FOREIGN EXCHANGE DERIVATIVES: Advanced Hedging and Trading Techniques

by
Dr. A. A. Kotzé
Financial Chaos Theory Pty. Ltd.
March 2011
http:\\www.quantonline.co.za

Philosophy is written in that great book whichever lies before our gaze — I mean the universe — but we cannot understand if we do not first learn the language and grasp the symbols in which it is written. The book is written in the mathematical language, and the symbols are triangle, circles and other geometrical figures, without the help of which it is impossible to conceive a single word of it, and without which one wonders in vain through a dark labyrinth.

Galileo Galilei (1564-1642)

Abstract

Instruments traded in the financial markets are getting more and more complex. This leads to more complex derivative structures that are harder to analyse and risk managed. These instruments cannot be traded or managed without the relevant systems and numerical techniques.

The global economy is becoming more and more interlinked with trading between countries skyrocketing. Due to the world trade, foreign exchange forwards, futures, options and exotics are becoming increasingly commonplace in today's capital markets.

The objective of these notes is to let the reader develop a solid understanding of the current currency derivatives used in international treasury management with an emphasis on the African continent. This will give participants the mathematical and practical background necessary to deal with all the products on the market.

Before I came here I was confused about the subject.
Having listened to your lecture I am still confused.
But on a higher level.

Enrico Fermi (1901-1954)

Financial Chaos Theory Pty. Ltd.
Illuminating OTC and listed Derivatives through:
consulting services
training workshops/seminars
In-house training
quantitative analysis
research
modeling complex optionality
model building
software products and development
risk management and analysis
structured products

Financial Chaos Theory
PO Box 16185
Doornfontein 2028
South Africa
consultant@quantonline.co.za

But the creative principle resides in mathematics.
In a certain sense, therefore, I hold it true that
pure thought can grasp reality, as the ancients dreamed.

Albert Einstein

Contents

List of Figures

List of Tables

Chapter 1

Introduction

More than 22.4 billion of derivative contracts were traded on exchanges worldwide in 2010 (11.2 billion futures and 11.1 billion options) against 17.8 billion in 2009, an increase of 25%. The number of futures traded increased faster - up 35% - than options - up 16%, according to statistics compiled by WFE, which annually conducts a survey for the International Options Markets Association[1] (IOMA).

"The strong volume in exchange-traded derivatives in 2010 indicate that reforms in regulation of over-the-counter derivatives markets are causing participants to shift some of their risk transfer activities to exchange-traded derivatives," commented Ronald Arculli, chairman of WFE and chairman of Hong Kong Exchanges and Clearing. During the same period, Arculli noted that, according to Bank for International Settlements (BIS) statistics, notional amounts outstanding of OTC Derivatives decreased by 13% between June 2009 and June 2010.

Currency derivatives remain the smallest section of the exchange-based derivatives markets, with 2.3 billion contracts traded in 2010. However, driven by the Indian exchanges that accounted for 71% of the volumes traded in 2010, they have experienced triple-digit growth (+144%). When Indian exchanges are removed from the statistics, the growth rate of currency volumes in 2010 was still very strong (+36%.)

Currency derivatives is a growing business, especially so for developing markets. Understanding the more complex nature of these markets is essential for all working in the capital markets. With increased volatility in global markets following the global credit and liquidity crisis, derivatives have again come to the forefront. Management of Currency risks have once more been seen as a critical tool in today's markets for offering a vehicle for hedging or for providing profit opportunities in difficult markets. This advanced course will provide a comprehensive analysis of Currency Options with focus on pricing and structuring.

We start by giving an overview of the basic concepts of options and option pricing. We will discuss the binomial model and volatility surfaces. In Chapter 2 we have an in-depth look at hedging. We discuss the Greeks and see what they really mean to a

[1] http://www.advancedtrading.com/derivatives/229300554

trader. We introduce the *Impact Delta* and *Gamma* which helps hedging when a skew is prevalent. In Chapter 3 we compare OTC derivatives to exchange traded contracts. We give a detailed analysis of how a derivatives exchange operates with all of its risk measures.

Chapter 4 is an overview of derivatives and Islamic finance. It is currently a topical issue. Due to the importance of volatility, Chapter 5 is devoted to volatility skews. We first look at how we can generate a skew from traded data and then look at the Vanna-Volga method widely used in FX markets.

The rest of the course is devoted to exotic options. We start start in Chapter 6 with the so-called 'first generation exotics' or vanilla exotic options. These include Barrier, Digital, Forward Start and Asian options. Chapter 7 introduce the more complex exotic options or so-called 'second generation exotic'. These include Timers, Variance Swaps, Range Accruals and Quantos.

In the last Chapter, Chapter 8 we discuss the pricing of exotic options under a volatility skew. We introduce the concepts of Local Volatility and Implied Binomial Trees.

1.1 Option Basics

There are two types of options: calls and puts. A *call* option gives the holder the right, but not the obligation, to buy the underlying asset by a certain date for a certain price.

A *put* option gives the holder the right, but not the obligation, to sell the underlying asset by a certain date for a certain price.

The price in the contract is known as the *exercise* price or *strike* price. The date in the contract is known as the expiration, *expiry*, *exercise* or *maturity* date.

An option can not be obtained free of charge. There is a price attached to it called a premium. The premium is the price paid for an option. The buyer of an option pays a price for the right to make a choice — the choice to exercise or not.

Call and put options are defined in one of two ways: American or European. A European option can only be exercised at the maturity date of the option whereas an American option can be exercised at any time up and including the maturity date[2]. If the holder of the option decides to exercise the option, this option becomes a simple FX contract. The holder of the option will only exercise the option when the market is in his favour, otherwise the option contract expires worthless.

[2]The terms European and American used here bear no relation to geographic considerations. European options trade on American exchanges and American option trade on European exchanges. We will later define some other option called Asian options, Bermudan options and Parisian options.

1.2 Why trade FX Options versus the Spot FX?

Currency options have gained acceptance as invaluable tools in managing foreign exchange risk. One of the primary benefits for trading FX Options versus Spot FX is that options provide investors with tremendous versatility including a wide range of strike prices and expiration months available for trading. Investors can implement single and multi-leg strategies, depending on their risk and reward tolerance. Investors can implement bullish, bearish and even neutral market forecasts with limited risk. FX Options also provide the ability to hedge against loss in value of an underlying asset. Options are attractive financial instruments to portfolio managers and corporate treasuries because of this flexibility.

Options can be a way for traders to limit their risk in a trade. For instance, if a trader believes the EURUSD will move upwards, he may purchase a call at a premium so that if the rate hits the option strike price he can exercise it. If the currency instead moves against the trader, all that is lost is the premium.

