Incomplete Markets, Volatility Smiles and Volatility Trading

David Hadyn Wayne

A thesis submitted for the Degree of Doctor of Philosophy of
the University of London

Centre for Quantitative Finance
Imperial College of Science, Technology and Medicine
University of London

September, 1999

Abstract

This thesis deals with the pricing of derivative products, both within the Black-Scholes framework and without key assumptions made under their methodology. Three topics are discussed, pricing options when the underlying asset is not liquid, pricing and hedging exotic derivatives in the presence of a volatility smile and the trading of volatility as an asset in itself.

We begin by applying a utility maximisation method to the pricing of European style contingent claims when the underlying asset may not be traded continuously. Prices and trading strategies are found under two different types of illiquidity. Secondly, we use a method of static replication to price new, exotic, option products with non-constant volatilities. Finally, we give a complete description of trading volatility as an asset class. We begin by showing how this may be achieved through standard products and progress to pure volatility derivatives known as Variance and Volatility Swaps. In the case of the Variance Swap, two methods of replication are implemented and compared for constant volatility and a volatility smile. Volatility Swaps are described in intuitive terms before a more exact value is found and risk parameters derived. Extensive computational results are presented throughout.

Acknowledgements

I would like to thank Professor Nicos Christofides and Mr. Gerry Salkin for the supervision they provided throughout the course of my studies. Their insight into the subjects of mathematics and finance has been invaluable.

My thanks are also due to Lehman Brothers, both for the financial assistance they provided and the wealth of knowledge available from my collegues at the bank. Special thanks are due to Robert Campbell, Tulug Temel, Sasi Digavalli and Adam Shapiro for the methods, tools and tricks they have shared with me. The skills I have learned from these professionals, and friends, have proved indispensable.

The advice and constructive criticism provided by Dr. Peter Whitehead was of enormormous value. I owe Peter a debt of gratitude, not only for his practical assistance, but also for his advice and encouragement which was of equal importance.

Lastly, thank you to my parents who have always supported and assisted me in continuing my studies and afforded me the opportunities which, over the years, I have been so happy to take advantage of.

Contents

1 Introduction 1

2 Pricing Options in an Illiquid Market 5

2.1 Introduction 5

2.2 The Writing Price Definition 6

2.3 Pricing Options in a Complete Market using Utility Maximisation 10

2.3.1 The Discrete Time Market Model 10

2.3.2 The Writing Price Under the Exponential Utility Function 13

2.4 Introducing Market Imperfections - Illiquid Stock 15

2.4.1 Results 16

2.5 Pricing Options when Trading Volume is Limited 19

2.5.1 Results 21

2.6 Conclusions 24

3 Volatility Smiles and Static Replication 26

3.1 Introduction 26

3.2 Implied Trinomial Trees 30

3.2.1 Constant Volatility Trees 30

3.2.2 Forward conditions and Arrow-Debreu Prices 31

3.2.3 Implying Node Probabilities and Constructing the Tree 34

3.2.4 Potential Difficulties 36

3.3 Improvements to the Basic Model 37

3.3.1 Foreign Exchange Barrier Options 37

3.4 Static Hedging and its Application to Volatility Smiles 39

3.4.1 Contract Replication Through Static Hedging 40

3.4.2 Static Hedging in the Continuous Black-Scholes World 40

3.4.3 Examples of Static Replication 41

3.4.4 Pricing Barrier Options on an Asset with Time Dependent Volatility 51

3.5 Conclusions 56

4 Volatility Trading **58**

4.1 Introduction 58

4.2 Volatility Trading through Delta-Hedged Vanilla Options 59

4.2.1 Profit and Loss due to Trading Delta-Hedged Vanilla Options 59

4.2.2 Vega, Gamma and Theta Associated with Trading Plain Vanilla Options 62

4.3 Variance Swaps 65

4.3.1 The Log Contract 66

4.3.2 Developing an Instrument for Speculating on Future Realised Variance 68

4.3.3 Numerical Results 77

4.4 Volatility Swaps 85

4.4.1 Description of the Volatility Swap 86

4.4.2 Approximation of the Volatility Swap by a Strip of Forward Starting Straddles 88

4.4.3 Pricing the Volatility Swap 90

4.4.4 Greeks and Hedging the Volatility Swap 95

4.4.5 Results 97

4.5 Conclusions 105

5 Conclusions **108**

A Arrow-Debreu Prices **118**

B Forward Starting Options **120**

C Derivation of Various Expectations and Greeks **121**

C.1 Proof of equation (4.42) for $\mathbb{E}_t^Q[\sqrt{\sigma_R^2}]$: 121

C.2 Derivation of b_1, b_2, b_3 and b_4 used in section (4.4.3): 122

C.3 Derivation of the Exact Greeks for $\hat{V}(t)$ 127

C.4 Derivation of the Approximate Greeks for $\hat{V}(t)$ 132

List of Tables

3.1 A Typical Volatility Smile for 3month USD/JPY Options. 27

3.2 Replication portfolio for an up and out call option to be revalued at time 0.6. 44

3.3 Replication portfolio for an up and out call option to be revalued at time 0.4. 44

3.4 Replication portfolio for an up and out call option. 44

3.5 Replication portfolio for an up and in call option. 47

3.6 Portfolio for the up and in call option at time 0.8 and spot 120. 47

3.7 Trigger levels for the moving double knockout option. 48

3.8 Replicating portfolio for the moving double knockout option. 49

3.9 Replicating portfolio for the moving double knockout option at time 0.625. 49

3.10 Replicating portfolio for the partial barrier option with a volatility curve. . 53

3.11 Replicating portfolio for the partial barrier option with a flat volatility curve. 54

3.12 Volatility smile for the capped power option. 56

3.13 Replication portfolio for a capped power option. 57

4.1 Replication using the method of Carr, with constant volatility for all strikes. 82

4.2 Replication using the method of DDKZ, with constant volatility for all strikes. 82

4.3 Convergence to fair variance with an increasing number of options. 83

4.4 Volatility smile used in the examples. 84

4.5 Carr's replication of a Variance Swap under a volatility smile. 84

4.6 DDKZ replication of a Variance Swap under a volatility smile. 85

List of Figures

2.1 Prices for an option on an asset which becomes illiquid before expiration. Volatility 0.15. 17

2.2 Prices for an option on an asset which becomes illiquid before expiration. Volatility 0.20. 18

2.3 Prices for an option on an asset with illiquidity periods of size 0.4 years starting at different time steps. Volatility 0.20. 19

2.4 Prices for an option on an asset with illiquidity periods of size 0.2 years starting at different time steps. Volatility 0.20. 20

2.5 Option prices when trading volume is limited. Volatility 0.10. 22

2.6 Option prices when trading volume is limited. Volatility 0.15. 23

2.7 Option prices when trading volume is limited. Drift 0.12. 24

2.8 Option prices when trading volume is limited. Drift 0.15. 25

3.1 The change in vega for a change in volatility for an up and out call option . 29

3.2 The change in vega for a change in spot for an up and out call option . . . 30

3.3 Notation for the implied trinomial trees. 32

3.4 Forward condition for the implied trinomial tree. 33

3.5 Construction of the option price. 34

3.6 Stock price node contributions to the value of a call option. 35

3.7 A general barrier option. 42

3.8 Boundary values for an up and out call option. 43

3.9 Replication step for an Up-and-Out Call option. 45

3.10 Covergence of the replication scheme to the analytic price of the barrier option. 46

3.11 Boundaries of the moving barrier option. 48

3.12 Partial barrier option price against front end volatility using one hundred time steps. 54

4.1 Vega against time for a plain vanilla option, strike price 100. 64

4.2 Vega against spot for a plain vanilla option, strike price 100. 65

4.3 Gamma against time for a plain vanilla option. 66

4.4 The convergence of values for a variance swap with no volatility smile. . . . 79

4.5 The convergence of values for a variance swap with a volatility smile. 81

4.6 Values for the Volatility Swap. 103

4.7 Vega for the Volatility Swap. 104

4.8 Delta for the Volatility Swap. 105

4.9 Gamma for the Volatility Swap. 106

A.1 Arrow-Debreu prices. 119

Chapter 1

Introduction

Throughout the last three decades, derivatives have grown to form a large part of the financial markets. With the proliferation of derivatives, their complexity has increased and the appetite for new ways to take a position, in order to hedge away risk or for speculation, continues to grow. This demand is being satisfied by the continuous creation of new derivative securities.

The equations derived by Black and Scholes [5] in the early 1970s for the pricing of European style options have been adopted as the basis for pricing and trading derivatives in the financial markets. However, the methodology relies upon certain assumptions in order to make the problem to be addressed tractable. Amongst those assumptions is the requirement that the asset upon which the contingent claim has been written has a liquid market itself, and that the prices for this asset are lognormally distributed. More specifically, the price of this asset moves continuously, randomly, and without gaps. A market participant is assumed to be able to transact as much of this asset as he would like, at any time, and without transaction costs. The volatility governing the random process by which the price of the underlying security moves is also assumed to be constant throughout the life of the option. In reality none of the above is wholly true.

In Chapter 2 the problem of illiquidity is addressed. Using a utility maximization scheme first introduced by Hodges and Neuberger [30] and later modified by Davis, Panas and Zariphopoulou [16] to calculate option prices under transaction costs, we describe an algorithm to price options on assets which may not be continuously traded. This scheme allows

us to examine the effect on option prices under different types of liquidity constraints and according to the risk preferences of the writer. A dynamic program is set up in order to maximise two value functions. Through the use of these functions the optimal trading strategies of the option writer are investigated under different circumstances. Numerical results are presented for solutions obtained under a special utility function.

The imperfections that exist in the marketplace for the underlying asset, and the flawed assumption of lognormality, are well known to practitioners. Nonetheless, the Black-Scholes equation continues to be widespread and provides the standard functions by which the financial services industry evaluates over-the-counter options. Market participants, well aware of these invalid assumptions, use this imperfect model and make adjustments to the inputs in order to recover the perceived fair value of the contract in question. Of the inputs used to value plain vanilla options, the only one not directly observable is the volatility of the underlying asset over the contract's life. It is by manipulation of this parameter that the values obtained from the standard Black-Scholes formulae for plain vanilla options are changed. This change in volatility for options of different strikes and maturities has become known as the "Volatility Smile", and leaves us with the curious situation of the same asset being assumed to have different volatilties over the same period of time. This is clearly incorrect but is the way that practitioners have been able to make allowances for the shortcomings of Black and Scholes whilst still having a relatively straightforward, and most importantly, universally accepted model. In fact vanilla option trades carried out between market counterparties are quoted on a volatility basis, each side of the bargain confident in the knowledge that the other uses precisely the same model when calculating the premium payment to be made.

The presence of a volatility smile enables plain vanilla options to be priced according to the market perception of fair value despite the assumptions of the model used. The prices for more exotic derivative products are also subject to the limitations of a lognormal distribution present in the Black-Scholes framework. In the case of these more complicated products it is generally not possible to use an implied volatility for pricing purposes. In practice, prices for non-vanilla products such as barrier options are quoted as a premium, often in terms of percentage of the notional amount. To price these products correctly and consistently with the prices of plain vanilla options, a premium must be added or subtracted to the value that simple, constant volatility, models give. The construction of a

model in order to price exotic derivatives consistent with the information contained within the volatility smile is discussed in Chapter 3.

Chapter 3 comprises two sections, the first of which implements a model described by Derman, Kani and Chriss [21]. This methodology for constructing an asset model that is consistent with plain vanilla option prices is implemented with specific reference to foreign exchange. The model is adapted to price a class of exotic options which is of particular relevance within foreign exchange and is known to have sensitivity to the shape and magnitude of the volatility smile. Using standard option products, the second part of Chapter 3 details a procedure for the replication of non-standard products. A framework proposed by Derman, Ergener and Kani [22] is described. This is then implemented for the same type of option singled out for attention in the first section of the chapter. Several, highly exotic structures are priced to demonstrate the fiexibility of the procedure and to generalise the original results. We extend the existing methodology to price exotic derivatives when the underlying asset is assumed to have time dependent volatility. Furthermore, we describe how the volatility smile model may be utilised to replicate products consistently with the full volatility surface. This is illustrated through the use of an example. Numerical results are presented throughout.

The existence of the volatility smile is to facilitate the trading of options with different strikes and maturities using a single model with an unrealistic assumption. Vanilla option transactions between market participants are conducted efficiently due to the practice of quoting option prices not in premium amounts, but in volatility. This volatility is used in the standard model to imply a premium payment to be made. In this way "Implied Premia" are traded rather than trading option prices that imply a volatility.

The practice of trading delta hedged options, and then continuously delta hedging them until expiration time, is discussed in the first part of Chapter 4. With the deep liquidity of many option markets, participants have become interested in taking positions purely in the level of volatility. In Chapter 4 we explore volatility trading in depth. After discussing the profit, loss and risk from the most simple form of volatility trading, delta hedged vanilla options, we look at new types of volatility products.

Contracts that have become known as Variance Swaps are described and the literature that exists reviewed. We show how Variance Swaps are related to delta hedged vanilla options

and another type of derivative contract that pays the natural logarithm of the return on the underlying asset over its lifetime. We investigate the methods of Carr and Madan [12] and Demerterfi, Derman, Kamal and Zou [18] to price and hedge these contracts. We then replicate the Variance Swap contract in two different ways. We present detailed numerical examples of the implementations under constant volatility and using a smile. Results are given and used to examine the comparative performance of the two hedging methods.

The advantage, as we will discuss, of Variance Swaps is that they are closely related to the payoff from a continuously delta hedged option. Due to the structure of the options market, speculators have views about the future level of implied volatility. This has led to the development of an alternative contract, the Volatility Swap, over the last two years. The majority of Chapter 4 is devoted to an examination of this new contract. We find an intuitive way of looking at the Volatility Swap and use this as a basis to examine the risks associated with it. An approximate value is found for the contract at a general point in time. "Greeks" for the purpose of hedging are also derived. We compare the different approximations, explain the nature of the swap and demonstrate the validity of our intuitive approach.

Full numerical results are given for the Volatility Swap valuation and "greeks". The behaviour of these "greeks" is also examined under certain circumstances and again our intuitive approach compared to more exact derivations.

Chapter 2

Pricing Options in an Illiquid Market

2.1 Introduction

The Black-Scholes [5] method for pricing derivatives relies upon what is known as a perfect market for the underlying security. Involved in this assumption is the requirement that the market participant is able to trade as much of the underlying asset as he desires, at no cost. In reality, market conditions and legislation affect the availability of assets and limit the size of transactions.

Black-Scholes uses perfect replication arguments in order to eliminate any risk assumed by the writer. Under perfect market conditions this is possible. However when perfection breaks down, so does this method. With preference-free pricing no longer possible, the attitudes of traders must be represented in some way. There exists an element of risk in the writing of an option and the writer requires compensation for this.

We represent attitudes towards risk by using utility functions to characterize the value of terminal wealth. A stochastic control approach is used to address the market incompleteness problems. The method used was applied by Hodges and Neuberger [30] to price an option with transaction costs incurred in implementing the hedging strategy. This work was later expanded and modified by Davis, Panas and Zariphopoulou [16] and it is their definition of writing price that is used here.

We begin with the description of two value functions giving the maximum utility of wealth of the investor. These functions are then used in the definition of the writing price of an option. The solution of these value functions is then described in a perfect market to price a European call option. A dynamic program is set up and implemented using a special utility function to speed up computation. We then go on to adapt this market model and dynamic program for use in an imperfect market where the underlying asset experiences periods of illiquidity. Prices are generated under several possible scenarios and the results discussed. The notion of a restriction on the volume of the underlying that an option writer is able to trade in, is then explored through the use of the value functions. The dynamic program is modified and implemented producing prices for different market conditions. The hedging strategy of the writer is investigated and the results analysed.

2.2 The Writing Price Definition

Risky assets are described as stochastic processes embedded in a probability space. To begin with we assume that the markets are perfect. In particular, the writer of the option is able to trade as much of the underlying stock and riskless bond as he needs to whenever he desires. It is possible, within this framework, to consider more than one risky asset. For our purposes we need only consider the one stock underlying the option to be written.

The price process of this stock, $S(t)$, is defined on the probability space $(\Omega, \mathbb{F}, \mathbf{P})$ over the finite interval $[0, T]$ as:

$$dS(t) = S(t)(\mu dt + \sigma dR(t)),$$

Where μ is the drift rate of the stock and σ its volatility per unit time. $R(t)$ is a one dimensional Brownian Motion.

The bond process, $B(t)$, is also described over $[0, T]$ and is the amount held in the bond at time t.

We also define $y(t)$ as the amount of stock held by the writer at time t where $t \in [0, T]$. Further, $c(S(t), y(t))$ is the cash value of the writer's holdings in the risky stock at time t. This is equal to the total number of units of stock held multiplied by the current stock price. The set of admissible trading strategies is denoted by A.

As in a standard perfect replication argument, the writer of the option forms a portfolio of stocks and bonds in order to hedge the option. In this case we will consider a call option. It is a trivial matter to modify the following to price a put option.

Define X to be the exercise price of the option and T its expiration time. Assume we are starting from initial time $s \in [0, T]$. At T the writer of the option must fulfil his obligations stated in the contract sold.

So, if $S(T) \leq X$ then the writer has wealth

$$B(T) + c(S(T), y(T)),$$

as the option will not be exercised.

If $S(T) > X$ then he has

$$B(T) + c(S(T), y(T) - 1) + X,$$

In other words the option is exercised and so he gives the purchaser one unit of stock in return for an amount of cash equal to the exercise price.

Now, to proceed further and define the value function, we use the following.

Definition:

A utility function is a strictly concave, strictly increasing function $U(x)$ such that

$$U : [0, \infty] \longrightarrow [0, \infty],$$

with $U(0) = 0$.

The concavity of the utility function implies that the writer is risk averse. See Davis, Panas and Zariphopoulou [16].

When considering imperfect market problems there is a fundamental need for some way to represent the attitudes of the trader towards risk. The market imperfection means that it is impossible to replicate the payoff of the option perfectly and thus eliminate risk. The writer of the option can no longer be sure of meeting his obligations at expiration. He may now only *expect* to have the wealth to settle up with the purchaser. It is this essential element of risk that forces us to characterize the attitude of the writer. There is no unique price for the option in an imperfect market, so the perception of 'fair price' varies from investor to investor.

We may now define a value function for the writer.

$$V_w(T,S,y,B) = \sup_{\pi \in A} \mathbb{E}\{U[B(T) + I_{S(T)\leq X}[c(S(T),y(T))] + I_{S(T)>X}[c(S(T),y(T)-1)+X]]\}, \quad (2.1)$$

where the supremum is taken over the set of all admissible trading strategies A and I is the indicator function with value 1 if its condition is satisfied, and value 0 otherwise. $\mathbb{E}$ denotes expectation in the relevant probability space. The set of admissible trading strategies A is defined in the following way. At time $t = 0$ the writer may start with any portfolio of stocks and bonds. A strategy π is a member of A if, and only if, $\forall t : 0 < t \leq T$, $S(t)y(t) + B(t) = S(t)y(t-\delta t) + B(t-\delta t)e^{r\delta t}$, where δt is an infinitessimally small step in time. This means that only transactions requiring no injection or removal of funds are admissible. All stock trades must be financed with the buying or selling of bonds.

The expectation and probabilities of up and down stock price movements are taken from a binomial lattice constructed in the normal way with the drift rate of the stock, μ, replacing the risk free rate of return in the calculations. This adjustment is due to the fact that we are not in the preference free probability space but the real probability space.

This value function gives the maximum expected utility of wealth given that the option has been written, the writer plays the market optimally within the set of admissible trading strategies and all obligations are fulfilled at expiration. The indicator functions denote whether the option is exercised or not. This affects the actual wealth of the writer at expiration time T. $V(B)$ is bounded $\forall B \in \mathbb{R}$, and is a continuous, monotonic increasing function of B.

Next we define B_w such that $V_w(0,S_0,0,B_w) = 0$, where S_0 is the spot price of one unit of stock at time 0.

Thus the writer, starting with no initial wealth, is indifferent between

(a) Doing nothing and

(b) Writing the option and charging B_w.

However this is not the writing price. The writer also has to play the market and construct the hedging portfolio. This must be taken into account when calculating a fair price for the

option. We therefore define a second value function.

$$V_u(T,S,y,B) = \sup_{\pi \in A} \mathbb{E}\left\{U[B(T) + c(S(T),y(T))]\right\}, \tag{2.2}$$

This function gives the maximum expected utility of wealth if the writer plays the market optimally within the set of admissible trading strategies. In this case there is no option written. The writer is playing the market by trading stock and bond purely to make a profit, not to hedge any instrument. The writer is again assumed to have started at time s with no wealth,

In a similar way to before we define B_u such that $V_u(0,S_0,0,B_u) = 0$.

We may then calculate the writing price of the call option given by

$$p_w = B_w - B_u,$$

In intuitive terms this call option price means that the writer is indifferent between

(a) Playing the market only and

(b) Charging the buyer p_w in order to hedge the option and play the market.

It is at this point that we return to the remark made earlier in this section about the notion of 'fair price' in an imperfect market. So far, care has been taken to refer to 'writing price' and not simply to 'price'. As a consequence of using utility functions , and a dynamic programming algorithm, the prices calculated for a buyer and a seller in an imperfect market do not coincide. This is in part due to the fact that the utility functions and therefore the risk preferences of the two parties may not coincide. More fundamental is the nature of buying and selling an option. In the real, non-Black-Scholes, world where risk may not be completely hedged away, the writing of an option corresponds to assuming a risk or liability and the purchase to a form of insurance. The difference in writing and purchasing price reflects this asymmetry.

2.3 Pricing Options in a Complete Market using Utility Maximisation

Before proceeding, a number of definitions are required.

$$\begin{aligned}\Phi_w(S,y,B) &= B(T) + I_{S(T)\leq X} c(S(T),y(T)) \\ &\quad + I_{S(T)>X}[c(S(T),y(T)-1)+X], \\ \Phi_u(S,y,B) &= B(T) + c(S(T),y(T)),\end{aligned}$$

The above describe the wealth of the writer firstly with the option written and secondly without. Again we have assumed it is a call option that has been written.

It has been shown by Davis, Panas and Zariphopoulou [16] that

$$\mathbb{E}\,|U(\Phi_i)| < k \quad \text{where} \quad k \in \mathbb{R}$$

In other words the value is initially bounded.

We now need to define a value function in order to maximise the writer's expected utility of wealth.

2.3.1 The Discrete Time Market Model

The market model describes the stock price, the number of shares, y, held by the writer and the size of the writer's holding in bonds.

We have the following discretisation parameters.

h_s - the discretised step in stock price

h_y - the discretised step in the writer's holding in stock

h_B - the discretised step in the writer's holding in bonds

δt - the discretised step in time.

The quantities of y (the stock), B (the bond) and S (the stock price) change only in discrete time steps. The changes are due only to the diffusion of the stock or trading of assets in

the course of maximising the writer's expected utility. We also define, for ease of notation, $t(n)$ to be $n\delta t$. That is, if we are on the n^{th} time step then the total amount of time elapsed is $t(n)$.

The change in stock price is

$$\delta S = \begin{cases} uS - S, & \text{with probability } p, \\ dS - S, & \text{with probability } 1 - p. \end{cases}$$

With $p \in [0, 1]$, m is the discrete time step and u and d are the factors given by the binomial lattice representing the price process.

As already defined, the smallest number of shares the writer is permitted to buy or sell is h_y, so the change in the writer's stock holding on the m^{th} time step is $\delta y(m)$. This quantity is a multiple of h_y, either positive or negative.

The admissible trading strategies are self financing, so for every amount of stock traded there must be a change in the bond holding corresponding to the price of one unit of stock multiplied by the number of stocks bought or sold. In other words

$$\delta B(m) = -S(m)\delta y(m).$$

A discrete time discount factor is given as follows

$$D(t, s) = e^{-r(t-s)}.$$

Now we are able to define our trial evaluation for the value function.

$$V_i(\tau, S, y, B) = \max_{\pi \in A} \left\{ \begin{array}{l} \mathbb{E}[V_i(\tau + \delta t, S + \delta S, y + \delta y, (B - \delta B)e^{r\delta t}] \\ \mathbb{E}[V_i(\tau + \delta t, S + \delta S, y, Be^{r\delta t})] \end{array} \right\} \quad i = u, w \qquad (2.3)$$

where $\delta y > 0$.

With boundary conditions

$$V_i(T, S, y, B) = U(\Phi_i(S, y, B)) \quad i = u, w \qquad (2.4)$$

Only buying stock is considered. The reason for this will become clear as the algorithm is described in more detail below.

At expiration time, T, the maximum expected utility of the writer's wealth is calculated according to the cash value of his portfolio. All possible portfolios that he may be holding

are considered. We then work recursively backwards through the underlying stock model finding the strategy that maximises the expected values. The value function is constructed for all the various discrete holdings of stock and bond until we reach the writing time.

The trial evaluation described above improves efficiency when calculating the value function in the following way.

At each time step, holding the minimum amount of stock, we initially compare the following.

(i) Buying h_y shares and allowing the stock price to diffuse.

(ii) Carrying out no trades and allowing the stock price to diffuse.

To make computation possible we must set a minimum and maximum number of units of stock the writer may own. The functions V_i are defined with a supremum over a set of admissible trading strategies A.

In a perfect market the writer of the option may, if he has the required funds, trade as much of the stock, as often as he likes. There is also no restriction on the volume or frequency of bond transactions so long as no injection of extra funds is required. This set of strategies must be reduced to allow numerical computation of the solution of the value function without the loss of any part of the optimal strategy. We therefore set a minimum and maximum number of units of stock and bond that the writer may own at any one time. It is crucial that the minimum is less than the least the writer would ever own under an optimal strategy and that the maximum is greater than the most he would ever optimally hold. If the parameters of the model have been set realistically then these bounds will be finite. The value function may then be solved in finite time. The value function is calculated for each of the points in this two dimensional array of stock and bond holdings. The array consists of all possible portfolios and the corresponding value for each if trading is carried out optimally. This array is different for each stock price and discrete point in time.

It is known that the buy and sell region and the no-trade line in the state space are well defined and distinct (i.e. there is no overlapping where buying and selling are both optimal policies). To begin the calculation of V_i we start by holding the minimum amount of stock permissible. We then compare the effect of buying the smallest amount of stock we are allowed to trade, δy, and leaving the stock price to diffuse, to the effect of leaving the

portfolio alone and the price diffusing. If the former is greater than the latter then we assume that we own $y_{min} + \delta y$ and repeat the procedure. This is carried out until the value function for trading is greater than for not trading. In the absence of transaction costs we may assume that whatever the stock and bond holdings of the writer, the optimal strategy for him to follow is to transact to y_{opt}. The value function may then be calculated for the portfolio he transacted from. This procedure means that the trial evaluation involving consideration of just buying stock and holding stock described above is sufficient to find the value function for every different portfolio in the state space of admissible strategies. The processing time is reduced as the optimisation loop is carried out for only one stock holding.

We now introduce the exponential utility function and show how it may be used to reduce the dimensionality of the dynamic programming problem.

2.3.2 The Writing Price Under the Exponential Utility Function

The dynamic programming algorithm is greatly simplified, and processing time significantly reduced if the utility function used is

$$U(x) = 1 - e^{-\gamma x}, \qquad \gamma \in \mathbb{R}$$

This function has the property that the function of risk averseness implied by it is independent of x, the wealth.

i.e.

$$-\frac{U''(x)}{U'(x)} = \gamma, \quad \forall x \in \mathbb{R}$$

This in turn implies that the writer is willing to borrow as much money, using the sale of bonds, as is necessary in order to purchase the optimal amount of stock. This eliminates the need to consider the bond axis in the calculation of the value function. The number of bonds traded by the writer is dictated only by the cost of the stocks he is transacting. A source of inaccuracy is removed as when bonds were being considered interpolation of points was necessary in order to calculate the value function. With the bond consideration removed the growth of their worth at the risk-free rate is dealt with by discounting stock values and the need for interpolation is gone.

We may use this special utility function to transform the optimisation in the following way.

$$\begin{aligned} V_i(S,y,B) &= \sup_{\pi\in A}\{U[\Phi_i(S,y,B)]\} \Rightarrow V_i(S,y,B) = \sup_{\pi\in A}\left\{1-e^{-\gamma\Phi_i(S,y,B)}\right\}, \\ &= 1-\inf_{\pi\in A}\left\{e^{-\gamma\Phi_i(S,y,B)}\right\}, \\ &= 1-\inf_{\pi\in A}\left\{e^{-\gamma B(T)}e^{-\gamma\chi_i}\right\}, \end{aligned}$$

where $\chi_i = \Phi_i(S,y,B) - B(T), \qquad i = u, w.$

Thus

$$V_w(S,y,B) = 1-\inf_{\pi\in A}\left\{\mathbb{E}[e^{-\gamma B(T)}e^{-\gamma\chi_w(S(T),y(T))}]\right\}, \tag{2.5}$$

and

$$V_u(S,y,B) = 1-\inf_{\pi\in A}\left\{\mathbb{E}[e^{-\gamma B(T)}e^{-\gamma\chi_u(S(T),y(T))}]\right\}. \tag{2.6}$$

For an initial time $s \in [0,T]$, we have

$$V_i(s,S,y,B) = 1-e^{\frac{\gamma B}{D(T,s)}}Q_i(s,S,y), \qquad i = u, w \tag{2.7}$$

where,

$$Q_i(s,S,y) = \inf_{\pi\in A}\left\{ e^{\gamma\chi_i(S(T),y(T))-\int_s^T \frac{S(\xi)}{D(T,\xi)}dy(\xi)} \right\}, \tag{2.8}$$

and,

$$D(T,\tau) = e^{-r(T-\tau)}.$$

We are, in other words, turning the problem into one of minimisation rather than maximisation and also reducing the dimensionality by removing the bond holding from the problem. The B in the above expressions refers to the initial holding that the writer has in bonds. In order to calculate the writing price of the option this should be assumed to be zero. In this way the writer has no initial advantage in hedging the option and the bonds are removed totally from calculations, not even occurring as a constant term. Q_i is now the function that must be discretised and solved. This is done in the following way.

$$Q_i(\tau,S,y) = \min_{\pi\in A}\left\{ \begin{array}{l} F(\tau,S,\delta y)\,\mathbb{E}[Q_i(\tau+\delta t, S+\delta S, y+\delta y)] \\ \mathbb{E}[Q_i(\tau+\delta t, S+\delta S, y)]\} \end{array} \right\} \qquad i = u, w \tag{2.9}$$

With boundary conditions

$$Q_i(T,S,y) = e^{-\gamma\chi_i(S(T),y(T))}, \qquad i = u, w \tag{2.10}$$

where,

$$F(\tau, S, \delta y) = e^{\frac{-\gamma \delta y}{D(T,\tau)}}.$$

Again it is best to proceed by starting with y_{min} and first compare buying h_y and doing nothing.

This time we assume that no matter what our holding in the stock, we transact to y_{opt} and may find the value function for any holding in the stock by using the following relationship.

$$Q_i(\tau, s, y) = F(\tau, S, y_{opt} - y) Q_i(\tau, S, y_{opt}), \qquad i = u, w$$

The writing price of the call option is then given by the following explicit function.

$$p_w = \frac{D(T,0)}{\gamma} \ln \left\{ \frac{Q_w(0,S,0)}{Q_u(0,S,0)} \right\}.$$

Where Q is evaluated at the end and for no initial holding of stock. If the writer were to hold stock at the writing time, s, then he would have an advantage when hedging the option. This would reduce the price of the option and so would not produce a fair result. In this way the preference free price for the writer is obtained.

In addition to the writing price of the option the dynamic program yields other information. The optimal holding in the underlying stock is found at each discrete time step. Under the exponential utility function this is given by

$$y^*(t, S(t)) = D(T,t) \frac{\mu - r}{\gamma \sigma^2 S(t)}.$$

It may be shown that this dynamic programming method converges to the Black-Scholes price in the perfect markets case. In addition to this the Black-Scholes deltas may be generated. To do this the optimal stock holding, y^*, is stored whenever the stock price is the same as the initial price at time s. If this is carried out when the option has been written and also for when it has not and the market is being played then the difference between the two, $y^*_{written} - y^*_{not}$ gives the delta hedge at each time step where recombination occurs.

2.4 Introducing Market Imperfections - Illiquid Stock

Using the method outlined above makes it possible to introduce certain market imperfections. The case we are interested in is that of a stock that it is not always possible to trade.

The way this particular problem has been tackled is by affecting the maximisation in the value function at times when trading is not possible. We consider the case when the period or periods of illiquidity are known before the option is written. It is then a case of identifying the points in discrete time when no trading is possible and imposing the condition that the only admissible trading strategy is the one where the stock is left alone and its price diffuses.

We now carry out the maximisations in the value function over a smaller set of admissible strategies, A^*. This set differs from the original set, A, in that there are certain times when trading is not possible so the size of the set is reduced. Then the writing price of the option may be calculated as before.

The trial evaluation for the value function (2.3) in discrete time is similarly modified and calculated as before.

The computation may, once again, be simplified using the exponential utility function. The resulting equations are similar to (2.7) and (2.8).

Expressions (2.5), (2.6) and (2.9) are modified by considering A^* instead of A, and the writing price may be calculated in the same way as before.

2.4.1 Results

In this section results are presented that were produced by the dynamic programming algorithm when pricing an option on a non-dividend paying stock and the market experiences periods of illiquidity. The periods of illiquidity are known before the option is written. In all cases the values used for calculation are as follows, $S = 10$, $X = 10$, $T = 1$, $r = 0.1$, $\mu = 0.12$ and the volatility of the stock is stated for each figure. The exponential utility function has been used throughout with a value of $\gamma = 1$ for the constant of risk averseness. Each time the options have been priced over one hundred time steps. The first set of figures (2.1) and (2.2), show the change in price when a period of illiquidity exists at the end of the option's life and this period becomes longer and longer. In figures (2.3) and (2.4), the length of time that the stock is untradeable is fixed but the location of this period within the life of the option is varied from expiry time backwards. The difference in price is also illustrated when the stock may not be bought or sold at the option's inception.

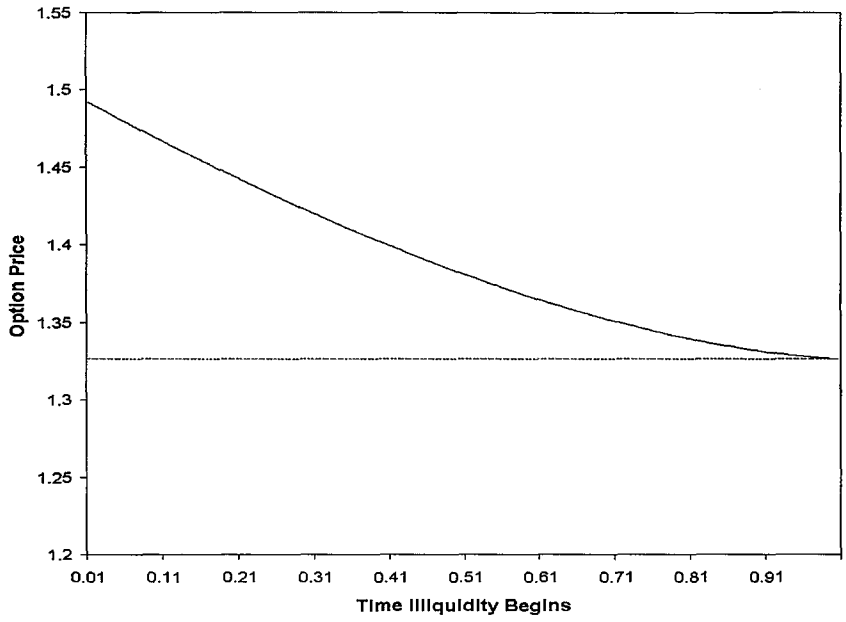


Figure 2.1: Prices for an option on an asset which becomes illiquid before expiration. Volatility 0.15.

As the length of time that the stock may not be traded is extended the cost of the option increases. This is what would be expected and goes some way to justify the approach used and the prices generated. The problem was solved by changing the set of admissible trading strategies from A to A^* and carrying out the optimization. As A^* is the set in which the liquidity limits are included $A^* \subseteq A$ and so the value function calculated over the set A^* is less than or equal to that calculated for the perfect markets case. In the figures the end of the liquidity period is fixed to the expiration date of the option. This produces the intuitive results mentioned above, as would fixing the end to any date during the life of the option.

It is a well known property that the higher the volatility of the stock price, the higher the price of an option written on it. Further to this, the relative increase in option price when the stock experiences periods of illiquidity is larger for more volatile stocks. This is a consequence of the greater likelihood of a large movement in stock price when there is no trading and the writer may not capitalise on this or adjust his hedge.