In general we note that options expanded the universe of tradable financial instruments. The consequence is that hedges can be tailored more precisely to the risk profile of the underlying and the risk can be managed more easily. Options allow an investor to construct different payoff profiles. You can mimic your actual exposure by trading in a portfolio of options. We will look at this in Chapter 8. All these benefits can be grouped together into six benefits [Sh 10]

Benefit One: The Ability to Leverage

Options provide both individuals and firms with the ability to leverage. In other words, options are a way to achieve payoffs that would usually be possible only at a much greater cost. Options can cause markets to become more competitive, creating an environment in which investors have the ability to hedge an assortment of risks that otherwise would be too large to sustain.

Benefit Two: Creating Market Efficiency

Options can bring about more efficiency in the underlying market itself. Option markets tend to produce information flow. Options enable investors to access and trade on information that otherwise might be unobtainable or very expensive. It is for instance difficult to short sell stock. This slows the process down in which adverse information is incorporated into stock prices and make markets less efficient. It is, however, easy to sell a future or put option.

Benefit Three: Cost Efficiency

Derivatives are cost efficient. Options can provide immense leveraging ability.An investor can create an option position that will imitate the underlying's position identically — but at a large cost saving.

Benefit Four: 24/7 Protection

Options provide relative immunity to potential catastrophic effects of gaps openings in the underlyings. Consider a stop-loss order put in place to prevent losses below a predetermined price set by the investor. This protection works during the day but what happens after market close. If the market gaps down on the opening the next day your stop-loss order might be triggered but at a price much lower than your stop-loss price. You could end up with a huge loss. Had you purchased a put option for downside protection you will also be protected against gap risk

Benefit Five: Flexibility

Options offer a variety of investment alternatives. You can hedge a myriad of risks under specific circumstances. We will look at structuring and constructing different payoff profiles later on in Chapter 8.

Benefit Six: Trading Additional Dimensions

Implementation of options opens up opportunities of additional asset classes to the investor that are embedded in options themselves. Options allow the investor not only to trade underlying movements, but to allow for the passage of time and the harnessing of volatility. The investor can take advantage of a stagnant or a range-bound market.

1.3 The *Black & Scholes* Environment

To obtain an understanding of what the *Black & Scholes* formula means, it is very important to know under what conditions the *Black & Scholes* formula hold. *Black & Scholes* (and researchers before them) understood very well that the market is complex. To be able to describe it mathematically and to enable them to obtain a useful model, they knew they had to simplify the market by making certain assumptions[3]. The following is a list of the more important assumptions *Black & Scholes* made in their analysis [MS 00]:

- *The underlying asset follows a lognormal random walk.* This was not a new assumption and was already proposed by Bachelier in 1900.
- *The efficient market hypothesis is assumed to be satisfied.* In other words, the markets are assumed to be liquid, have price-continuity, be fair and provide all

[3]What inevitably happens after such a "simple" model has been proposed and understood well, is that people start to relax certain assumptions to move closer to a more realistic model.

players with equal access to available information. This implies that there are no transaction costs.

- *We live in a risk-neutral world i.e., investors require no compensation for risk.* The expected return on all securities is thus the riskfree interest rate with the consequence that there are no arbitrage opportunities. This was one of *Black & Scholes* ' insights and is known as risk-neutral valuation.
- The stock's *volatility is known* and does not change over the life of the option. In statistical talk we say the means and variances of the distribution or process is "stationary".
- The *short-term interest never changes.*
- *Short selling of securities with full use of the proceeds is permitted.*
- *There are no dividends.*
- *Delta hedging is done continuously.* This is impossible in a realistic market but makes their analysis possible.

The above describes the so-called *Black & Scholes* environment within which they did their analysis.

In general we can say that *Black & Scholes* assumed that the financial market is a system that is in equilibrium. With equilibrium we mean that, if there are no outside or exogenous influences, then the system is at rest - everything thus balances out; supply equals demand. Any distortion or perturbation is thus quickly handled by the market players so as to restore the equilibrium situation. More so, the systems reacts to the perturbation by reverting to equilibrium in a *linear fashion.* The system reacts immediately because it wishes to be at equilibrium and abhors being out of balance [Pe 91]. An option's price is thus the value obtained under this equilibrium situation.

These assumptions are very restrictive, as a matter of fact *Black* went on to say that "Since these assumptions are mostly false, we know the formula must be wrong" [Bl 88]. But, he might not be far from the truth when he further stated that "But we may not be able to find any other formula that gives better results in a wide range of circumstances."

1.4 The Seminal Formula

The seminal formula is[4]

$$V(S,t) = \phi\left(Se^{-d\tau}N(\phi x) - Ke^{-r\tau}N(\phi y)\right) \tag{1.1}$$

[4] *Clark* gives an extensive review of the derivation of the *Black & Scholes* equation [Cl 11]

where

$$\begin{aligned} x &= \left[\ln\frac{S}{K} + \left(r - d + \frac{1}{2}\sigma^2\right)\tau\right]\frac{1}{\sigma\sqrt{\tau}} \\ y &= x - \sigma\sqrt{\tau}. \end{aligned}$$

where ϕ is a binary parameter defined as

$$\phi = \begin{cases} 1 & \text{for a call option} \\ -1 & \text{for a put option.} \end{cases} \tag{1.2}$$

Here, S is the current spot stock price, K is the strike price, d is the dividend yield, σ^2 is the variance rate of the return on the stock prices, $\tau = T - t$ is the time to maturity and $N(x)$ is the cumulative standard normal distribution function[5]. Note that Eq. 1.1 in general format is; it holds for both calls and puts.

Equation (1.1) is not the original formula as published by *Black & Scholes* . They considered stock that do not pay dividends, i.e. $d = 0$. This formula is called the *modified Black & Scholes* equation adapted by *Merton* in 1973 to include a continuous dividend yield d [Me 73]. He did this by correctly assuming that an option holder does not receive any cash flows paid by the underlying instrument. This fact should be reflected in a lower call price or a higher put price. The *Merton* model provides a solution by subtracting the present value of the continuous cash dividend flow from the price of the underlying instrument. The original equation had $d = 0$.

The beauty of this formula lies in the fact that one does not need to estimate market expectation or risk preferences. This was a revolutionary improvement over its predecessors. There are 3 parameters that needs to be estimated: the riskfree interest rate; the dividend yield and the variance or volatility. Note that the volatility needed in the *Black & Scholes* formula is the volatility of the underlying security that will be observed in the future time interval τ - *volatility thus needs to be predicted* [MS 00]. *Black & Scholes*, however, assumed that the variance is known and that it is constant.

1.5 A Currency Option Model

Garman and *Kohlhagen* provided a formula for the valuation of foreign currency options [GK 83]. They followed the *Black & Scholes* lines of thought but set their riskless hedge portfolio up by investing in foreign bonds, domestic bonds and the option. They had some more assumptions though

[5] Note that

$$N(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-a^2/2}da$$

but it can be determined numerically.