If we assume that trading is possible at the time of writing, it is observed that the most expensive time for the underlying to become illiquid is at the end of the option's life. It seems that if the writer is able to trade after a period of illiquidity has occurred then the

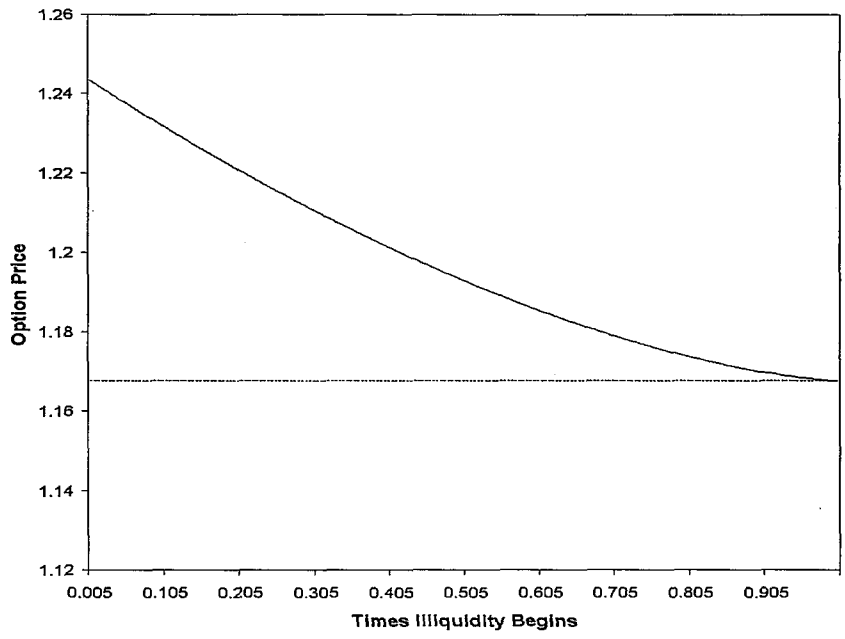


Figure 2.2: Prices for an option on an asset which becomes illiquid before expiration. Volatility 0.20.

hedge he requires in order to replicate the option payoff may be recovered. The extra cost of the option is to compensate the writer for the fact that the hedge may be expensive to re-establish due to the market movement when the option position is not covered. The nearer to expiration that the liquidity problem occurs, the less chance he has to make back the money lost by playing the market. If he cannot trade until the option has expired then the option is at its most expensive as the position is not adequately hedged when the purchaser is demanding settlement. This observation is backed up by Grannan and Swindle [26] who sought to find the most efficient times to hedge given the existence of transaction costs. Even though the problem here is considered in the absence of transaction costs, the conclusion drawn by Grannan et al [26] that it is most important to hedge more often towards expiration helps back up the higher premium shown here.

By far the most serious situation for the writer of an option is when he may not buy the security underlying the option until some point in the future after writing. When this is the case, the dynamic programming algorithm produces prices which are much larger than for the previous situations discussed. The main thing to notice is that the option price is no longer bounded by the spot price of the stock at writing. The writer holds none of

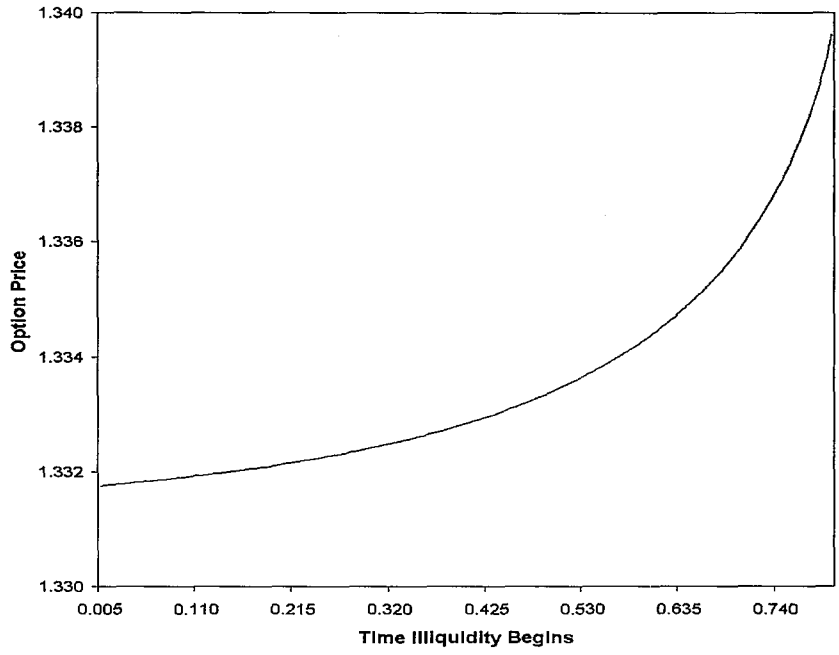


Figure 2.3: Prices for an option on an asset with illiquidity periods of size 0.4 years starting at different time steps. Volatility 0.20.

the stock yet is exposed to its price movement in the form of the option. The price could move a large amount before liquidity returns and the writer has no insurance against it. Therefore the writer may only have an expectation of his wealth and the purchaser must pay a high price. Again, the more volatile the stock price the higher the cost of the option and the more the writer should charge for this liquidity problem.

2.5 Pricing Options when Trading Volume is Limited

We now consider a market for the underlying asset where a maximum transaction size is imposed and define our trial evaluation for the value function.

$$V_i(\tau, S, y, B) = \max_{\pi \in A} \left\{ \begin{array}{l} \mathbb{E}[V_i(\tau + \delta t, S + \delta S, y + \delta y, (B - \delta B)e^{r\delta t}] \\ \mathbb{E}[V_i(\tau + \delta t, S + \delta S, y - \delta y, (B + \delta B)e^{r\delta t}] \\ \mathbb{E}[V_i(\tau + \delta t, S + \delta S, y, Be^{r\delta t})] \end{array} \right\} \quad i = u, w \qquad (2.11)$$

With boundary conditions

$$V_i(T, S, y, B) = U(\Phi_i(S, y, B)) \quad i = u, w \qquad (2.12)$$

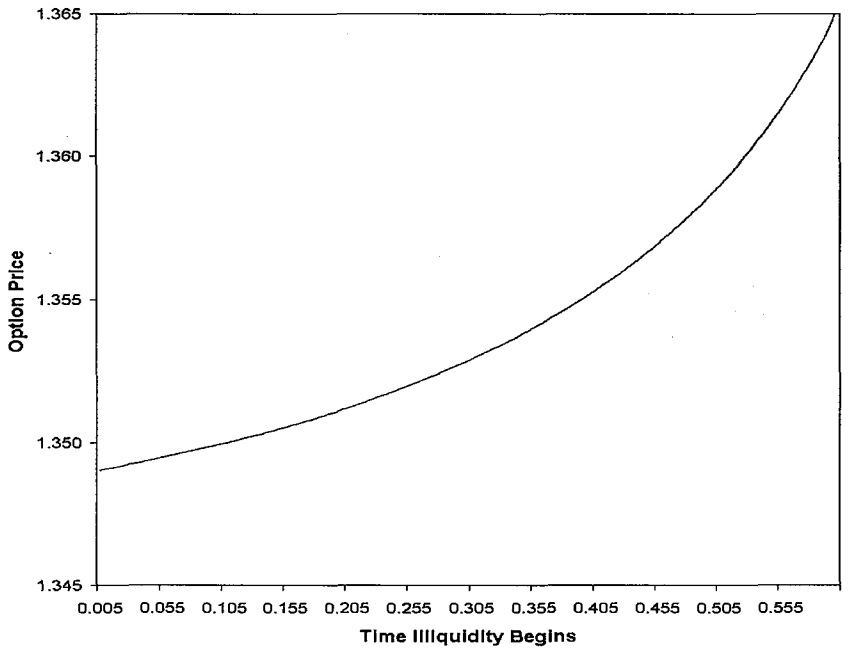


Figure 2.4: Prices for an option on an asset with illiquidity periods of size 0.2 years starting at different time steps. Volatility 0.20.

At each time step, and for each possible holding of stock, we initially compare the following.

(i) Buying h_y shares and allowing the stock price to diffuse.

(ii) Selling h_y shares and allowing the stock price to diffuse.

(iii) Carrying out no trades and allowing the stock price to diffuse.

If buying h_y shares produces a higher expected utility than selling or holding then we know that buying is the optimal strategy. The comparison of the strategy buy and the strategy hold is then repeated for increasing stock holdings until no improvement in value is found, or the transaction size limit is reached. The optimal holding y_{opt} is then stored and the value function calculated. If, through the initial comparison, selling is found to be the correct course of action an analogous procedure is followed until the optimal amount to sell is found or the trade limit is reached.

Under these market conditions the optimization loop must be repeated for every single possible portfolio in the set of admissible strategies, A^*. It is not possible to assume that, after performing one part of the procedure on the minimum stock holding, the figure y_{opt}

is the optimal holding point to transact to no matter what the portfolio's present position may be. The restriction on transaction size means that the optimal portfolio given a perfect market may not be obtainable in the imperfect market. It is not necessarily the case that the best course of action is to transact as closely to this value as possible.

We reduce the dimensionality of the dynamic programming problem using the exponential utility function to obtain the following value function for solution

$$Q_i(\tau, S, y) = \min_{\pi \in A} \left\{ \begin{array}{l} F(\tau, S, \delta y)\, \mathbb{E}[Q_i(\tau + \delta t, S + \delta S, y + \delta y)] \\ F(\tau, S, -\delta y)\, \mathbb{E}[Q_i(\tau + \delta t, S + \delta S, y - \delta y)] \\ \mathbb{E}[Q_i(\tau + \delta t, S + \delta S, y)]\} \end{array} \right\} \tag{2.13}$$

and calculate the writing price of the option as previously described.

2.5.1 Results

Here we present results produced by the dynamic program for options written on a non-dividend paying stock where the volume of the stock that may be transacted is limited. More precisely, the writer may trade only once in each discrete time step and the absolute size of this trade is limited to a certain amount known at the time of writing. The total amount of stock the writer may hold is not limited apart from in a non-intrusive way to make numerical computation simpler. The set of parameters taken as a basis for evaluation are as follows, $S = 10$, $X = 10$, $T = 1$, $r = 0.1$, $\mu = 0.12$, $\sigma = 0.05$. Any variation in the above is stated clearly at the top of the particular figure. The exponential utility function has been used throughout with a value of $\gamma = 1$ for the function of risk averseness. Each time the option was priced using one hundred time steps.

Irrespective of the parameters given for the option, the price curve produced as the maximum trade size is increased appears to display certain characteristics. When trade volume is severely limited, for instance the maximum trade is a single discrete stock division, the price is badly affected. If the option is then re-priced with the maximum trade increased to two stock divisions the price is reduced significantly. As the maximum size is increased this improvement in price becomes less and less dramatic until no improvement is found. It is at this point that the desired trades are not limited by the imperfection in the market and the Black-Scholes price is produced.

The reason for the shape of this curve is probably due to the cumulative effect of a trade limit. The hedge error is carried over to the next time step and so the strategy is changed all the way through the life of the option. The relative improvement in hedging opportunity is most exaggerated at the small end of the scale. An improvement in market conditions from a trade size of 1 to 1.1 is not as significant as from 0.1 to 0.2 even though in absolute terms the improvement is the same. The prices found reflect this difference.

In figures (2.5) and (2.6) the effect of trading limits on option price was investigated for differing volatilities. As expected, the higher the volatility the higher the option price. However the change in price as the market imperfection eases is more pronounced as volatility rises. As with illiquidity problems, the mis-hedge caused by problems in the market is more harshly penalised when stock price volatility is higher. The potential loss, which is what the writer is charging the purchaser of the option for, is greater for an option with greater volatility. Figures (2.7) and (2.8) represent the change in price when the drift

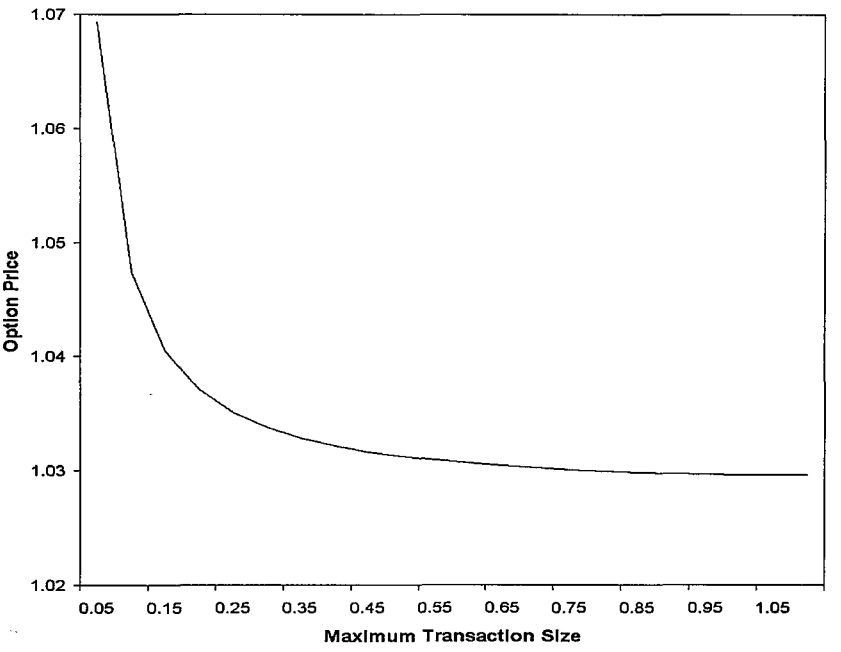


Figure 2.5: Option prices when trading volume is limited. Volatility 0.10.

parameter, μ, is changed. The drift price of the stock is removed from consideration in the Black-Scholes framework as a consequence of a change of probability measure and the existence of a replicating portfolio. In the method described above drift plays a crucial role in the modelling of stock price movement. A perfect market results in the Black-Scholes

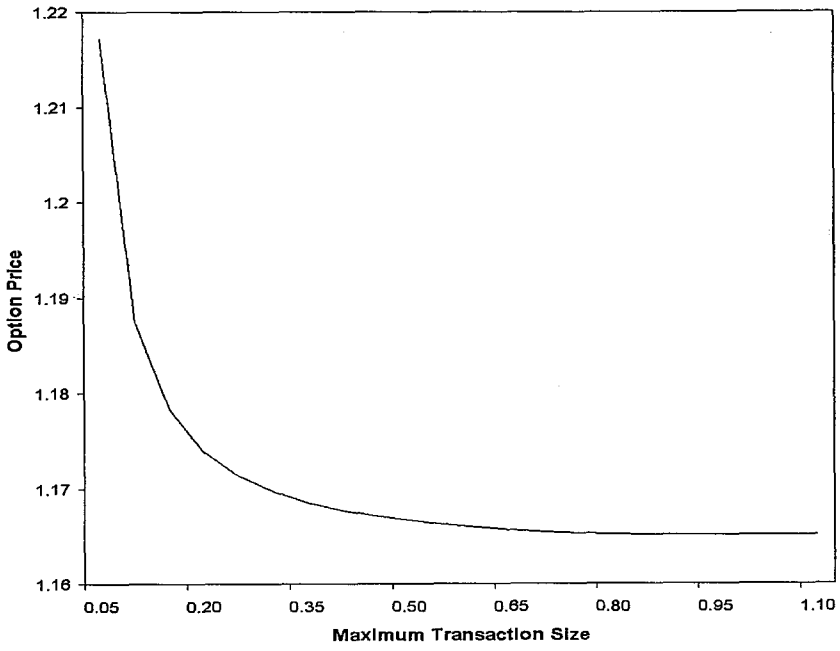


Figure 2.6: Option prices when trading volume is limited. Volatility 0.15.

price for the option. If a limit is placed on the size of transactions carried out then for a higher drift rate a higher price is produced. The reason for this is that a risk averse investor, as the option writer is assumed to be, will want to purchase more of the risky stock if the drift rate or expected return is high. The risk and return of the stock if the drift is significantly higher than the risk free interest rate is a far more attractive proposition than if the difference is marginal. The limit on volume, therefore, affects both the hedge for the option and the expected utility generated through optimal playing of the stock market. Both of these factors combine to produce the high writing price. A different algorithm was used in the preparation of this Chapter to test the validity of a naive approach to hedging. In this approach the optimal stock holding given a perfect market was found in the trial evaluation. Then it was assumed that for each possible portfolio, if this position was possible given the market imperfection then the writer transacted to this optimum. If it was not possible then he transacted to a holding as close as possible to this point. The value function is then calculated. This strategy when applied to the hedging of a vanilla call option gives the same writing price as the globally optimal.

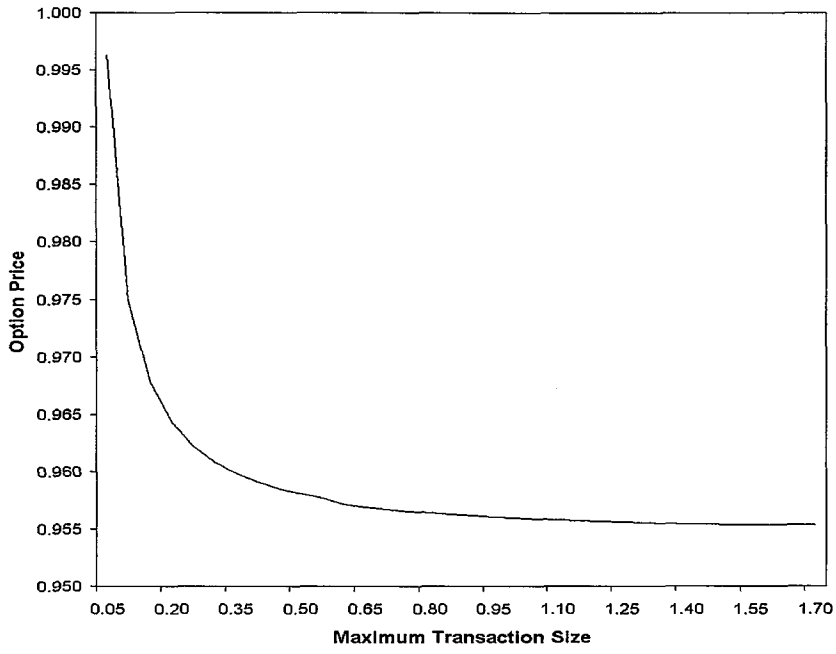


Figure 2.7: Option prices when trading volume is limited. Drift 0.12.

2.6 Conclusions

A general approach to option pricing has been described and used to investigate the effect that market imperfections have on option prices.

The prices computed for options on illiquid assets confirm the general intuitive ideas one might have. The main problem not addressed here is the random nature of market illiquidity and the inability of participants to predict it. However, the random nature of illiquidity may be addressed by simulation, using the current model to evaluate the scenario simulations.

The drift of the underlying asset plays a major role in the price of an option when transactions of the asset are restricted in size. The beauty of Black-Scholes is that this drift rate is eliminated from the equation making the result risk neutral with no need to carry out the tricky procedure of estimating the drift. However under the assumption that trade sizes are limited this drift is essential if pricing is to be done accurately. The presence of this drift is a drawback but unavoidable in an imperfect market setting where risk preferences are being specifically described and no risk is assumed by the writer for free.

The two different algorithms used for the trade limit problem yield the same writing price

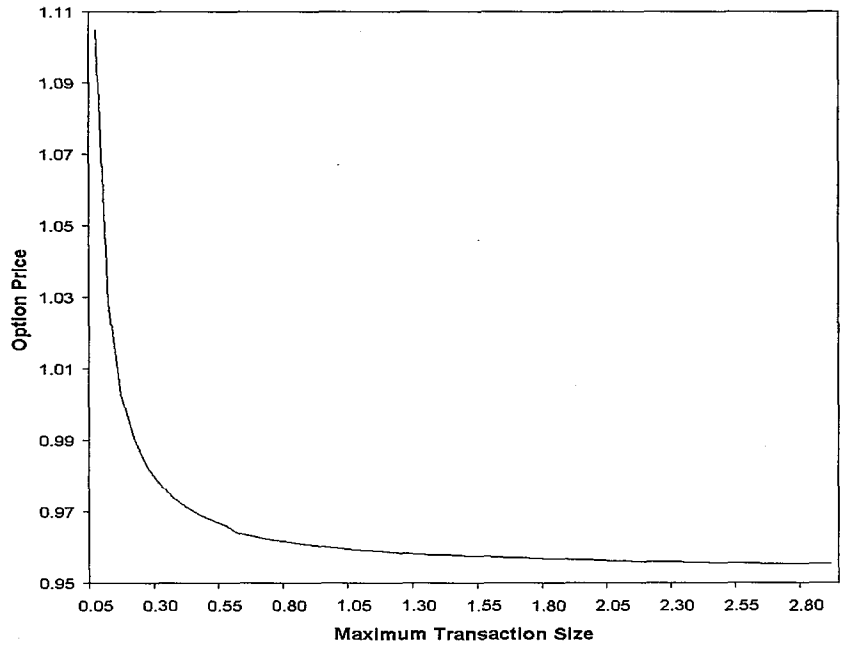


Figure 2.8: Option prices when trading volume is limited. Drift 0.15.

showing that for a plain vanilla option, the optimal strategy is to transact as near as possible to the stock holding you would otherwise want if the trading limitation was not in place.

Chapter 3

Volatility Smiles and Static Replication

3.1 Introduction

The Black-Scholes equation for the pricing of derivative securities assumes a constant volatility for the underlying asset. In reality the market prices of plain vanilla options imply volatilities that differ according to the time to expiration and strike price of the option. The variation of implied volatility with time to expiration may be extremely pronounced. The lack of trading in the underlying currency pair leads to a "Monday Effect" in the foreign exchange options market. The lack of trading leads to a lack of volatility on a Saturday and Sunday. Therefore an option written on Friday with expiration on Monday will be priced using a volatility such that the premium is similar to that for an identical option with only one day until expiration. The option expiring on Tuesday will then have a very different implied volatility. Supply and demand for options, which is a market just like any other, also leads to variation of volatilities for at-the-money options across expirations. This may also be exaggerated by the existence of perceived "good" and "bad" dates when economic numbers may be released. In addition to this time dependence, the situation also exists where options with the same tenor and underlying security have different implied volatilities due to their different strike prices. This difference may be explained as the market's modification of the lognormal assumption made by Black-Scholes [5]. The market for options may believe that out-of-the-money options are undervalued by the Black-Scholes model and

increase the volatility used in order to price such options. This change in volatility may also be viewed as a supply and demand issue. If, for instance, intervention was expected from a central bank in order to strengthen their currency, a hedger may wish to purchase out-of-the-money call options as protection from such an event. This would result in the short supply of these options and implied volatility would rise accordingly. This time and strike dependence of the volatility has become known as the "smile". Table (3.1) provides an example of a typical volatility smile for 3month USD/JPY options. It is convention

Call/Put	Delta	Volatility	over ATM
c	0.10	15.55%	1.55%
c	0.25	14.45%	0.45%
c	0.35	14.25%	0.25%
c	0.50	14.00%	0.00%
p	-0.50	14.00%	0.00%
p	-0.35	14.10%	0.10%
p	-0.25	14.30%	0.30%
p	-0.10	15.15%	1.15%

Table 3.1: A Typical Volatility Smile for 3month USD/JPY Options.

in the over-the-counter foreign exchange plain vanilla options market to refer to "deltas" (where the delta of an option is defined as the first derivative of the option's price with respect to spot price, $\frac{\partial C}{\partial S}$) rather than a strike price. With an underlying asset which moves so much, it would be inconvenient to adjust volatility prices for each strike price when spot changes. To avoid these difficulties, interbank counterparties trade options complete with a spot transaction equal to the delta hedge for that particular contract. With this the case, as long as the writer and purchaser agree on a spot price, volatility and delta ratio for the option, the strike price may be implied. This leads to the commoditisation of options according to their delta, for example, 25Δ calls. This results in a number of benchmark options, typically 50 (approximately at-the-money forward), 35, 25 and 10-Δ calls and puts. Each of these is quoted as a volatility price by market makers. It is that array of prices which is illustrated in table (3.1).

Work by Dupire [25] and Derman, Kani and Chriss [21] addresses the problem of constructing a market model that is consistent with current market prices for plain vanilla options.

It is the work of Derman, Kani and Chriss that is described here, and then expanded for use in pricing exotic foreign exchange options. Rather than assuming a lognormal process for asset prices, constructing a trinomial tree, then using this tree to price an option, the converse is attempted. The price of liquid vanilla options are well known in the market. These prices may be written down in terms of a payoff multiplied by the discounted path probabilities until expiration implied by option market prices. By inverting these relationships for call and put options a price distribution is recovered and a tree is constructed. This tree may then be used to price other derivatives written on the same asset giving a value we are confident is consistent with the liquid products in the market.

Over recent years barrier options have become important products in the foreign exchange market. Their dependence on hitting a trigger to knock in or out, makes knowledge of the smile crucial for accurate pricing. Volatility throughout the entire life of the barrier affects the expected first hitting time and so must be considered in any pricing and hedging. In this section the work of Derman, Kani and Chriss [21] is adapted specifically to improve convergence of barrier prices and make foreign exchange product pricing possible.

In common with an out-of-the-money vanilla option, the price of a knockout option becomes more sensitive to changes in volatility as volatility increases. That is, the vega (defined as the first derivative of options price with respect to volatility) of the option increases as the volatility increases. Also, the vega of the option decreases as volatility decreases. This is a favourable situation that we would be willing to pay money for. However, traditional valuation assumes a constant volatility thus missing this extra piece of value. In addition, we see that as the spot price of the knockout option increases or decreases the vega of this option increases. This reminds us of being long a vanilla option with a strike different the forward price. If such options are given a volatility higher than the at-the-money volatility we use to price the exotic, then we are clearly under-pricing the barrier option. Through purchasing the knockout option using at-the-money volatility to value it, we are obtaining the properties of an option which, in volatility terms, is more expensive. This clearly represents some kind of arbitrage opportunity.

Figures (3.1) and (3.2) show the variation in the vega for an up and out call option with maturity 3months, spot 110.00, volatility 14 percent, foreign rate 0.05 and domestic interest rate zero. The strike price is 108.00 and the barrier is 117.00. Prices and vega were found using a constant volatility, analytic formula. As may be observed from figure (3.1) the

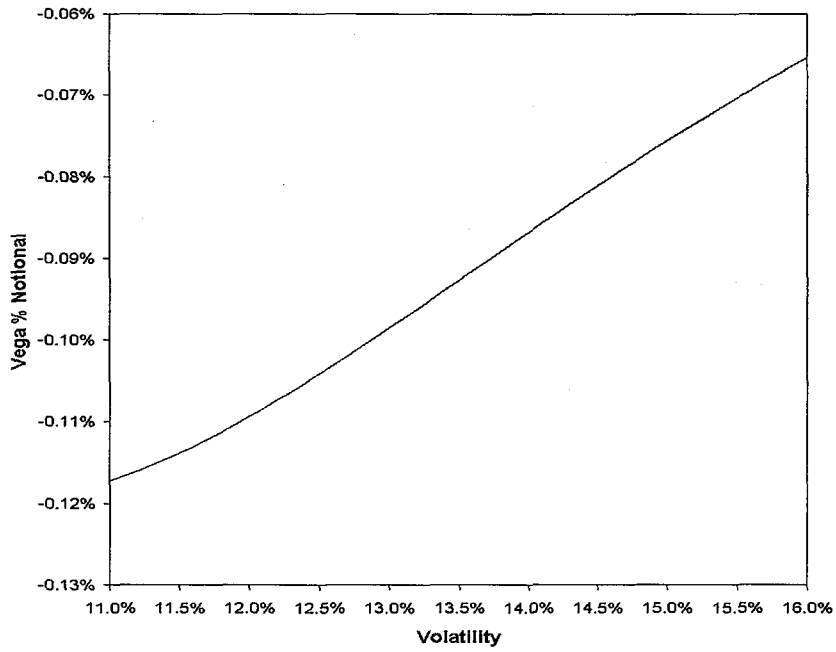


Figure 3.1: The change in vega for a change in volatility for an up and out call option

exotic option has larger vega as volatility increases and smaller vega as volatility decreases. This is the same characteristic as for an out-of-the-money vanilla option (as is discussed in section (4.2.2)). The barrier option therefore behaves in some way like an out-of-the-money option. From the smile effect of the market given in table (3.1) we note that all out-of-the-money options are assigned a higher volatility than the at-the-money options. This implies that the barrier option, priced using at-the-money volatility, should trade at a premium to its theoretical value. We also note that in figure (3.2) the vega increases as spot increases. This is similar to that of an out-of-the-money option with strike higher than the forward price. From table (3.1) we see that options with strikes greater than the forward price have implied volatilities greater than the at-the-money options. Again this shows that the exotic option has properties priced at a premium by the vanilla option market.

In the last section of the chapter the static hedging of exotic options is examined. Again, particular consideration is given to the hedging of barrier options. Present methodology is reviewed and applied to various foreign exchange instruments. The work is then extended to price exotic options under a time dependent volatility. Further examples are provided detailing the replication of exotic options under a full volatility smile. Results are presented and possible further work is discussed.

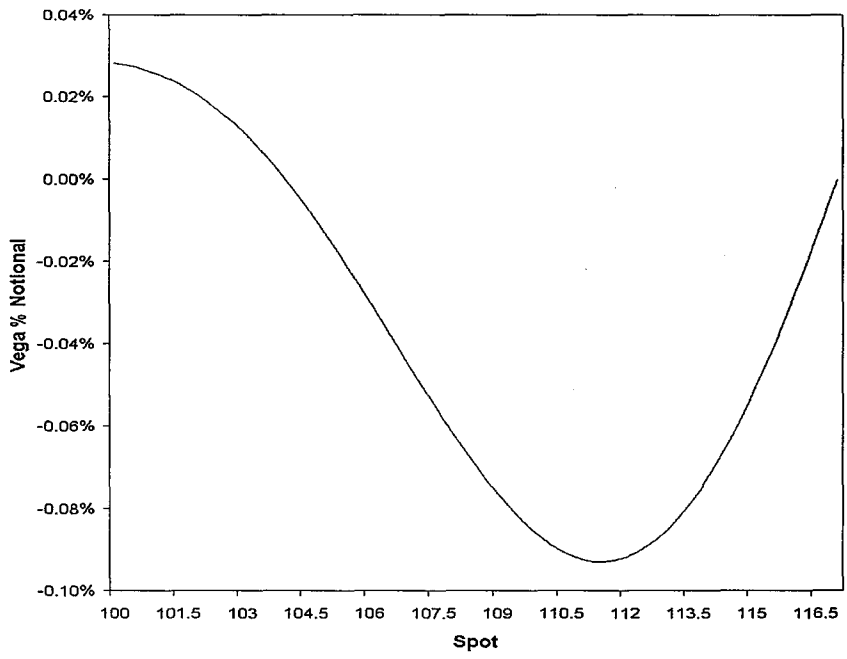


Figure 3.2: The change in vega for a change in spot for an up and out call option

3.2 Implied Trinomial Trees

3.2.1 Constant Volatility Trees

The traditional trinomial tree is constructed assuming constant volatility for the underlying asset. Several parameters require calculation for the model to be complete. A basic tree centred about the initial spot price, S, could be specified by the following,

$$S_{up} = Se^{\sigma\sqrt{2\Delta t}}, \tag{3.1}$$

$$S_{mid} = S, \tag{3.2}$$

$$S_{down} = Se^{-\sigma\sqrt{2\Delta t}},$$

and the constant probabilities at each node given by,

$$p = \left(\frac{e^{\frac{r\Delta t}{2}} - e^{-\sigma\sqrt{\frac{\Delta t}{2}}}}{e^{\sigma\sqrt{\frac{\Delta t}{2}}} - e^{-\sigma\sqrt{\frac{\Delta t}{2}}}} \right)^2, \tag{3.3}$$

$$q = \left(\frac{e^{\sigma\sqrt{\frac{\Delta t}{2}}} - e^{\frac{r\Delta t}{2}}}{e^{\sigma\sqrt{\frac{\Delta t}{2}}} - e^{-\sigma\sqrt{\frac{\Delta t}{2}}}} \right)^2, \tag{3.4}$$

$$m = 1 - p - q.$$

S is the spot price, Δt the time step, σ the volatility and r the riskfree rate. p and q represent the up and down probabilities at each node, whilst m represents the probability that spot remains unchanged.

The aim of Derman, Kani and Chriss' [21] work is to construct a trinomial tree that correctly prices the liquid vanilla options used to build it. Volatility will not be constant and the probabilities of up and down moves will be different at each node of the tree. As we will see, the probabilities are implied by the market prices so that the options used are priced by the tree exactly as in the market.

3.2.2 Forward conditions and Arrow-Debreu Prices

We use the same notation as Derman, Kani and Chriss [21]. Referring to figure (3.3), the following will be used:

n = no. of the time step in the trimomial tree,

i = no. of the stock price level in the trimomial tree,

(n, i) = node of the the trimomial tree at time step n and stock level i,

r = risk – free interest rate,

s_i = known stock price at node (n, i),

F_i = forward price at level $n + 1$ of the stock price s_i at level n,

S_i = known stock price at node $(n + 1, i)$,

and also the strike for options expiring at level $n + 1$,

λ_i = known Arrow – Debreu price at node (n, i),

p_i = unknown risk – neutral transition probability,

from node (n, i) to node $(n + 1, i + 2)$,

q_i = unknown risk – neutral transition probability,

from node (n, i) to node $(n + 1, i)$,

m_i = unknown risk – neutral transition probability,

from node (n, i) to node $(n+1, i+1)$.

Assume that we have already specified such parameters as the number of time steps over

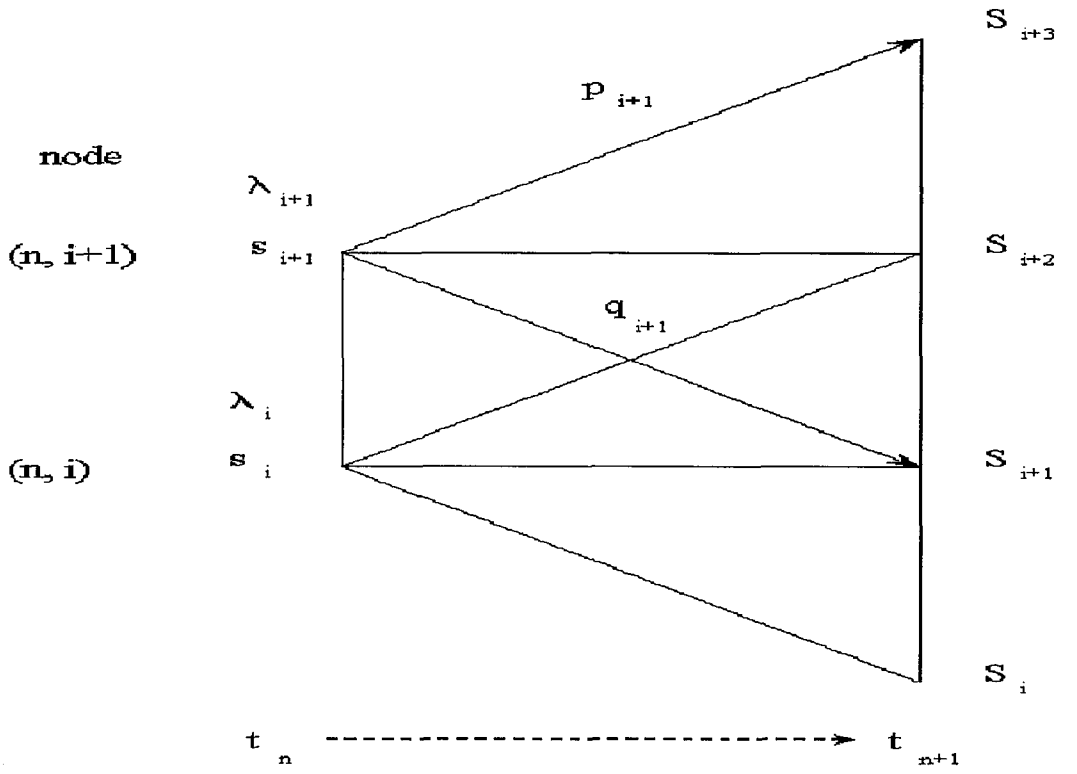


Figure 3.3: Notation for the implied trinomial trees.

which we are calculating the option value, the initial stock price and the state-space for the trinomial tree. Tree nodes are denoted (n, i) where n refers to the index time step t_n, and i is the stock price level index.

The tree to be constructed using these stock prices (n, i) and the probabilities p_i, q_i must be arbitrage free. The following conditions must be satisfied. Firstly the stock prices must be chosen so that the forward price of s_i denoted F_i lies between the prices S_i and S_{i+2}:

$$S_i \leq F_i \leq S_{i+2}.$$

The next relationship to be satisfied also concerns the forward price, but restricts the node probabilities. The forward price of the stock at node (n, i) must equal the expected value of the stock at the next time level t_{n+1}. In other words

$$p_i S_{i+2} + (1 - p_i - q_i) S_{i+1} + q_i S_i = F_i. \tag{3.5}$$

This is illustrated in figure (3.4).

If this is not the case and the forward price is larger than the left hand side of (3.5) then an arbitrageur may sell the forward and buy the stock and on average make a profit. If the opposite is true and the forward price is the smaller, the arbitrageur would buy the forward and sell the spot. Equation (3.5) rules out such opportunities. Also denote vanilla put and

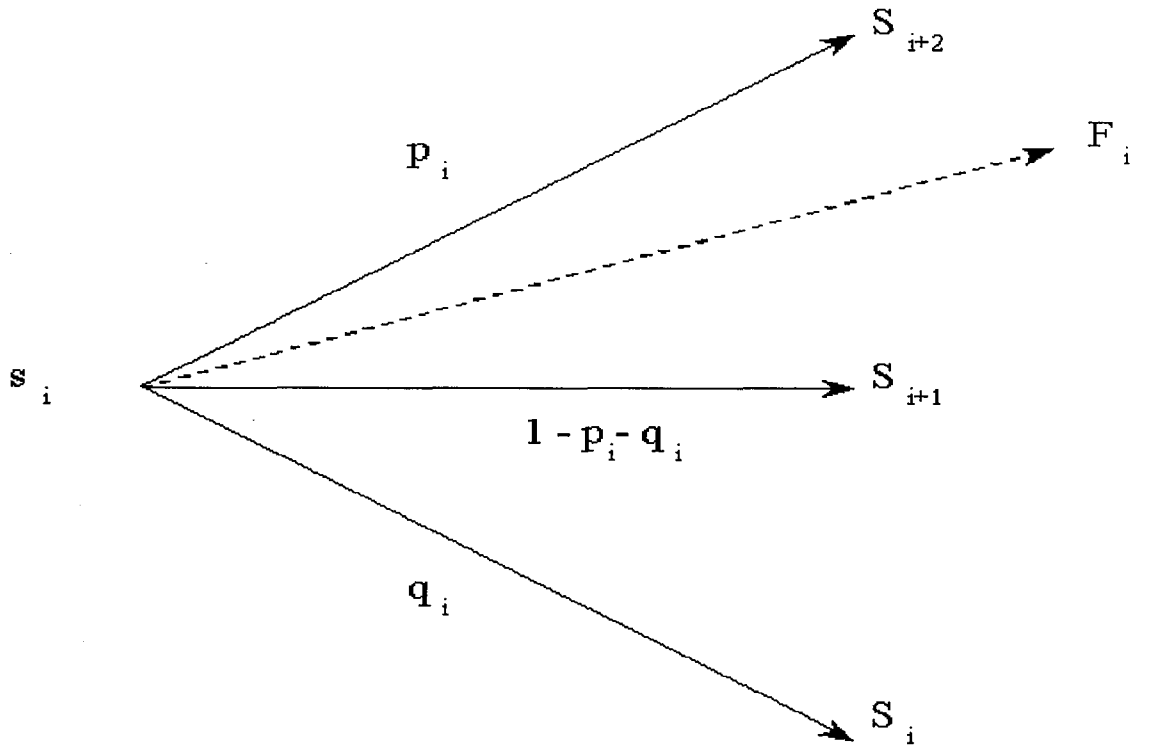


Figure 3.4: Forward condition for the implied trinomial tree.

call prices recovered from the market as in Derman and Kani. Let $C(S_{i+1}, t_{n+1})$ represent the price today of a call with strike price S_{i+1} and expiration at time t_{n+1}. The put price is denoted in a similar way replacing C with P. Prices for these options are taken from a "volatility surface" where known prices are interpolated in order to find values of less liquid options. It is necessary to find either a put or a call price for options coinciding in strike and expiration with every node of the tree.