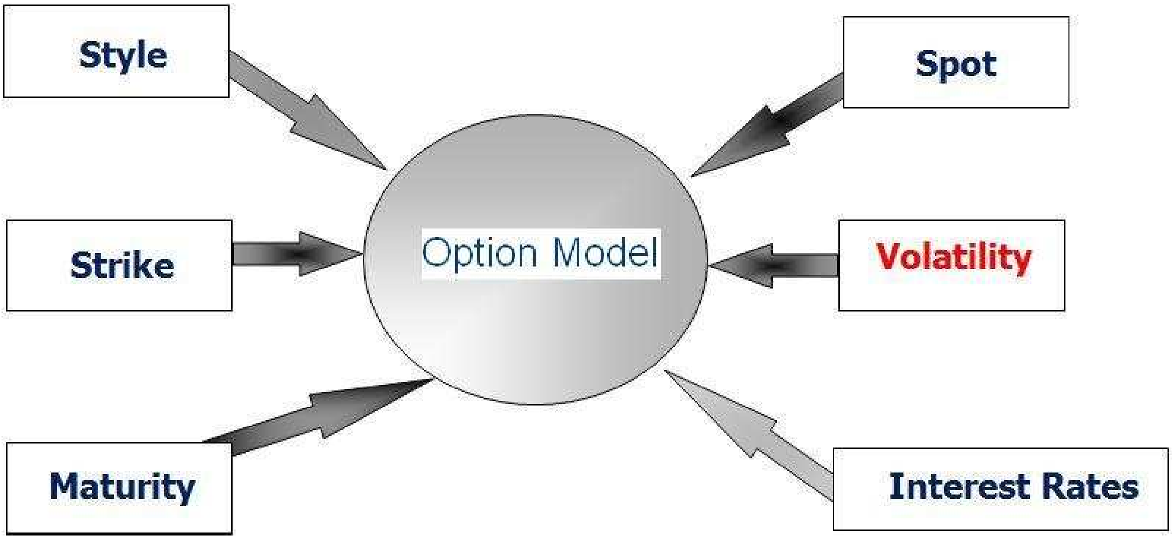

Figure 1.1: Information necessary to price an option.

- It is easy to convert the domestic currency into the foreign currency;
- We can invest in foreign bonds without any restrictions.

Their analysis led them to the following *Black & Scholes* -type equation

$$V(S,K,\tau,\sigma,r_d,r_f,\phi) = \phi\left(Se^{-r_f\tau}N(\phi x) - Ke^{-r_d\tau}N(\phi y)\right) \tag{1.3}$$

where

$$\begin{aligned} x &= \left[\ln\frac{S}{K} + \left(r_d - r_f + \frac{1}{2}\sigma^2\right)\tau\right]\frac{1}{\sigma\sqrt{\tau}} \\ y &= x - \sigma\sqrt{\tau} \end{aligned} \tag{1.4}$$

with ϕ defined in Eq. 1.2. Note that S is the current spot exchange rate, K is the strike price, r_f is the foreign interest rate, r_d is the domestic interest rate σ^2 is the variance rate of the return on the exchange rate, $\tau = T - t$ is the time to maturity and $N(x)$ is the cumulative standard normal distribution function. We note that Eq. 1.3 is exactly the *Black & Scholes* equation given in Eq. 1.1 where we have substituted $d = r_f$ and $r = r_d$. To price an option we thus need six quantities as depicted in Fig. 1.1

Note that $V(S,K,t,\sigma,r_d,r_f,\phi)$ will sometimes be shortened to $V(S,\tau)$ or just $V(S)$. V denotes the value of a Call or Put. In future we sometimes also have

$$\begin{aligned} C(S,K,\tau,\sigma,r_d,r_f) &= V(S,K,t,\sigma,r_d,r_f,+1) \Rightarrow \text{Call} \\ P(S,K,\tau,\sigma,r_d,r_f) &= V(S,K,t,\sigma,r_d,r_f,-1) \Rightarrow \text{Put.} \end{aligned} \tag{1.5}$$

1.6 Options on Forwards and Futures

In 1976 *Fischer Black* presented a model for pricing commodity options and options on forward contracts [Bl 76]. Define a forward contract's value to be

$$F_T = F = Se^{(r_d - r_f)T} \tag{1.6}$$

with S the spot currency exchange rate at the start of the forward contract. *Black* showed that a futures contract can be treated in the same way as a security providing a continuous dividend yield equal to the riskfree interest r. This means that we can use Eq. 1.1 where we have $r = d$ and we use the forward price F as the price of the underlying instead of the cash price S. Turning to currencies and the *Garman* and *Kohlhagen* model we put $r_d = r_f$.

The easiest way in obtaining *Black's* formula is to invert Eq. 1.6 to obtain S and substitute this into Eq. 1.3 to give

$$V(F, t) = \phi e^{-r_d \tau} \left[FN(\phi x) - KN(\phi y) \right] \tag{1.7}$$

if $T_d = T_f$.

1.7 Settlement Adjustments

Foreign exchange spot transactions generally settle in two business days. There are thus four dates of importance for option contracts: today, spot, expiry and delivery [Cl 11]. The delivery date is usually set to the expiry spot, i.e., so that delivery bears the same spot settlement relationship to expiry. This means that if

$$\text{spot} = \text{today} + 2 \text{ (T+2), then delivery} = \text{expiry} + 2 \text{ (T+2)}$$

as well. In the *Black & Scholes* equation given in Eq. 1.1 we defined the time to expiry by stating $\tau = T - t$ where t is the trade date and T is the expiry date. Now that we have four dates, which ones are the correct ones to use when we price options? On a time line we have

$$T_{today} \rightarrow T_{spot} \longrightarrow T_{exp} \rightarrow T_{es}$$

and we usually have the delivery time $T_{del} = T_{es}$.

To understand why this is important, we refer back to the actual cash flows. The premium will only be received on T_{spot} although we enter into the contract at T_{today}. The premium should thus reflect the premium today at T_{today}. Also, if the option is in the money at expiry, the profit will flow on T_{del} only. Now it becomes tricky because the volatility applicable is the volatility over the period $T_{today} \leq t \leq T_{exp}$ because that is the real terms of the agreement. To price the option correctly, we price it from cash flow to cash flow thus from T_{spot} to T_{del}. We then discount the premium back from T_{del} to T_{today} and then forward value to T_{spot}.

This changes the *Black & Scholes* formula somewhat to

$$V(S,t) = Se^{-r_f(T_{es}-T_{spot})}N(\phi x) - Ke^{-r_d(T_{es}-T_{spot})}N(\phi y) \tag{1.8}$$

where

$$\begin{aligned} x &= \left[\ln\frac{S}{K} + (r_d - r_f)(T_{es} - T_{spot}) + \frac{1}{2}\sigma^2\,(T_{exp} - T_{today})\right]\frac{1}{\sigma\sqrt{T_{exp} - T_{today}}} \\ y &= x - \sigma\sqrt{T_{exp} - T_{today}}. \end{aligned}$$

If $\tau = T_{es} - T_{spot} = T_{exp} - T_{today}$ Eq. 1.8 is exactly the same as Eq. 1.3.

If one extract the relevant interest rates from zero coupon yield curves, note the following: the interest rates should be the forward interest rates that hold from T_{spot} to T_{es}. Graphically we depict this in Fig. 1.2.

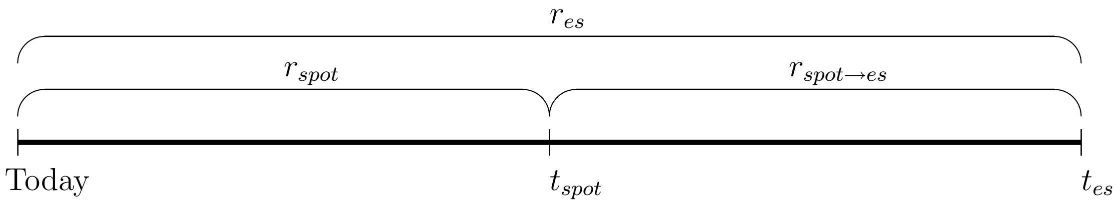

Figure 1.2: The correct interest rates for a currency option.

1.8 Put-Call-Parity

In the previous chapter we looked at the formulas available for trading European options. Can we say anything about a relationship between puts and calls?

Put-Call-Parity is a very important relationship that is distribution-free. It relationship that exists between the prices of European put and call options where both have the same underlier, strike price and expiry date. It is given by

$$C + Ke^{-r\tau} = P + Se^{-d\tau}. \tag{1.9}$$

Note, we used the same notation as in Eq. 1.1 where the strike price $K = X$ and we included a dividend yield d.

This is a fundamental arbitrage relationship which forces call and put prices to be tied to their underlying market and to each other. Note that it is not based on any option pricing model. It was derived purely using arbitrage arguments. Put-call parity offers a simple test of option pricing models. Any option pricing model that produces put and call prices that do not satisfy put-call parity must be rejected as unsound. Such a model will suggest trading opportunities where none exist. Put-call-parity is used to create synthetic securities.

From Eq. 1.9 and the *Garman* and *Kohlhagen* model given in Eq. 1.3 we obtain Put-Call-Parity for currency options to be

$$C + Ke^{-r_d\tau} = P + Se^{-r_f\tau}. \tag{1.10}$$

From Eq. 1.7 we deduce Put-Call-Parity for a currency option on a forward or futures contract is given by

$$\begin{aligned} C + Ke^{-r_d\tau} &= P + Fe^{-r_d\tau} \qquad (1.11) \\ \Rightarrow C - P &= (F - K)\, e^{-r_d\tau}. \end{aligned}$$

1.9 Option Dynamics and Risk Managements

If you have traded a few options but are relatively new to trading them, you are probably battling to understand why some of your trades aren't profitable. You start to realise that trying to predict what will happen to the price of a single option or a position involving multiple options as the market changes can be a difficult undertaking. You experience the forces of the market and see that an option price does not always appear to move in conjunction with the price of the underlying asset or share. If you want to trade in any financial instrument, it is important to understand what factors contribute to the movement in the price of that instrument, and what effect they have.

Option prices are influenced by six quantities or variables:

- the current FX rate;
- the strike price;
- the time to expiration;
- the local riskfree interest rate;
- the foreign riskfree interest rate;
- the volatility.

If you want to manage the risk associated with an option you need to understand the dynamics of option values in relation to these quantities.

Futures traders are almost exclusively interested in the direction of the market. Option traders, on the other hand, must also take note of how fast the market is moving or will move. If both futures and options traders take a long market position and the market does move higher, the futures trader is assured of a profit while the option trader may show a loss. To maintain a profit margin the option trader must analyze manage the risk associated with an option at once. He needs to understand the dynamics of option values in relation to these six quantities.

Because this is a difficult task we resort to theoretical models. The goal of theoretical (and numerical) evaluation of option prices is to analyze an option based on current market conditions as well as expectations about future conditions. This evaluation then *assists the trader in making an intelligent decision on an option* [Na 88]. Such an analysis can be done; all that is needed is information that characterizes the probability distribution of future FX rates and interest rates!

Here we consider what happens to options prices when one of these quantities changes while the others remain fixed/constant. We draw the relevant graphs for puts and calls and deduct from there the option's behaviour. The following have to be remembered: the FX rate, strike and time to expiry are known quantities. The riskfree interest rates and the volatility are mostly unknown. These have to be estimated. We will return to these later.

Please note that we use the same notation as for the currency option model in Eq. 1.3.

1.9.1 The Greeks

When we talk about these six variables in relation to option pricing, we call them the Greeks. You might have heard terms such as Delta or Vega and you immediately thought option trading is too difficult or risky.

However, what you will learn in this lesson is that learning things the 'Greek' way is like knowing the baby steps towards potential gains. While many traders focus on spot prices and trends, options pricing and its unpredictability seems to be a bigger problem. For one, the value of options is so uncertain that sometimes trends and factors provide no help at all. If you know about technical analysis of shares, try some of those analyses on option values. You will quickly realise that momentum or stochastics are of no use at all.

Further to the above, we pile on the fact that the Greeks cannot simply be looked up in your everyday option tables nor will you see them on screen where you see the option bids and offers. They need to be calculated which means you will need access to a computer or electronic calculator that calculates them for you.

The Delta

The delta is a measure of the ratio of option contracts to the underlying asset in order to establish a neutral hedge. We can also state that the Δ is a measure of how fast an option's value changes with respect to changes in the price of the underlying asset.

From 1.3 we have

$$\Delta = \frac{\partial V}{\partial S} = \phi e^{-r_f \tau} N(\phi x). \tag{1.12}$$

From this we see that if the option is far out-of-the-money $\Delta \simeq 0$, however if the option is deep in-the-money we have $\Delta \simeq \phi$. Δ is thus the probability that the option will end in-the-money.

If you do not have a *Black & Scholes* calculator giving the Delta, you can easily calculate the numerical Delta as follows

$$\Delta_{num} = (V(S + 0.0001) - V(S)) \times 10000.$$

Just add one pip to the price. Here we multiply by 10,000 because of the pip size.