The probabilites p_i and q_i are found such that the option used to produce them will be exactly priced by the resulting tree. In order to do this the price of the option must be written in terms of discounted path probabilities - Arrow-Debreu prices - and the payoff at

expiration. We have the following for a call.

$$C(K, t_{n+1}) = e^{-r\Delta t} \sum_j \{\lambda_{j-2} p_{j-2} + \lambda_{j-1}(1 - p_{j-1} - q_{j-1}) + \lambda_j q_j\} \max(S_j - K, 0), \quad (3.6)$$

where the λ_i are the discounted path probabilities or Arrow-Debreu prices. See Appendix (A) for further details.

Graphically the option price has been constructed as in figure (3.5).

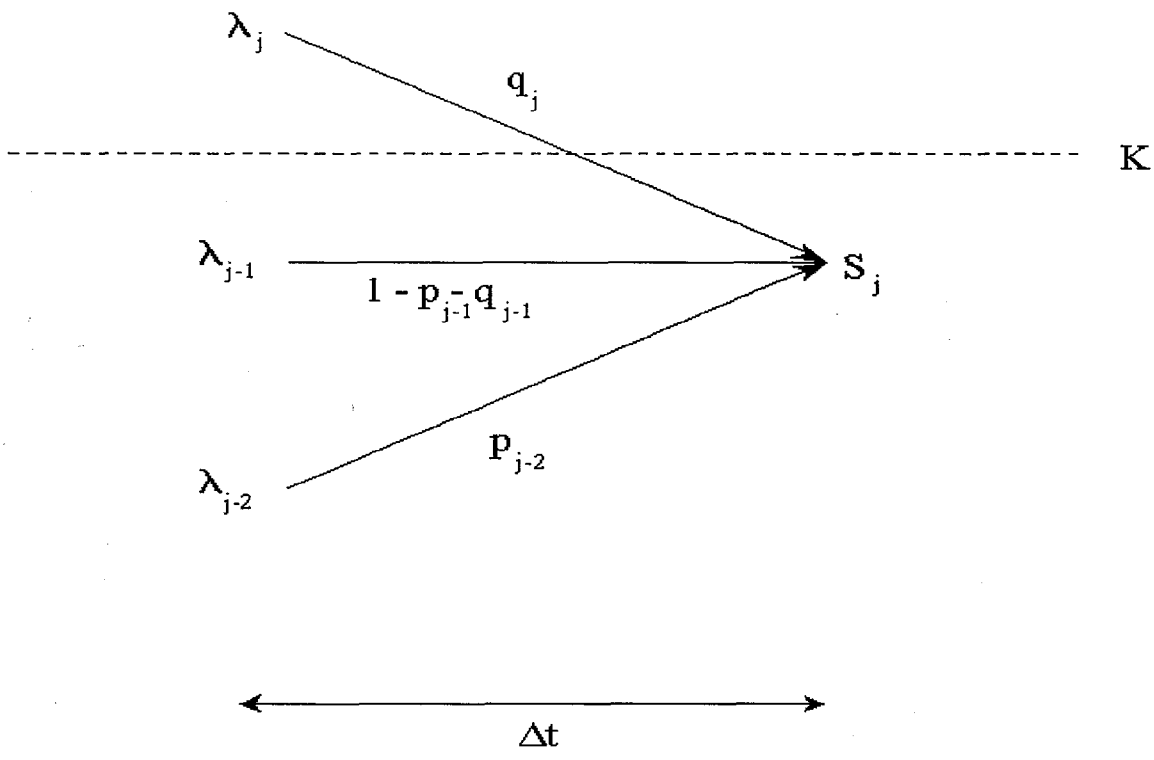


Figure 3.5: Construction of the option price.

3.2.3 Implying Node Probabilities and Constructing the Tree

The expression for the prices given above may now be broken in two by substituting $K = S_{i+1}$, so that the $C(K, t_{n+1})$ term is split.

This operation gives

$$e^{r\Delta t} C(S_{i+1}, t_{n+1}) = \lambda_i p_i (S_{i+2} - S_{i+1}) + \sum_{j=i+1}^{2n} \lambda_j (F_j - S_{i+1}), \quad (3.7)$$

Figure (3.6) shows the contributions that the stock price nodes make to the value of the option and why the call price may be represented as described above. The maximise operator is split nicely in two by considering only the upper portion of the diagram which contributes the non-zero part of the price. In this expression the only unknown is the

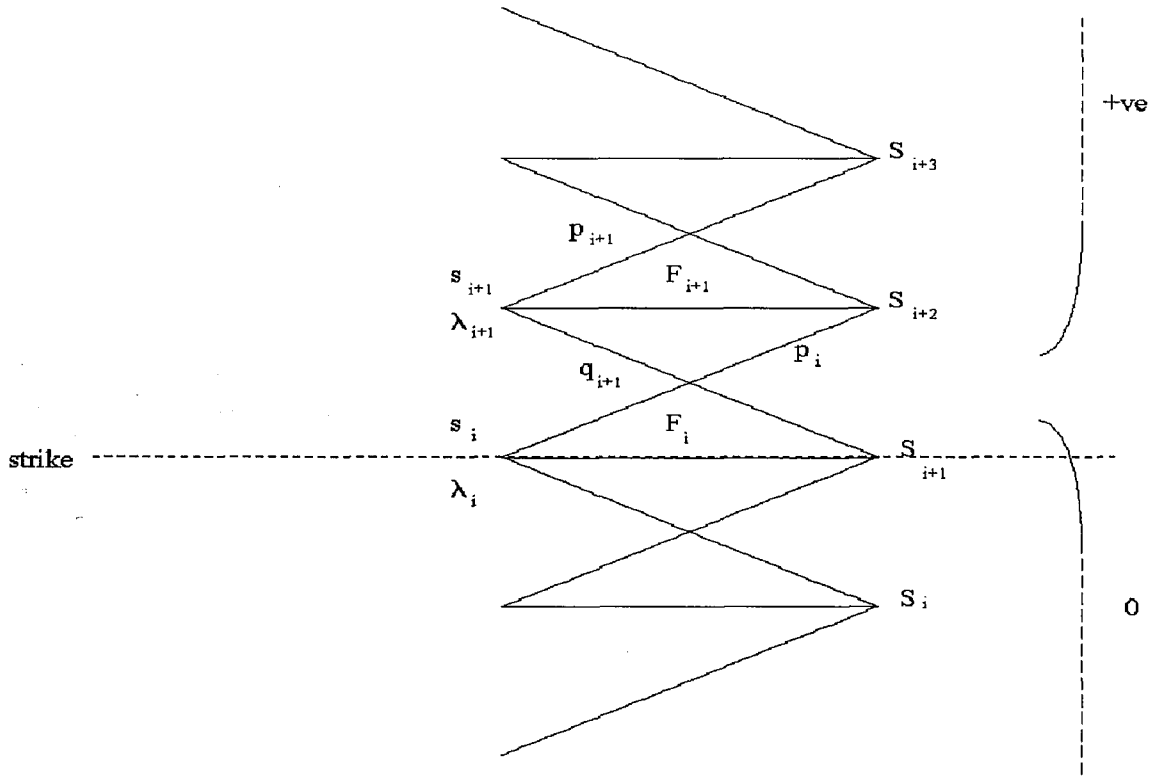


Figure 3.6: Stock price node contributions to the value of a call option.

probability p_i. Solving for p_i gives

$$p_i = \frac{e^{r\Delta t}C(S_{i+1}, t_{n+1}) - \sum_{j=i+1}^{2n} \lambda_j (F_j - S_{i+1})}{\lambda_j (S_{i+2} - S_{i+1})}. \tag{3.8}$$

We use the no-arbitrage condition (3.5) that forward prices must equal the expectation of the future spot price to derive q_i from p_i.

$$q_i = \frac{F_i - p_i(S_{i+2} - S_{i+1}) - S_{i+1}}{S_i - S_{i+1}}. \tag{3.9}$$

Using market call prices the probabilities at nodes above the centre of the tree (the initial stock price in our case) have been calculated. An analogous method requires the probabilities on and below the central node to be implied by put option prices.

$$P(K, t_{n+1}) = e^{-r\Delta t} \sum_j \{\lambda_{j-2} p_{j-2} + \lambda_{j-1}(1 - p_{j-1} - q_{j-1}) + \lambda_j q_j\} \max(K - S_j, 0), \tag{3.10}$$

Letting $K = S_{i+1}$,

$$e^{r\Delta t}P(S_{i+1}, t_{n+1}) = q_i\lambda_i(S_{i+1} - S_i) + \sum_{j=0}^{i-1}\lambda_j(S_{i+1} - F_j). \tag{3.11}$$

Solving for q_i,

$$q_i = \frac{e^{r\Delta t}P(S_{i+1}, t_{n+1}) - \sum_{j=0}^{i-1}\lambda_j(S_{i+1} - F_j)}{\lambda_i(S_{i+1} - S_i)}, \tag{3.12}$$

and using the forward condition,

$$p_i = \frac{F_i + q_i(S_{i+1} - S_i) - S_{i+1}}{S_{i+2} - S_{i+1}}. \tag{3.13}$$

The entire tree has now been implied.

3.2.4 Potential Difficulties

In constructing our implied trinomial tree we have already imposed some restrictions on the state-space which may be chosen to represent the values of the stock price throughout the options life. We were also careful to ensure the presence of a forward condition on the probabilities so that the forward price at a time step is equal to the expected value at the next time step. These measures do not however ensure that things always run smoothly.

The values of p_i, q_i and m_i must always lie between 0 and 1 due to the fact that they are probabilities. It is not always possible to fit these probabilities to the option prices and have them remain within these limits. The local volatility from one step to the next may just be too high or alternatively the smile of the option volatility too extreme for our chosen state-space to cope. Generally these problems arise at the extremes of the tree and when outlying prices have occurred for some reason or other. Derman, Kani and Chriss [21] suggest the the first type of problem may be avoided in some cases by the judicious choice of a state-space that exhibits the main characteristics of the smile. For instance, if the smile is mainly temporal then the stock prices may be specified to mimic this temporal bias to some degree.

If redefining the state-space is unsuccessful we must discard the option price that has produced the unsatisfactory results and overwrite it with some price of our choosing. A simple way of doing this given by Derman and Kani [20] is to set

$$p_i = \frac{(F_i - S_i)}{(S_{i+2} - S_i)},$$

and $q_i = 1 - p_i$, with the central probability set to zero.

3.3 Improvements to the Basic Model

The model outlined above provides a good basis for the pricing of exotic options consistent with a volatility smile. However modifications are needed to price many of these exotics to within an acceptable error and make the model computationally tractable.

3.3.1 Foreign Exchange Barrier Options

Barrier options have become very popular instruments in recent years with a large volume being traded in the foreign exchange market. The key feature of a barrier option is that there is an asset price level in addition to the strike price that governs the value and payoff. An "out" barrier option knocks out if a trigger price is hit leaving the option worthless. Similarly, an "in" barrier option becomes live if a certain price is hit by the underlying.

This feature of barrier options makes the issue of "fat tails" of the distributions of prices particularly relevant. The payoff depends not only on the asset price at expiration but at every other time during the life of the option. The volatility smile is therefore of particular interest to a trader of barrier options. In pricing a barrier we must be sure that the smile is considered not only for the expiration and strike of the constituent call or put but for every time during the life of the option. The following section addresses these problems.

Choosing the State-Space for Pricing a Barrier Option

Despite the fact that options are routinely priced using both binomial and trinomial trees, the use of these for barrier options presents major problems. The first problem is convergence of any n-nomial model when used for pricing a barrier option. Prices found for barrier options using a binomial tree converge in a saw-tooth pattern. Even when ten thousand time steps are used the price produced may not be acceptable. The reason for this is the discrete nature of the tree and the continuous process it attempts to describe. If the value of the trigger price in the barrier option does not coincide exactly with nodes of the tree then the actual derivative priced is a barrier with some trigger other than the one desired.

It seems logical to increase the number of time steps used in order to get a more accurate model for asset price evolution. In fact this may lead to a worse price as the nodes could be further from the barrier than before.

Various methods have been employed to solve this problem. In our case the degrees of freedom provided by the trinomial tree are exploited and the state-space chosen carefully to aid the convergence of the price. The aim is to construct a state-space which still makes mathematical sense but also has nodes coincident with the barrier specified in the option contract. Firstly we remind ourselves how to build a trinomial tree with constant volatility.

A general constant volatility trinomial tree may be constructed using the following values for S_u, S_m and S_d the up, middle and down ratios for the movement of stock price.

$$S_u = Se^{\pi\Delta t+\phi\sigma\sqrt{\Delta t}}, \tag{3.14}$$

$$S_m = Se^{\pi\Delta t}, \tag{3.15}$$

$$S_d = Se^{\pi\Delta t-\phi\sigma\sqrt{\Delta t}}, \tag{3.16}$$

for $\phi > 1$ and any reasonable value of π.

In practice the values used most often are $\phi = \sqrt{2}$ and $\pi = 0$. This is the starting point for building a state-space for the barrier option.

Denote the factor by which the stock price moves up in one time step by u and the factor by which it may move down by d.

$$u = e^{\phi\sigma\sqrt{\Delta t}}, \tag{3.17}$$

$$d = e^{-\phi\sigma\sqrt{\Delta t}}. \tag{3.18}$$

Furthermore, m, the factor by which the stock price moves if the middle branch is taken is 1, to produce a recombining tree.

We want to adjust the value of ϕ so that, after a number of steps, the nodes become coincident with the trigger price. We do this in the following way. Assume an up barrier, with trigger price H.

(i) Calculate S_u.

(ii) Starting with the initial stock price S_0, calculate the prices for the top edge of the tree. This is done by calculating $S = S_0u^j$ where j represents the discrete time step.

(iii) When the trigger price satisfies $S_0u^{j-1} \leq H \leq S_0u^j$ stop the calculations and note the value of j. This is the number of steps at which we want to force the trigger and stock price to be equal.

(iv) Invert the formula for generating the tree to find the value of ϕ that causes the equality of trigger and node price on the top edge of the tree after $j-1$ discrete time steps.

(v) Using this value ϕ instead of $\sqrt{2}$ generate the stock price nodes for the entire tree.

This algorithm removes the inaccuracies caused by a trigger and tree that do not have compatible nodes. The inversion mentioned in step (iv) is given by,

$$\phi = \frac{1}{(j-1)\sigma\sqrt{\Delta t}} \ln\left(\frac{H}{S_0}\right).$$

The trigger of a down barrier option may be fitted in a very similar way by replacing u with d in the steps described above. Fitting double barriers is a slightly more tricky proposition but not more conceptually difficult. As long as the forward conditions are satisfied, thus ruling out arbitrage, we may choose any convenient state-space.

3.4 Static Hedging and its Application to Volatility Smiles

Black-Scholes [5] requires the continuous rebalancing of a replicating portfolio in order to hedge derivatives. This portfolio consists exclusively of risk-free bonds and the asset underlying the product to be priced. Under the perfect market assumptions of Black and Scholes this works well; however when real markets are considered, problems start to arise. A major practical drawback comes from the implementation of the hedging portfolio. The infinite trades involved in rebalancing incur transaction costs. It is also impossible to carry out an infinite number of trades so the practitioner must be selective about when he hedges.

These problems have been addressed by a number of researchers all seeking to simplify the replication procedure. A focus for their attention has been the use of standard "vanilla" options in set quantities to synthesise more exotic products. Work carried out by Derman, Ergener and Kani [22] is described and generalised in this section.

Firstly their general methodology is described in a continuous setting. Examples are given of the replication of some important foreign exchange products. The work is then extended to address problems encountered in the pricing of options with non-constant volatility.

3.4.1 Contract Replication Through Static Hedging

In this section we investigate the use of plain vanilla options to create new exotic derivative products. The types of derivatives currently on offer in the financial markets contain features that make their risk management distinctly different to that of plain vanilla options. For example, in-the-money knockout barrier options (known as reverse barriers), have an infinite gamma close to the trigger level near expiration. In addition, such products experience large positive theta when expiration is close and the barrier is near to current spot. These features are difficult to hedge dynamically but may be hedged using plain vanilla options as we will discuss here. Static replication will enable us to consider the effect of non-constant volatilities on the value of new exotic products and in some cases allow the use of a volatility smile for pricing. This scheme makes the assumption that, with the exception of the spot price, all parameters remain the same or behave in the manner that was prescribed at the time of pricing.

The options portfolios produced for the static hedges may not be unique. Selection of the portfolio may by carried out under further conditions additional to the mimicking of the exotic. For example, vega, gamma or some other risk management parameter may be minimised.

Under the scheme described here, in order to be perfectly hedged an infinite number of options must be purchased and sold. There are exceptions and some of these may be found in Derman, Ergener and Kani [5]. As we will demonstrate, a modest number of options results in reasonably accurate replication.

3.4.2 Static Hedging in the Continuous Black-Scholes World

To statically hedge an option we must know its payoff along all boundaries. We may then replicate these values in the relevant places without any injection of extra funds or introduction of payoffs within the boundaries so that the value of the portfolio is equal to the value of the target option.

To replicate the option perfectly at every single point along its boundary would require an infinite number of vanillas. We therefore specify points in time at which we want to be hedged. We may choose points at which we are expecting particular events or simply

specify arbitrary points in time.

To replicate the target option we use liquid products, taking care to choose those options which do not introduce extra, unwanted, cashflows or payoffs. At the expiration boundary we use a product or products that replicate the payoff precisely and use only out-of-the-money calls for upper boundaries (above the current asset price) and out-of-the-money puts to replicate payoffs at lower boundaries (below the current asset price) throughout the life of the option. In trying to replicate an up and out call option with barrier 120, for instance, we would use call options with strikes of 120 or higher. Any strikes less than 120 would introduce some intrinsic value in addition to the replication already performed for the expiration boundary. This would result in an incorrect replicating portfolio.

Consider the general target option with payoff boundaries as in figure (3.7). We begin the replication by choosing a vanilla option that pays off in the region between the upper and lower boundaries to the same value as the target option. Generally, when pricing barrier options, a single put or call with expiration time the same as the target option held long or short will suffice. Rolling backwards to the penultimate hedging point it is clear that we require another two vanillas to replicate the target option at the upper and lower boundaries. For the upper boundary we add a number of calls struck at, or above, the boundary asset price. For the lower boundary add to the position with out-of-the-money puts. Using out-of-the-money vanillas leaves the value within the boundaries. which was matched at the expiration time. untouched and so no additional payoffs are introduced. Rolling back to the next time point leaves us unhedged again so we must calculate the put and call positions needed to maintain the replication. Repeating this procedure until we are back at the present day, when the target option is to be written or bought, provides the replicating portfolio.

3.4.3 Examples of Static Replication

We replicate a barrier option with the following specifications,

Type: Up and out call option, a reverse knock out = UOC,

- spot: $S = 100$,

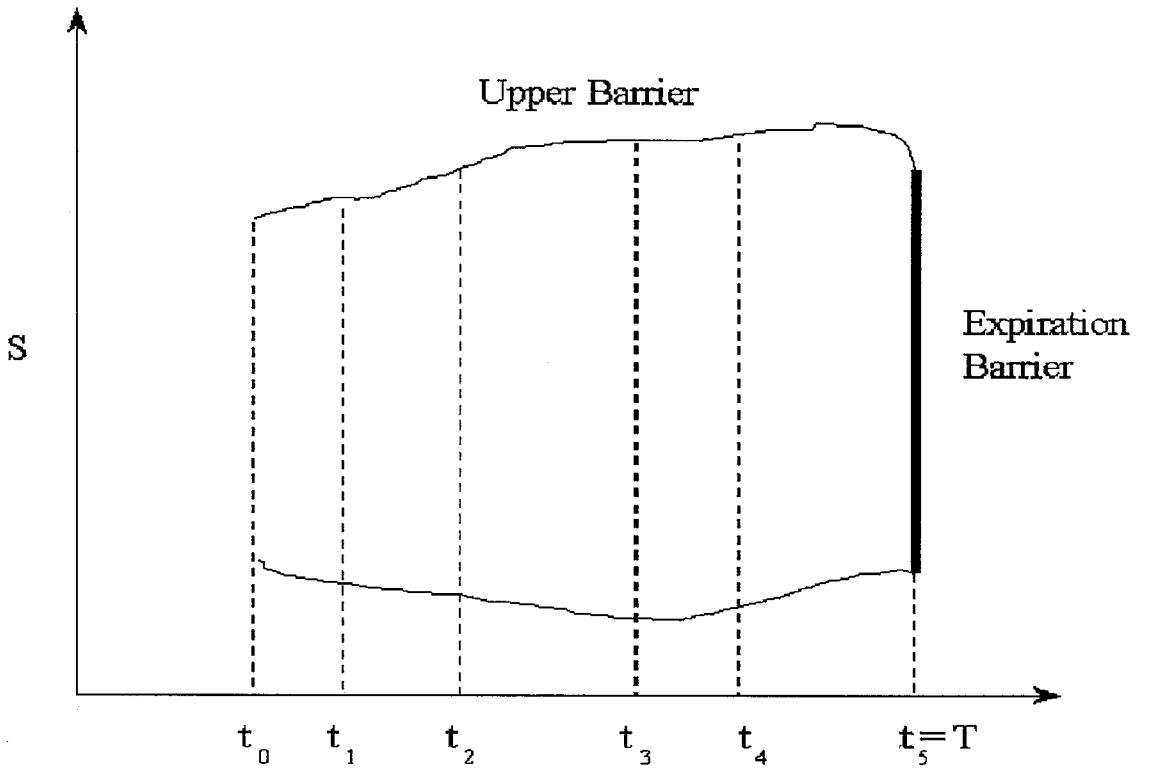


Figure 3.7: A general barrier option.

- strike: $K = 100$,
- barrier: $B = 120$,
- expiration: $T = 1$,
- volatility: $\sigma = 0.15$,
- domestic rate: $r_d = 0.08$,
- foreign rate: $r_f = 0.05$.

In order to hedge this option the payoff must be replicated at all the boundaries. See figure (3.8). We will use the following notation in all examples. $C(S, K, t_p, t_e)$ refers to the price of a call option at pricing time t_p and spot level S with strike K and expiration time t_e. The time left before such an option expires is therefore given by $t_e - t_p$.

Starting at expiration we find an option that fits the boundary of the UOC with strike 100. A call option with strike 100 that expires at time T matches the payoff precisely on the final boundary if the barrier at 120 is never hit. We therefore have one 100 call of expiration

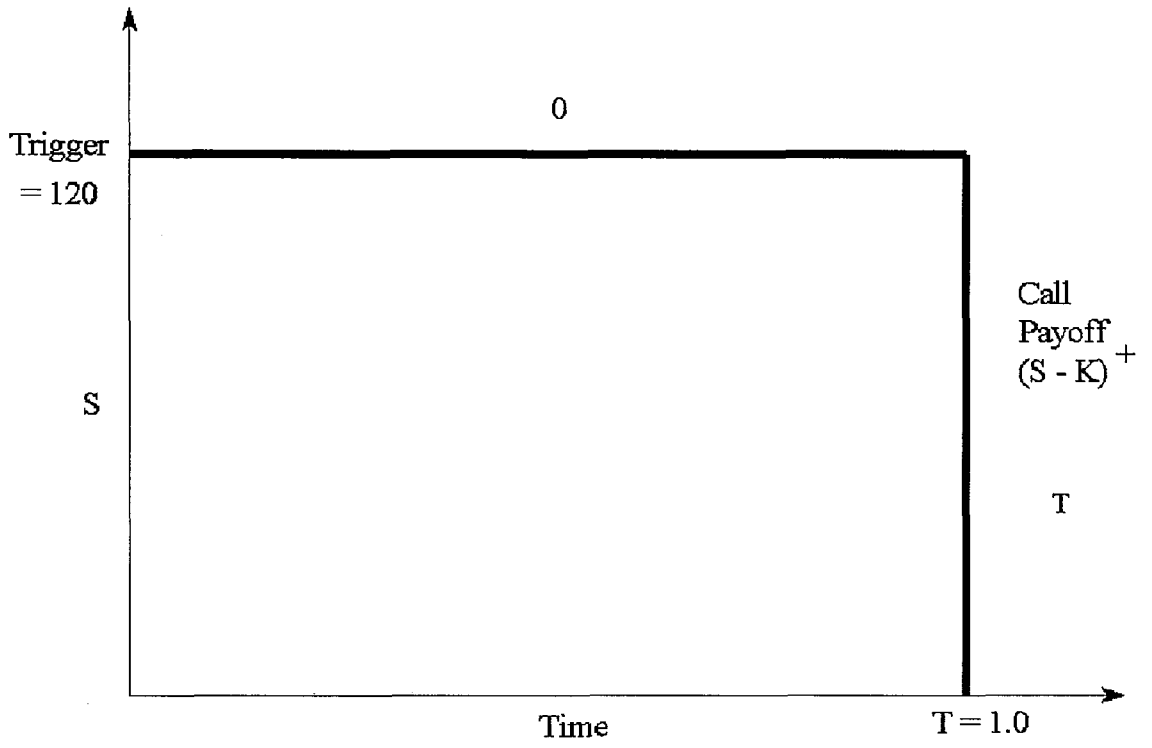


Figure 3.8: Boundary values for an up and out call option.

time 1.0 in our portfolio. We now go back from time 1.0 to time 0.8. Here we must fit our portfolio to the value the UOC has at the boundary. At the boundary, 120, the call option knocks out and so has a value of zero.

Revaluing our current portfolio at time 0.8, which consists of one unit of the 100 call option expiring at time 1.0 we find $C(120, 100, 0.8, 1.0) = 20.39863$. We require a portfolio worth zero at spot 120.00 so we add, or subtract, a quantity of call options of strike 120 and expiration at $T = 1.0$. At time 0.8 this one unit of this option is worth $C(120, 120, 0.8, 1.0) = 3.537377$. We go short 5.76659 of these 120 calls.

$$\frac{20.399}{3.537} = 5.767. \tag{3.19}$$

The replicating portfolio now stands as in table (3.2).

Rolling back another 0.2 years to time 0.6 the portfolio in table (3.2) is worth,

$$1.0 \times C(120, 100, 0.6, 1.0) - 5.77 \times C(120, 120, 0.6, 1.0) = -8.91. \tag{3.20}$$

Now both options are 0.4 years from expiration. See figure (3.9).

Amount	Strike	Expiration	Type
1.000	100	1.0	Call
-5.767	120	1.0	Call

Table 3.2: Replication portfolio for an up and out call option to be revalued at time 0.6.

The portfolio is required to have value zero on the boundary, 120. We then add x 120 calls valued at time 0.6, that expire at time 0.8 such that the following is satisfied,

$$1.0 \times C(120, 100, 0.6, 1.0) - 5.77 \times C(120, 120, 0.6, 1.0) + x \times C(120, 120, 0.6, 0.8) = 0. \quad (3.21)$$

Rearranging gives $x = 2.519$.

We then roll back another time step to time 0.4 and revalue the portfolio shown in table (3.3).

Amount	Strike	Expiration	Type
1.000	100	1.0	Call
-5.767	120	1.0	Call
2.519	120	0.8	Call

Table 3.3: Replication portfolio for an up and out call option to be revalued at time 0.4.

Subsequent operations, until the replication is complete for every required time point, are similar. See table (3.4) for a complete listing of the portfolio.

Amount	Strike	Expiration	Type	Value
1.000	100	1.0	Call	7.120
−5.767	120	1.0	Call	−7.169
2.519	120	0.8	Call	2.050
0.804	120	0.6	Call	0.349
0.373	120	0.4	Call	0.055
0.211	120	0.2	Call	0.002

Table 3.4: Replication portfolio for an up and out call option.

Summing the values at time zero of the options in the portfolio gives a reasonable approx-

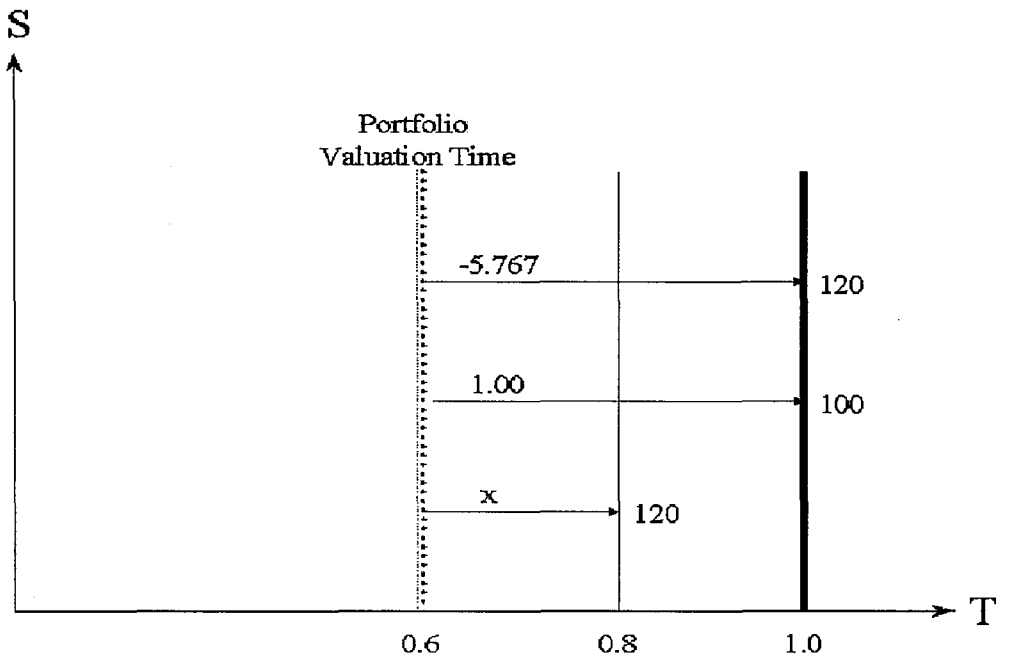


Figure 3.9: Replication step for an Up-and-Out Call option.

imation to the value of the target option considering the small number of options we have chosen to replicate it with. Increasing the number of options used in the replication produces a better price. For instance, if we used fifty options in the example above the price given for the barrier by adding the value of the options portfolio would be 1.985, close to the analytic value of 1.936. See figure (3.10). A common way to price an up-and-in call (or put) option is to use the following relationship,

$$UIC + UOC = \text{Call}.$$

An interesting experiment is to see whether the static replication portfolio for an UOC may be modified in such a simple way and provide a replicating portfolio for a UIC of the same parameters. The portfolio should consist of a long call option with expiration and strike the same as the barrier option, and short the entire options array constructed for the knockout option. The portfolio in table (3.5) is produced for an UIC with parameters the same as in our UOC example above.

By examining a couple of boundary points we find that the replication is correct and was simple to derive from the knockout barrier option.

For example, at time 0.8 and stock price 120 the call should have knocked in and be worth

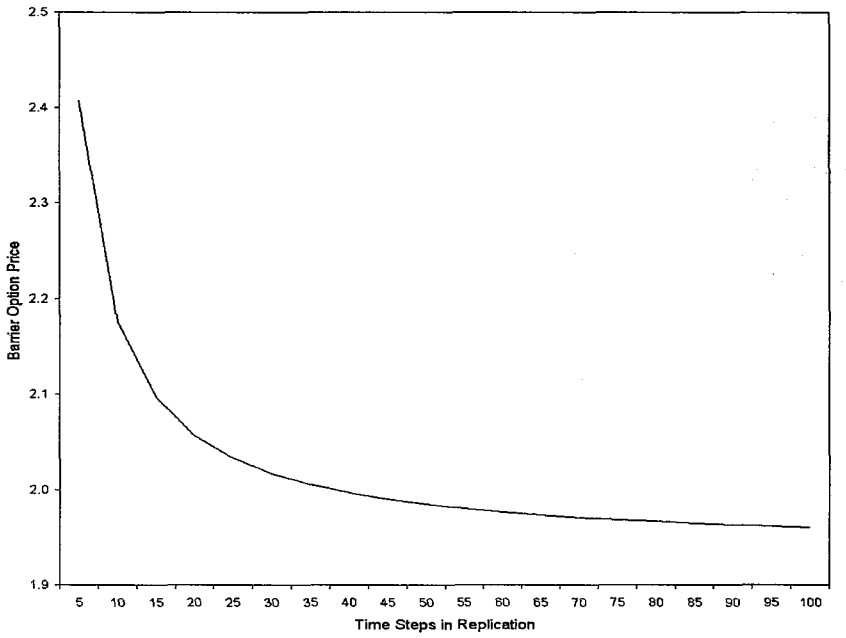


Figure 3.10: Covergence of the replication scheme to the analytic price of the barrier option.

$C(120, 100, 0.2) = 20.39$. The portfolio consists of the options shown in table (3.6) that are still live at $T = 0.8$.

Now $5.766 \times 3.537 = 20.39$ so the UIC has been replicated.

Moving Double Barrier Options

We now consider a more exotic style of barrier option and demonstrate the flexibility of the static replication scheme. The contract we describe and price is known as a moving double knockout call option. The terms of such a contract is as follows. The option has payoff $\max(S_T - K, 0)$ at time T. If the barriers depicted in figure (3.11) are touched whilst active then the option knocks out and the contract expires worthless. See table (3.7) and figure (3.11).

We, find that the replicating portfolio for such a call option with volatility 12%, zero interest rates, spot 120.00 and strike 120.00 is as shown in table (3.8)

As with the more standard barrier option, the Up-and-Out Call, we begin by choosing an option, or options, with the same payoff as the exotic at expiration time. In this case

Amount	Strike	Expiration	Type	Value
−0.211	120	0.2	Call	−0.002
−0.373	120	0.4	Call	−0.055
−0.804	120	0.6	Call	−0.349
−2.519	120	0.8	Call	−2.050
5.766	120	1.0	Call	7.169
1.0 − 1.0 = 0.0	100	1.0	Call	0

Table 3.5: Replication portfolio for an up and in call option.

Amount	Strike	Expiration	Type	Value
5.766	120	0.2	Call	3.537
1.0 − 1.0 = 0.0	100	0.2	Call	0

Table 3.6: Portfolio for the up and in call option at time 0.8 and spot 120.

we require a call option struck at 120. This is first option in the replicating portfolio. We continue by rolling back one time step, in this case we choose 0.125 of a year, and calculating the value of the current portfolio (containing just one option) at the boundaries. In the case of the moving double barrier option, there are two boundariers at each time. At time 0.125 from the expiration (1.0 - 0.125 = 0.875) the barriers are set at 115 and 145. We replicate at 115 by calculating the value of our portfolio at time 0.875 and spot 115 and buying x units of the 115 put option expiring at time 1.0 such that the total portfolio value is zero.

$$1.0 \times C(115, 120, 0.875, 1.0) + xP(115, 115, 0.875, 1.0) = 0, \tag{3.22}$$

where $P(S, K, t_p, t_e)$ denotes the value of a put option with strike K and expiration time t_e at spot price S and time t_p. Rearranging gives $x = -0.0212$. We also need to replicate the exotic at its other boundary. At spot 145 the exotic has zero value. We therefore find y such that:

$$1.0 \times C(145, 120, 0.875, 1.0) + yC(145, 145, 0.875, 1.0) = 0.$$

Rearranging gives $y = -10.187$. Now assume that we have successfully found the replicating portfolio from 1.0 until time 0.625. We find the options required to replicate the exotic at time 0.5. The portfolio currently assembled is given in table (3.9).

Time	Low Barrier	High Barrier
0.25	115	127
0.50	115	131
0.75	115	137
1.00	115	145

Table 3.7: Trigger levels for the moving double knockout option.