The Premium-Included Delta

If a Nairobi based FX trader wants to hedge his USDKES book in Shillings, he will use the Delta given in Eq. 1.12. However, if a trader has a USDKES book, but the trader sits in New York, his profits and losses will be computed in Dollars and what he really aim at is hedging the option values converted into Dollars. Hence, if V is the option value in Shillings and S is the USDKES spot exchange rate, V/S is the option value converted into Dollars. What the New York trader wants to hedge is

$$\frac{\partial \frac{V}{S}}{\partial S}.$$

This is called the *premium-included Delta* and is given by [Ca 10]

$$\Delta_{pi} = \phi \frac{K}{S} e^{-r_d \tau} N(\phi y). \tag{1.13}$$

Elasticity

The elasticity (denoted by Λ) is the elasticity of an option and shows the percentage change in its value that will accompany a small percentage change in the underlying asset price such that

$$\begin{aligned}
\Lambda_c = \frac{\partial C}{\partial S}\frac{S}{C} &= \frac{S}{C}\Delta_C = e^{-r_f \tau}\frac{S}{C}N(x) \;\; > 1 \\
\Lambda_p = \frac{\partial P}{\partial S}\frac{S}{P} &= \frac{S}{P}\Delta_P = e^{-r_f \tau}\frac{S}{P}N(y) \;\; < 0.
\end{aligned}$$

The elasticity increases when the FX rate decreases. Λ also increases as time to expiration decreases. A call will thus be more sensitive to FX rate movements 'in percentage terms', the shorter the time remaining to expiration [Ko 03].

Gamma

The Γ is the rate at which an option gains or loses deltas as the underlying asset's price move up or down. It is a measure of how fast an option is changing its market characteristics and is thus a useful indication of the risk associated with a position. We have

$$\Gamma = \frac{\partial^2 V}{\partial S^2} = \frac{\partial \Delta}{\partial S} = e^{-r_f \tau} \frac{N'(x)}{S\sigma\sqrt{\tau}} \tag{1.14}$$

where $N'(x)$ is the standard normal probability density function and the cumulative normal's derivative given by

$$N'(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right). \tag{1.15}$$

If the option is far out-of-the-money or far in-the-money, $\Gamma \simeq 0$. If Γ is small, Δ changes slowly and adjustments to keep the portfolio Δ-neutral need only be made relatively infrequently. If Γ is large, however, changes should be made frequently because the Δ is then highly sensitive to the price of the underlying asset. This happens when the FX rate is close to the strike price with very little time to expiry — such as in the morning of the option expires in the afternoon.

Speed

This quantity measures how fast the Γ is changing. It is given by

$$\frac{\partial^3 V}{\partial S^3} = \frac{\partial \Gamma}{\partial S} = -e^{r_f \tau} \frac{N'(x)}{S^2 \sigma\sqrt{\tau}} \left(\frac{x}{\sigma\sqrt{\tau}} + 1\right) \tag{1.16}$$

with $N'(x)$ define in Eq. 1.15.

Theta

Both puts and calls lose value as maturity approaches. The Θ is the "time decay factor" and measures the rate at which an option loses its value as time passes such that

$$\Theta = \frac{\partial V}{\partial \tau} = -e^{-r_f \tau} \frac{SN'(x)\sigma}{2\sqrt{\tau}} + \phi\left[r_f S e^{-r_f \tau} N(\phi x) - \phi r_d K e^{-r_d \tau} N(\phi y)\right] \tag{1.17}$$

The size of the Γ correlates to the size of the Θ position where a large positive Γ goes hand in hand with a large negative Θ. A large negative Γ correlates with a large positive Θ. This means that every option position is a trade-off between market movement and time decay. Thus if Γ is large market movement will help the trader but time decay will hurt him and vice versa.

Some market participants calculates the numerical time decay defined as

$$\text{Time Decay} = V(t+i) - V(t) \tag{1.18}$$

where i is a day count parameter. If we put $i = 1$, $V(t+1)$ means the value of the option tomorrow keeping all the other parameters the same. The time decay is thus just the value of the option tomorrow minus the value today. The time decay over a weekend can be obtained by putting $i = 3$.

Charm

Charm is the change of Δ with time

$$\frac{\partial^2 V}{\partial S \partial \tau} = \frac{\Delta}{\partial \tau} = \phi e^{-r_f \tau}\left[N'(x)\frac{2(r_d - r_f)\tau - y\sigma\sqrt{\tau}}{2\tau\sigma\sqrt{\tau}} - r_f N(\phi x)\right] \tag{1.19}$$

Color

Color is the change of Γ with time

$$\frac{\partial^3 V}{\partial S^2 \partial \tau} = \frac{\Gamma}{\partial \tau} = -e^{-r_f \tau}\frac{N'(x)}{2S\tau\sigma\sqrt{\tau}}\left[2r_f\tau + 1 + \frac{2(r_d - r_f)\tau - y\sigma\sqrt{\tau}}{2\tau\sigma\sqrt{\tau}}x\right] \tag{1.20}$$

Vega

The Vega measures the change in the option's price as volatility changes such that

$$Vega = \frac{\partial V}{\partial \sigma} = Se^{-r_f \tau}\sqrt{\tau}N'(x) \tag{1.21}$$

"At-the-money" options are the most sensitive to volatility changes. Some market participants calculates the numerical Vega defined as

$$\text{Vega Numerical} = V(\sigma_{up}) - V(\sigma)$$

where $\sigma_{up} = \sigma + 1\%$ — the current volatility adjusted upwards by 1%.

Volga and Vanna

The Volga measures the speed of the change of the Vega

$$Volga = \frac{\partial^2 V}{\partial \sigma^2} = Se^{-r_f \tau}\sqrt{\tau}N'(x)\frac{xy}{\sigma}. \tag{1.22}$$

The Volga is also called the Vomma or Volgamma — due to the definition of the Γ. The Vanna measures how the Vega changes if the spot price changes

$$Vanna = \frac{\partial^2 V}{\partial \sigma \partial S} = \frac{\partial Vega}{\partial S} = -e^{r_f \tau}N'(x)\frac{y}{\sigma}. \tag{1.23}$$

Rho

The ρ measures the change of an option's price to changes in interest rates such that

$$\rho_d = \frac{\partial V}{\partial r_d} = \phi K \tau e^{-r_d \tau} N(\phi y) \quad (1.24)$$

$$\rho_f = \frac{\partial V}{\partial r_f} = -\phi S \tau e^{-r_f \tau} N(\phi x). \quad (1.25)$$

An increase in interest rates will decrease the value of an option by increasing carrying costs. This effect is, however, outweighed by considerations of volatility, time to expiration and the price of the underlying asset.