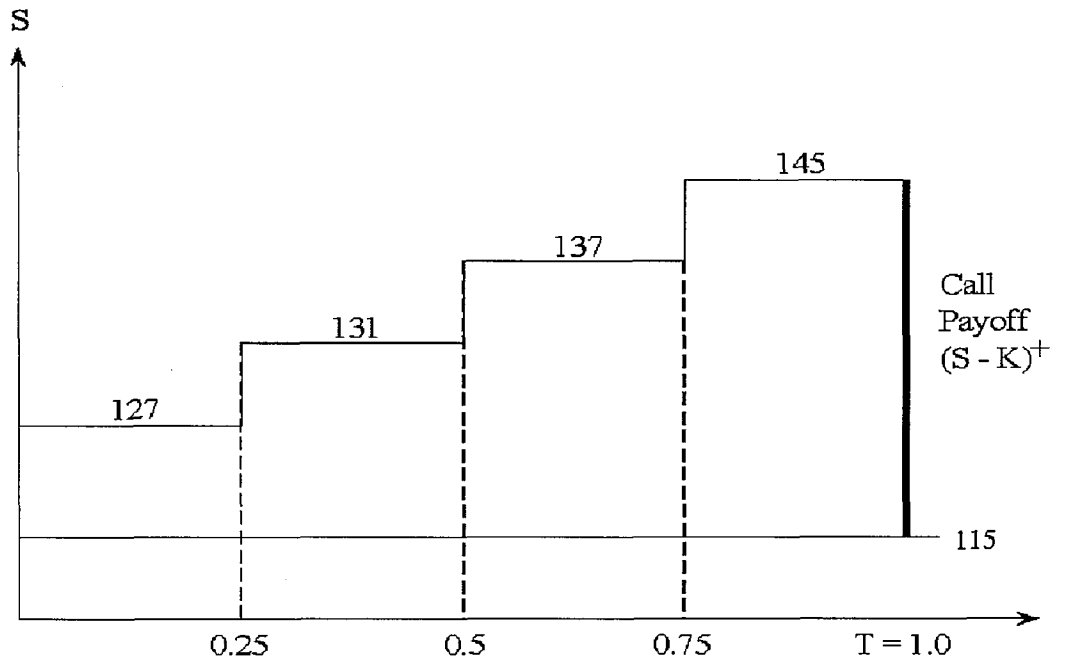


Figure 3.11: Boundaries of the moving barrier option.

At time 0.5 we require replication at the boundaries 115 and 131 where the exotic has zero value. We begin with lower boundary and find the value of the current portfolio at time 0.5 and spot 115. A quantity, x, of put options with expiration time 0.625 and strike 115 is then added such that the total value of the portfolio is zero.

$$\begin{aligned} 1.0 \times C(115,120,0.5,1.0) - 10.187C(115,145,0.5,1.0) - 0.212P(115,115,0.5,1.0) \\ +4.218C(115,145,0.5,0.875) - 0.209P(115,115,0.5,0.875) \\ -3.066C(115,137,0.5,0.75) - 0.100P(115,115,0.5,0.75) \\ +xP(115,115,0.5,0.625) = 0. \quad (3.23) \end{aligned}$$

Rearranging (3.23) gives $x = -0.026$. At the higher barrier, 131, we find y such that the

Expiration	Weight	Strike	Type	Weight	Strike	Type
1.000	1.000	120	Call	0.000	0	Put
1.000	-10.187	145	Call	-0.212	115	Put
0.875	4.218	145	Call	-0.209	115	Put
0.750	-3.066	137	Call	-0.100	115	Put
0.625	2.425	137	Call	-0.026	115	Put
0.500	-1.188	131	Call	0.044	115	Put
0.375	1.285	131	Call	0.108	115	Put
0.250	-0.215	127	Call	0.122	115	Put
0.125	0.611	127	Call	0.124	115	Put

Table 3.8: Replicating portfolio for the moving double knockout option.

Expiration	Weight	Strike	Type	Weight	Strike	Type
1.000	1.000	120	Call	0.000	0	Put
1.000	-10.187	145	Call	-0.212	115	Put
0.875	4.218	145	Call	-0.209	115	Put
0.750	-3.066	137	Call	-0.100	115	Put
0.625	2.425	137	Call	-0.026	115	Put

Table 3.9: Replicating portfolio for the moving double knockout option at time 0.625.

following is satisfied:

$$\begin{aligned} 1.0 \times C(131,120,0.5,1.0) - 10.187C(131,145,0.5,1.0) - 0.212P(131,115,0.5,1.0) \\ +4.218C(131,145,0.5,0.875) - 0.209P(131,115,0.5,0.875) \\ -3.066C(131,137,0.5,0.75) - 0.100P(131,115,0.5,0.75) \\ +yC(131,137,0.5,0.625) = 0. \quad (3.24) \end{aligned}$$

Rearranging (3.24) gives $y = 2.425$.

European Style Digital Options

We demonstrate the flexibility of this scheme by pricing a well known exotic, the Digital European Call Option. This contract requires the specification of a strike price, whether it

is a call or a put, an expiration time and, in the case of foreign exchange, a payoff currency. If the contract is specified to be a call, the owner receives one unit of the face currency at expiration if the spot price is above the strike price. This contract is sometimes known as a European Bet.

In the case of this digital option, we have no boundaries to replicate other than the expiration boundary. Let us consider the case of a T maturity call option paying the domestic currency if spot is above strike K at expiration. In general there exists a minimum tick size, that is a minimum amount by which the spot price of an asset may move. We take advantage of this in our replication procedure. We require a combination of options which pay one unit of the domestic currency (ccy) above K. A strategy resulting in a similar payoff is a call spread. A call spread is simply a combination of options where a lower strike call option is purchased and a higher strike call is sold in the same notional amount and with the same expiration date. The resulting payoff is limited to the range $[0, K_u - K_l]$, where K_u is the upper strike and K_l the lower strike. We therefore choose a call spread with $K_u = K + \delta$ and $K_l = K$ where δ is the minimum tick size. The individual legs of the call spread have notional amounts that result in a payoff of $1ccy$ if spot is above K at expiration.

We now give the example of a USD call, JPY put digital on the USD/JPY exchange rate with expiration of one year and strike 100.00. The contract pays 100m JPY. The minimum tick size for USD/JPY rates is 0.01 of a JPY.

The USD call spread we require is therefore of strikes 100.00 and 100.01 and expiration one year. The notional amount, N, of these options is given by:

$$N = \frac{100,000,000}{0.01} = 1 \times 10^{10}\text{USD}.$$

Evaluating the call spread above we find that the replicated value of the digital call is 47.6028%jpy or 47,602,800jpy. The analytic value of the digital is known to be 47.6045%jpy showing just how good our replication is. In fact, if we are able to buy and sell options with strikes within the minimum tick size, we may make the replication almost perfect. Using strikes of 100.00 and 100.00001 and a notional amount of the call spread of 1×10^{13}USD we find a value of 47.6045% jpy, exactly the same, to four decimal places, as the analytic valuation.

3.4.4 Pricing Barrier Options on an Asset with Time Dependent Volatility

Rather than pricing an exotic option, then finding a portfolio of standard options that replicates the payoffs on its boundaries perfectly we now construct the replicating portfolio from the known payoffs in order to price the target option. Ideally we would want to do this in a continuous setting picking only the vanilla option prices relevant to our replicating portfolio from the volatility surface. It is possible to do this with no loss of information if the volatility has only a time dependence.

We follow the notation of Derman, Ergener and Kani [22] with adaptations for non-constant volatility.

The spot price process is given as follows,

$$dS = S(r_d - r_f)dt + S\sigma(t)dZ_t,$$

where S is the stock price, r_d the domestic risk-free rate, r_f the foreign risk-free rate and $\sigma(t)$ the function for the volatility of the stock price.

We let $\Phi(S, T-t)$ denote the value of the target option at time t contingent on the stock price S. We assume that the spot prices $S_{\partial\Omega}$ on the boundary where the payoff is defined may be locally parameterised in terms of some sufficiently regular function $B(t)$. We wish to work in continuous time and obtain some representation of the target option's value $\Phi(S, T-t)$ at any instant in time and spot price S. We find this in terms of a weighted average of values of a set of vanilla (or some other standard) options.

$$\Phi(S, T-t) = \int_t^T \alpha(t,u) C\left(S, K(u), \sigma(u-t), u-t\right) du,$$

where $C(S, K, \sigma(u), u)$ is the value of a European style contingent claim at spot price S, strike K, volatility $\sigma(u)$ and time to expiration u.

We want a static replication strategy and so all weights $\alpha(t,\tau)$ must be independent of initial time t. i.e.

$$\frac{\partial}{\partial t}\alpha(t,\tau) = 0.$$

Thus we use simply the parameter $\alpha(T)$ to denote the weights of the standard options.

Although it is not strictly necessary as any out-of-the-money options may be used, for simplicity assume that all options in the replicating portfolio have strikes $K(\tau)$ which coincide with the stock value at the boundary $B(\tau)$.

$$K(\tau) = B(\tau) \quad \forall \tau > t.$$

This condition (and the weaker assumption mentioned) guarantees that no cashflows or payoffs are introduced within the boundary $\partial\Omega$.

Thus we have the conditional relationship,

$$\Phi(S, T-t) = \int_t^T \alpha(u) C(S, B(u), \sigma(u-t), u-t) du.$$

In many cases the payoff at expiration time T is discontinuous. In these cases the weight $\Phi(T)$ takes the form of a Dirac delta function and so we are able to separate the terminal weights and re-express our relation from above.

$$\Phi(S, T-t) = \int_t^T \alpha(u) C\left(S, B(u), \sigma(u-t), u-t\right) du + \alpha_T C_T(S, B(T), \sigma(T), T-t), \quad (3.25)$$

where $C_T(S, B(T), \sigma(T), T-t)$ represents the combination of options required to replicate the terminal payoff of the target option and $\alpha(T)$ represents their collective weights.

Let ξ represent the known payoff of the target option on the boundary point at time t.

$$\xi(t) = \Phi(B(t), T-t)$$

Evaluating equation (3.25) on the boundary gives the identity,

$$\xi(t) = \int_t^T \alpha(u) C\left(S, B(u), \sigma(u-t), u-t\right) du + \alpha_T C_T(B(t), B(T), \sigma(T), T-t). \quad (3.26)$$

Solving (3.26) recursively for the weights is done by finding the weights at expiration α_T. We match the terminal payoff of the target option with some standard options and roll backwards through time along the boundary.

To do this we must discretise the time interval (t, T). Do this in the following way with $t = t_0, t_1, t_2, ..., t_N = T$. The weights are denoted similarly $\alpha_0, ..., \alpha_N$.

Evaluating (3.26) gives at time t_{N-2},

$$\begin{aligned}\xi_{N-2} = \alpha_{N-1} C\left(B_{N-2}, B_{N-1}, \sigma_f(t_{N-2}, t_{N-1}), t_{N-1} - t_{N-2}\right) \\ + \alpha_N C_N\left(B_{N-2}, B_N, \sigma_f(t_{N-2}, t_N), t_N - t_{N-2}\right),\end{aligned}$$

where $\sigma_f(i,j)$ denotes the forward volatility from time i until j, $i < j$ given by

$$\sigma_f(i,j) = \sqrt{\frac{\sigma^2(j)j - \sigma^2(i)i}{j-i}}.$$

We solve to find,

$$\alpha_{N-1} = \frac{\xi_{N-2} - \alpha_N C_N\left(B_{N-2}, B_N, \sigma_f(t_{N-2}, t_N), t_N - t_{N-2}\right)}{C\left(B_{N-2}, B_{N-1}, \sigma_f(t_{N-2}, t_{N-1}), t_{N-1} - t_{N-2}\right)},$$

A general step i is given by,

$$\alpha_i = \frac{\xi_{i-1} - \alpha_{i+1} C_{i+1}\left(B_i, B_{i+1}, \sigma_f(t_i, t_{i+1}), t_{i+1} - t_i\right)}{C\left(B_{i-1}, B_i, \sigma_f(t_{i-1}, t_i), t_1 - t_{i-1}\right)},$$

and the solution is complete.

Partial Barrier Options with Time Dependent Volatility

An increasingly popular product is a variant upon the standard knockout option. The partial barrier option is similar to the knockout option in that a level is specified which, if touched, causes the termination of the option contract. The difference is that the partial barrier option may only knockout during a specified period of the option's life. We will consider an example of such a contract and show why the use of a volatility curve in its pricing is crucial for accurate results.

Consider the following derivative contract. Call option, strike 100, expiration 1.0, barrier 110 active from time zero until time 0.25. Using the following parameters we price the contract using flat volatility and a volatility curve. We use time steps of size 0.125, spot 100, zero interest rates, volatility to 1.0 of 12% and volatility to 0.25 of 16%. The values we find under the two different volatility assumptions are quite different. Using eight time steps and a constant volatility of 12% gives a price of 3.918, while using a volatility curve gives 2.928. Tables (3.10) and (3.11) give the replicating portfolios.

Expiration	Weight	Strike	Type
1.000	1.000	100	Call
0.250	-5.909	110	Call
0.125	4.948	110	Call

Table 3.10: Replicating portfolio for the partial barrier option with a volatility curve.

Expiration	Weight	Strike	Type
1.000	1.000	100	Call
0.250	-6.067	110	Call
0.125	2.388	110	Call

Table 3.11: Replicating portfolio for the partial barrier option with a flat volatility curve.

Figure (3.12) demonstates just how sensitive this option is to the value of volatility at specific times. We plot the price of the option, using one hundred time steps, for different levels of volatility to the end of the barrier period while keeping the volatility to the expiration date constant at 12%. By purchasing such a partial barrier option, a market participant is getting exposure to the shape of the volatility curve. We can see that he is getting "short"" front end volatility and "long" backend volatility. We could view this as buying volatility forward (see Appendix (B)). Clearly, the owner is speculating that volatility will be small in the near term and increase in the long term. That is, the barrier will not be touched whilst alive but the option will expire in-the-money.

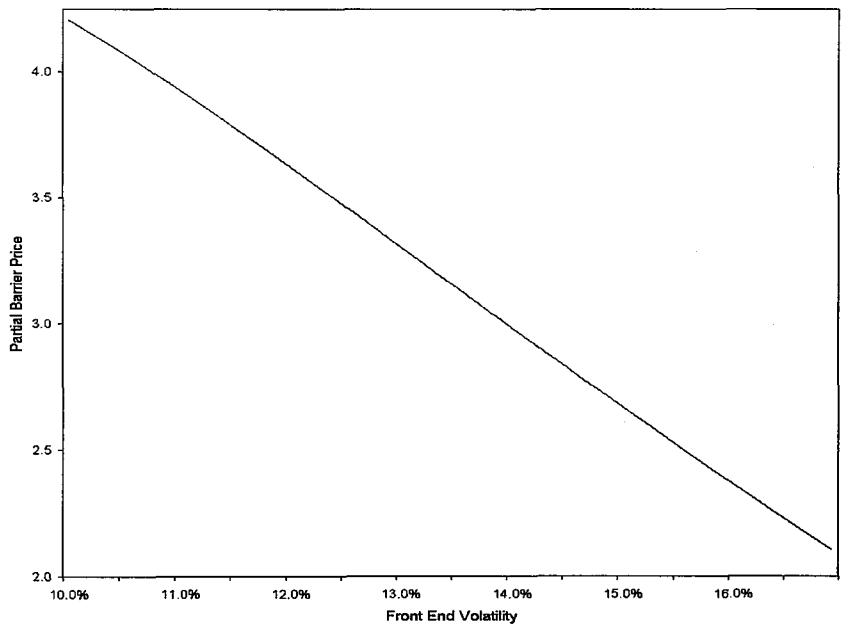


Figure 3.12: Partial barrier option price against front end volatility using one hundred time steps.

A Full Volatility surface

We have shown that the framework for static replication may be generalised for pricing under a time dependent volatility. So far, we have concentrated upon a continuous time setting in which the replication takes place. It should be noted, however, that for a large number of time steps the replication provides good results when used in conjunction with a binomial or trinomial tree model. If we use the implied tree model described in section (3.2) then it is possible to replicate derivative contracts consistently with the volatility surface at the time of pricing.

We refer back to our example of an up and out call option. The initial replication portfolio would be given in the same way, purchasing one unit of the call option. Then rolling back one time step as in equation (3.19), the value of the call options of strikes 120 and 100 would be found using the implied tree we have grown from the volatility surface given by market prices. The ratio of these new options would then be found as in (3.20). Rolling back recursively through time we carry on in a similar fashion, finding option prices from our implied tree and using those prices to find the number of options to write or purchase. At the time of writing, t, we then have a portfolio of options that replicates and prices our target security consistently with the implied volatility surface. We give a simple example below for a contract known as a "Capped Power Option".

Consider the contract that pays the square of the payoff of the corresponding plain vanilla option up to a certain maximum. The contract we will price is as follows. The expiration is one year, interest rates are zero, spot price 110, strike price 114 and cap 119. This contract pays $\max(\max((S(T)-K)^2, 25), 0)$ at expiration. Replication of this contract only requires options of maturity 1. However we define the volatility smile shown in table (3.12).

The portfolio of options in table (3.13) replicates the payoff of the capped power option. We use options with strikes 0.5 apart and price the contract using a volatility smile and flat volatility.

From table (3.13) we see how the value of the exotic is affected by the presence of the volatility smile, and compare the replicating portfolio using flat, at-the-money volatility. In this case the derivative is far cheaper under a volatility smile. This is due to the large number of out-of-the-money options that must be sold. Use of a constant volatility model

Strike	Volatility
114	12.0
115	12.5
116	13.0
117	13.5
118	14.0
119	14.5

Table 3.12: Volatility smile for the capped power option.

would have resulted in mispricing.

3.5 Conclusions

We have shown how the work by Derman, Ergener and Kani [22] may be used to price a large array of new products. The framework provides a flexible way of hedging standard products and we have, through practical examples, shown how this may be used to price particularly exotic derivative contracts. The existing methodology was generalised in this Chapter for use with a time dependent volatility function so that we may price contracts consistently with the current volatility curve. It was also shown that, under certain conditions, it is possible to take account of the change of volatility with strike price. Numerical examples were given throughout.

			Smile			Flat	
Strike	**Volatility**	**Weight**	**Value**	**Total**	**Volatility**	**Value**	**Total**
114.00	12.00%	0.5	3.594	1.797	12.00%	3.594	1.797
114.50	12.25%	1.0	3.523	3.523	12.00%	3.417	3.417
115.00	12.50%	1.0	3.457	3.457	12.00%	3.247	3.247
115.50	12.75%	1.0	3.395	3.395	12.00%	3.084	3.084
116.00	13.00%	1.0	3.338	3.338	12.00%	2.927	2.927
116.50	13.25%	1.0	3.284	3.284	12.00%	2.776	2.776
117.00	13.50%	1.0	3.234	3.234	12.00%	2.631	2.631
117.50	13.75%	1.0	3.187	3.187	12.00%	2.493	2.493
118.00	14.00%	1.0	3.144	3.144	12.00%	2.360	2.360
118.50	14.25%	1.0	3.103	3.103	12.00%	2.233	2.233
119.00	14.50%	-9.5	3.066	-29.123	12.00%	2.112	-20.065
			Price	2.339		**Price**	6.900

Table 3.13: Replication portfolio for a capped power option.

Chapter 4

Volatility Trading

4.1 Introduction

Under the Black-Scholes framework we may evaluate the fair price of plain vanilla options. In order to price these contracts six inputs are required - spot price, time to maturity, strike price, interest rate for the foreign currency, interest rate for the domestic currency and volatility. All of the above inputs are either specific to the contract in question or may be observed directly, except for one: the volatility. This volatility is the annualised standard deviation of returns on the asset over the life of the derivative contract. This forward looking volatility is not yet known and is effectively set by the options market. This volatility assigned to the asset by the market is known as implied volatility and is subject to change in the same way that the spot price of the underlying asset may change. Indeed, options traders are often described as being volatility traders or market makers in volatility.

Due to the variable nature of implied volatility and the existence of a volatility market a position in options could, or should, be viewed as a position in volatility. A perfectly delta hedged options portfolio will still show a profit or loss on a mark-to-market basis due the sensitivity of its value to changes in implied volatility. This sensitivity is known as vega. Given this movement it is natural that there is interest in products that allow market participants to speculate on levels of volatility in a more direct fashion.

In the following chapter several different types of volatility contracts are investigated. We

will begin by examining the simplest form of volatility speculation, namely positions in delta-hedged portfolios of options. Secondly, a type of contract known as a variance swap is investigated and the limited amount of work published on this subject is reviewed and implemented. Thirdly, a form of contract known as a volatility swap is described, and an intuitive approach is taken to discovering the nature of this derivative security. We then demonstrate a more accurate methodology for pricing and hedging these contracts which corresponds very closely to the intuitive approximation under perfect market conditions. The difficulties involved in the pricing, and therefore hedging, of this contract under typical market conditions is illustrated. Numerical examples are given throughout.

4.2 Volatility Trading through Delta-Hedged Vanilla Options

The value of plain vanilla options depends on several factors, one of which being implied volatility. Therefore, if we assume that all other exposures have been hedged away through construction of the portfolio of the option, zero coupon bonds and the underlying asset delta neutral, we are left with a portfolio whose value is sensitive only to volatility. Therefore, in the simplest sense, a delta-neutral options portfolio may be viewed as a derivative written on volatility.

4.2.1 Profit and Loss due to Trading Delta-Hedged Vanilla Options

Analysis of the Black-Scholes equation for valuing European style put and call options shows that the vega of such contracts is always greater than or equal to zero. In other words, an increase in implied volatility leads to an increase in the value of the derivative. The vega of a plain vanilla call or put option on a foreign exchange rate valued at time t and expiring at time T is given by:

$$Vega_{BS} = \frac{\partial C}{\partial \sigma} = S\sqrt{T-t}N'(d_1),$$

where,

$$d_1 = \frac{\ln(S/K) + (r_d - r_f + \frac{1}{2}\sigma^2)(T-t)}{\sigma(T-t)},$$

and,

S = spot price at time t,

K = strike price of option,

σ = annualised volatility of underlying asset to time T,

r_d = domestic continuously compounded risk – free interest rate,

r_f = foreign continuously compounded risk – free interest rate,

and,

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

is the normal probability density function.

Therefore, if implied volatility increases for the remainder of the option's life, then the value of the option will increase. Correspondingly, a decrease in implied volatility leads to a decrease in the value of the derivative. Referring back to our derivative security consisting of a delta hedged vanilla option, what contributes to the profit and loss of this strategy over the lifetime of the contract? The sensitivity of the contract's value to implied volatility has already been derived, but what of the delta hedging that has gone on before? We are left to answer the question - what is the profit or loss arising from a short position in a continuously delta hedged vanilla option held until expiration?

Theorem 1

A European style contingent claim sold at price $C(S_t, t, \sigma_I)$ at time t which expires at T is delta hedged continuously using σ_I to calculate the delta ratio. If $\sigma_I > \sigma_r(\tau) \forall t \leq \tau \leq T$, such a portfolio will result in a profit for the holder.

Moreover, the profit or loss realised by the holder of this portfolio is given by,

$$\frac{1}{2} \int_t^T \left({\sigma_I}^2 - \sigma_r(\tau)^2 \right) {S_\tau}^2 \frac{\partial^2 C}{\partial S^2} d\tau. \tag{4.1}$$

where σ_I is the volatility at which the hedge is calculated and executed.

Proof

Let us assume, without loss of generality, that all interest rates are zero. In other words, the spot price of the underlying asset is equal to the forward price.

The investor's portfolio consists, at any time t, of the following,

Short one unit of the derivative security, long $\Delta_t = \frac{\partial C}{\partial S_t}$ of the underlying asset.

The initial value of the contingent claim as well as the delta hedge ratio are, under Black-Scholes conditions, given by:

$$\frac{\partial C}{\partial t} + \frac{1}{2}{\sigma_I}^2 {S_t}^2 \frac{\partial^2 C}{\partial S^2} = 0. \tag{4.2}$$

The value of the hedging portfolio, Π_t, is given by,

$$\Pi_t = \Delta_t S_t.$$

Then the change in value of this portfolio is,

$$d\Pi_t = \Delta_t dS_t.$$

Integrating this expression over time we obtain,

$$\Pi_T = \Pi_t + \int_t^T \Delta_\tau dS_\tau,$$

where Π_t is the initial value of the derivative.

Now assume that the underlying asset follows a lognormal process given by,

$$\frac{dS_t}{S_t} = \sigma_r(t) dW_t,$$

where $\sigma_r(t)$ is the volatility realised by the underlying asset over time step dt.

Using Itô's lemma on $C(S_t, t, \sigma_I)$, we find,

$$\begin{aligned} dC(S_t, t, \sigma_I) &= \frac{\partial C}{\partial t} dt + \frac{\partial C}{\partial S} dS_t + \frac{1}{2} \frac{\partial^2 C}{\partial S^2} {dS_t}^2, \\ &= \frac{\partial C}{\partial t} dt + \frac{\partial C}{\partial S} dS_t + \frac{1}{2} {\sigma_r}^2(t) {S_t}^2 \frac{\partial^2 C}{\partial S^2} dt. \end{aligned}$$

Integrating the above expression and making use of equation (4.2), we find,

$$\begin{aligned} C(S_T, T, \sigma_I) &= C(S_t, t, \sigma_I) + \int_t^T \frac{\partial C}{\partial S} dS_\tau + \int_t^T \left(\frac{\partial C}{\partial \tau} + \frac{1}{2} {\sigma_r}^2(\tau) {S_\tau}^2 \frac{\partial^2 C}{\partial S^2} \right) d\tau, \\ &= C(S_t, t, \sigma_I) + \int_t^T \frac{\partial C}{\partial S} dS_\tau + \frac{1}{2} \int_t^T \left({\sigma_I}^2 - {\sigma_r}^2(\tau) \right) {S_\tau}^2 \frac{\partial^2 C}{\partial S^2} d\tau \end{aligned} \tag{4.3}$$

Rearranging we find,

$$C(S_T, T, \sigma_I) - \left(C(S_t, t, \sigma_I) + \int_t^T \frac{\partial C}{\partial S} dS_\tau \right) = \frac{1}{2} \int_t^T \left({\sigma_I}^2 - {\sigma_r}^2(\tau) \right) {S_\tau}^2 \frac{\partial^2 C}{\partial S^2} d\tau. \tag{4.4}$$

The left-hand side of (4.4) clearly represents the difference between the initial premium received by the writer of the European style contingent claim and the value of the security at time T and the loss arising from following a dynamic delta-hedging strategy from the time of writing, t, to expiration time, T.

The right-hand side of (4.4) is then equal to the profit or loss from the portfolio consisting of a delta hedged short option position. Recognise that the term $\frac{\partial^2 C}{\partial S^2}$ is the gamma of the contingent claim at time τ, $t \leq \tau \leq T$. It is a well known property of European style contingent claims with convex payoffs that this gamma is nonnegative, see Hull [28].

Thus we deduce that if $\sigma_I > \sigma_\tau(\tau)$ $(\forall \tau : t \leq \tau \leq T)$ the holder of the portfolio consisting of a continuously delta hedged short position in a European style contingent claim will make a profit. Moreover, the profit or loss on this position is given by

$$PnL = \frac{1}{2} \int_t^T \left({\sigma_I}^2 - {\sigma_\tau}^2(\tau) \right) {S_\tau}^2 \frac{\partial^2 C}{\partial S^2} d\tau. \tag{4.5}$$

End of proof.

As Theorem (1) shows, a market participant may trade European style options and, by neutralising any sensitivity to the underlying asset, effectively take positions in the volatility realised by the underlying over the life of the derivative security. The profit and loss on this strategy is also known and given by (4.5).

4.2.2 Vega, Gamma and Theta Associated with Trading Plain Vanilla Options

As we have seen, a speculator wishing to take a position in the volatility of an asset may do so through plain vanilla options. Theorem (1) shows the profit and loss experienced by an investor holding, and delta hedging, European style contingent claims. We noted that this PnL was dependent upon the path taken by the underlying and, more precisely, the gamma associated with the contingent claim throughout its life. It is interesting that through seeking a position in volatility we have found ourselves with a fundamental dependence upon the gamma. In fact, as we will see below, there is a direct link between the gamma and vega of a plain vanilla option. In other words, $\frac{\partial^2 C}{\partial S^2}$ is related to $\frac{\partial C}{\partial \sigma}$. The relationship

for plain vanilla call and put options is as follows:

$$\frac{\partial^2 C}{\partial S^2} = \frac{\partial C}{\partial \sigma} \frac{1}{(T-t)S^2\sigma}. \tag{4.6}$$

Where σ is the volatility of the spot price S over the life of the option. Ignoring interest rates for convenience and taking the specific example of a call option we find this in the following way.

The Black-Scholes formula gives us,

$$C = SN(d_1) - KN(d_2),$$

where S is the spot price at the time of pricing, K is the strike price of the option and $N()$ denote the standard cumulative normal distribution. In addition:

$$d_1 = \frac{\ln \frac{S}{K} + \frac{1}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}}. \tag{4.7}$$

$$d_2 = \frac{\ln \frac{S}{K} - \frac{1}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}}. \tag{4.8}$$

$$\begin{aligned} \frac{\partial C}{\partial \sigma} &= SN'(d_1)\left(-\frac{\ln\left(\frac{S}{K}\right)}{\sigma^2\sqrt{(T-t)}} + \frac{1}{2}\sqrt{(T-t)}\right) \\ &- KN'(d_2)\left(-\frac{\ln\left(\frac{S}{K}\right)}{\sigma^2\sqrt{(T-t)}} - \frac{1}{2}\sqrt{(T-t)}\right) \\ &= SN'(d_1)\left(\frac{-d_1}{\sigma}\right) - KN'(d_2)\left(\frac{-d_1}{\sigma}\right) \\ &= -\frac{1}{\sigma}\left[SN'(d_1)d_2 - KN'(d_2)d_1\right]. \end{aligned} \tag{4.9}$$

However, we notice that:

$$N'(d_1) = \frac{K}{S}N'(d_2).$$

Substituting into the equation above we find:

$$\begin{aligned} \frac{\partial C}{\partial \sigma} &= -\frac{1}{\sigma}SN'(d_1)(d_2 - d_1), \\ &= \sqrt{T-t}SN'(d_1). \end{aligned} \tag{4.10}$$

Now we find the gamma:

$$\begin{aligned} \frac{\partial^2 C}{\partial S^2} &= \frac{\partial}{\partial S}\left(N(d_1)\right) \\ &= N'(d_1)\left(\frac{1}{S\sigma\sqrt{T-t}}\right). \end{aligned} \tag{4.11}$$

Then by inspection we see that:

$$\frac{\partial^2 C}{\partial S^2} = \frac{\partial C}{\partial \sigma} \frac{1}{(T-t)S^2\sigma}. \tag{4.12}$$

Using put-call parity, $C - P = F$, where F is the price of a forward contract with strike price identical to that of the call and put option, we deduce that the same relationships hold for a put option.

As we see from the expressions and graphically from figures (4.1) and (4.2), the vega of an option is dependent upon, amongst other things, the proximity of the spot price to the strike price and the time to expiration. By simply trading a plain vanilla option in order to capture its sensitivity to implied volatility, the vega, a market participant may find that all of his desired exposure is lost due to a large move in the underlying asset. In addition to

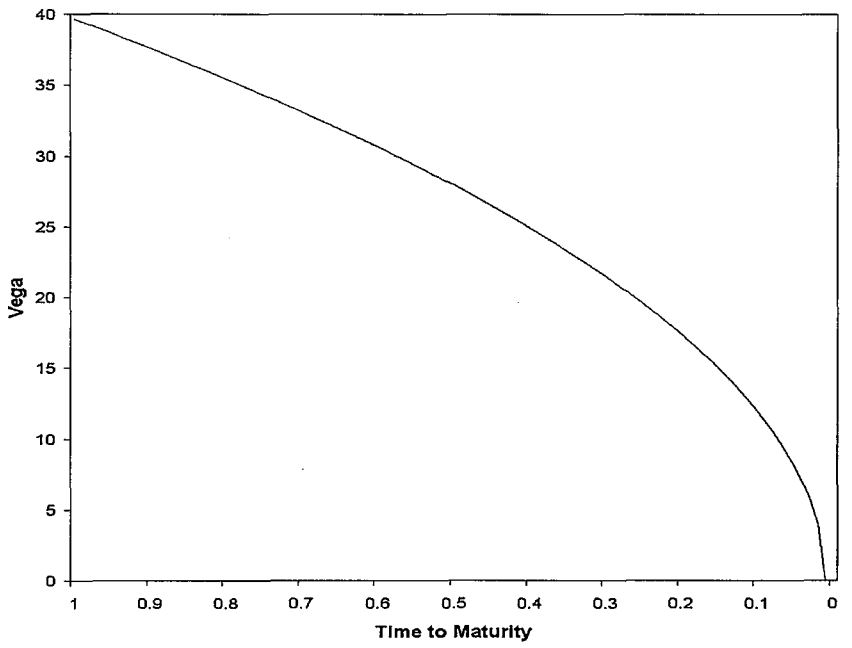


Figure 4.1: Vega against time for a plain vanilla option, strike price 100.

the problems that may be experienced due to the loss of vega, the analysis above exposes another pitfall of trading in this way. The gamma of the at-the-money option, given all else remains the same, increases as time to expiration decreases. See figure (4.3) for an illustration. Indeed if spot is at the strike at expiration time the gamma is infinite. This leads to a large exposure to the spot price path that a volatility trader is unlikely to want to have. These are the kinds of problems that have led to the development of "pure" volatility

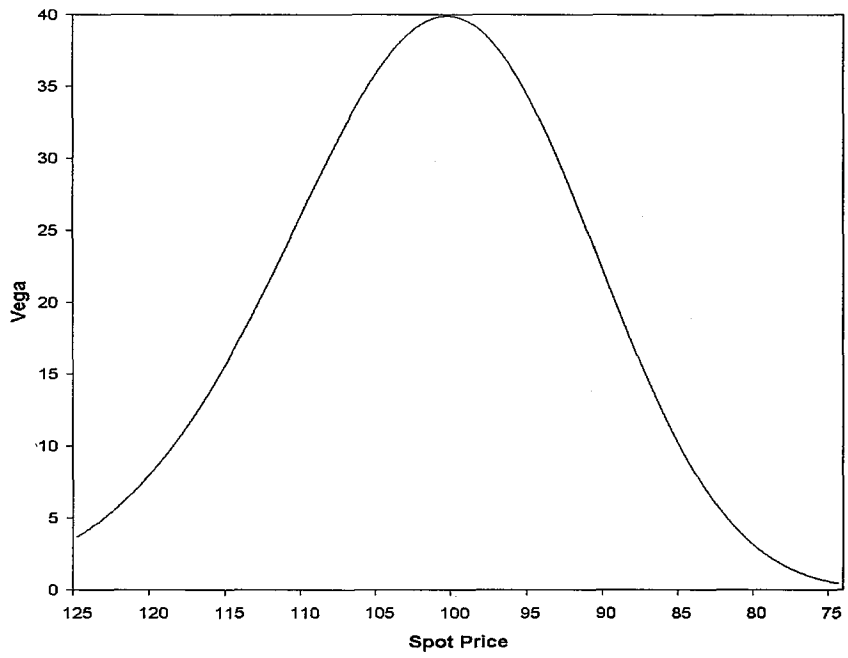


Figure 4.2: Vega against spot for a plain vanilla option, strike price 100.

products.

4.3 Variance Swaps

A desire to trade volatility, first as an aid to hedging vanilla options portfolios and subsequently as an asset in itself, has led to the development of various volatility related contracts. The simplest of these contracts, from a theoretical point of view, is the Log Contract first suggested by Neuberger [40]. This contract was approached with the specific aim of finding an aid for plain vanilla option traders wishing to hedge away the volatility risk arising from incorrectly estimating the volatility of an asset over its lifetime.

A different point of view is taken by Carr and Madan [12] who developed the idea of a contract paying the holder an amount proportional to the variance of an underlying asset measured over some fixed time. This is the contract that has become known as a Variance Swap. Later on, Demeterfi, Derman, Kamal and Zou [18] investigated the Variance Swap contract further, rederiving the results of Carr and Madan [12] and suggesting a more stable method for replicating and hedging these contracts. Work was also carried out by Demeterfi, Derman, Kani and Zou [18] on the effects of a market for the underlying and

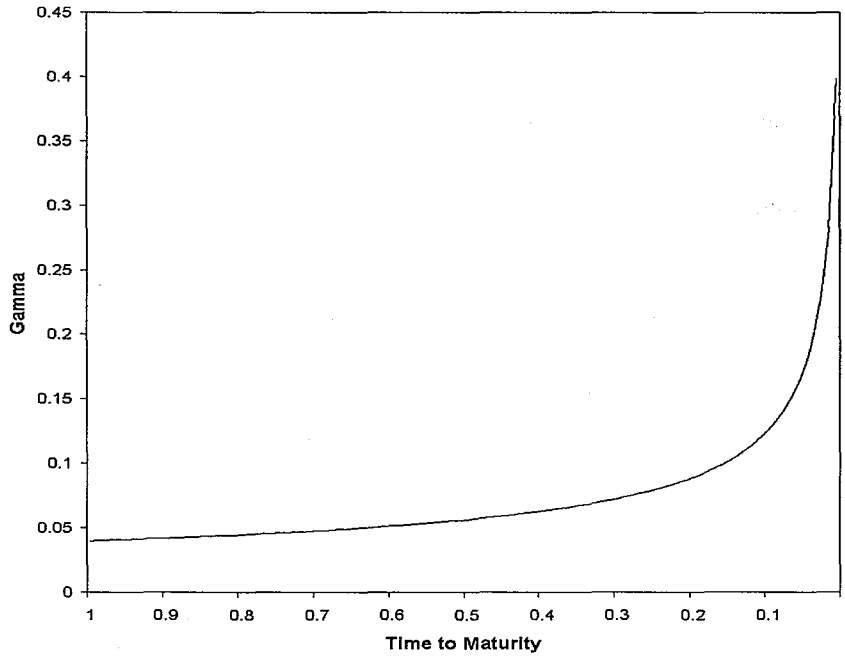


Figure 4.3: Gamma against time for a plain vanilla option.

options on the underlying which does not obey standard Black-Scholes assumptions.

In this section we will review the existing work that concerns variance swaps. Two methods for hedging these contracts, one suggested by Carr and Madan [12] and the other by Demerterfi, Derman, Kani and Zou [18], are implemented and compared. These implementations also include the effects of a volatility skew for plain vanilla options on the fair variance of the swap contract.

4.3.1 The Log Contract

Among the first attempts made to construct a contract that paid-off some kind of volatility related quantity was that by Neuberger [40]. It had been noticed that the classic replication strategy of Black-Scholes relies upon the hedger of an option correctly estimating the future average volatility, or variance, of the underlying asset when initially pricing the contract. It also requires that this volatility remains constant throughout the life of the option and that this volatility is the one used in the Black-Scholes calculation of the delta ratio when rehedging is carried out.