Dual Delta and Gamma

These quantities measure the changes in option prices as the strike changes

$$Dual\ \Delta = \Delta_K = \ = \frac{\partial V}{\partial K} = -\phi e^{-r_d \tau} N(phiy) \quad (1.26)$$

$$Dual\ \Gamma = \Gamma_K = \frac{\partial^2 V}{\partial K^2} = e^{-r_d \tau} \frac{N'(y)}{K \sigma \sqrt{\tau}}. \quad (1.27)$$

1.9.2 The Gamma and Theta Relationship

From the *Black & Scholes* differential equation we have [Hu 06]

$$\frac{1}{2}\sigma^2 S^2 \Gamma + (r_d - r_f) S \Delta + \Theta = (r_d - r_f) V.$$

For a delta-neutral position we have have

$$\frac{1}{2}\sigma^2 S^2 \Gamma + \Theta = (r_d - r_f) V.$$

1.10 Useful Relationships

1.10.1 Identities

The following identities are very useful when we do differentiation in calculating the theoretical Greeks for *Black & Scholes* -type formulas [Wy 06]

$$\begin{aligned}
\frac{\partial x}{\partial \sigma} &= -\frac{y}{\sigma} \\
\frac{\partial y}{\partial \sigma} &= \frac{x}{\sigma} \\
\frac{\partial x}{\partial r_d} = \frac{\partial y}{\partial r_d} &= \frac{\sqrt{\tau}}{\sigma} \\
\frac{\partial x}{\partial r_f} = \frac{\partial y}{\partial r_f} &= -\frac{\sqrt{\tau}}{\sigma} \\
S e^{-r_f \tau} N'(x) &= K e^{-r_d \tau} N'(y)
\end{aligned} \tag{1.28}$$

Some identities also hold for the cumulative normal distribution function

$$\begin{aligned}
\frac{\partial N(\phi x)}{\partial S} &= \frac{\phi N'(\phi x)}{S \sigma \sqrt{\tau}} \\
\frac{\partial N(\phi y)}{\partial S} &= -\frac{\phi N'(\phi y)}{S \sigma \sqrt{\tau}} \\
\frac{\partial N'(x)}{\partial x} &= -x N'(x).
\end{aligned} \tag{1.29}$$

1.10.2 Symmetries

In Eq. 1.10 we wrote down the Put-Call-Parity relationship. There is also a Put-Call value symmetry for puts and calls with different strikes such that (where we use the same notation as set out in Eq. 1.6)

$$\begin{aligned}
C(S, K, \tau, \sigma, r_d, r_f) &= \frac{K}{S e^{(r_d - r_f)\tau}} P\left(S, \frac{\left(S e^{(r_d - r_f)\tau}\right)^2}{K}, \tau, \sigma, r_d, r_f\right) \\
\Rightarrow C(S, K, \tau, \sigma, r_d, r_f) &= \frac{K}{F} P\left(S, \frac{F^2}{K}, \tau, \sigma, r_d, r_f\right).
\end{aligned} \tag{1.30}$$

The symmetry for puts and calls on a forward or futures contract is quite simple

$$C(F, K, \tau, \sigma, r_d, r_f) = P(F, K, \tau, \sigma, r_d, r_f). \tag{1.31}$$

Another useful symmetry between puts and calls is given by [Ha 07]

$$C(S, K, \tau, \sigma, r_d, r_f) = P(-S, -K, \tau, -\sigma, r_d, r_f). \tag{1.32}$$

Let's say we wish to measure the value of the underlying in a different unit. That will effect the option value. The new option price can be calculated if we use the following state space transformation (also known as *space homogeneity*) [Ha 07, Wy 06]

$$a \times V(S, K, \tau, \sigma, r_d, r_f, \phi) = V(a \times S, a \times K, \tau, \sigma, r_d, r_f, phi) \; \forall \; a > 0 \tag{1.33}$$

where a is some constant.

From Eqs. 1.32 and 1.33 we deduce the follwing

$$\begin{aligned} C(S, K, \tau, \sigma, r_d, r_f) &= -P(S, K, \tau, -\sigma, r_d, r_f) \\ P(S, K, \tau, \sigma, r_d, r_f) &= -C(S, K, \tau, -\sigma, r_d, r_f). \end{aligned} \tag{1.34}$$

This is also know as "put-call-supersymmetry".

The symmetries mentioned here simplify coding and implementation of option pricing calculators.

1.10.3 Put-Call-Delta Parity

If we differentiate the put-call-parity relationship in Eq. 1.10 with respect to S we get

$$\Delta_C = \Delta_P + e^{-r_f \tau}. \tag{1.35}$$

For an option on a future we get from Eq. 1.12

$$\Delta_C = \Delta_P + e^{-r_d \tau}. \tag{1.36}$$

Now, look at the *space homogeneity* relationship in Eq. 1.33, and we differentiate both sides with respect to a and we then set $a = 1$, we get [Wy 06]

$$V = S\Delta + K\Delta_K \tag{1.37}$$

where Δ is the ordinary Delta and Δ_K is the dual delta defined in Eq. 1.27. This again can help in simplifying coding and option calculators. They also help in double checking or verifying the Greek numbers. These homogeneity methods can easily be extended to other more complex options.

1.10.4 Symmetries for Currency Options

By combining the Rho Greeks given in Eqs. 1.25 and 1.25 we obtain the *rates symmetry*

$$\begin{aligned} \frac{\partial V}{\partial r_d} + \frac{\partial V}{\partial r_f} &= -\tau V \\ \Rightarrow \rho_d + \rho_f &= -\tau V. \end{aligned} \tag{1.38}$$

We also have *foreign-domestic-symmetry* given by

$$\frac{1}{S}V(S,K,\tau,\sigma,r_d,r_f,\phi) = KV\left(\frac{1}{S},\frac{1}{K},\tau,\sigma,r_d,r_f,-\phi\right) \tag{1.39}$$

or

$$\frac{1}{S}C(S,K,\tau,\sigma,r_d,r_f) = KP\left(\frac{1}{S},\frac{1}{K},\tau,\sigma,r_d,r_f\right). \tag{1.40}$$

This equality is one of the faces of put-call-symmetry. The reason is that the value of an option can be computed both in a domestic as well as foreign scenario.

1.11 The Volatility Skew

The Black and Scholes model assumes that volatility is constant. This is at odds with what happens in the market where traders know that the formula misprices deep in-the-money and deep out-the-money options. The mispricing is rectified when options (on the same underlying with the same expiry date) with different strike prices trade at different volatilities — traders say volatilities are skewed when options of a given asset trade at increasing or decreasing levels of implied volatility as you move through the strikes. The empirical relation between implied volatilities and exercise prices is known as the "volatility skew".

The volatility skew can be represented graphically in 2 dimensions (strike versus volatility). The volatility skew illustrates that implied volatility is higher as put options go deeper in the money. This leads to the formation of a curve sloping downward to the right. Sometimes, out-the-money call options also trade at higher volatilities than their at-the-money counterparts. The empirical relation then has the shape of a smile, hence the term "volatility smile". This happens most often in the currency markets.

1.11.1 Universality of the Skew

The skew is a universal phenomenon. It is is seen in most markets around the globe. One of the best and comprehensive studies to confirm this was done by *Tompkins* in 2001 [To 01]. He looked at 16 different options markets on financial futures comprising four asset classes: equities, foreign exchange, bonds and forward rate agreements (FRA's). He compared the relative smile patterns or shapes across markets for options with the same time to expiration. His data set comprised more than 10 years of option prices spanning 1986 to 1996.

Tompkins concluded that regularities in implied volatility surfaces exist and are similar for the same asset classes even for different exchanges. A further result is that the shapes of the implied volatility surfaces are fairly stable over time. We show his results in Fig. 1.3 for currencies. South Africa's stock indices and currencies exhibit similar shapes as those determined in this study. Note the 'smiles' in all markets.

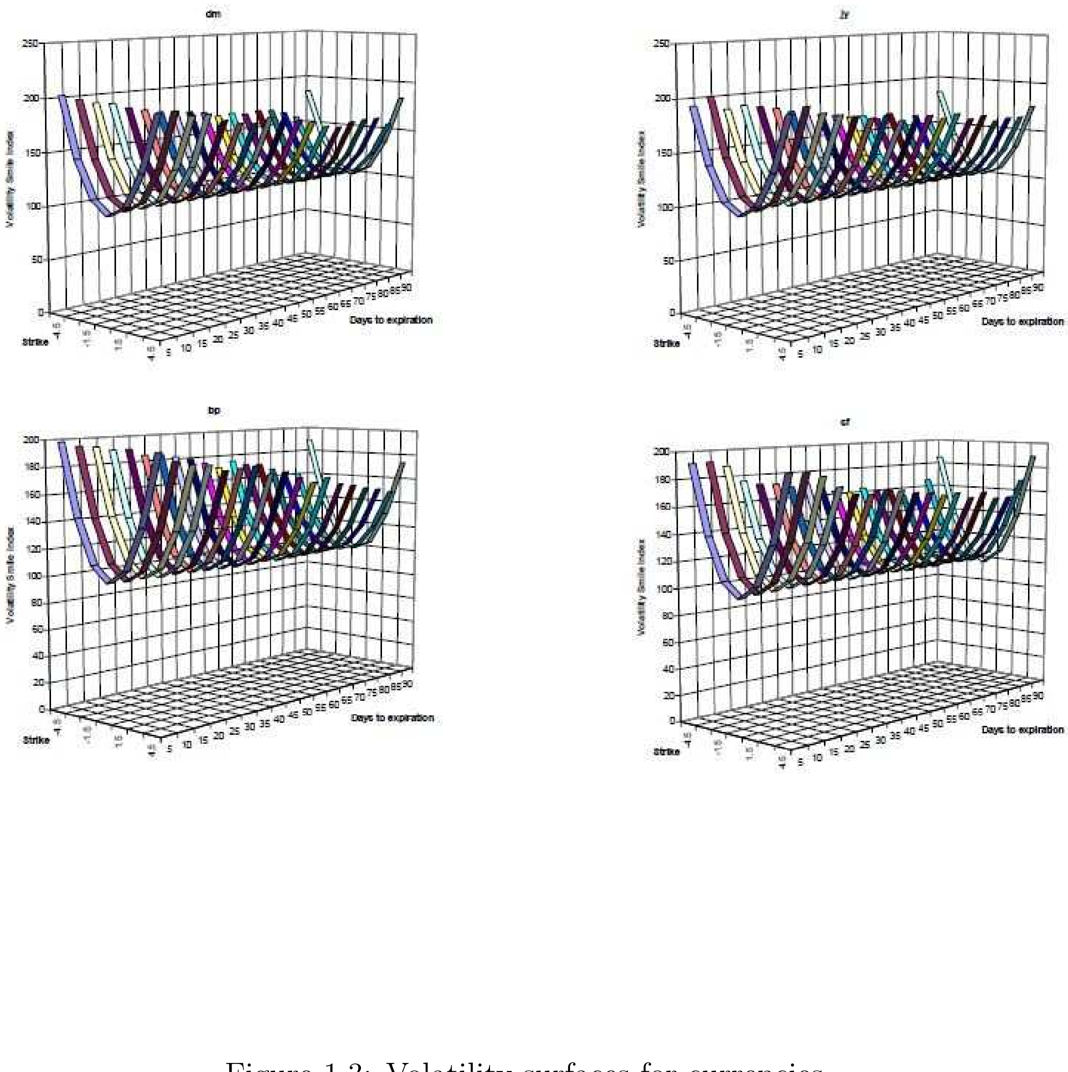

Figure 1.3: Volatility surfaces for currencies.

1.11.2 Why do we observe a Skew?

In 1972 Black and Scholes mentioned in a paper "the historical estimates of the variance include an attenuation bias - the spread of the estimated variance is larger than the true variance" [BS 72]. This would imply that for securities with a relatively high variance (read volatility), the market prices would imply an underestimate in the variance, while using historical price series would overestimate the variance and the resulting model option price would be too high; the converse is true for relative low variance securities. Black and Scholes further showed that the model performed very well, empirically, if they use the right variance. In 1979, *Macbeth* and *Mervile* extended this empirical research of Black and Scholes and also showed that the skew existed [MM 79]. In this paper, *Macbeth* and *Merville* reported that the *Black & Scholes* model undervalues in-the-money and overvalues out-the-money options. At that point in time the skew wasn't pronounced but the market crash of October 1987 changed all of that.

If one looks at option prices before and after October 1987, one will see a distinct break. Option prices began to reflect an "option risk premium" — a crash premium that comes from the experiences traders had in October 1987. After the crash the demand for protection rose and that lifted the prices for puts; especially out-the-money puts. To afford protection, investors would sell out-the-money calls. There is thus an over supply of right hand sided calls and demand for left hand sided puts — alas the skew. A skew represents the market's bias toward calls or puts.

The skew tells us there exists multiple implied volatilities for a single underlying asset. This should be somewhat disconcerting. How can the market be telling us that there is more than one volatility for the asset? The real phenomenon underlying volatility skews is that either [Ma 95]

- market imperfections systematically prevent prices from taking their true *Black & Scholes* values or
- the underlying asset price process differs from the lognormal diffusion process assumed by the *Black & Scholes* model[6].

These two points show us there is something wrong with the Black-Scholes model, which is that it fails to consider all of the factors that enter into the pricing of an option. It accounts for the stock price, the exercise price, the time to expiration, the dividends, and the risk-free rate. The implied volatility is more or less a *catch-all term*, capturing whatever variables are missing, as well as the possibility that the model is improperly specified or blatantly wrong. The volatility skew is thus the market's way of getting around Black and Scholes's simplifying assumptions about how the market behaves.

[6]Stochastic volatility models closer to the truth.