In practice implied volatilities change the whole time. As mentioned before, there is a

market in implied volatility. The assumption that the correct volatility will be used to price the option initially and that this volatility will not change is thus a rather restrictive and unrealistic one. Neuberger was first to seek a contract to hedge this exposure.

Neuberger [40] considered the contract which has a terminal payoff of $L_T = \ln(S_T)$, where T is the time of expiration and S_t denotes the spot price at time $t, 0 \leq t \leq T$. The payoff is simply the natural logarithm of the terminal spot price.

It may be shown, using standard Black-Scholes assumptions, that the fair price of the Log Contract, L_t is given by,

$$L_t = \ln S_t - \frac{1}{2}\sigma_R^2(T-t), \tag{4.13}$$

for any $t \ : \ 0 \leq t \leq T$ and where

$$\sigma_R^2(T-t) = \int_t^T \sigma_\tau^2(\tau)d\tau.$$

with σ_τ the realised volatility defined in theorem (1). See Neuberger [40].

Note that this implies that the delta of the log contract at any time, t, throughout its life is given by $\frac{1}{S_t}$. If we use the example of a Log Contract on cable, the exchange rate for GBP/USD quoted in terms of USD per one GBP, we see that the holder of one unit of the Log Contract should sell one USD worth of GBP against the Dollar in order to be delta neutral. Suppose that, as the owner of one unit of the Log Contract, you are initially delta hedged and therefore long 1 USD worth of GBP against USD. If GBP strengthens versus the USD then the price of a Log Contract will also rise. The value of the delta hedge position in the underlying will therefore be worth less than 1 USD meaning that, to rebalance the hedging portfolio, more of the underlying must be purchased.

An important feature of the Log Contract is that its delta is, as shown above, independent of the volatility, σ, that may be estimated for use in pricing options. This independence from σ is what makes this contract a suitable hedging instrument for plain vanilla options. Recall theorem (1), where the profit or loss from continuously delta hedging a European style contingent claim was derived. The PnL from that particular strategy was found to be:

$$\frac{1}{2}\int_t^T \left({\sigma_I}^2 - {\sigma_\tau}^2(\tau)\right) {S_\tau}^2 \frac{\partial^2 C}{\partial S^2} d\tau.$$

If we substitute C, our general European style option, with the Log Contract L, and calculate the Log Contract's gamma

$$\frac{\partial^2 L}{\partial S^2} = -\frac{1}{S_t^2},$$

we find that the continuously delta hedged Log Contract has a PnL equal to the following:

$$PnL = \frac{1}{2}\left({\sigma_R}^2(T-t) - {\sigma_I}^2(T-t)\right), \tag{4.14}$$

where σ_I is the volatility implied by the price of the Log Contract and σ_R is the total volatility realised by the underlying asset over the entire life of the Log Contract.

Neuberger [40] provides numerical results and shows the efficacy of this contract in hedging the volatility risk inherent in positions in plain vanilla options.

The delta of the Log Contract does not depend on the initial forecast of volatility. Thus the hedger does not need to forecast the volatility accurately in order to delta hedge the contract with zero error. The outcome of the Log Contract depends only on the volatility realised by the underlying asset and not on any estimate of future volatility. It is also important to note that no assumptions have been made about the process that the volatility of the returns on the underlying asset follow. It is this, as will be shown later, that enables us to consider pricing variance swaps under more realistic conditions than standard option pricing theory would otherwise allow.

4.3.2 Developing an Instrument for Speculating on Future Realised Variance

The Log Contract was born of a desire to find a way for derivatives traders and risk managers to hedge away the error inherent in their own estimates of future volatility. Neuberger [40] demonstrated how this might be done. It has transpired that the greatest interest in these types of realised volatility contracts has come from those wishing to take risk rather than hedge it away. With the development of a liquid options market and, as its proxy an implied volatility market, speculators using derivatives have sought a way to take positions purely in volatility.

As discussed in the first section of this chapter, the pricing of a European style contingent claim (or a European style plain vanilla call option more specifically) relies on several

parameters that are directly observable and one that is not - the volatility over the life of the contract. More accurately, the value of an option depends on the *variance* of the underlying asset over the lifetime of the contract. If, as previously, we denote the annualised implied volatility of the asset to be σ_I and the time to expiration T, then the variance of the underlying asset over the lifetime of the contract is given by $\sigma_I^2 T$.

A paper by Carr and Madan [12] first described several different types of instruments for trading "volatility". The main contribution of this work was to describe a strategy for hedging a contract which pays the owner the variance of an asset over a predetermined time period. This contract was termed a "Variance Swap". The payoff from this swap is the following:

$$\text{Payoff} = \sigma_R^2 T,$$

where $\sigma_R^2 T$ is the total realised variance of asset S over the time interval $[0, T]$.

In practice the payoff of the Variance Swap will often be changed with the inclusion of a constant strike variance $\sigma_K^2 T$ so that the contract is initially zero cost. The contract then resembles a forward contract on the realised variance struck at $\sigma_K^2 T$. In this case "swap" is somewhat of a misnomer - realised variance forward contract is a more accurate description.

The Variance Swap, Carr's Approach to Pricing and its Relation to the Log Contract

Recall that the profit and loss experienced by an investor employing a strategy where a plain vanilla style option is sold and then delta hedged at a volatility σ_I that is different to the realised volatility $\sigma_r(\tau)$ is given by:

$$\frac{1}{2}\int_0^T \left({\sigma_I}^2 - {\sigma_r}^2(\tau)\right) {S_\tau}^2 \frac{\partial^2 C}{\partial S^2} d\tau.$$

Considering a contract C with a payoff of $2\ln S_T$ we notice that the gamma of such a contract is given by $\frac{\partial^2 C}{\partial S^2} = -\frac{2}{S_t^2}$ at any time t, $0 \leq t \leq T$. Substituting this into the equation above we get the PnL for this type of contract which is being continuously hedged at volatility σ_I whilst the underlying asset S_t exhibits volatility $\sigma_r(\tau)$ as:

$$\int_0^T \left({\sigma_r}^2(\tau) - {\sigma_I}^2\right) d\tau. \tag{4.15}$$

Notice that this is equivalent to a contract which pays the difference, over a time interval $[0, T]$, between the realised variance and the implied variance set at time zero.

This is the observation we shall use as a basis for constructing a Variance Swap contract. However we rely upon some other results to follow Carr's description of the hedging strategy for this new product.

Static Replication of a European Style Payoff

We consider the same general setting as that described in earlier sections for our Variance Swap. Explicitly, we consider the interval $[0, T]$ and assume that there exists a market for a risky stock S. Further, we make the assumption that there exists a liquid market for plain vanilla options on the asset S and that this market allows the sale or purchase of options with any strike price. However we do impose the restriction that market participants are unable to trade at any time other than 0 and T.

A result of Breeden and Litzenberger [10] and Green and Jarrow [27] shows that any European style contingent claim with payoff $f(S_T)$, where $f(.)$ is a smooth function of spot prices at the expiration time T, may be replicated using only positions in options, written or purchased, at time 0. This position is static.

If we denote the Dirac delta function by $\delta(.)$ we may then write, by definition, any function $f(S)$ as:

$$f(S) = \int_0^\infty f(K)\delta(S-K)dK \tag{4.16}$$

$$= \int_0^\kappa f(K)\delta(S-K)dK + \int_\kappa^\infty f(K)\delta(S-K)dK, \tag{4.17}$$

for any non-negative κ. We will refer to this κ more informally as our reference price. We now integrate the right-hand side of (4.17) using integration by parts to find,

$$f(S) = f(\kappa)1_{(S<K)} - f'(K)(K-S)^+|_0^\kappa + \int_0^\kappa f''(K)(K-S)^+dK \tag{4.18}$$

$$+ \; f(\kappa)1_{(S\geq K)} - f'(K)(K-S)^+|_\kappa^\infty + \int_\kappa^\infty f''(K)(S-K)^+dK.$$

$$= f(\kappa) + f'(\kappa)[(S-\kappa)^+ - (\kappa-S)^+] \tag{4.19}$$

$$+ \int_0^{\kappa} f''(K)(K-S)^+ dK + \int_{\kappa}^{\infty} f''(K)(F-S)^+ dK.$$

Note that the last line in (4.19) is a combination of weighted call and put options - contingent claims with payoffs at time T of $(S-K)^+$ for a call option and $(K-S)^+$ for the put. Any twice differentiable payoff, f, maybe written in the above way. We interpret these terms in the following way, assuming zero interest rates for simplicity. The first term is simply a cash flow in the amount $f(\kappa)$ whilst the second and third terms are simply weighted amounts of call and put options. Rewriting equation (4.19), we obtain the initial cost of the payoff $f(S_T)$:

$$\begin{aligned} V_0^f = f(\kappa) &+ f'(\kappa)[C_0(\kappa) - P_0(\kappa)] \\ &+ \int_0^{\kappa} f''(K)P_0(K)dK + \int_{\kappa}^{\infty} f''(K)C_0(K)dK. \end{aligned} \tag{4.20}$$

In particular, note that this replication with option prices (bond prices also in the case where interest rates are considered) makes no assumptions about the process followed by the underlying asset or its volatility. This lack of assumptions is important as it enables us, later on in this Chapter, to price these contracts under very general, realistic settings.

Application of Static Replication and Delta Hedging to the Variance Swap Problem

Firstly, let us refer once again to the payoff from a continuously delta hedged European style contingent claim. Suppose that we set the implied volatility used to calculate the delta hedge, σ_I, to zero. We then have from equation (4.4):

$$\int_0^T \frac{S_t^2}{2} f''(S_t)\sigma_t^2 dt = f(S_T) - f(S_0) - \int_0^T f'(S_t)dS_t. \tag{4.21}$$

Looking at the above equation we see that the left hand side is dependent only upon the path of the price S_t, and that the dependence on f is only due to the second derivative f''. In accordance with Carr and Madan [12], we will therefore only concern ourselves with payoffs f whose value and delta at a point κ are zero.

Carr and Madan observe that the right hand side of (4.21) depends on, and only on, the sum of the following two payoffs.

(i) A static position in options maturing at time T, paying $f(S_T)$ at time T.

(ii) The payoff from a hedging strategy where a position of $f'(S_t)$ underlying asset is maintained throughout the time interval $[0, T]$.

Note that this may easily be generalised to include interest rates and also a forward starting time interval $[T_{start}, T]$.

In particular let us look once again at expression (4.15), the payoff from a delta hedged Log Contract. This contract, as we have observed before, has a gamma equal to $-\frac{1}{S_t^2}$, written in our current notation:

$$\frac{\partial^2 f}{\partial S^2} = -\frac{1}{S_t^2},$$

where $f(S_T) = \ln(S_T)$.

Carr and Madan [12] consider the following function:

$$\Psi(S) = 2\left[\ln\left(\frac{\kappa}{S}\right) + \frac{S-\kappa}{\kappa}\right] \tag{4.22}$$

The function $\Psi(S)$ is a combination of Log Contracts and cash settled forward contracts struck at our reference price κ. Differentiating once with respect to S and a second time we find the delta and gamma of this function:

$$\Psi'(S) = 2\left[\frac{1}{\kappa} - \frac{1}{S}\right] \tag{4.23}$$

$$\Psi''(S) = \frac{2}{S^2} \tag{4.24}$$

The gamma is the same as that for the Log Contract shorted twice, meaning that we will obtain the desired payout of the realised variance $\sigma_R^2 T$ if we follow the replication strategy described above. Substituting for delta and gamma into equation (4.20) and remembering that $\Psi(S)$ is equal to zero when $S = \kappa$, yields:

$$V_0^f = \int_0^\kappa \frac{2}{K^2} P_0(K)dK + \int_\kappa^\infty \frac{2}{K^2} C_0(K)dK. \tag{4.25}$$

Returning now to our contract paying future realised variance, we consider the case where $f(S) = \Psi(S)$. Substituting Ψ for C in equation (4.3), and assuming σ_I is equal to zero we find the relation:

$$\sigma_R^2 T = \int_0^T \sigma_r(t)^2 dt = 2\left[\ln\left(\frac{\kappa}{S_T}\right) + \frac{S_T}{\kappa} - 1\right] - 2\int_0^T \left[\frac{1}{\kappa} - \frac{1}{S_t}\right] dS_t. \tag{4.26}$$

Then we note that the first term in the relationship above is just that which we examined with a view to static replication with plain vanilla option contracts. We may then, after substituting for $2\left[\ln\left(\frac{\kappa}{S_T}\right)+\frac{S_T}{\kappa}-1\right]$ in (4.26), write the initial cost of the variance contract in terms of plain vanilla options:

$$\int_0^T \sigma_R(t)^2 dt = \int_0^\kappa \frac{2}{K^2} P_0(K) dK + \int_\kappa^\infty \frac{2}{K^2} C_0(K) dK. \tag{4.27}$$

This summation of at and out-of-the-money call and put options then provides us with the initial fair value of the future variance of the asset S over the time interval $[0,T]$. By initially establishing this portfolio of options we may price the contract, and because the assumptions we made about the process for S did not include any about the volatility over the interval in question, we may use the market prices of options in order to find the market price of the variance contract. This fair value provides us with the strike variance, σ_K^2, that makes the Variance Swap zero cost.

The replication strategy after the initial hedge is established is simply to delta hedge the portfolio of options on a continuous basis. We made the assumption, when deriving the replicating portfolio for the variance contract, that the hedging volatility σ_I is equal to zero. It is therefore under this assumption that the continuous rebalancing of delta will take place. This should seem natural given the contracts demonstrated relationship to the Log Contract. As in the case of the Log Contract this results in a replication strategy whose value is totally independent of implied volatility. It only depends upon the path taken by the underlying asset and is equal in value at time T to the total variance exhibited over interval $[0,T]$.

A More Accurate Replication Strategy

Most recent work extended and explained the original papers of Carr and Madan [12] and that of Neuberger [40]. Demeterfi, Derman, Kamal and Zou [18] (hereafter referred to as DDKZ for convenience) published a paper in which they rederive the valuation methodology described above and introduce an improved strategy for static replication of the Variance Swap contract. Below, the alternative derivation is outlined and numerical results obtained for the price of the Variance Swap. These results, in both the constant volatility case and the case of an options market with a volatility smile, are compared with the values obtained from the method described in the previous section.

As well as giving a thorough explanation of the Variance Swap contract, the paper of DDKZ [18] also attempts to derive some results concerning the pricing of variance contracts under imperfect market conditions. We will look in particular at their approximation for a value when a volatility smile is exhibited by the options implied volatility market. This is compared to the results obtained through replication using a finite number of out-of-the-money options.

A Re-Derivation of the Valuation of a Variance Swap

DDKZ [18] also noted that a portfolio of options used to hedge a Variance Swap must have a constant exposure to the variance. This quantity they term "Variance Vega". Noting that a portfolio of plain vanilla options, weighted inversely proportionally to the square of their strikes has constant variance vega, they constructed the portfolio in the same way as Carr and Madan [12] where the portfolio valued at any time $(t:\ 0 \leq t \leq T)$ with volatility σ_R, is denoted $\Phi(S, \sigma_R\sqrt{t})$, and is equal to:

$$\Phi(S, \sigma_R\sqrt{t}) = \int_0^{\kappa} \frac{1}{K^2} P(S, K, t) dK + \int_{\kappa}^{\infty} \frac{1}{K^2} C(S, K, t) dK.$$

At the expiration time T they found that the value of $\Phi(S_T, 0)$, is given by:

$$\Phi(S_T, 0) = \frac{S_T - \kappa}{\kappa} - \ln\left(\frac{S_T}{\kappa}\right).$$

Note that this is simply a position in a forward contract with strike κ, the reference spot price, and a short position in a Log Contract. Thus the value of the portfolio $\Phi(S, \sigma_R\sqrt{t})$ at any time t before expiration is given by the result from equation (4.13).

$$\Phi(S, \sigma_R\sqrt{t}) = \frac{S - \kappa}{\kappa} - \ln\left(\frac{S}{\kappa}\right) + \frac{\sigma_R^2}{2}. \tag{4.28}$$

Thus, if delta hedged continuously, the only difference in price over the total life of the portfolio's life will be one half the total variance of the underlying asset price. We then see that in order to replicate a contract paying the total variance of the price path from 0 to T we hold and delta hedge continuously, two times the contents of portfolio Φ. DDKZ is actually concerned with, as we wil be, a contract whose exposure to the variance may be measured in terms of USD (or any other currency) per volatility point squared. This is used simply because market participants wish to speculate on volatility in general and must be able to easily understand their exposure to it's movement. We denote this new

portfolio as:

$$\Pi(S, \sigma_R\sqrt{t}) = \frac{2}{T}\left[\frac{S-\kappa}{\kappa} - \ln\left(\frac{S}{\kappa}\right)\right] + \frac{t}{T}\sigma_R^2. \tag{4.29}$$

Note, that this function will result in the annualised variance rather than the un-normalised variance found in Carr and Madan [12]. Multiplying the whole of equation (4.29) by T would give the non-annualised variance.

Thus it has been shown that a Variance Swap may be constructed from a portfolio containing of a forward contract and this new derivative security, the Log Contract. In fact these Log Contracts may not be traded directly. The last section presenting the work of Carr and Madan described an initial attempt at synthesising this payoff, but as a practical issue it is important that we are able to do this efficiently. It is upon replication of this Log Contract that our attention must then be focused.

An Alternative Replication Strategy for the Log Contract

We have found, through two different methods, that in order to replicate a Variance Swap we require a combination of two products. The first product, a forward contract, is a freely traded product. The second, on the other hand, is not. As we have already discussed, there is no market in existence for the trading of Log Contracts and it is therefore left to those wishing to trade them to replicate the contracts synthetically.

Our initial attempt to replicate the Log Contract was to purchase (or sell if the position is reversed) out-of-the-money call and put options expiring at the same time as the Variance Swap. These options were found to be weighted proportionally to the inverse of the square of their strikes. This theoretical static hedge requires the existence of options struck at every level. To perfectly replicate the Log Contract an infinite number of strikes must be used, covering every conceivable spot price at expiration time. In practice, the strike prices of options available in the market place may be limited. Due to transaction costs and liquidity constraints, a practitioner might choose not to trade a large number of different options. The results of the original static hedge used to find the initial fair value of the variance show that, with a limited number of strikes available, the error may be large. See the results section (4.3.3). Noticing this led DDKZ to suggest an alternative method. Establishing portfolio II at time zero as described in equation (4.29), using implied volatily

σ_I leads the fair value of the contract to be:

$$\Pi_0 = \frac{2}{T}\left[\frac{S_0 - \kappa}{\kappa} - \ln\left(\frac{S_0}{\kappa}\right)\right] + \sigma_I^2$$

Continuous delta hedging of the portfolio is done at a realised volatility σ_R. Using this volatility the fair variance would have been given by the same equation as above with σ_I replaced with σ_R. As a result the net payoff would be $(\sigma_R^2 - \sigma_I^2)$. This is the payoff of the Variance Swap with non-zero strike price $\sigma_K^2 = \sigma_I^2$. Note that in our previous descriptions we considered the contract paying the realised variance σ_R^2 rather than its difference from a predetermined strike level. The problem remains the same however - to find the fair variance price. In this case we find the fair annualised variance, $\frac{\sigma_K^2}{T}$.

Setting $\kappa = S_0$ in the above simplifies our problem. In doing this we are effectively using today's spot price as our reference price for creating the replication portfolio. Any value of κ may be used, but this simply introduces an extra forward term.

The problem is then one of replicating the function $f(S_T)$ at time T. That is:

$$f(S_T) = \frac{2}{T}\left[\frac{S_T - S_0}{S_0} - \ln\left(\frac{S_T}{S_0}\right)\right]. \tag{4.30}$$

Instead of the infinite number of options available for every strike above and below our reference price S_0 that we assumed existed before, we now only consider discretely spaced liquid strikes for call and put options. These are

$$K_0 = S_0 < K_{1c} < K_{2c} < ...$$

for the calls, and for the puts

$$K_0 = S_0 > K_{1p} > K_{2p} > ...$$

The payoff $f(S_T)$ is then approximated, in a linear fashion, using weighted amounts of call and put options. The initial call option weight is given by

$$w_c(K_0) = \frac{f(K_{1c}) - f(K_0)}{K_{1c} - K_0}. \tag{4.31}$$

We find the second weighting that must be added to the existing options, and fit the next part of the payoff function.

$$w_c(K_1) = \frac{f(K_{2c}) - f(K_{1c})}{K_{2c} - K_{1c}} - w_c(K_0). \tag{4.32}$$

This process is continued using the general relation for the call option weights:

$$w_c(Kn, c) = \frac{f(K_{n+1,c}) - f(K_{n,c})}{K_{n+1,c} - K_{n,c}} - \sum_{i=0}^{n-1} w_c(K_{i,c}).$$

For the put options we have a similar relation:

$$w_p(Kn, p) = \frac{f(K_{n+1,p}) - f(K_{n,p})}{K_{n+1,p} - K_{n,p}} - \sum_{i=0}^{n-1} w_p(K_{i,p}).$$

This replication is implemented in the section (4.3.3) and compared to the Carr and Madan methodology.

4.3.3 Numerical Results

We compare the results obtained through the use of the two replication schemes. This is done through a series of worked examples, beginning with the case where the implied volatility of spot price, S, is constant across all strikes at the time of writing the Variance Swap. We consider the following Variance Swap:

$$\begin{aligned} \text{Time to Maturity}, T &= 1, \\ \text{Spot Price}, S &= 100.0, \\ \text{Domestic Rate} &= 0.0, \\ \text{Foreign Rate} &= 0.0, \\ \text{Implied Volatility} &= 12.80\%. \end{aligned}$$

Firstly, we use the method of Carr and Madan to replicate this Variance Swap with an initial static options hedge. Remembering that this is done by finding the integral in equation (4.27), we discretise the possible range of strikes we wish to use to replicate the payoff. In practice this may be done by choosing the most liquid strike prices in the market. In our example we will simply divide a chosen range into equal pieces and use these divisions as the strike prices for the options in our replicating portfolio. We choose the maximum strike price to be 140.00 and the minimum to be 60.00. Taking 5 each of call and put options we have the strikes $100, 110, ..., 140$ for the calls and $60, 70, ..., 100$ for the put options.

The first piece of the replication we calculate is the weighting of the lowest delta, or most out-of-the-money, put option. This option has a strike price of 60. According to the integral (4.27) we take a weighting of $\frac{2}{K^2}$. Thus we find the weight is 5.56×10^{-4}. The value of

this option, P_0, is found using the Black-Scholes equation for plain vanilla options. A one year, 60 strike put, with volatility 12.8% and spot price 100 has value 7.3×10^{-5}. The difference between one strike price and the next is 10, so the contribution of this put option to the total fair variance is found by multiplying the weight by the unit cost of the option and the difference between this strike and the next. In our example this gives 4.06×10^{-7}. Subsequent put option weights are found in a similar way. The process is then repeated for the selected call options. Summing all of these contributions up produces the fair variance for the Variance Swap. Table (4.1) gives all of the weights and their sum. Taking the square root of this sum and dividing through by the square root of time gives us a figure to compare to the flat, annualised, implied volatility we used in order to price the Variance Swap. Market participants think in terms of implied volatility points and will usually request the strike of derivative contracts such as these to be quoted in terms of volatility points squared. Given Black-Scholes type market conditions, we know that the fair variance of the contract is equal to $\sigma_I^2 T$. We may then assess the accuracy of our replication strategy by comparing the fair variance (or its associated volatility) with the true value. In this case we see there is a significant difference in the two quantities. Increasing the number of options in the replication increases the accuracy of the strategy as shown in figure (4.4).

Table (4.2) illustrates the same example as above using the DDKZ method for replicating the derivative contract. Under this alternative scheme we are unable to start at the extreme of our chosen range of strike prices. Instead we must begin with the call or put option struck at the reference spot price κ. In this case κ is chosen to be equal to 100. We start the replication on the call side of the reference price and using equation (4.31) find the weighting for the call option struck at 100 to be 0.004787. $K_{1c} = 110$ and $K_0 = 100$, giving $f(K_{1c}) = 2(0.1 - \ln(1.1))$ and $K_0 = 0$. Substituting these values into equation (4.31) for $w_c(K_0)$ we get a weight of 0.00094. Again we multiply this by the unit value of the option which is 5.1030 to get a contribution to the fair variance of 0.004787. Now, moving up the strike array, we solve for the weight of our second call option struck at 110. $K_{2c} = 120$, and we have just found $w_c(K_0)$ so substituting into equation (4.32) we find:

$$w_c(K_1) = \frac{(0.0354 - 0.00938)}{10} - 0.00094 = 0.00166.$$

Multiplying by the value of the 110 call option we find the contribution of this option to the fair variance is 0.002945. All other values and weights are given in table (4.2). The fair variance is found by summing up all of these contributions. Given that the volatility

used to price these options was constant, the fair variance of the swap is equal to the square of the volatility. In this case the volatility used was 12.8%. The two examples demonstrate, as does figure (4.4) and table (4.3), that the method of replication devised by DDKZ results in a better approximation to the fair variance than Carr's methodology. Using only a relatively small number of options DDKZ provides a good estimate of the fair variance. Carr and Madan's method takes many more strikes to converge. Both methods

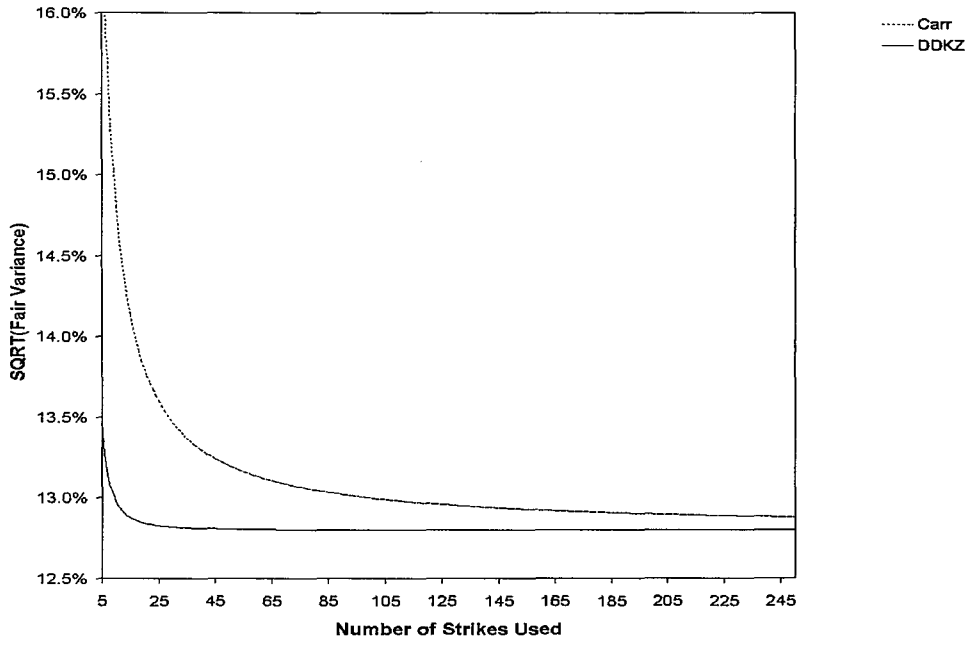


Figure 4.4: The convergence of values for a variance swap with no volatility smile.

allow the inclusion of a volatility smile when finding the value of the Variance Swap. This is achieved by using the market value of options when calculating the cost of the hedging strategy. Note that in the calculation of the weights in both approaches no input for the value of volatility was required. Implied volatility only played a part in determining the ultimate cost of replication, not its construction. The examples illustrated in tables (4.5) and (4.6) find the fair variance for swaps under the same conditions, and using the same array of strikes, as for the flat smile case just described. The one difference is that we use a volatility smile to evaluate the option prices. The smile is given in table (4.4). This smile is typical for the USD/JPY OTC options market. It is not symmetric about the 50-delta strike, showing some skew - or bias - to the market. We assume that the smile is linear in delta in order to price options of a specific strike if its implied volatility is not given. The results of pricing using an increasing number of options is demonstrated in figure (4.5)

showing that, under a volatility smile, the two different methods converge.

We note that, under the influence of a volatility smile where the out-of-the-money option implied volatilities are always greater than the at-the-money volatility, the fair variance we find using both methods of replication is greater than the at-the-money implied variance. Taking the Carr and Madan method of replication, we find the fair variance by adding option values together. Increasing the out-of-the-money implied volatilies increases the value of options compared to using fiat, at-the-money, volatility. In particular, the most effect will be found when put options struck out-of-the-money have a significantly higher volatility than at-the-money options. This is due to the weighting of contributing options by $\frac{1}{K^2}$ meaning more put than call options must be sold when implementing the replicating portfolio. In theory an option may have zero value. In practice it is unlikely that a trader would sell an option for nothing. Assuming that all options have value of at least one basis point (one hundreth of a percentage point), we may see that under market conditons, if implemented over a large range of strikes, both replication methods result in very large values for implied variance. If the range of strikes is not limited and every option assumed to have value of one basis point or more, then every additional option in our replication adds to the fair variance. An option with a Black-Scholes value of zero (using the volatility of the lowest delta option quoted directly in the market) sold for one basis point will have a very large implied volatility. The cumulative effects of these large implied volatilities dramatically increases the calculated fair variance.

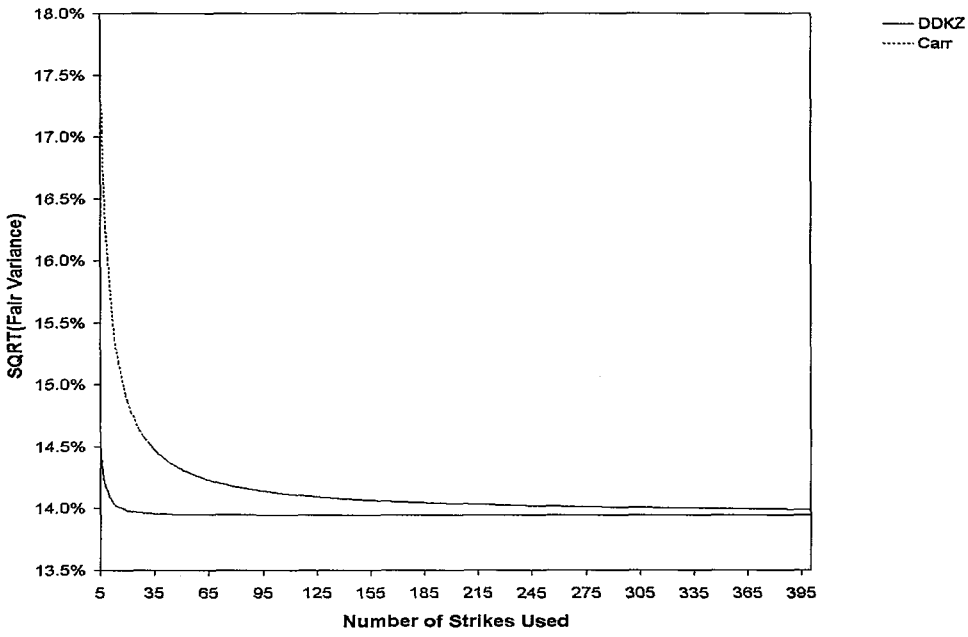


Figure 4.5: The convergence of values for a variance swap with a volatility smile.

Call/Put	Strike	Weight	Unit Cost	Total Cost
c	140	0.001020	0.020189	0.000020593
c	130	0.001183	0.108175	0.000127971
c	120	0.001389	0.486445	0.000675672
c	110	0.001653	1.775241	0.002934473
c	100	0.002000	5.102992	0.010205984
p	100	0.002000	5.102992	0.010205984
p	90	0.002469	1.399417	0.003455161
p	80	0.003125	0.187971	0.000587409
p	70	0.004082	0.00851	0.000034738
p	60	0.005556	0.000073	0.0000004056

Table 4.1: Replication using the method of Carr, with constant volatility for all strikes.

Call/Put	Strike	Weight	Unit Cost	Total Cost
c	140	0.00102	0.0202	0.0000206533
c	130	0.00119	0.1082	0.0001284037
c	120	0.00139	0.4864	0.0006781043
c	110	0.00166	1.7752	0.0029469001
c	100	0.00094	5.1030	0.0047866065
p	100	0.00107	5.1030	0.0054704074
p	90	0.00249	1.3994	0.0034775512
p	80	0.00315	0.1880	0.0005921087
p	70	0.00412	0.0085	0.0000350952
p	60	0.00563	0.0001	0.0000004113

Table 4.2: Replication using the method of DDKZ, with constant volatility for all strikes.

No. Calls and Puts	Carr	DDKZ
10	14.800%	12.969%
20	13.798%	12.842%
30	13.465%	12.818%
40	13.298%	12.810%
50	13.198%	12.806%
60	13.132%	12.804%
70	13.084%	12.803%
80	13.049%	12.802%
90	13.021%	12.802%
100	12.999%	12.801%
110	12.981%	12.801%
120	12.966%	12.801%
130	12.953%	12.800%
140	12.942%	12.800%
150	12.932%	12.800%
160	12.924%	12.800%
170	12.917%	12.800%
180	12.910%	12.800%
190	12.904%	12.800%
200	12.899%	12.800%

Table 4.3: Convergence to fair variance with an increasing number of options.

Call/Put	Delta	Strike	Volatility	over ATM
c	0.10	118.80	14.15%	1.35%
c	0.25	109.91	13.05%	0.25%
c	0.35	105.92	12.95%	0.15%
c	0.50	100.00	12.80%	0.00%
p	-0.50	100.00	12.80%	0.00%
p	-0.35	95.97	13.20%	0.40%
p	-0.25	94.48	13.60%	0.80%
p	-0.10	85.57	15.15%	2.35%

Table 4.4: Volatility smile used in the examples.

Call/Put	Weight	Strike	Unit Cost	Volatility	Totals
c	0.0010200	140	0.070754	14.846%	0.000072
c	0.0011840	130	0.249069	14.711%	0.000295
c	0.0013880	120	0.740332	14.246%	0.001028
c	0.0016520	110	1.859329	13.064%	0.003072
c	0.0020000	100	5.102992	12.800%	0.010206
p	0.0020000	100	5.13008	12.868%	0.010260
p	0.0024700	90	1.802158	14.246%	0.004451
p	0.0031260	80	0.505561	15.818%	0.001580
p	0.0040820	70	0.064595	16.161%	0.000264
p	0.0055560	60	0.002716	16.183%	0.000015
				Sum	**0.031243**
				SQRT(Variance)	**17.68%**

Table 4.5: Carr's replication of a Variance Swap under a volatility smile.

Call/Put	Weight	Strike	Unit Cost	Volatility	Totals
c	0.0010230	140	0.070754	14.846%	0.000072
c	0.0011870	130	0.249069	14.711%	0.000296
c	0.0013940	120	0.740332	14.246%	0.001032
c	0.0016600	110	1.859329	13.064%	0.003086
c	0.0009380	100	5.102992	12.800%	0.004787
p	0.0010720	100	5.130080	12.868%	0.005499
p	0.0024850	90	1.802158	14.246%	0.004478
p	0.0031500	80	0.505561	15.818%	0.001593
p	0.0041240	70	0.064595	16.161%	0.000266
p	0.0056340	60	0.002716	16.183%	0.000015
				Sum	**0.021125**
				SQRT(Variance)	**14.53%**

Table 4.6: DDKZ replication of a Variance Swap under a volatility smile.

4.4 Volatility Swaps

The last year or two has seen a great increase in demand from speculators for volatility related products. The variance swap described in the section above is a product with characteristics that make it attractive from the point of view of a liquidity provider. As a market maker in these products you are dealing with a contract which may be replicated, under certain conditions, using plain vanilla options and liquid products. The methodology for replicating these contracts is also free from assumptions regarding the process followed by the implied volatility of the underlying asset, allowing the smile exhibited by the options market to be taken into account and locked in through the static hedges described and implemented above.

The major drawback of the variance swap contract is that speculators do not generally wish to take positions in variance. Convention amongst those in the options market is to trade via implied volatilities, that is implied annualised volatility. Speculators therefore have a much better understanding of volatility, not variance, and it is volatility in which they wish to take a position. This set of circumstances has led to the development of what is known as a Volatility Swap. This differs significantly from a variance swap and presents many

more difficulties in its pricing and hedging.

It is impractical to write contracts based upon continuous, realised, variance or volatility. In practice such contracts are written upon a quantity that is predefined and corresponds to an estimate of the realised volatility from a number of price fixings for the underlying asset. Typically these fixings are taken on a daily basis at a predetermined time and from a predetermined source. The number of these fixings may be as few as nineteen in the case of a contract of duration one calendar month. It is convention to ignore non-business days when taking estimates of realised volatility.

In this section, a full description of the Volatility Swap is given. An intuitive approach is initially taken in order to gain an insight into the nature of this *swap* and to enable us to devise a hedging strategy. A second, more direct, methodology is then employed to derive a price and hedge ratios for the Volatility Swap. Several, distinctively different, pieces are found to contribute towards the value of the Volatility Swap and these are each described in turn. Both exact and approximate greeks are found for the contract. The approximate greeks are useful in helping to check the values we recover versus the intuitive approach we first employ. Results are given for the value of the Volatility Swap and the greeks that characterise it under particular conditions and the three methods for finding then are compared. Finally, numerical results are given for a practical example and the accuracy of the approximations tested against the exact greeks for the Volatility Swap. Full derivations for all quantities given in the following section are provided in appendices (C.1), (C.2), (C.3) and (C.4).

4.4.1 Description of the Volatility Swap

The type of contract we deal with here has become known as a *swap*. In fact it is not a swap at all. A Volatility Swap is a forward contract written on future realised volatility of the returns on an underlying asset. We will consider the case of a *volswap* written on a foreign exchange rate.

Let $S(t)$ denote the spot exchange rate at time t for a currency pair $ccy1/ccy2$. Then a volswap maturing at time T written on S has the following payoff,

$$\text{payoff} = 100 \times (\sigma_R - \sigma_K),$$

where, σ_K is a constant fixed at time 0 (the time the volswap contract was written),

and

$$\sigma_R = \sqrt{\frac{1}{(N-1)\Delta} \sum_{i=0}^{N-1} (\ln S(t_{i+1}) - \ln(S_i) - m_R\Delta)^2},$$

where, $\{t_i, i = 0, t_1, ..., t_N\}$ are known fixings that are assumed to be equidistant with $\Delta \equiv t_{i+1} - t_i$, $\forall i$ and with $t_0 = 0$ and $t_N = T$. Note that there are $N + 1$ spot fixings in total giving N returns.

Moreover, the annualised mean ,m_R, is defined to be,

$$m_R = \sum_{i=0}^{N-1} \frac{(\ln S(t_i + 1) - \ln S(t_i))}{N\Delta} \equiv \frac{(\ln S(T) - \ln S(0))}{N\Delta}.$$

It may be seen from the definition above that σ_R is the estimate of (annualised) standard deviation of the changes in the natural logarithm of the spot exchange rate $S(t)$ based upon N observations equally spaced Δ years apart. Clearly, $N\Delta = T$. For small values of N the error resulting from this estimate may be significant.

The volswap contract has a payoff of one unit of face currency per percentage point difference between the strike volatility ,σ_K, and the realised volatility over the life of the contract. If this difference is positive then the holder of the contract receives the corresponding amount, if this difference is negative then the holder of the contract makes a payment to the writer of the Volatility Swap.

In practice, these fixing times, t_i, are generally taken on a daily basis from a published page of exchange rates agreed, between the counterparties, at the time of writing.

The *strike volatility*, σ_K, is a constant. Common market practice is to set σ_K so that the contract is initially zero cost. Of course the volswap may be written such that there is an initial debit or credit. This adjustment would simply introduce a constant into the problem. The problem we address is how to determine the value of the volatility strike such that the contract is initially zero cost, and how to hedge the contract.

Volatility swaps offer an investor a direct way to speculate on future levels of market volatility. They also make it possible to express an opinion on the levels of future volatilities implied by the options market. Volswaps are pure volatility instruments which offer a payoff intuitively understood by speculators. Another attractive feature of these contracts when

compared to the delta-hedged options strategy discussed in the first section of this chapter, is that it is low maintenance requiring no rebalancing during its life.

We first demonstrate that it is intuitive to think of the volswap as a strip of forward starting straddles written on the spot price $S(t)$.

4.4.2 Approximation of the Volatility Swap by a Strip of Forward Starting Straddles

In this section we show that it is intuitive, and useful, to think of the volatility swap contract as a strip of forward starting straddles and that the volswap is bounded below by this strip.

Definition 1

Let $C(S,K,T)$ denote a call option on an asset S with strike price K and expiration T. Let $P(S,K,T)$ denote the corresponding put option. Then a straddle on asset S with strike price K and maturity T is the sum of long positions in one unit of $C(S,K,T)$ and one unit of $P(S,K,T)$.

From here onwards we will concentrate only on the risky part of the volswap payoff, i.e.

$$\sigma_R = \sqrt{\frac{\sum_{i=0}^{N-1} \left(\ln S(t_{i+1}) - \ln(S_i) - m_R\Delta\right)^2}{(N-1)\Delta}}. \tag{4.33}$$

In other words, we focus our efforts on finding the fair value of σ_K so that the swap is initially zero cost.

Note that the risky piece of the payoff (4.33) may be written as,

$$\sqrt{\frac{N}{N-1}}\sqrt{\frac{\sum_{i=0}^{N-1} \left[\ln S(t_{i+1}) - \ln S(t_i)\right]^2}{N\Delta} - {m_R}^2\Delta}.$$

and that this payoff is approximated by,

$$\sqrt{\frac{N}{N-1}}\sqrt{\frac{\sum_{i=0}^{N-1} \left[\ln S(t_{i+1}) - \ln S(t_i)\right]^2}{N\Delta}}. \tag{4.34}$$

By Jensen's inequality, the expression (4.34) is greater than,

$$\sqrt{\frac{N}{N-1}}\left[\frac{\sum_{i=0}^{N-1}|[\ln S(t_{i+1})-\ln S(t_i)]|}{N\sqrt{\Delta}}\right]. \tag{4.35}$$

Upon approximating log returns by discrete returns, expression (4.35) is approximately equal to,

$$\sqrt{\frac{N}{N-1}}\sum_{i=0}^{N-1}\left[\frac{|[S(t_{i+1})-S(t_i)]|}{S(t_i)N\sqrt{\Delta}}\right]. \tag{4.36}$$

By inspection, we see that the i^{th} term in (4.36) is the payoff on $K(t_i)$ straddles purchased at time t_i, expiring at time t_{i+1} and where the strike of the call and put option making up the straddle is $S(t_i)$. Note that the payoffs of these straddles are compounded to the ultimate expiration of the volswap, T. If we have a volswap written on a foreign exchange rate S between two currencies, $ccy1$ and $ccy2$, where the rate is quoted in terms of $ccy2$ per $ccy1$ (written $ccy1/ccy2$ under market conventions), then the straddles are written each on one unit of $ccy2$ and the face currency of the swap is $ccy1$. Furthermore, we see that the face amount of each straddle $K(t_i)$ is given by,

$$K(t_i)=e^{-r_d(T-t_{i+1})}\sqrt{\frac{N}{N-1}}\left[\frac{1}{N\sqrt{\Delta}S(t_i)}\right],$$

where r_d is the continuously compounded, risk-free interest rate for the face currency of the volswap. This is due to the fact that foreign exchange options written on a currency pair $ccy1/ccy2$ pays off in $ccy2$. As our straddle is written on $ccy2$ the payoff will be in terms of $ccy1$ and it is this which must be compounded to T, the expiration date of the volswap.

We have shown that the payoff from a volatility swap is bounded below by the strip of straddles mentioned above. We will now look briefly at what this implies about the characterstics of a volswap and its hedge.

The first thing we notice is that the approximation of the volswap by a strip of forward starting straddles has ignored the realised drift of the underlying asset. We are also using at-the-money-spot straddles rather than those struck at-the-money-forward. This means that the delta of these option combinations will not be precisely zero at each fixing time. However, as long as the interest differential, $r_d - r_f$, is not large and that the time between fixes, Δ is sufficiently small, the delta of the strategy should be very close to zero. Also note that forward starting options have zero delta and zero gamma. See Appendix (B) for a proof of this. The only sensitivity that a forward starting option has, ignoring the interest

rates, is to volatility. The only greek to hedge for an option which is forward starting is therefore vega.

From this analysis we expect the Volatility Swap to have a delta close to zero for when a fixing has just been taken, and at any point in time this delta (and the gamma) is due only to the next fixing that is to take place. Other fixings manifest themselves as drift or realised volatility in the case of the past fixes or, in the case of future fixes, as vega. Therefore, wherever spot may go we still expect the volswap to have some vega, and once the next fixing has been taken we expect the delta to reset to a value close to zero and have a non-zero gamma. The volswap, when looked at as a strip of forward starting straddles, retains sensitivity to implied volatility wherever spot moves. In the case of a plain vanilla option, this sensitivity is lost should spot move sufficiently far away from the current strike price. The volswap does not lose this sensitivity and could therefore be seen to have some elements of out-of-the-money options embedded within it that move towards the money as spot moves helping to maintain the non-zero vega. Recall that a variance swap may be replicated by using options with an infinite number of strikes. The discrete volswap appears to have similar characteristics when looked at in this manner.

4.4.3 Pricing the Volatility Swap

We examine this contract under the Black-Scholes model. Let us begin with some preliminary assumptions.

The spot price $S(t)$ evolves in continuous time as a geometric Brownian motion under the risk-neutral probability measure as follows:

$$\frac{dS}{S} = (r_d - r_f)dt + \sigma dz,$$

where $\{z(t)\}$ is a standard Brownian motion, r_d is the domestic, continuously compounded, risk-free interest rate and r_f is the corresponding foreign interest rate. We assume that all parameters are deterministic.

Under these assumptions,

$$\ln S(t_{i+1}) - \ln S(t_i),$$

the return over an infinitesimally small time step for the spot price, is normally distributed

and has mean,

$$(r_d - r_f - \frac{1}{2}\sigma^2)\Delta,$$

and variance $\sigma^2\Delta$.

The strike volatility, σ_K, is to be set such that the swap has an initial value of *zero*.

The price of the volatility swap is, at any instant t , $0 \leq t \leq T$, given by $V(t)$ where,

$$V(t) = A(t)\left[\mathbb{E}_t^Q[\sigma_R] - \sigma_K\right],$$

where,

$$A(t) = 100e^{-r_d(T-t)}, \tag{4.37}$$

the present value of one percentage point in terms of units of the face currency.

An Upper Bound on the Price of the Volatility Swap

Now, through the use of an approximation, we compute an upper bound on the mark-to-market value of the volatility swap at any point, t, during the life of the contract.

Looking at the expression given above for $V(t)$ we see that the only non-constant piece is $\mathbb{E}_t^Q[\sigma_R]$ so it is on this expectation that we will concentrate.

Jensen's inequality implies that,

$$\mathbb{E}_t^Q \sigma_R \leq \sqrt{\mathbb{E}_t^Q [{\sigma_R}^2]}.$$

We therefore concentrate on this upper bound and compute its value.

Let

$$\hat{V}(t) = A(t)\sqrt{\mathbb{E}_t^Q {\sigma_R}^2}.$$

As discussed earlier, the realised volatility of the asset S is calculated on the basis of discrete fixings. We therefore write $\hat{V}(t)$ as,

$$\hat{V}(t) = \sqrt{\mathbb{E}_t^Q {s_R}^2} = \sqrt{\mathbb{E}_t^Q \sum_{i=0}^{N-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i) - m\Delta]^2}{(N-1)\Delta}}, \tag{4.38}$$

where m is the mean of the log changes in $S(t)$,

$$m = \frac{\ln S(T) - \ln S(0)}{N\Delta}$$

Rewriting the above yields,

$$\hat{V}(t) = \sqrt{\frac{N}{N-1}}\,\mathbb{E}_t^Q \sqrt{\sum_{i=0}^{N-1} \frac{\{\ln S(t_{i+1}) - \ln S(t_i)\}^2}{N\Delta} - \left\{\frac{\ln S(T) - \ln S(0)}{N}\right\}^2 \frac{1}{\Delta}}.$$

Now, suppose that we are at time t, $0 \le t \le T$, when i_0 fixes have already been observed. More formally, $t_{i_0} \le t < t_{i_0+1}$, $0 \le i_0 \le N-1$.

We compute the value of $\mathbb{E}_t^Q\, {s_R}^2$ by splitting ${s_R}^2$ into four distinct pieces and calculating the expectation of each piece in turn. The decomposition is as follows:

$$\frac{(N-1)}{N} s_R^2 = s_1^2 + s_2^2 + s_3^2 - s_4^2$$

where,

$$s_1^2 = \sum_{i=0}^{i_0-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i)]^2}{N\Delta} = \frac{i_0}{N} v_H^2,$$

and,

$$v_H^2 = \sum_{i=0}^{i_0-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i)]^2}{i_0\Delta}$$

also,

$$s_2^2 = \frac{[\ln S(t_{i_0+1}) - \ln S(t_{i_0})]^2}{N\Delta},$$

$$s_3^2 = \sum_{i=i_0+1}^{N-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i)]^2}{N\Delta},$$

and finally

$$s_4^2 = \frac{1}{\Delta}\left\{\frac{\ln S(T) - \ln S(0)}{N}\right\}^2.$$

Notice that the payoff has been broken into four different pieces. The first corresponds to the contribution from the volatility of the fixings taken so far. The second piece, s_2^2, corresponds to the contribution that the return for the next fixing coupled with the last fixing will make to the volswap valuation. s_3^2 corresponds to the future fixings and their expected contribution. Finally, s_4^2 is the contribution to the value of the volswap arising from the realised drift of the underlying asset over the entire life of the contract.

If i_0 is equal to zero then there is no contribution from previous fixes and so s_1^2 is taken to be zero. If $i_0 = N-1$ then there are no more fixes to be taken and s_3^2 is taken to be zero.

Now, in order to calculate the value of $\hat{V}(t)$ the expectation of each piece is taken. We find that:

$$b_1 = \mathbb{E}_t^Q[s_1^2] = s_1^2 = \frac{i_0}{N} v_H^2.$$

as all components are known at time t.

Also

$$b_2 = \mathbb{E}_t^Q[s_2^2] = \frac{1}{N\Delta}\left[\left\{\ln\left[\frac{S(t)}{S(t_{i_0})}\right] + (\alpha - \frac{1}{2}\sigma_G^2)\hat{\Delta}\right\}^2 + \sigma_G^2\hat{\Delta}\right],$$

and

$$b_3 = \mathbb{E}_t^Q[s_3^2] = \left(\frac{N - i_0 - 1}{N}\right)\left[(\alpha - \frac{1}{2}\sigma_I^2)^2\Delta + \sigma_I^2\right],$$

finally

$$\begin{aligned} b_4 &= \mathbb{E}_t^Q[s_4^2] \\ &= \frac{1}{N^2\Delta}\left[\left\{(\alpha - \frac{1}{2}\sigma_I^2)(T - t_{i_0+1}) + (\alpha - \frac{1}{2}\sigma_G^2)\hat{\Delta} + \ln\left(\frac{S(t)}{S(0)}\right)\right\}^2 + \sigma_I^2(T - t_{i_0+1}) + \sigma_G^2(\hat{\Delta})\right], \end{aligned}$$

where,

$$\begin{aligned} N &= \text{total no. of returns,} \\ i_0 &= \text{no. fixes expired,} \\ t &= \text{time now,} \\ t_{i_0} &= \text{time of the previous fixing,} \\ \hat{\Delta} &= \text{time from t until the next fix,} \\ \Delta &= \text{time between fixes,} \\ T &= (t_N - t_0) = \text{total expiration of swap,} \\ S(t) &= \text{spot price at time t,} \\ \sigma_I &= \text{implied volatility for interval } [t_{i_0+1}, T], \\ \sigma_G &= \text{implied volatility for interval } [t, t_{i_0+1}], \\ \alpha &= r_d - r_f, \text{ the interest rate differential.} \end{aligned}$$

See Appendix (C.3) for calculations of the expectations stated above.

So, combining all of the components computed above and denoting $H(t)$ as

$$H(t) = \sqrt{(b_1 + b_2 + b_3 - b_4)},$$

we find

$$\hat{V}(t) = A(t)H(t)\sqrt{\left(\frac{N}{N-1}\right)}, \tag{4.39}$$

for $t \in [t_{i_0}, t_{i_0+1})$ with $0 \leq i_0 \leq N - 1$.

An Exact Expression for $V(t)$.

We recall that the expression derived in (4.39) for $\hat{V}(t)$ is an upper bound on the value of the volatility swap at time t. For large N, the number of fixings to be taken, this is likely to be very close to the true value of $V(t)$. In fact it is possible to calculate the exact value of $V(0)$ under the Black-Scholes model.

A well-known result of the geometric Brownian motion assumption being used in the Black-Scholes model to price the volatility swaps is that the changes, $\ln S(t_{i+1}) - \ln S(t_i)$, are normally distributed with means $\mu_i = (r_d - r_f - \frac{1}{2}\sigma^2)\ \Delta$ and variances $\sigma_i^2 = \sigma^2\Delta$, where $\Delta = (t_{i+1} - t_i)$. Furthermore, for non-overlapping time intervals, these changes are pairwise independent. The sum of the squares of N independent random variables X_i, each having a standard normal distribution, i.e. $X_i \sim N(0,1)$, is known to possess a chi-squared distribution. Therefore, the random variable X defined by

$$X \triangleq \sum_{i=1}^{N} \left(\frac{X_i - \mu_i}{\sigma_i}\right)^2, \tag{4.40}$$

has a chi-squared distribution with N degrees of freedom. Recalling that σ_R^2 can be written as

$$\sigma_R^2 \times \frac{(N-1)}{\sigma^2} = \sum_{i=0}^{N-1} \left(\frac{\ln S(t_{i+1}) - \ln S(t_i) - \mu_i}{\sigma_i}\right)^2, \tag{4.41}$$

we get the result that $\sigma_R^2 \times \frac{(N-1)}{\sigma^2} \sim \chi^2(N)$. Thus, the following required result is obtained:

$$\mathbb{E}[\sqrt{\sigma_R^2}] = \frac{\sigma}{\sqrt{N-1}}\,\mathbb{E}\left[\sqrt{\left(\sigma_R^2\frac{(N-1)}{\sigma^2}\right)}\right] = \sigma\sqrt{\frac{2}{N-1}}\,\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)}. \tag{4.42}$$

See Appendix (C.1) for a full derivation of this result.

Therefore the exact value of $V(0)$ is given by:

$$V(0) = A(0)\,\overset{Q}{\mathbb{E}}_0\sqrt{\sigma_R^2} - A(0)\sigma_K = A(0)\left[\sigma\sqrt{\frac{2}{N-1}}\,\frac{\Gamma\left(\frac{N}{2}\right)}{\Gamma\left(\frac{N-1}{2}\right)} - \sigma_K\right] \tag{4.43}$$

Now, we estimate the difference between the approximate price $\hat{V}(0)$ of the volswap and the exact price $V(0)$. We look at the difference between $\mathbb{E}_t^Q\,\sigma_R$ and $\sqrt{\mathbb{E}_t^Q\,\sigma_R^2}$. Using a second order Taylor series expansion of $\sqrt{\sigma_R}$ around $\mathbb{E}_t^Q\,\sigma_R^2$ we obtain:

$$\sigma_R = \sqrt{\sigma_R^2} \simeq \sqrt{\mathbb{E}_t^Q\,\sigma_R^2} + \frac{1}{2}[\mathbb{E}_t^Q\,\sigma_R^2]^{-\frac{1}{2}}[\sigma_R^2 - \mathbb{E}_t^Q\,\sigma_R^2] - \frac{1}{8}[\mathbb{E}_t^Q\,\sigma_R^2]^{-\frac{3}{2}}[\sigma_R^2 - \mathbb{E}_t^Q\,\sigma_R^2]^2.$$

Then, by taking expectations under Q of both sides of the equation,

$$\mathbb{E}_t^Q = V(0) \simeq \sqrt{\sigma_R^2} \simeq \sqrt{\mathbb{E}_t^Q \sigma_R^2} - \frac{1}{8}[\mathbb{E}_t^Q \sigma_R^2]^{-\frac{3}{2}} Var_0^Q[\sigma_R^2].$$

Substituting for $Var_0^Q[\sigma_R^2] = \frac{2\sigma^4}{N-1}$ and $\mathbb{E}_t^Q \sigma_R^2 = \sigma^2$, we find that,

$$V(0) \simeq \hat{V}(0) - \frac{\sigma}{4(N-1)}.$$

The above relation shows that as the number of fixes to be taken during the life of the volswap N approaches infinity, the error in pricing will tend towards zero. In other words $V(0)$ and $\hat{V}(0)$ will be equal and given by $A(0)(\sigma - \sigma_K)$

We have therefore found the error bound for $V(0)$ and $\hat{V}(0)$ at time $t = 0$. Note, however that as we progress further through the life of the volswap and more fixings have expired the randomness of σ_R decreases. This suggests that the error in pricing arising from our approximation $\hat{V}(t)$ is likely to become smaller as t increases.

4.4.4 Greeks and Hedging the Volatility Swap

We have derived an approximate value for the volswap at any point, t, during the life of the swap. Given this approximate value $\hat{V}(t)$ we may then find the Black-Scholes hedge ratios associated with this contract. We calculate the delta, gamma and vega by differentiating with respect to the current spot price once for the delta, twice for the gamma and finding the derivative with respect to the implied volatility σ for the vega. See Appendix (C.2) for the derivations. The expressions found are as follows.

$$\begin{aligned}\Delta_e &= \frac{A(t)}{H(t)S(t)} \frac{1}{N^2\Delta} \sqrt{\frac{N}{N-1}} \left[(N-1) \left(\ln \frac{S(t)}{S(t_{i_0})} + \left[\alpha - \frac{1}{2}\sigma_G^2 \right] \hat{\Delta} \right) \right. \\ &\qquad \left. - \left(\ln \frac{S(t_{i_0})}{S(0)} + \left[\alpha - \frac{1}{2}\sigma_I^2 \right] [T - t_{i_0+1}] \right) \right]. \qquad (4.44)\end{aligned}$$

$$\begin{aligned}\Gamma_e &= \sqrt{\frac{1}{N(N-1)}} \frac{1}{N\Delta S^2(t)} \frac{A(t)}{H^2(t)} \left\{ H(t) \left[(N-1) R_1(t, t_{i_0}) + R_2(t, t_{i_0}, t_{i_0+1}) \right] \right. \\ &\qquad \left. - \frac{1}{N^2 \Delta H(t)} \left[(N-1) \left(1 - R_1(t, t_{i_0}) \right) - R_2(t, t_{i_0}, t_{i_0+1}) \right]^2 \right\}, \qquad (4.45)\end{aligned}$$

where

$$R_1(t, t_{i_0}) = 1 - \left(\alpha - \frac{1}{2}\sigma_G^2 \right) \hat{\Delta} - \ln \frac{S(t)}{S(t_{i_0})}; \text{ and} \qquad (4.46)$$

$$R_2(t, t_{i_0}, t_{i_0+1}) = \ln \frac{S(t_{i_0})}{S(0)} + \left(\alpha - \frac{1}{2}\sigma_I^2 \right) (T - t_{i_0+1}). \qquad (4.47)$$

$$
\begin{aligned}
vega_e &= \sqrt{\frac{1}{N(N-1)}}\frac{A(t)}{H(t)}\left[\sigma_I(N-i_0-1)[1-(\alpha-\frac{1}{2}\sigma_I^2)\Delta]\right. \\
&+ \sigma_G\frac{\hat{\Delta}}{\Delta}\left[1-\ln\frac{S(t)}{S(t_{i_0})}-\left(\alpha-\frac{1}{2}\sigma_G^2\right)\hat{\Delta}\right] \\
&- \frac{1}{N\Delta}\left[\sigma_G\hat{\Delta}+\sigma_I(T-t_{i_0+1})\right]\left[1-\ln\frac{S(t)}{S(0)}\right. \\
&\qquad \left.\left.-\left(\alpha-\frac{1}{2}\sigma_G^2\right)\hat{\Delta}-\left(\alpha-\frac{1}{2}\sigma_I^2\right)(T-t_{i_0+1})\right]\right]. \qquad (4.48)
\end{aligned}
$$

Greeks Arising from an Approximation for $\hat{V}(t)$.

We use an approximation for the term $\hat{V}(t)$ in order to see the qualitative properties of the volatility swap and its hedge more easily. Recall that $V(t)$ is the exact volswap price and $\hat{V}(t)$ an exact upperbound to it.

If, instead of writing the exact expression for $\hat{V}(t)$ in equation (4.38), we write it as

$$
\hat{V}(t) \approx \hat{W}(t)\sqrt{\mathbb{E}_t^Q\, s_R{}^2} = \sqrt{\mathbb{E}_t^Q \sum_{i=0}^{N-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i) - z\Delta]^2}{N\Delta}}, \qquad (4.49)
$$

where $z = r_d - r_f - \frac{1}{2}\sigma^2$ is the drift of our Brownian motion. In a similar fashion to the exact calculation for $\hat{V}(t)$ we rewrite q_R^s as the sum of three quantities.

$$
q_R^2 = q_1^2 + q_2^2 + q_3^2
$$

where,

$$
q_1^2 = \sum_{i=0}^{i_0-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i) - m\Delta]^2}{N\Delta} = \frac{i_0}{N} w_H^2,
$$

and w_H^2 is an estimate of σ^2 given by the i_0 fixes that have expired up until time t and is equal to,

$$
w_H^2 = \sum_{i=0}^{i_0-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i) - m\Delta]^2}{i_0\Delta}.
$$

In addition,

$$
q_2^2 = \frac{[\ln S(t_{i_0+1}) - \ln S(t_{i_0})] - m\Delta^2}{N\Delta},
$$

$$
q_3^2 = \sum_{i=i_0+1}^{N-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i) - m\Delta]^2}{N\Delta}.
$$

If i_0 is equal to zero then q_1^2 is taken to be zero. If $i_0 = N-1$ then q_3^2 is taken to be zero.

Taking expectations at time t under measure Q gives,

$$
\hat{W}(t) \approx A(t)J(t),
$$

where,

$$J(t) = \sqrt{\frac{i_0}{N} w_H^2 + \frac{1}{N} w_c^2 + \left[\frac{N - i_0 - 1}{N}\right] \sigma^2}$$

and,

$$w_c^2 = \frac{1}{\Delta} \left[\frac{[S(t) - S(t_{i_0})(1 + \alpha(\Delta - \hat{\Delta}))]^2}{S^2(t_{i_0})} + \sigma^2 \hat{\Delta} \right].$$

As with the exact expression for $\hat{V}(t)$ we differentiate the expressions above to find the approximate greeks of the volatility swap.

We find,

$$\Delta_a = A(t) \left[\frac{R(t, t_{i_0})}{N \Delta J(t) S(t_{i_0})} \right], \tag{4.50}$$

$$\Gamma_a = A(t) \frac{\left[1 - \left(\frac{R^2(t, t_{i_0})}{N \Delta J^2(t)}\right)\right]}{J(t) S^2(t_{i_0}) N \Delta}, \tag{4.51}$$

$$Vega_a = A(t) \frac{\sigma}{J(t)} \left[\frac{N - i_0 - 1}{N} + \frac{\hat{\Delta}}{\Delta} \frac{1}{N} \right]. \tag{4.52}$$

where $J(t)$ is the estimate of σ^2 given above and,

$$R(t, t_{i_0}) = \left[\frac{S(t)}{S(t_{i_0})} \right] - (1 + \alpha(\Delta - \hat{\Delta})).$$

4.4.5 Results

We now examine, in detail, the values for the greeks implied by the model and its approximations. Here we check that the model agrees with any intuitive ideas we may have had prior to the construction of the model and that, under particular conditions, the different expressions yield similar answers.

Values of the Greeks at time Zero

At time zero, we notice the following:

$$\begin{aligned} t &= 0, \\ S(t) &= S(t_{i_0}), \\ \hat{\Delta} &= \Delta. \end{aligned} \tag{4.53}$$

Thus, substituting the above into equation (4.50) for the approximate delta Δ_a and assuming that $\sigma_G = \sigma_I = \sigma$ for simplicity, we find that $R(t, t_{i_0}) = 0$ and so,

$$\Delta_a(0) = 0. \tag{4.54}$$

Substituting the same values into equation (4.44) for the exact value of the delta gives us the same result. That is:

$$\Delta_e(0) = 0. \tag{4.55}$$

Given our previous discussion in section (4.4.2), this is precisely the result we expect. The delta of an at-the-money forward straddle is zero, so if the approximation is a good one then the greeks derived directly from the expression for $\hat{V}$ should also be very close, or equal to zero. Our intuition tells us that at the time of the first fixing in the volswap there is no reference price to affect the value of the contract. Thus, rather like an at-the-money straddle, at the instant that the strike is set the contract has zero delta. When the spot price of the underlying asset starts to move we do expect to see a change in this delta to something non-zero. In other words, we expect the gamma to be non-zero at the start of the volswap's life.

Again making the substitutions from (4.53) only this time into equations (4.45) and (4.51), we find that:

$$\Gamma_a(0) = \frac{A(0)}{T\sigma S^2(0)}, \tag{4.56}$$

and,

$$\Gamma_e(0) = \frac{A(0)}{TS^2(0)\sqrt{(\alpha - \frac{1}{2}\sigma^2)^2 + \sigma^2}}. \tag{4.57}$$

Note that, if $(\alpha - \frac{1}{2}\sigma^2)^2$ in equation (4.57) is small compared to σ^2 the two approximations for the gamma are nearly equal, the only difference being the risk-neutral drift term. We also note the dependence of gamma on the maturity of the volswap. In both cases, exact and approximate, the gamma is inversely proportional to the term of the swap, T. In common with plain vanilla options, short maturity volswaps have larger gamma than long maturity volswaps. As we will see below, the vega of these contracts does not behave in the same way as that of a plain vanilla.

Once more we make use of (4.53), this time in finding the initial vega of the volswap. Substituting these into equations (4.52) and (4.48) we find that:

$$Vega_a(0) = A(0), \tag{4.58}$$

and,

$$Vega_e(0) = \frac{A(0)\sigma}{\sqrt{(\alpha - \frac{1}{2}\sigma^2)^2 + \sigma^2}}. \tag{4.59}$$

Again, we observe that if the value of $(\alpha - \frac{1}{2}\sigma^2)^2$ is small compared to σ^2, we have almost equal values arising from each calculation. The value of the vega in each case corresponds to our intuitive expectations, which gives us confidence in our methodology. A volswap is a contract based purely upon volatility. It is, in essence, a forward contract written on realised volatility. If the face amount is of the swap is 1 as we have assumed throughout this chapter then a vega of $A(0)$ corresponds to the present value of one volatility percentage point. If implied volatility moves by one percent, and all else remains equal, then we approximately expect the value of our swap to change by one point at time T. The extra term in (4.59) occurs due to the careful treatment of the risk-neutral drift term in our expectation calculations.

Values of the Greeks at any Fixing Time t_{i_0}

In the following, we will assume that $\sigma_G = \sigma_I = \sigma$ for the sake of simplicity. Substituting this into equations (4.50), (4.51) and (4.52) for the approximate greeks and into (4.44), (4.45) and (4.48) for the exact greeks at general time, t, we have:

$$\Delta_a = \frac{1}{N\Delta}\frac{A(t)}{J(t)}\frac{R(t,t_{i_0})}{S(t_{i_0})}. \tag{4.60}$$

$$\Gamma_a = \frac{1}{N\Delta}\frac{1}{S^2(t_{i_0})}\frac{A(t)}{J(t)}\left[1 - \frac{R^2(t,t_{i_0})}{N\Delta J^2(t)}\right]. \tag{4.61}$$

$$vega_a = \frac{\sigma}{N}\frac{A(t)}{J(t)}\left[\frac{\hat{\Delta}}{\Delta} + (N - i_0 - 1)\right]. \tag{4.62}$$

where,

$$R(t,t_{i_0}) \triangleq \frac{S(t)}{S(t_{i_0})} - \left[1 + \alpha\left(\Delta - \hat{\Delta}\right)\right]. \tag{4.63}$$

and for the exact greeks,

$$\begin{aligned}\Delta_e \;=\; & \frac{A(t)}{H(t)S(t)}\frac{1}{N^2\Delta}\sqrt{\frac{N}{N-1}}\left[N\left(\ln\frac{S(t)}{S(t_{i_0})} + \left[\alpha - \frac{1}{2}\sigma^2\right]\hat{\Delta}\right)\right. \\ & \left. - \left(\ln\frac{S(t)}{S(0)} + \left[\alpha - \frac{1}{2}\sigma^2\right][T-t]\right)\right].\end{aligned} \tag{4.64}$$

$$\Gamma_e = \sqrt{\frac{1}{N(N-1)}}\frac{A(t)}{H(t)}\frac{1}{N\Delta S^2(t)}\left\{N - 1 - R_3(t,t_{i_0},t_{i_0+1}) - \frac{R_3^2(t,t_{i_0},t_{i_0+1})}{N^2\Delta H^2(t)}\right\}. \tag{4.65}$$

where,

$$R_3(t, t_{i_0}, t_{i_0+1}) = N \ln \frac{S(t)}{S(t_{i_0})} - \left(\alpha - \frac{1}{2}\sigma^2\right)\left(T - t - N\hat{\Delta}\right) - \ln \frac{S(t)}{S(0)}, \tag{4.66}$$

$$\begin{aligned} vega_e &= \sigma\sqrt{\frac{1}{N(N-1)}}\frac{A(t)}{H(t)}\left\{(N - i_0 - 1) + \frac{\hat{\Delta}}{\Delta} - \frac{(T-t)}{N\Delta}\right. \\ &+ \left[\alpha - \frac{1}{2}\sigma^2\right]\left[\frac{(T-t)^2}{N\Delta} - \frac{\hat{\Delta}^2}{\Delta} - (N - i_0 - 1)\Delta\right] \\ &+ \left.\frac{(T - t - N\hat{\Delta})}{N\Delta}\ln S(t) - \frac{(T-t)}{N\Delta}\ln S(0) + \frac{\hat{\Delta}}{\Delta}\ln S(t_{i_0})\right\}. \end{aligned} \tag{4.67}$$

We look first at the delta of the volswap at a time t_{i_0} when a fixing for the spot price of the underlying is taken. We notice that $\hat{\Delta} = \Delta$, $t = t_{i_0}$ and $t = i_0\Delta$. Substituting these into equation (4.60) we find that the approximate delta at a general fixing time is zero. This corresponds to the intuitive approach taken when thinking of the volswap as a strip of forward starting straddles. As an at-the-money straddle has a delta very close to zero, we would hope that our approximated delta would be close to zero too.

Making the same substitutions into equation (4.64) we find:

$$\Delta_e(t_{i_0}) = \frac{A(t_{i_0})}{H(t_{i_0})S(t_{i_0})}\frac{1}{N^2\Delta}\sqrt{\frac{N}{N-1}}\left(i_0\Delta(\alpha - \frac{1}{2}\sigma^2) - \ln\frac{S(t_{i_0})}{S(0)}\right). \tag{4.68}$$

Again, the difference in our two approximations is, as we expect, in the treatment of the drift of the underlying asset's spot price. Here we are left with, instead of zero, the difference between the expected drift and the actual drift so far.

Making the same substitutions, only this time into equation (4.62), we find the approximate vega of the volswap at a fixing time t_{i_0}:

$$Vega_a(t_{i_0}) = \frac{N - i_0}{N}\sigma\frac{A(t_{i_0})}{J(t_{i_0})}. \tag{4.69}$$

and for the exact vega from equation (4.67):

$$Vega_e(t_{i_0}) = \frac{N - i_0}{N}\sigma\frac{A(t_{i_0})}{H(t_{i_0})}\sqrt{\frac{1}{N(N-1)}}\left((N - i_0)\Delta - 1 + \ln\frac{S(t_{i_0})}{S(0)}\right). \tag{4.70}$$

Note that if the fixing in question, i_0, is the last fixing of the volswap contract to take place then in both the approximate and exact expressions we find a vega of exactly zero. Furthermore, during the life of the contract the vega is proportional to the fraction of the swap's remaining life. It is also dependent upon the ratio of the implied volatility to $\sqrt{\frac{N}{N-1}}H(t_{i_0})$, the effective volatility found by combining implied and realised volatilities.

In other words, as long as the realised volatility is close to the implied volatility our initial vega hedge should work well. Large deviation from this results in rebalancing being required which can be costly and is not priced into the contract at time zero. A similar analysis of the gamma resulting from the volatility swap contract yields the following for the approximate gamma:

$$\Gamma_a(t_{i_0}) = \frac{A(t_{i_0})}{J(t_{i_0})} \frac{1}{N\Delta} \frac{1}{S^2(t_{i_0})}. \tag{4.71}$$

In fact for gamma, both approximate and exact, values change little between t_{i_0} and t_{i_0+1}, for large values of N. We may see from equation (4.71) that gamma is inversely proportional to the effective volatility $\sqrt{\frac{N}{N-1}}H(t_{i_0})$. It is the deviation in the effective volatility from the implied volatility that plays the largest role in altering the gamma of the swap during its life. If $\sqrt{\frac{N}{N-1}}H(t)$ remains close to the initial implied volatility then the value of gamma remains close to the initial value. This again gives us confidence due to the similarity to a strip of forward starting straddles. The straddles will all have the same gamma when initiated at each fixing time as long as the volatility at that fixing time is the same. Also, recall from equation (4.6) that the gamma of a vanilla option is inversely proportional to the volatility. The exact gamma yields the same qualitative results.

$$\begin{aligned} \Gamma_e(t_{i_0}) &= \frac{A(t_{i_0})}{H(t_{i_0})} \frac{1}{N\Delta} \frac{1}{S^2(t_{i_0})} \sqrt{\frac{1}{N(N-1)}} \\ &\times \left((N-1) - R_3(t_{i_0}, t_{i_0}, t_{i_0+1}) - \frac{R_3(t_{i_0}, t_{i_0}, t_{i_0+1})}{N^2 \Delta H^2(t_{i_0})} \right), \end{aligned} \tag{4.72}$$

where,

$$R_3(t, t_{i_0}, t_{i_0+1}) = N \ln \frac{S(t)}{S(t_{i_0})} - \left(\alpha - \frac{1}{2}\sigma^2\right)\left(T - t - N\hat{\Delta}\right) - \ln \frac{S(t)}{S(0)}. \tag{4.73}$$

Noting that $N\Delta = T$, the total term of the volswap from writing time until expiration, we see that the gamma is inversley proportional to the term of the contract. As with the exact and approximate gammas at time zero, the gammas at any fixing time are inversely proportional to the term of the swap. In fact, examination of equations (4.51) and (4.45) shows this to be true for general time t.

Numerical Examples

The following figures illustrate the values of the greeks for a volatility swap with four fixes already taken. In all cases the following inputs were used.

$$spot = 97.00, \tag{4.74}$$

$$
\begin{aligned}
T &= \frac{21}{252}, \\
t &= \frac{17}{252}, \\
\Delta &= \frac{1}{252}, \\
\hat{\Delta} &= \frac{1}{252}, \\
\sigma &= 0.14, \\
r_d &= 0.005, \\
r_f &= 0.035, \\
N &= 21, \\
i_0 &= 4.
\end{aligned}
$$

It may be seen from figures (4.6), (4.7), (4.8) and (4.9) that the approximate and exact greeks for the volatility swap correspond closely, and that the greeks are also qualitatively similar to that of a plain vanilla contingent claim. As with a long position in a European call or put option, a long position in a volatility swap results in the owner's portfolio being long gamma and long vega. The difference comes in the magnitude of these sensitivities and the property that the volatility swap's delta resets at a fixing time to a value near to zero. Consider the volswap described in our example immediately before the fourth fixing time. We assume that the spot price is currently 97 and the previous three spot fixings were 100, 100.30 and 99.20. Here $\hat{\Delta}$ is almost zero, so our expected fourth fixing is 97. The value of such a swap with a notional amount of 1 unit of the foreign currency and zero strike volatility is 20.435. The delta of the swap just before the fixing is taken is 135.135.

We now revalue the swap immediately after the fixing has been taken and assume that this fixing is 97.00. We find that the volswap has exactly the same value after the fix as it had immediately before. Tedious calculation using equation (4.39) shows that this is true for a general fixing. Evaluation of the delta shows that it is no longer large and positive, but small and negative. The delta immediately after the fixing has been taken is -9.823. The act of taking the fix has reset the delta close to zero. The hedger of a Volatiltiy Swap must then delta-hedge at the same time as the fixing is taken to remain delta-neutral. Failure to do so results in a large discrepancy in deltas, particularly when using plain vanilla options to hedge the volswap. A hedge executed at a spot rate different to the fixing price also results in mis-hedging. In our example, if we assume the contract was written on a notional

of 100,000 USD per volatility percentage point, the hedger must sell 14,495,800 USD at a spot rate of 97.00 to replicate the volswap. If the dollars were sold at 96.80 rather than 97.00 a loss of 29,888 USD would result. In volatility terms this is 0.299%, a significant sum. Delta hedging at a rate equal to the fixing is crucial.

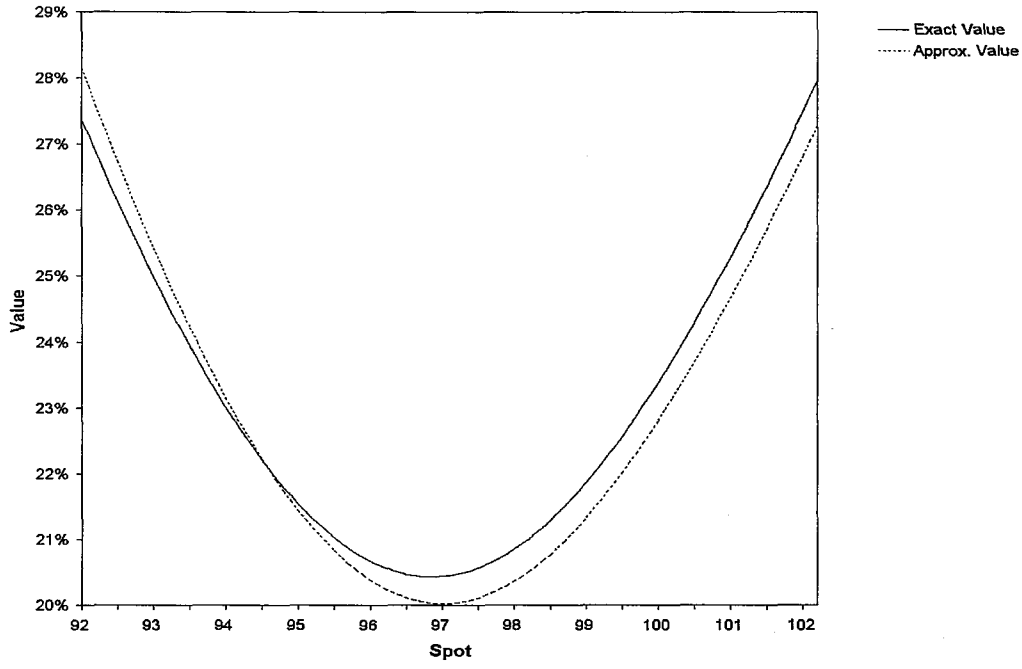


Figure 4.6: Values for the Volatility Swap.

Hedging in Practice

We have derived the approximate price of the volatility swap at any time, t, during its life as well as its delta, gamma and vega. Thus at any time, we are able to value the swap and know how the value of the swap will change given a range of circumstances. We may readjust our hedges to create a delta, gamma and vega neutral position at any time. Delta hedging presents us with no problem. In the foreign exchange market, the underlying asset is typically liquid and transaction costs associated with trading minimal. The rebalancing of the gamma and vega positions present us with additional problems.

The simplest way that a market maker may chose to attempt to replicate the payoff of a volswap may be through delta hedging an at-the-money-forward straddle of expiration T, and vega $A(t)\sigma$ on the fixing points $t_0, t_1, ..., t_{N-1}, t_N$. At t_0 this results in a vega, delta

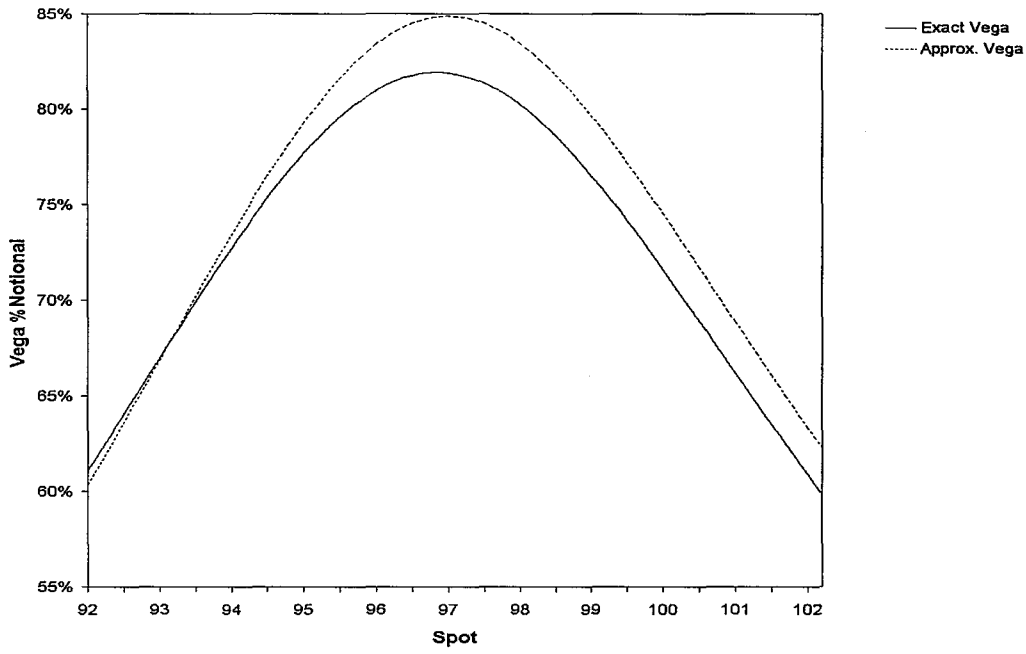


Figure 4.7: Vega for the Volatility Swap.

and gamma neutral position. However, as equation (4.1) shows, the ultimate payoff will not be that of the volswap. It should also be noted that, as the expiration time nears, the value of spot will effect the gamma of the straddle enormously. If we are very near to the strike of the straddle just before expiration then its gamma will be very large, much larger than that of the volswap which has no spot strike. If spot is very far away from the strike of the straddle then the resulting gamma will be near to zero. The same is not true for the volswap so a mismatch will arise. Clearly, from two perspectives, what appears to be a good initial hedge quickly turns into a bad one.

We may choose to use forward starting at-the-money straddles to hedge the volswap. This is only an approximation to the value of the volswap as we have already shown. Although useful in aiding our intuition about these contracts, it will not provide a good practical hedge for them. Regardless of the effectiveness of such a hedge, these forward starting options are not liquid products and so trading, or replicating, them would be prohibitively expensive.

The vega of the contract has been shown to be linear in time, so long as the effective volatility, $\sqrt{\frac{N}{N-1}}H(t)$, is close to the implied volatility σ. When these deviate the vega may be very different to that which we were expecting at that time given that the dependence

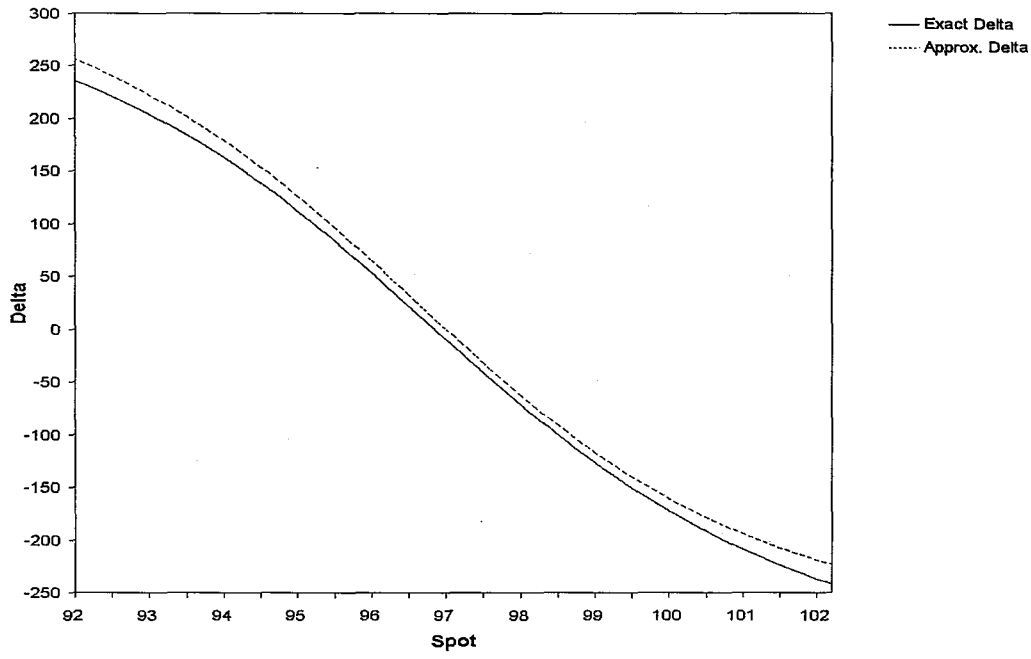


Figure 4.8: Delta for the Volatility Swap.

on σ was assumed to decay linearly in time. This dependence on the realised volatility (not variance as with the variance swap) presents us with a dynamic vega hedging problem. This need to rebalance the vega hedge, and the significant transaction costs involved in trading options on the underlying asset, should not be ignored when initially calculating the fair volatility, σ_K at which the volswap contract is struck to make it zero cost.

Considering the problem of gamma hedging leads us to similar conclusions. We saw that the gamma of a volswap remains constant as long as the effective and implied volatilities do not deviate from one another. In the case where there is a large difference, there will be the need to transact plain vanilla options on the underlying security in order to maintain a gamma neutral position. Again, this rebalancing involves significant transaction costs and should not be ignored in practice.

4.5 Conclusions

Volatility trading is a concept that has been understood by marketmakers in options for a long while. However, the instruments to enable speculators to participate easily in this market have only recently become available and the literature on this subject limited. We

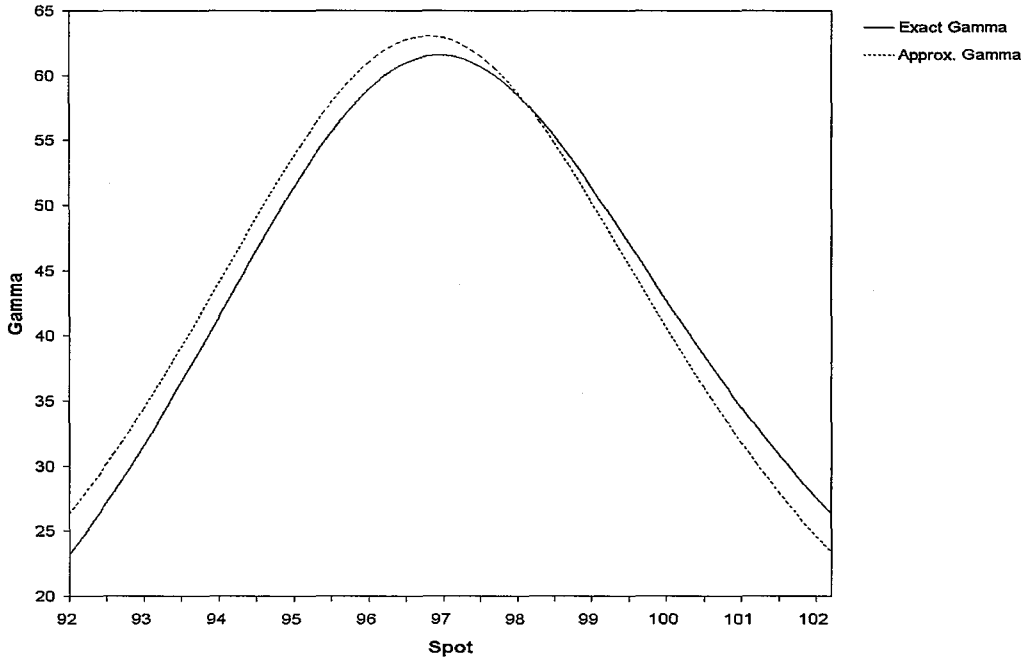


Figure 4.9: Gamma for the Volatility Swap.

have examined three particular areas of volatility trading. We began by taking the most simple approach to trading volatility, that being vanilla options. These have a dependence on implied volatility, or vega, which plain vanilla option traders have long been using to trade option books. The payoff from continuously delta-hedged plain vanilla options was derived.

A further step was then taken. Existing literature deals almost exclusively with contracts known as Variance Swaps. This literature was reviewed and the connections between these contracts and the most simple form of volatility trading, the delta-hedged option, explained. Two different schemes for replicating the Variance Swap contract were implemented and compared, both in terms of performance and practicality. The two schemes were implemented under a constant implied volatility and a more general volatility smile.

Finally, we described a different volatility contract known as a Volatility Swap. This contract has proved to be the most popular amongst market participants, particularly speculators. We derived an approximate value for this contract under the definition of realised volatility used in practice. The volatility was not found on a continuous basis, but by the use of an estimate based upon discrete time fixings. An intuitive way of viewing this new contract was introduced that allowed us to draw some elementary conclusions about the

nature of the contract and how its greeks may look. A value was found at a general time, t, by decomposing its payoff into several more easily handled components, each with their own separate contribution to the volswap. We were then able to derive greeks for this contract in order to define a hedging strategy. Exact values for delta, gamma and vega were found, but approximate values for these were also presented in order to help with the analysis of the risk characteristics of the Volatility Swap. The greeks were analysed under particular conditions and compared to those results that our intuitive approach led us to expect. Finally full numerical examples and values were given for both the exact and approximate greeks and the values compared.

Chapter 5

Conclusions

The differences that exist between the perfect markets assumed under the Black-Scholes framework and the financial markets create difficulties for practitioners in terms of pricing and hedging derivative products. At the same time, the inconsistencies between traditional option pricing theory and practice have led to the development of new and innovative products to enable participants to exploit these differences. This thesis investigated the effect that market imperfections have on contingent claim prices and the way that practitioners attempt to make allowances for these effects.

We began, in Chapter 2, by looking at one specific violation of the Black-Scholes assumption. Using a method used by Davis, Panas and Zariphopoulou [16] for the pricing of contingent claims under proportional transaction costs, we found the prices of plain vanilla options when the market for the underlying asset is not liquid. A utility maximisation method was used to find the writing prices of options under different types of illiquidity conditions. Trading strategies under a specific utility function were found. It was found that, under illiquid conditions, we may not work under the risk-neutral measure and require the real drift of the asset to be known. The level of drift and the risk-averseness of the writer was found to play a large role in determining the price of the option. The writing and purchasing prices of an option do not coincide for these same reasons.

In reality, options markets attempt to cope with the shortcomings of Black-Scholes and its assumptions by adjusting the volatility used to price options of different strikes and maturities. The first part of Chapter 3 introduced the notion of a volatility smile or surface, then

reviewed and implemented the work of Derman, Kani and Chriss [21] to grow a trinomial tree consistent with market prices of plain vanilla options. This model was then extended to price foreign exchange barrier options effectively. In the second part of Chapter 3 we continued our focus on exotic options and the volatility smile by looking at static replication. By extending the methods of Derman, Ergener and Kani [22] we found portfolios of vanilla options that replicate the payoffs for several non-standard products, including barrier options. This work was then generalised so that a time dependent volatility function may be used when pricing non-standard exotic foreign exchange products.

With the growth of an over-the-counter options market with deep liquidity, has come the existence of a volatility market. The presence of such a market has led to demand for pure volatility products allowing speculators to take positions in the level of future volatility. Chapter 4 began by discussing the dependence of plain vanilla option products on the level of volatility used when determining their value. The variation of this dependence, known as vega, with changes in other parameters such as spot price and time to maturity was shown. We went on to describe the most simple form of volatility trading, the delta hedged vanilla option. We derived the profit and loss from such a strategy and showed that this may not be satisfactory for the volatility speculator as it has a dependence on the gamma of the option being hedged. A new product known as a Variance Swap was then introduced. We reviewed previous work on this product and discussed the connection between Variance Swaps and another exotic product known as Log Contract. Two different methods for replicating the Variance Swap were described and implemented. Their relative performance was investigated and we found that the method of Demeterfi, Derman, Kamal and Zou [18] gave more accurate replication with far fewer options needed than the method of Carr and Madan [12]. Both methods were implemented under a volatility smile and the effects found. It was found that the use of a smile where out-of-the-money options have higher implied volatility than at-the-money-options results in larger values for the fair variance than that we would find using only at-the-money volatility for our estimate.

Finally, we took the idea of volatility trading a step further by introducing the Volatility Swap. The Volatility Swap is a new product designed specifically with the speculator in mind and pays the difference between a predetermined strike value and an estimate of the realised volatility over the life of the contract. We began by demonstrating that it is intuitive to think of the volswap contract as being bounded below by a strip of forward starting

straddles each of maturity Δ, the time between each spot price fixing taken for the estimate of volatility. We used this to investigate the behaviour of the volswap, finding that at a fixing point the delta of such a contract resets to near zero. We then found that the volswap value may be split into several distinct pieces, each with its own contribution and that each piece had an intuitive description. This re-expression of the contract allowed us to find an approximate value for the volswap at a general point in time. Using this, and another approximation, we were able to derive the greeks for this contract. The three approaches used were examined under particular conditions and they were found to yield similar answers, each approximation providing more accurate results than the last. Volswaps, it was found, have zero delta at the first spot fixing time and this delta resets to near zero value at each subsequent fixing time. Vega is dependent upon both time to expiration and, importantly, the ratio between implied and effective volatility. Volswaps were also found to exhibit large values for gamma particularly when the time to expiration is small. This gamma is a major factor in effective hedging of the swap contract and remains reasonably constant as long as implied and effective volatilities are similar.

Through the work carried out in this thesis several interesting areas for further research have come to light. The utility maximisation method used in Chapter 2 to price european style options on a non-dividend paying stock could be used to price other types of products under incomplete market conditions. American style options as well as the barrier type products discussed throughout this thesis provides a extensive area of application for this methodology. In Chapter 3, where stock price models for pricing exotic options were implied from plain vanilla option prices, work still remains to be carried out on the efficient implementation of these methods. We have contributed towards the rapid convergence of results for certain types of options priced using this model, however the construction of the tree is initially quite slow. Further work carried out on this topic would be valuable. Finally, Chapter 4 dealt with trading volatility as an asset in itself. The most popular of the volatility products is the volatility swap and it is in this area where most work remains to be done. We used a constant volatility model to price and hedge the volaitlity swap contract. An imporatant extension to our work is to consider the volswap contract within a stochastic volatility model. Results from this may be comapared with the work presented here.

Glossary

ATM: at-the-money option. An option with strike price the same as the current spot price.

Barrier Option: option contract, normally with the same terminal payoff as a plain vanilla option, but which, in the case of an Out Barrier, is cancelled or, in the case of an In Barrier, comes alive if a predetermined barrier level is reached by the price of the underlying asset at any time during the life of the contract.

Butterfly: options strategy consisting of buying a strangle and selling a straddle with the same expiration.

Call Option: contract giving the owner the right, but not the obligation, to buy an asset at a set price at a set time in the future.

Call Spread: option strategy consisting of buying a call option with a lower strike price and selling a call with a higher strike price and the same expiration.

Delta: first derivative of an option's price with respect to the underlying asset's spot price.

Digital Option: option contract paying the owner a fixed amount if it is in-the-money at expiration and zero otherwise.

Exotic: non-standard option contract with conditions which differ from a standard, plain vanilla, option.

Gamma: second derivative of an option's price with respect to the underlying asset's spot price.

Greeks: collective name for the Delta, Gamma and Vega of an option.

Implied Volatility: volatility of an asset implied by the price for a plain vanilla option.

ITM: in-the-money option. An option with strike price far from the current spot price and which will be exercised at expitration if the current spot price does not change.

Log Contract: european style contract which pays the owner the natural logarithm of the underlying spot price at expiration.

OTM: out-of-the-money option. An option with strike price far from the current spot price and which will not be exercised at expiration if the current spot price does not change.

Plain Vanilla Option: european style put or call option.

Power Option: european style put or call option which pays the square of the payoff from a plain vanilla.

Put Option: contract giving the owner the right, but not the obligation, to sell a certain asset at a certain price at a set time in the future.

Put Spread: option strategy consisting of buying a put option with a higher strike price and selling a put with a lower strike price and the same expiration.

Risk Reversal: option strategy consisting of buying a call option and selling a put option with the same expiration.

Straddle: option strategy consisting of buying a call and a put with the same strike and expiration.

Strangle: option strategy consisting of buying a call and a put with out-of-the-money strikes and the same expiration.

UIC: Up and In Call option. A barrier call option with a knock-in level above the current market.

UOC: Up and Out Call option. A barrier call option with a knock-out level above the current market.

UIP: Up and In Put option. A barrier put option with a knock-in level above the current market.

UOP: Up and Out Put option. A barrier put option with a knock-out level above the current market.

Variance Swap: contract which pays the owner the difference between the realised variance of the underlying asset's spot price over the term of the contract and a predetermined strike price.

Vega: first derivative of an option's price with respect to volatility.

Volatility Swap: contract which pays the owner the difference between the realised volatility of the underlying asset's spot price over the term of the contract and a predetermined strike price.

Volatility Smile: difference observed in the options market, for the the volatility impled for an asset by options with the same expiration but different strike prices.

Wing: deep out-of-the-money option. An option with a delta close to zero.

Bibliography

[1] Abramowitz, M. and Stegun, I.E.,"Handbook of Mathematical Functions," 1^{st} Ed. Dover 1970.

[2] Barle S. and N. Cakici, "Growing a Smiling Tree," *Risk Magazine*, Vol. 8, October 1995.

[3] Bensaid, B., J. P. Lesne, H. Pagès, and J. Scheinkman, "Derivative Asset Pricing with Transaction Costs," *Mathematical Finance*, Vol. 2, pp. 63-86, 1992.

[4] Benson, R., and N. Daniel, "Up, Over and Out," *Risk*, Vol. 4, pp. 17-19, June 1991.

[5] Black, F., and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, pp. 637-654, 1973.

[6] Bowie J. and P. Carr, "Static Simplicity," *Risk Magazine*, Vol. 7, August 1994.

[7] Boyle, P. P., and S. Lau, "Bumping Up Against the Barrier with the Binomial Method," *Journal of Derivatives*, 1, 4, pp. 6-14, 1994.

[8] Boyle, P. P., and K. S. Tan, "Lure of the Linear," *Risk*, Vol. 7, pp. 43-46, 1994.

[9] Boyle, P. P., and T. Vorst, "Option Replication in Discrete Time with Transaction Costs," *Journal of Finance*, Vol. 47, pp. 271-293, 1992.

[10] Breeden D. and R. Litzenberger, "Prices of State Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, pp.621-651, 1978.

[11] Calderón, P., and T. Temel, "Pricing of Barrier Options with Term Structure of Interest Rates and Volatilities," Lehman Brothers working paper, 1993.

[12] Carr P. and D. Madan, "Towards a Theory of Volatility Trading," *Working Paper*, December 1997.

[13] Carr P. and D. Madan, "Currency Covariance Contracting," *Working Paper*, March 1999.

[14] Cox, J. C., S. A. Ross, and M. Rubinstein, "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, Vol. 7, pp. 229-263, 1979.

[15] Cox, J., and M. Rubinstein, *Options Markets*, Prentice-Hall, 1985.

[16] Davis M.H.A., V.G. Panas and T. Zariphopoulou, "European Option Pricing with Transaction Costs," *SIAM J. Con. & Opt.*, Vol. 31, pp.470-493, 1993.

[17] Davis, M.H.A., and V.G. Panas, "The Writing of a European Contingent Claim Under Proportional Transaction Costs," *Comp. Appl. Math.*, Vol.13, pp.115-157, 1994.

[18] Demeterfi K., E. Derman, M. Kamal, J. Zou, "More Than You Ever Wanted to Know About Volatility Swaps," *Goldman Sachs Quantitative Strategies Research Notes*, March 1999.

[19] Derman, E., and I. Kani, "The Ins and Outs of Barrier Options," Goldman Sachs working paper, June 1993.

[20] Derman E. and I. Kani, "The Volatility Smile and its Implied Tree," *Quantitative Strategies Research Notes, Goldman Sachs*, January 1994.

[21] Derman E., I. Kani and N. Chriss, "Implied Trinomial Trees of the Volatility Smile," *Quantitative Strategies Research Notes, Goldman Sachs*, February 1996.

[22] Derman E., D. Ergener and I. Kani, "Static Options Replication," *Quantitative Strategies Research Notes, Goldman Sachs*, May 1994.

[23] Dothan, M. U., *Prices in Financial Markets*, Oxford University Press, 1990.

[24] Duffie, D., *Dynamic Asset Pricing Theory*, Princeton, 1992.

[25] Dupire B., "Pricing and Hedging with Smiles," *Paribas Capital Markets, Swaps and Options Research Team.*, April 1993.

[26] Grannan E.R., and G.H. Swindle, "Minimising Transaction Costs of Option Hedging Strategies."

[27] Green R. and R.A. Jarrow, "Spanning and Completeness in markets with Contingent Claims," *Journal of Economic Theory*, 41, pp.202-210, 1987.

[28] Hull J.C., "Options, Futures and other Derivative Securities," 2^{nd} Ed. Prentice Hall, 1993.

[29] Heynen, R., and H. Kat, "Crossing Barriers," *Risk*, Vol. 7, No. 6, pp. 46-51, June 1994.

[30] Hodges S.D., and A. Neuberger, "Optimal Replication of Contingent Claims Under Transaction Costs," *Rev. Futures Markets.* Vol. 8, pp.222-239, 1989.

[31] Hoggard, T., A. E. Whalley, and P. Wilmott, "Hedging Option Portfolios in the Presence of Transaction Costs," *Advances in Futures and Options Research*, Vol. 7, pp. 21-35, 1994.

[32] Hudson, M., "The Value in Going Out," *Risk*, Vol. 4, pp. 29-33, March 1991.

[33] Hull, J., *Options, Futures, and Other Derivatives*, 3^{rd} Ed., Prentice-Hall, 1997.

[34] Karatzas I., and S.E. Shreve, "Brownian Motion and Stochastic Calculus," 2^{nd} Ed., Springer Verlag, 1991.

[35] Kunitomo, N., and M. Ikeda, "Pricing Options with Curved Boundaries," *Mathematical Finance*, Vol. 2, pp. 275-298, 1992.

[36] Leland, H. E., "Option Pricing and Replication with Transactions Costs," *Journal of Finance*, Vol. 40, pp. 1283-1301, 1985.

[37] Merton R. C., "Continuous Time Finance," (revised ed.) Blackwell, 1992

[38] Merton, R. C., "Option Pricing when Underlying Stock Returns are Discontinuous," *Journal of Financial Economics*, Vol. 3, pp. 124-144, 1976.

[39] Mood, A. M., F. A. Graybill and D. C. Boes, "Introduction to the Theory of Statistics," 3^{rd} Ed., McGraw-Hill International Editions, 1974.

[40] Neuberger A., "The Log Contract: A new instrument to hedge volatility," *Journal of Portfolio Management*, pp.74-80, Winter 1994.

[41] Oksendal B., "Stochastic Differential Equations," 4^{th} Ed., Springer Verlag, 1995.

[42] Rubinstein M., "Implied Binomial Trees," *Journal of Finance*, 69, pp.771-818, 1994.

[43] Rubinstein, M., and E. Reiner, "Breaking Down the Barriers," *Risk*, Vol. 4, pp. 28-35, September 1991.

[44] Wayne D. and V. Naik, "Pricing and Hedging Volatility Swaps," *Lehman Borthers Technical Note*, September 1998.

[45] Wilmott, P., J. N. Dewynne, and S. Howison, *Option Pricing: Mathematical Models and Computation*, Oxford Financial Press, 1993.

Appendix A

Arrow-Debreu Prices

The notion of an Arrow-Debreu price is something that may not be familiar. All that this name refers to is the discounted path probability of reaching a certain node. The definition is best made clear through an example.

Let us consider a binomial tree for simplicity. Assume a value of 0.1 for the risk-free rate. $\Delta t = 1$. Assign to this tree an up probability of $\frac{3}{5}$ and a down probability of $\frac{2}{5}$. We may then calculate the Arrow-Debreu prices as follows.

For a tree such as in figure (A.1) where probabilities are known for every node the Arrow-Debreu price may be found at every point. In the case of non-constant volatilities the Arrow-Debreu prices must be calculated recursively and implied up a level as the tree is built. The calculations are carried out in this recursive manner below. We require a starting point. This start is provided at the root node where no time has elapsed. Thus the probability of reaching this node is 1. The probability of one is then discounted over zero time at the riskfree rate giving us a starting value of 1 for the Arrow-Debreu prices.

Moving up one time level we have the up node firstly. Here the A-D price is given by,

$$\lambda_{11} = \lambda_{00} \times p \times e^{-r\Delta t} = 1.0 \times 0.6 \times e^{-0.1 \times 1.0} = 0.543.$$

The lower A-D price, λ_{10}, is found as above with q replacing p. All we have done is taken the probability of getting to that node and discounted it at the risk-free rate over the period of time elapsed. We will now select certain nodes that illustrate our method.

Take, for example, the calculation of the A-D price at node (2,2). We could do this in two

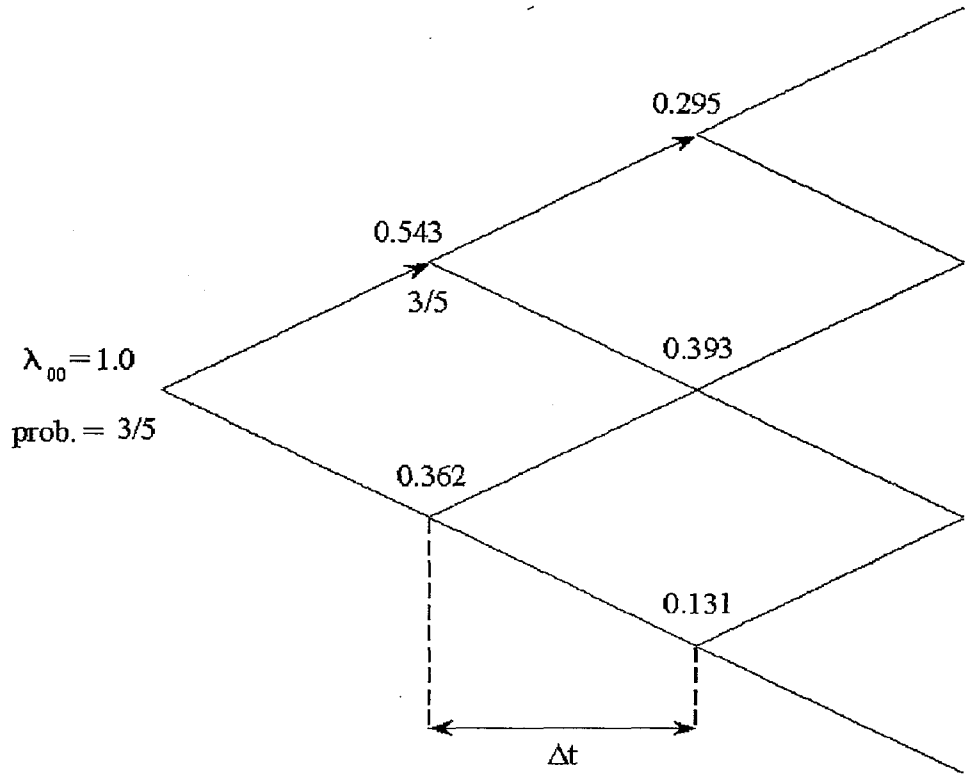


Figure A.1: Arrow-Debreu prices.

ways. Firstly by recursion. We take the A-D price at (0,1) denoted λ_{11} and multiply it by the probability of getting up to point (2,2). The result is then discounted over the time step Δt.

$$\lambda_{22} = \lambda_{11} p e^{-r\Delta t} = 0.295.$$

An alternative way of seeing things is the following. There is only one path leading from the root node at (0,0) to the node (2,2). The probability of this path being taken is $p \times p$ and the time elapsed is $2\Delta t$. The results is as follows,

$$\lambda_{22} = p^2 e^{-2r\Delta t} = 0.295.$$

Notice that $\lambda_{11} = pe^{-r\Delta t}$ so the two ways of finding λ_{22} are precisely the same.

Finally we shall find the value of λ_{21}. In this case there are two possible paths leading to the node (2,1) which have probabilities attached,

$$\lambda_{21} = (p\lambda_{10} + q\lambda_{11})e^{-r\Delta t} = 0.395.$$

We may carry on in a similar way as long as the path probabilities are known.

Appendix B

Forward Starting Options

Consider a plain vanilla option which is paid for now, but does not have its strike price set until some time in the future. The option is paid for at time t_0, struck at-the-money at time t_1, and expires at time T. Using similar notation to before, where S_t denotes the spot price at time t and σ denotes the volatility of the underlying asset, we derive the value of this option. Without loss of generality we consider a call option and assume that all interest rates are zero.

From the Black-Scholes formula for a call option struck at-the-money we note that the value of the option is proportional to the spot price. Therefore the value of a call option at time t_1 is $\frac{CS_{t_1}}{S_{t_0}}$, where C denotes the value, at time t_0, of an at-the-money option that expires in time $T - t_1$. Using risk-neutral valuation, the value of the forward starting option at time t_0 is

$$\mathbb{E}^q\left[\frac{CS_{t_1}}{S_{t_0}}\right], \tag{B.1}$$

where $\mathbb{E}^q$ denotes expectation in a risk-neutral world. It follows that the price of the forward starting option is C as S_{t_0} and C are known at time t_0 and the expectation of S_{t_1} at time t_0 is S_{t_0}.

From the valuation of the forward starting option it follows that the delta of such a contract is zero, as is the gamma. The vega on the other hand, is simply equal to the vega of the at-the-money option C.

Appendix C

Derivation of Various Expectations and Greeks

C.1 Proof of equation (4.42) for $\mathbb{E}_t^Q[\sqrt{\sigma_R^2}]$:

Consider a random variable X which has a chi-squared distribution with ν degrees of freedom, i.e. $X \sim \chi^2(\nu)$. Its probability density function is given by[1]

$$f(x) = \left(\frac{1}{2}\right)^{\frac{\nu}{2}} \frac{1}{\Gamma\left(\frac{\nu}{2}\right)} x^{\frac{\nu}{2}-1} e^{-\frac{1}{2}x}, \; x > 0. \tag{C.1}$$

Note that for a χ^2-distributed random variable X,

$$\mathbb{E}[X] = \nu, \text{ and } \mathrm{Var}[X] = 2\nu. \tag{C.2}$$

We are interested in calculating the following expectation:

$$\mathbb{E}\left[\sqrt{X}\right] = \frac{2^{-\frac{\nu}{2}}}{\Gamma\left(\frac{\nu}{2}\right)} \int_0^\infty dx \; x^{\frac{\nu-1}{2}} e^{-\frac{1}{2}x} = \frac{2^{-\frac{\nu}{2}}}{\Gamma\left(\frac{\nu}{2}\right)} I, \tag{C.3}$$

in which the density given in (C.1) is being used to evaluate the expectation. Substituting $x = 2y$, the integral I can be rewritten as

$$I = 2^{\frac{\nu+1}{2}} \int_0^\infty dy \; y^{\frac{\nu-1}{2}} \; e^{-y},$$

and putting $q = \frac{\nu+1}{2}$, this can be expressed as

$$I = 2^{\frac{\nu+1}{2}} \int_0^\infty dy \; y^{q-1} e^{-y}. \tag{C.4}$$

[1] See, for example, Mood et al. [39], page 241-242.

However, the integral in (C.4) can easily be recognised as that defining the gamma function[2] $\Gamma(q)$, so that on substituting for q in this function and back substituting for I in (C.3), the following result is obtained:

$$\mathbb{E}\left[\sqrt{X}\right] = \sqrt{2}\,\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)}. \tag{C.5}$$

A well-known result of the geometric Brownian motion assumption being used in the Black-Scholes model to price the volatility swaps is that the changes, $\ln S(t_{i+1}) - \ln S(t_i)$, are normally distributed with means $\mu_i = (r_d - r_f - \frac{1}{2}\sigma^2)\,\Delta$ and variances $\sigma_i^2 = \sigma^2\Delta$, where $\Delta = (t_{i+1} - t_i)$. Furthermore, for non-overlapping time intervals, these changes are pairwise independent. The sum of the squares of N independent random variables X_i, each having a standard normal distribution, i.e. $X_i \sim N(0,1)$, is known to possess a chi-squared distribution[3]. Therefore, the random variable X defined by

$$X \stackrel{\Delta}{=} \sum_{i=1}^{N}\left(\frac{X_i - \mu_i}{\sigma_i}\right)^2, \tag{C.6}$$

has a chi-squared distribution with N degrees of freedom. Recalling that σ_R^2 can be written as

$$\sigma_R^2 \times \frac{(N-1)}{\sigma^2} = \sum_{i=0}^{N-1}\left(\frac{\ln S(t_{i+1}) - \ln S(t_i) - \mu_i}{\sigma_i}\right)^2, \tag{C.7}$$

we get the result that $\sigma_R^2 \times \frac{(N-1)}{\sigma^2} \sim \chi^2(N)$. Using (C.5) with $\nu = N-1$, the following required result is finally obtained:

$$\mathbb{E}[\sqrt{\sigma_R^2}] = \frac{\sigma}{\sqrt{N-1}}\,\mathbb{E}\left[\sqrt{\left(\sigma_R^2\frac{(N-1)}{\sigma^2}\right)}\right] = \sigma\sqrt{\frac{2}{N-1}}\,\frac{\Gamma\left(\frac{N}{2}\right)}{\Gamma\left(\frac{N-1}{2}\right)}. \tag{C.8}$$

C.2 Derivation of b_1, b_2, b_3 and b_4 used in section (4.4.3):

From section (4.4.3), we have that

$$b1 = \mathbb{E}_t^Q[s_1^2] = \frac{i_0}{N}\,\mathbb{E}_t^Q\left[v_H^2\right], \tag{C.9}$$

where

$$v_H^2 = \sum_{i=0}^{i_0-1}\frac{[\ln S(t_{i+1}) - \ln S(t_i)]^2}{i_0\Delta}, \tag{C.10}$$

$$\Delta = t_{i+1} - t_i, \text{ is the equidistant time interval between fixings.} \tag{C.11}$$

[2]See, for example, Abramowitz and Stegun [1] equation(6.1.1) page 255.

[3]See, for example, Mood et al. [39] theorem 7, page 242.

However, as the expectation in (C.9) is taken with respect to time t, all the terms in v_H^2 are known and can therefore be taken out of the expectation. We then have that

$$b1 = \frac{i_0}{N} v_H^2. \tag{C.12}$$

b_2 is defined in a similar manner to b_1, namely

$$b_2 = \mathbb{E}_t^Q[s_2^2], \tag{C.13}$$

where

$$s_2^2 = \frac{[\ln S(t_{i_0+1}) - \ln S(t_{i_0})]^2}{N\Delta}. \tag{C.14}$$

Writing

$$\begin{aligned} \ln S(t_{i_0+1}) - \ln S(t_i) &= \ln\left[\frac{S(t_{i_0+1})}{S(t_{i_0})}\right] = \ln\left[\frac{S(t_{i_0+1})}{S(t)} \times \frac{S(t)}{S(t_{i_0})}\right], \\ &= \ln\left[\frac{S(t_{i_0+1})}{S(t)}\right] + \ln\left[\frac{S(t)}{S(t_{i_0})}\right], \end{aligned}$$

we then have that

$$\begin{aligned} b_2 N\Delta &= \mathbb{E}_t^Q\left[\left(\ln\left[\frac{S(t_{i_0+1})}{S(t)}\right] + \ln\left[\frac{S(t)}{S(t_{i_0})}\right]\right)^2\right], \\ &= \mathbb{E}_t^Q\left[\left(\frac{\ln S(t_{i_0+1})}{S(t)}\right)^2 + \left(\ln\frac{S(t)}{S(t_{i_0})}\right)^2 + 2\ln\left(\frac{S(t_{i_0+1})}{S(t)}\right)\ln\left(\frac{S(t)}{S(t_{i_0})}\right)\right]. \end{aligned}$$

Since we have that $t_{i_0} \leq t \leq t_{i_0+1}$ and the expectation is taken with respect to time t, $S(t_{i_0})$ and $S(t)$ are known quantities and play no part in the expectation. We then have that

$$\begin{aligned} b_2 N\Delta &= \left(\ln\frac{S(t)}{S(t_{i_0})}\right)^2 + \mathbb{E}_t^Q\left[\left(\ln\frac{S(t_{i_0+1})}{S(t)}\right)^2\right] \\ &+ 2\ln\left(\frac{S(t)}{S(t_{i_0})}\right)\mathbb{E}_t^Q\left[\ln\left(\frac{S(t_{i_0+1})}{S(t)}\right)\right]. \end{aligned} \tag{C.15}$$

For a generic time T, consider the following two expectations:

$$\begin{aligned} E_1 &= \mathbb{E}_t^Q\left[\ln\left(\frac{S(T)}{S(t)}\right)\right], \text{ and} \\ E_2 &= \mathbb{E}_t^Q\left[\left(\ln\left[\frac{S(T)}{S(t)}\right]\right)^2\right]. \end{aligned}$$

From our stochastic differential equation for $S(t)$, one has

$$\ln S(T) = \ln S(t) + \left(r_d - r_f - \frac{1}{2}\sigma_{tT}^2\right)(T-t) + \int_t^T \sigma_{tT} dz_Q(u), \tag{C.16}$$

in which we have placed explicit subscripts on the diffusion coefficient. Then, using the result that[4]

$$\mathbb{E}_t^Q \left[\int_t^T f(u,\omega)\ dZ_Q(u) \right] = 0, \tag{C.17}$$

we obtain

$$E_1 = \left(r_d - r_f - \frac{1}{2}\sigma_{tT}^2 \right) (T-t). \tag{C.18}$$

Substituting for $\ln S(T) - \ln S(t)$ from (C.16) in (C.16) and expanding the square, we then have

$$\begin{aligned} E_2 &= 2\left(r_d - r_f - \frac{1}{2}\sigma_{tT}^2 \right)(T-t)\,\mathbb{E}_t^Q \left[\int_t^T \sigma_{tT} dz_Q(u) \right] \\ &+ \left(r_d - r_f - \frac{1}{2}\sigma_{tT}^2 \right)^2 (T-t)^2 + \mathbb{E}_t^Q \left[\left(\int_t^T \sigma_{tT} dz_Q(u) \right)^2 \right]. \end{aligned} \tag{C.19}$$

The first tem on the RHS of equation (C.19) vanishes as a result of equation (C.17) whereas the second expectation can be calculated using the Itô isometry[5], namely

$$\mathbb{E}_t^Q \left[\left(\int_t^T f(u)\ dZ_Q(u) \right)^2 \right] = \mathbb{E}_t^Q \left[\int_t^T f^2(\omega,u) du \right], \tag{C.20}$$

to yield

$$\mathbb{E}_t^Q \left[\left(\int_t^T \sigma_{tT} dz_Q(u) \right)^2 \right] = \sigma_{tT}^2 (T-t). \tag{C.21}$$

Substituting these results back into the equation for E_2 leads to

$$E_2 = \left(r_d - r_f - \frac{1}{2}\sigma_{tT}^2 \right)^2 (T-t)^2 + \sigma_{tT}^2 (T-t). \tag{C.22}$$

Finally, considering explicitly the case where $T \to t_{i_0+1}$, we have

$$\mathbb{E}_t^Q \left[\ln\left(\frac{S(t_{i_0+1})}{S(t)} \right) \right] = \left(\alpha - \frac{1}{2}\sigma_G^2 \right) \hat{\Delta}, \tag{C.23}$$

$$\mathbb{E}_t^Q \left[\left(\ln \frac{S(t_{i_0+1})}{S(t)} \right)^2 \right] = \left(\alpha - \frac{1}{2}\sigma_G^2 \right)^2 \hat{\Delta}^2 + \sigma_G^2 \hat{\Delta}, \tag{C.24}$$

where

$$\begin{aligned} \alpha &= r_d - r_f, \\ \hat{\Delta} &= (t_{i_0+1} - t) \text{ is the time until the next fixing date,} \\ \sigma_G &= \text{the implied volatility for the interval } [t, t_{i_0+1}]. \end{aligned} \tag{C.25}$$

[4] See, for example, Oksendal [41] theorem 3.7(iii) p22.

[5] See, for example, Oksendal [41] equation 3.13, p21.

The choice of value for the volatility is the one unknown for the Black-Scholes model. The subscripts tT placed on the diffusion coefficient reflect the fact that, in practice, a constant volatility is not used and different choices of σ will be made depending on the time t and the time span considered. The best market guess for the volatility is then simply the implied volatility at time t for the period in question.

Substituting equations (C.23) and (C.24) back into equation (C.15), one obtains

$$\begin{aligned} b_2 N\Delta &= \left(\ln\frac{S(t)}{S(t_{i_0})}\right)^2 + \left(\alpha - \frac{1}{2}\sigma_G^2\right)^2 \hat{\Delta}^2 + \sigma_G^2 \hat{\Delta} \\ &+ 2\ln\left(\frac{S(t)}{S(t_{i_0})}\right)\left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta}, \end{aligned}$$

which can be simplified to give

$$b_2 = \frac{1}{N\Delta}\left[\left(\ln\frac{S(t)}{S(t_{i_0})} + \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta}\right)^2 + \sigma_G^2\hat{\Delta}\right]. \tag{C.26}$$

b_3 is given by

$$b_3 = \mathbb{E}_t^Q\left[s_3^2\right], \tag{C.27}$$

where

$$s_3^2 = \sum_{i=i_0+1}^{N-1} \frac{\left[\ln S(t_{i+1}) - \ln S(t_i)\right]^2}{N\Delta} = \frac{1}{N\Delta}\sum_{i=i_0+1}^{N-1}\left[\ln\frac{S(t_{i+1})}{S(t_i)}\right]^2. \tag{C.28}$$

By the linearity property of the expectation operator, we have that

$$b_3 = \frac{1}{N\Delta}\sum_{i=i_0+1}^{N-1}\mathbb{E}_t^Q\left[\left(\ln\frac{S(t_{i+1})}{S(t_i)}\right)^2\right]. \tag{C.29}$$

Using the result (C.22) for E_2 with $T \to t_{i+1}$ and $t \to t_i$ and upon simplifying, b_3 is finally given by

$$b_3 = \frac{1}{N}\sum_{i=i_0+1}^{N-1}\left(\alpha - \frac{1}{2}\sigma_i^2\right)^2\Delta + \sigma_i^2, \tag{C.30}$$

where α is given by (C.25) and Δ by (C.11). We need to discuss the diffusion coefficient σ_i appearing in (C.30). This is the annualised volatility appropriate for the interval from t_i to t_{i+1}. We will assume that these "gap" volatilities σ_i, $i = i_0 + 1, \ldots, N-1$ are constant and equal to σ_I, the implied volatility for the period from t_{i_0+1} to T. This is equivalent to assuming a flat forward volatility term structure from t_{i_0+1} to T. Note that σ_I can be calculated from σ_G and the $\sigma(T)$, the implied volatility at time t for maturity T. Replacing σ_i by a constant σ_I means that all the terms in the sum in (C.30) can be taken out of the sum to yield

$$b_3 = \frac{(N - i_0 - 1)}{N}\left[\left(\alpha - \frac{1}{2}\sigma_I^2\right)^2\Delta + \sigma_I^2\right]. \tag{C.31}$$

Finally, for b_4, we have that

$$b_4 = \mathbb{E}_t^Q \left[s_4^2 \right], \tag{C.32}$$

where

$$s_4^2 = \frac{1}{\Delta} \left(\frac{\ln S(T) - \ln S(0)}{N} \right)^2. \tag{C.33}$$

Writing s_4^2 as

$$\begin{aligned} s_4^2 N^2 \Delta &= \left(\ln \frac{S(T)}{S(0)} \right)^2 = \left(\ln \left[\frac{S(T)}{S(t_{i_0+1})} \times \frac{S(t_{i_0+1})}{S(t)} \times \frac{S(t)}{S(0)} \right] \right)^2, \\ &= \left(\ln \frac{S(T)}{S(t_{i_0+1})} \right)^2 + \left(\ln \frac{S(t_{i_0+1})}{S(t)} \right)^2 + \left(\ln \frac{S(t)}{S(0)} \right)^2 \\ &+ 2 \ln \frac{S(T)}{S(t_{i_0+1})} \ln \frac{S(t_{i_0+1})}{S(t)} + 2 \ln \frac{S(T)}{S(t_{i_0+1})} \ln \frac{S(t)}{S(0)} \\ &+ 2 \ln \frac{S(t_{i_0+1})}{S(t)} \ln \frac{S(t)}{S(0)}, \end{aligned}$$

and recalling that all terms involving $S(t)$ and $S(0)$ are known at time t and can be taken out of the expectations leads

$$\begin{aligned} b_4 N^2 \Delta &= \mathbb{E}_t^Q \left[\left(\ln \frac{S(T)}{S(t_{i_0+1})} \right)^2 \right] + \mathbb{E}_t^Q \left[\left(\ln \frac{S(t_{i_0+1})}{S(t)} \right)^2 \right] + \left(\ln \frac{S(t)}{S(0)} \right)^2 \\ &+ 2 \mathbb{E}_t^Q \left[\ln \frac{S(T)}{S(t_{i_0+1})} \ln \frac{S(t_{i_0+1})}{S(t)} \right] + 2 \ln \frac{S(t)}{S(0)} \mathbb{E}_t^Q \left[\ln \frac{S(T)}{S(t_{i_0+1})} \right] \\ &+ 2 \ln \frac{S(t)}{S(0)} \mathbb{E}_t^Q \left[\ln \frac{S(t_{i_0+1})}{S(t)} \right]. \end{aligned} \tag{C.34}$$

The second and the last expectation in (C.34) are given respectively by equations (C.24) and (C.23). The first and penultimate expectations in (C.34) can be obtained from equations (C.22) and (C.18) by putting $t \to t_{i_0+1}$ to give

$$\mathbb{E}_t^Q \left[\left(\ln \frac{S(T)}{S(t_{i_0+1})} \right)^2 \right] = \left(\alpha - \frac{1}{2} \sigma_I^2 \right)^2 (T - t_{i_0+1})^2 + \sigma_I^2 (T - t_{i_0+1}), \tag{C.35}$$

$$\mathbb{E}_t^Q \left[\ln \left(\frac{S(T)}{S(t_{i_0+1})} \right) \right] = \left(\alpha - \frac{1}{2} \sigma_I^2 \right) (T - t_{i_0+1}), \tag{C.36}$$

where σ_I is once again the implied volatility for the period from time t_{i_0+1} to time T. Consider the expectation

$$E_3 = \mathbb{E}_t^Q \left[\ln \frac{S(T)}{S(t_{i_0+1})} \ln \frac{S(t_{i_0+1})}{S(t)} \right]. \tag{C.37}$$

From equation (C.16) with respectively $t \to t_{i_0+1}$ and then $T \to t_{i_0+1}$, we have that

$$E_3 = \mathbb{E}_t^Q \left[\left(\left[\alpha - \frac{1}{2} \sigma_I^2 \right] [T - t_{i_0+1}] + \int_{t_{i_0+1}}^{T} \sigma_I \, dz_Q(u) \right) \right.$$

$$\times \quad \left(\left[\alpha - \frac{1}{2}\sigma_G^2 \right] \hat{\Delta} + \int_t^{t_{i_0+1}} \sigma_G dz_Q(u) \right) \Bigg], \tag{C.38}$$

$$= \left[\alpha - \frac{1}{2}\sigma_G^2 \right] \left[\alpha - \frac{1}{2}\sigma_I^2 \right] [T - t_{i_0+1}] \hat{\Delta}. \tag{C.39}$$

Upon multiplying out the terms in the expectation in (C.38), only one term does not vanish in the final result (C.38). The two expectations involving terms of the form $(\alpha - \frac{1}{2}\sigma)(T-t)$ vanish as a result of equation (C.17). Finally, the expectation of the product of the Brownian motion integrals is equal to the product of the expectations of the individual integrals since these are independent of one another (they have no time overlap). However, each one of these terms vanishes as a result of (C.17) and their product also vanishes.

Substituting equations (C.24), (C.35), (C.23), (C.36) and (C.39) into equation (C.34) and rearranging, the following expression is obtained for b_4:

$$\begin{aligned} b_4 &= \frac{1}{N^2\Delta} \left[\sigma_G^2 \hat{\Delta} + \sigma_I^2 (T - t_{i_0+1}) \right] \\ &+ \frac{1}{N^2\Delta} \left(\left[\alpha - \frac{1}{2}\sigma_G^2 \right] \hat{\Delta} + \left[\alpha - \frac{1}{2}\sigma_I^2 \right] (T - t_{i_0+1}) + \ln \frac{S(t)}{S(0)} \right)^2. \end{aligned} \tag{C.40}$$

C.3 Derivation of the Exact Greeks for $\hat{V}(t)$

Recall the exact expression (4.39) for the upper bound to the volswap price $V(t)$, namely

$$\hat{V}(t) = A(t) \left(\sqrt{\frac{N}{N-1}} H(t) - \sigma_K \right), \tag{C.41}$$

where

$$H(t) = \sqrt{b_1 + b_2 + b_3 - b_4}, \tag{C.42}$$

and b_1, b_2, b_3 and b_4 are respectively given by equations (C.12), (C.26), (C.31) and (C.41).

For a general variable X, we have that

$$\frac{\partial \hat{V}(t)}{\partial X} = \hat{V}(t) \frac{\partial \ln A(t)}{\partial X} + A(t) \frac{\partial}{\partial X} \left[\frac{\hat{V}(t)}{A(t)} \right]. \tag{C.43}$$

From equation (4.37) for $A(t)$,

$$\ln A(t) = \ln 100 - r_d(T - t),$$

so that

$$\frac{\partial \ln A(t)}{\partial X} = \begin{cases} -(T-t), \text{ if } X \equiv r_d; \\ r_d, \text{ if } X \equiv t; \\ -r_d, \text{ if } X \equiv T; \text{ and} \\ 0, \text{ otherwise.} \end{cases} \tag{C.44}$$

Furthermore, from (C.42),

$$\frac{\partial}{\partial X}\left[\frac{\hat{V}(t)}{A(t)}\right] = \begin{cases} \sqrt{\frac{N}{N-1}}\frac{\partial H(t)}{\partial X}, \text{ if } X \neq \sigma_K; \text{ and} \\ \\ \sqrt{\frac{N}{N-1}}\frac{\partial H(t)}{\partial X} - 1, \text{ if } X = \sigma_K. \end{cases} \tag{C.45}$$

From equations (C.43)-(C.45) with $X \equiv S(t)$, the exact delta Δ_e is given by

$$\Delta_e \triangleq \frac{\partial \hat{V}(t)}{\partial S(t)} = \sqrt{\frac{N}{N-1}} A(t) \frac{\partial H(t)}{\partial S(t)}, \tag{C.46}$$

and, by applying these results again to the RHS of (C.46), the exact gamma Γ_e is

$$\Gamma_e \triangleq \frac{\partial^2 \hat{V}(t)}{\partial S^2(t)} = \sqrt{\frac{N}{N-1}} A(t) \frac{\partial^2 H(t)}{\partial S^2(t)}. \tag{C.47}$$

Now for any general variable X, from equation (C.42),

$$\frac{\partial H(t)}{\partial X} = \frac{1}{2}\frac{1}{H(t)}\left(\frac{\partial b_1}{\partial X} + \frac{\partial b_2}{\partial X} + \frac{\partial b_3}{\partial X} - \frac{\partial b_4}{\partial X}\right). \tag{C.48}$$

We therefore require the partial derivatives of the b coefficients with respect to $S(t)$. From (C.12), b_1 is a known constant at time t and from (C.31), b_3 does not depend on $S(t)$, and thus

$$\frac{\partial b_1}{\partial S} = 0, \tag{C.49}$$

$$\frac{\partial b_3}{\partial S} = 0. \tag{C.50}$$

We then have that

$$\frac{\partial H}{\partial S} = \frac{1}{2}\frac{1}{H(t)}\left(\frac{\partial b_2}{\partial S(t)} - \frac{\partial b_4}{\partial S(t)}\right), \tag{C.51}$$

and then, from the quotient rule, we obtain

$$\frac{\partial H^2(t)}{\partial S^2(t)} = \frac{1}{2}\frac{1}{H^2(t)}\left[H(t)\left(\frac{\partial^2 b_2}{\partial S^2(t)} - \frac{\partial^2 b_4}{\partial S^2(t)}\right) - \frac{\partial H(t)}{\partial S(t)}\left(\frac{\partial b_2}{\partial S(t)} - \frac{\partial b_4}{\partial S(t)}\right)\right]. \tag{C.52}$$

However, from equation (C.26),

$$\frac{\partial b_2}{\partial S(t)} = \frac{2}{N\Delta}\frac{1}{S(t)}\left(\ln\left[\frac{S(t)}{S(t_{i_0})}\right] + \left[\alpha - \frac{1}{2}\sigma_G^2\right]\hat{\Delta}\right), \tag{C.53}$$

and then

$$\frac{\partial^2 b_2}{\partial S^2(t)} = \frac{2}{N\Delta}\frac{1}{S^2(t)}\left(1 - \left[\alpha - \frac{1}{2}\sigma_G^2\right]\hat{\Delta} - \ln\left[\frac{S(t)}{S(t_{i_0})}\right]\right), \tag{C.54}$$

in which the result $\frac{d}{dx}\left[\frac{1}{x}\ln x\right] = \frac{(1-\ln x)}{x^2}$ has been used. On the other hand, from (C.41), we have that

$$\frac{\partial b_4}{\partial S(t)} = \frac{2}{N^2\Delta}\frac{1}{S(t)}\left(\ln\left[\frac{S(t)}{S(0)}\right] + \left[\alpha - \frac{1}{2}\sigma_G^2\right]\hat{\Delta} + \left[\alpha - \frac{1}{2}\sigma_I^2\right][T - t_{i_0+1}]\right), \tag{C.55}$$

and

$$\frac{\partial^2 b_4}{\partial S^2(t)} = \frac{2}{N^2 S^2(t)\Delta}\left(1 - \left[\alpha - \frac{1}{2}\sigma_G^2\right]\hat{\Delta} - \left[\alpha - \frac{1}{2}\sigma_I^2\right][T - t_{i_0+1}] - \ln\left[\frac{S(t)}{S(0)}\right]\right). \tag{C.56}$$

From equations (C.53) and (C.55), we obtain upon rearranging terms that

$$\begin{aligned}\frac{\partial b_2}{\partial S(t)} - \frac{\partial b_4}{\partial S(t)} &= \frac{2(N-1)}{N^2\Delta S(t)}\left[\ln\frac{S(t)}{S(t_{i_0})} + \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta}\right] \\ &- \frac{2}{N^2\Delta S(t)}\left[\ln\frac{S(t_{i_0})}{S(0)} + \left(\alpha - \frac{1}{2}\sigma_I^2\right)(T - t_{i_0+1})\right],\end{aligned} \tag{C.57}$$

and from (C.54) and (C.56) that

$$\begin{aligned}\frac{\partial^2 b_2}{\partial S^2(t)} - \frac{\partial^2 b_4}{\partial S^2(t)} &= \frac{2}{N\Delta S^2(t)}\left(\frac{N-1}{N}\right)\left[1 - \ln\frac{S(t)}{S(t_{i_0})} - \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta}\right] \\ &+ \frac{2}{N^2\Delta S^2(t)}\left[\ln\frac{S(t_{i_0})}{S(0)} + \left(\alpha - \frac{1}{2}\sigma_I^2\right)(T - t_{i_0+1})\right].\end{aligned} \tag{C.58}$$

From equations (C.46), (C.51) and (C.57), we finally obtain the following result for the delta of the volswap:

$$\begin{aligned}\Delta_e &= \frac{A(t)}{H(t)S(t)}\frac{1}{N^2\Delta}\sqrt{\frac{N}{N-1}}\left[(N-1)\left(\ln\frac{S(t)}{S(t_{i_0})} + \left[\alpha - \frac{1}{2}\sigma_G^2\right]\hat{\Delta}\right)\right. \\ &\left. - \left(\ln\frac{S(t_{i_0})}{S(0)} + \left[\alpha - \frac{1}{2}\sigma_I^2\right][T - t_{i_0+1}]\right)\right].\end{aligned} \tag{C.59}$$

This can be simplified if we consider the case where $\sigma_G = \sigma_I = \sigma$ and by noting that $\hat{\Delta} + [T - t_{i_0+1}] = T - t$ to give

$$\begin{aligned}\Delta_e &= \frac{A(t)}{H(t)S(t)}\frac{1}{N^2\Delta}\sqrt{\frac{N}{N-1}}\left[N\left(\ln\frac{S(t)}{S(t_{i_0})} + \left[\alpha - \frac{1}{2}\sigma^2\right]\hat{\Delta}\right)\right. \\ &\left. - \left(\ln\frac{S(t)}{S(0)} + \left[\alpha - \frac{1}{2}\sigma^2\right][T - t]\right)\right].\end{aligned} \tag{C.60}$$

The exact gamma Γ_e is given by equation (C.47). Substituting for $\frac{\partial^2 H(t)}{\partial S^2(t)}$ from equation (C.52), we have that

$$\Gamma_e = \frac{1}{2}\sqrt{\frac{N}{N-1}}\frac{A(t)}{H^2(t)}\left[H(t)\left(\frac{\partial^2 b_2}{\partial S^2(t)} - \frac{\partial^2 b_4}{\partial S^2(t)}\right) - \frac{1}{2}\frac{1}{H(t)}\left(\frac{\partial b_2}{\partial S} - \frac{\partial b_4}{\partial S}\right)^2\right]. \tag{C.61}$$

Substituting for the various partial derivatives from (C.57) and (C.58), we finally obtain

$$\begin{aligned}\Gamma_e \;=\; & \sqrt{\frac{1}{N(N-1)}}\frac{1}{N\Delta S^2(t)}\frac{A(t)}{H^2(t)}\left\{H(t)\left[(N-1)R_1(t,t_{i_0}) + R_2(t,t_{i_0},t_{i_0+1})\right]\right. \\ & \left. -\frac{1}{N^2\Delta H(t)}\left[(N-1)\left(1-R_1(t,t_{i_0})\right) - R_2(t,t_{i_0},t_{i_0+1})\right]^2\right\},\end{aligned} \tag{C.62}$$

where

$$R_1(t,t_{i_0}) \;=\; 1-\left(\alpha-\frac{1}{2}\sigma_G^2\right)\hat{\Delta} - \ln\frac{S(t)}{S(t_{i_0})};\ \text{and} \tag{C.63}$$

$$R_2(t,t_{i_0},t_{i_0+1}) \;=\; \ln\frac{S(t_{i_0})}{S(0)} + \left(\alpha-\frac{1}{2}\sigma_I^2\right)(T-t_{i_0+1}). \tag{C.64}$$

For the simple case of $\sigma_G = \sigma_I = \sigma$ and defining

$$R_3(t,t_{i_0},t_{i_0+1}) = N\ln\frac{S(t)}{S(t_{i_0})} - \left(\alpha-\frac{1}{2}\sigma^2\right)\left(T-t-N\hat{\Delta}\right) - \ln\frac{S(t)}{S(0)}, \tag{C.65}$$

the exact gamma can be written as

$$\Gamma_e = \sqrt{\frac{1}{N(N-1)}}\frac{A(t)}{H(t)}\frac{1}{N\Delta S^2(t)}\left\{N-1-R_3(t,t_{i_0},t_{i_0+1}) - \frac{R_3^2(t,t_{i_0},t_{i_0+1})}{N^2\Delta H^2(t)}\right\}. \tag{C.66}$$

In order to calculate the exact vega, we define a composite volatility $\overline{\sigma} \triangleq \sigma_G + \sigma_I$ such that for any function Y of σ_G and σ_I,

$$\frac{\partial Y}{\partial \overline{\sigma}} = \frac{\partial Y}{\partial \sigma_G}\frac{\partial \sigma_G}{\partial \overline{\sigma}} + \frac{\partial Y}{\partial \sigma_I}\frac{\partial \sigma_I}{\partial \overline{\sigma}} = \frac{\partial Y}{\partial \sigma_G} + \frac{\partial Y}{\partial \sigma_I}. \tag{C.67}$$

The exact vega of the volswap is therefore

$$vega_e \triangleq \frac{\partial \hat{V}(t)}{\partial \overline{\sigma}} = \frac{\partial \hat{V}(t)}{\partial \sigma_G} + \frac{\partial \hat{V}(t)}{\partial \sigma_I}. \tag{C.68}$$

Using equations (C.43), (C.44) and (C.45), this is then given by

$$vega_e = \sqrt{\frac{N}{N-1}}A(t)\left[\frac{\partial H(t)}{\partial \sigma_G} + \frac{\partial H(t)}{\partial \sigma_I}\right], \tag{C.69}$$

which, in turn, can be expressed in terms of the coefficients b_1, b_2, b_3 and b_4 as

$$vega_e = \frac{1}{2}\sqrt{\frac{N}{N-1}}\frac{A(t)}{H(t)}\left[\frac{\partial b_1}{\partial \sigma_G} + \frac{\partial b_1}{\partial \sigma_I} + \frac{\partial b_2}{\partial \sigma_G} + \frac{\partial b_2}{\partial \sigma_I} + \frac{\partial b_3}{\partial \sigma_G} + \frac{\partial b_3}{\partial \sigma_I} - \frac{\partial b_4}{\partial \sigma_G} - \frac{\partial b_4}{\partial \sigma_I}\right]. \tag{C.70}$$

From equations (C.12), (C.26), (C.31) and (C.41)for b_1, b_2, b_3 and b_4, the following results are easily obtained:

$$\begin{aligned}
\frac{\partial b_1}{\partial \sigma_G} &= 0; \\
\frac{\partial b_1}{\partial \sigma_I} &= 0; \\
\frac{\partial b_2}{\partial \sigma_G} &= \frac{2\sigma_G}{N}\frac{\hat{\Delta}}{\Delta}\left[1 - \ln\frac{S(t)}{S(t_{i_0})} - \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta}\right]; \\
\frac{\partial b_2}{\partial \sigma_I} &= 0; \\
\frac{\partial b_3}{\partial \sigma_G} &= 0; \\
\frac{\partial b_3}{\partial \sigma_I} &= 2\sigma_I\frac{(N - i_0 - 1)}{N}\left[1 - \left(\alpha - \frac{1}{2}\sigma_I^2\right)\Delta\right]; \\
\frac{\partial b_4}{\partial \sigma_G} &= \frac{2\sigma_G}{N^2}\frac{\hat{\Delta}}{\Delta}\left[1 - \ln\frac{S(t)}{S(0)} - \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta} - \left(\alpha - \frac{1}{2}\sigma_I^2\right)(T - t_{i_0+1})\right]; \\
\frac{\partial b_4}{\partial \sigma_I} &= \frac{2\sigma_I(T - t_{i_0+1})}{N^2\Delta}\left[1 - \ln\frac{S(t)}{S(0)} - \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta} - \left(\alpha - \frac{1}{2}\sigma_I^2\right)(T - t_{i_0+1})\right].
\end{aligned}$$

Substituting these results back into equation (C.70) and rearranging, we obtain

$$\begin{aligned}
vega_e &= \sqrt{\frac{1}{N(N-1)}}\frac{A(t)}{H(t)}\left[\sigma_I(N - i_0 - 1)[1 - (\alpha - \frac{1}{2}\sigma_I^2)\Delta]\right. \\
&+ \sigma_G\frac{\hat{\Delta}}{\Delta}\left[1 - \ln\frac{S(t)}{S(t_{i_0})} - \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta}\right] \\
&- \frac{1}{N\Delta}\left[\sigma_G\hat{\Delta} + \sigma_I(T - t_{i_0+1})\right]\left[1 - \ln\frac{S(t)}{S(0)}\right. \\
&\qquad \left.\left. - \left(\alpha - \frac{1}{2}\sigma_G^2\right)\hat{\Delta} - \left(\alpha - \frac{1}{2}\sigma_I^2\right)(T - t_{i_0+1})\right]\right].
\end{aligned} \tag{C.71}$$

For the case in which $\sigma_G = \sigma_I = \sigma$, this simplifies to

$$\begin{aligned}
vega_e &= \sigma\sqrt{\frac{1}{N(N-1)}}\frac{A(t)}{H(t)}\left\{(N - i_0 - 1) + \frac{\hat{\Delta}}{\Delta} - \frac{(T-t)}{N\Delta}\right. \\
&+ \left[\alpha - \frac{1}{2}\sigma^2\right]\left[\frac{(T-t)^2}{N\Delta} - \frac{\hat{\Delta}^2}{\Delta} - (N - i_0 - 1)\Delta\right] \\
&+ \left.\frac{(T - t - N\hat{\Delta})}{N\Delta}\ln S(t) - \frac{(T-t)}{N\Delta}\ln S(0) + \frac{\hat{\Delta}}{\Delta}\ln S(t_{i_0})\right\}.
\end{aligned} \tag{C.72}$$

C.4 Derivation of the Approximate Greeks for $\hat{V}(t)$

We shall now derive the greeks which result from the approximation $\hat{W}(t)$ to the upper bound $\hat{V}(t)$ for the volswap price $V(t)$. Recall the equations for this approximation, namely

$$\hat{W}(t) = A(t)\,[J(t) - \sigma_K]\,, \tag{C.73}$$

where

$$J(t) = \sqrt{\frac{i_0}{N}w_H^2 + \frac{1}{N}w_c^2 + \frac{(N - i_0 - 1)}{N}\sigma^2}; \tag{C.74}$$

$$w_H^2 = \sum_{i=0}^{i_0-1} \frac{[\ln S(t_{i+1}) - \ln S(t_i) - m\Delta]^2}{i_0\Delta}; \tag{C.75}$$

$$\text{and } w_c^2 = \frac{1}{\Delta}\left[\frac{[S(t) - S(t_{i_0})(1 + \alpha(\Delta - \hat{\Delta}))]^2}{S^2(t_{i_0})} + \sigma^2\hat{\Delta}\right]. \tag{C.76}$$

Then, from equations (C.43), (C.45), (C.46) and (C.47) with $\hat{V}(t) \to \hat{W}(t)$ and $\sqrt{\frac{N}{N-1}}H(t) \to J(t)$ and from the results (C.44), the approximate delta Δ_a and gamma Γ_a are given by

$$\Delta_a \triangleq \frac{\partial\hat{W}(t)}{\partial S(t)} = A(t)\frac{\partial J(t)}{\partial S(t)}, \text{ and} \tag{C.77}$$

$$\Gamma_a \triangleq \frac{\partial^2\hat{W}(t)}{\partial S^2(t)} = A(t)\frac{\partial^2 J(t)}{\partial S^2(t)}. \tag{C.78}$$

The partial derivative of $J(t)$ with respect to an arbitrary variable X is given by

$$\frac{\partial J(t)}{\partial X} = \frac{1}{2N}\frac{1}{J(t)}\left[i_0\frac{\partial[w_H^2]}{\partial X} + \frac{\partial[w_c^2]}{\partial X} + 2\sigma(N - i_0 - 1)\frac{\partial\sigma}{\partial X}\right]. \tag{C.79}$$

For the case where $X \equiv S(t)$, the following two results hold:

$$\frac{\partial\sigma}{\partial S(t)} = 0, \tag{C.80}$$

$$\frac{\partial[w_H^2]}{\partial S(t)} = 0, \tag{C.81}$$

where the latter results from the fact that all quantities in w_H^2 are known prior to time t. Substituting equations (C.80) and (C.81) into (C.79), we obtain

$$\frac{\partial J(t)}{\partial S(t)} = \frac{1}{2N}\frac{1}{J(t)}\frac{\partial[w_c^2]}{\partial S(t)}, \tag{C.82}$$

and differentiating this once again with respect to $S(t)$ through an application of the quotient rule leads to

$$\frac{\partial^2 J(t)}{\partial S^2(t)} = \frac{1}{2N}\frac{1}{J(t)}\left[\frac{\partial^2[w_c^2]}{\partial S^2(t)} - \frac{1}{2N}\left(\frac{1}{J(t)}\frac{\partial[w_c^2]}{\partial S(t)}\right)^2\right], \tag{C.83}$$

in which equation (C.82) has also been used. Differentiating equation (C.76) with respect to $S(t)$ leads to

$$\frac{\partial[w_c^2]}{\partial S(t)} = \frac{2}{\Delta S(t_{i_0})} R(t, t_{i_0}), \tag{C.84}$$

where

$$R(t, t_{i_0}) \triangleq \frac{S(t)}{S(t_{i_0})} - \left[1 + \alpha\left(\Delta - \hat{\Delta}\right)\right]. \tag{C.85}$$

Differentiating (C.84) with respect to $S(t)$ then leads to

$$\frac{\partial^2[w_c^2]}{\partial S^2(t)} = \frac{2}{\Delta S^2(t_{i_0})}. \tag{C.86}$$

Substituting (C.82) into (C.77) and using equation (C.84), we finally obtain the following equation for the approximate delta:

$$\Delta_a = \frac{1}{N\Delta} \frac{A(t)}{J(t)} \frac{R(t, t_{i_0})}{S(t_{i_0})}, \tag{C.87}$$

where $R(t, t_{i_0})$ is defined by equation (C.85). The approximate gamma is obtained by substituting equations (C.86) and (C.84) into equation (C.83) and then substituting the resultant equation back into (C.78) to give

$$\Gamma_a = \frac{1}{N\Delta} \frac{1}{S^2(t_{i_0})} \frac{A(t)}{J(t)} \left[1 - \frac{R^2(t, t_{i_0})}{N\Delta J^2(t)}\right]. \tag{C.88}$$

The approximate vega is given by

$$vega_a = A(t) \frac{\partial J(t)}{\partial \sigma}. \tag{C.89}$$

Furthermore, from equation (C.79) with $X \equiv \sigma$, we have that

$$\frac{\partial J(t)}{\partial \sigma} = \frac{1}{2N} \frac{1}{J(t)} \left[i_0 \frac{\partial[w_H^2]}{\partial \sigma} + \frac{\partial[w_c^2]}{\partial \sigma} + 2\sigma(N - i_0 - 1)\right]. \tag{C.90}$$

The following two results follow from equations (C.75) and (C.76):

$$\frac{\partial[w_H^2]}{\partial \sigma} = 0, \tag{C.91}$$

$$\text{and } \frac{\partial[w_c^2]}{\partial \sigma} = 2\sigma \frac{\hat{\Delta}}{\Delta}. \tag{C.92}$$

Substituting equations (C.91) and (C.92) into (C.90) and then back substituting this result into equation (C.89), we finally obtain the approximate vega as

$$vega_a = \frac{\sigma}{N} \frac{A(t)}{J(t)} \left[\frac{\hat{\Delta}}{\Delta} + (N - i_0 - 1)\right]. \tag{C.93}$$