1.11.3 Shapes of the skew

There are three distinct shapes

- **Supply Skew**: The supply skew is defined by higher implied volatility for lower strikes and lower volatility for higher strikes. Supply refers to the natural hedging activity for the major players in the market who have a supply of something they need to hedge. Stock Index and interest rate markets have supply skews. The natural hedge for stock owners is to buy puts in order to protect the value of their "supply" of stock and sell calls to offset the puts' prices — collars. This natural action in the marketplace determines the structure of the skew. This type of skew is also known as a 'smirk'.

- **Demand Skew**: Demand skews have higher implied volatility at higher strikes and lower implied volatility at lower strikes. The natural hedger in demand markets is the end user. The "collar" for a demand hedger is to buy calls and sell puts. The grain markets and energy markets are good examples of demand skews. This type of skew is also known as a 'hockey stick'.

- **Smile Skew**: The third type of skew is called a smile skew. The smile skew is illustrated by higher implied volatility as you move away from the at-the-money strike. The at-the-money strike would have the lowest implied volatility while the strikes moving up and down would have progressively higher implied volatility. The smile skew is generally only observed in the currency markets. The natural hedger has to hedge currency moves in either direction depending of whether they have accounts payable or receivable in the foreign currency.

We show these shapes in Figs. 1.4. Note that the EURUSD smile is symmetric around the ATM. Both currencies are perceived to be stable with similar risks. However, due to the currency being a pair, the risk is that either currency of the exchange rate can collapse. This means that the smile can be skewed to one side due to the country risk or stability of a particular exchange rate, e.g. the markets perceive the Kenyan country risk as being higher than the USD or Euro country risk. Look at the shape of the smile of the Brazilian Real and Euro in Fig. 1.5 [DW 08]. It is not symmetrical around the ATM. This was also shown in Fig. 1.4. This shows that while indices across the globe mostly have similar shaped skews, the skews for different currency pairs can be vastly different.

1.11.4 Delta Hedging and the Skew

Another view on the skew is the fact that if the markets go down they tend to become more volatile. Equity markets crash downward but hardly ever 'crash' or gap upward. However, the currency market with its smile do 'crash' upward and downward. This alone does not explain the skew as realised volatility is the same regardless of any

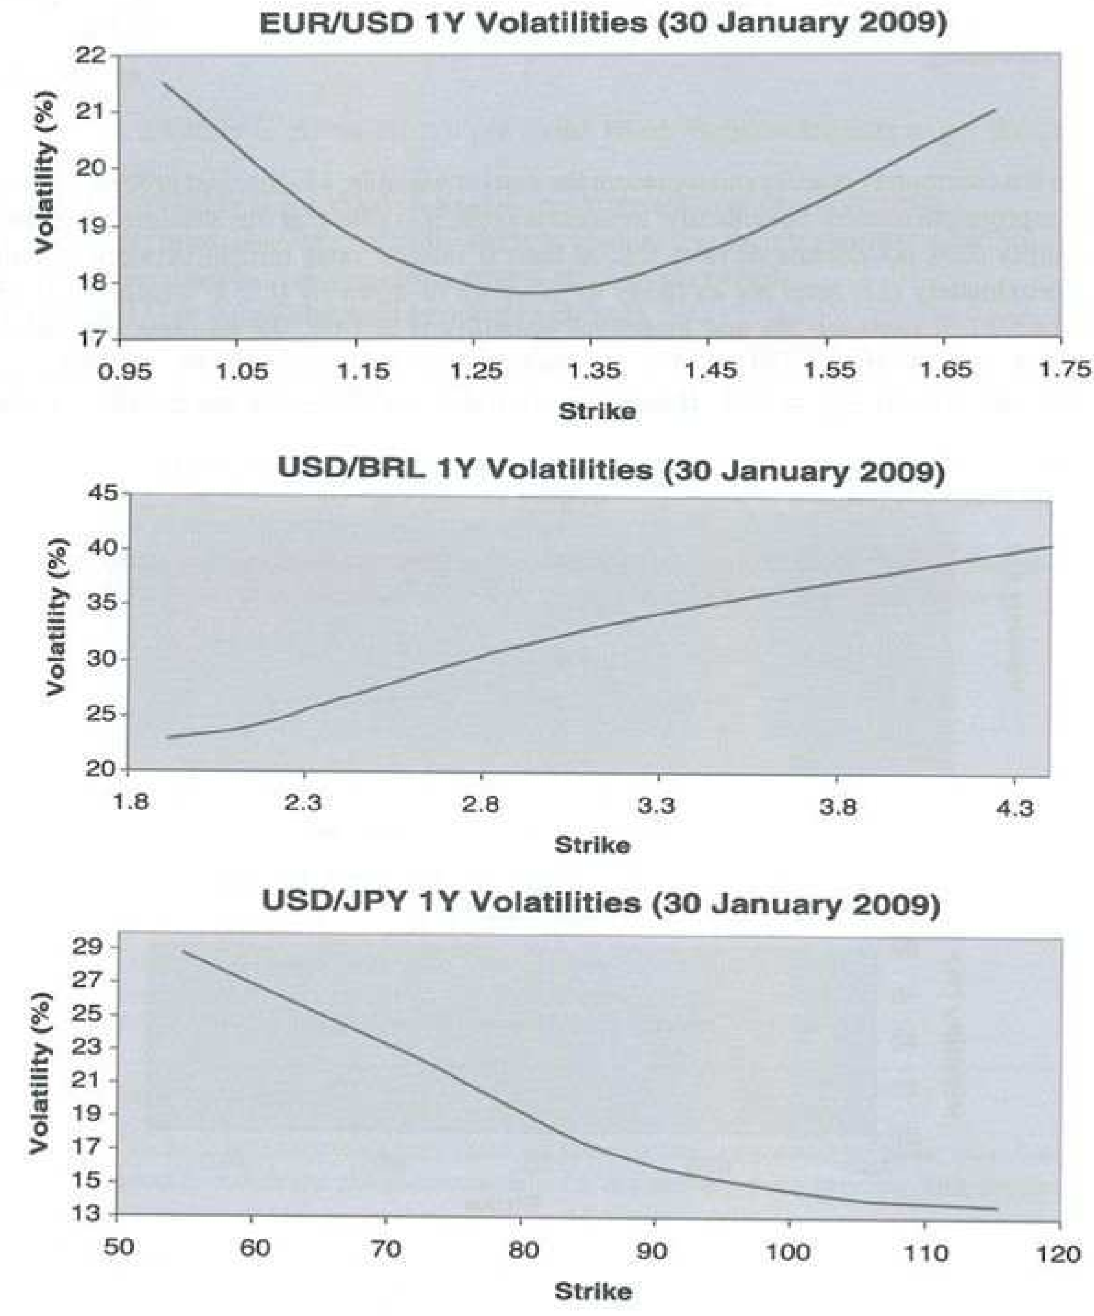

Figure 1.4: Different shape currency skews.

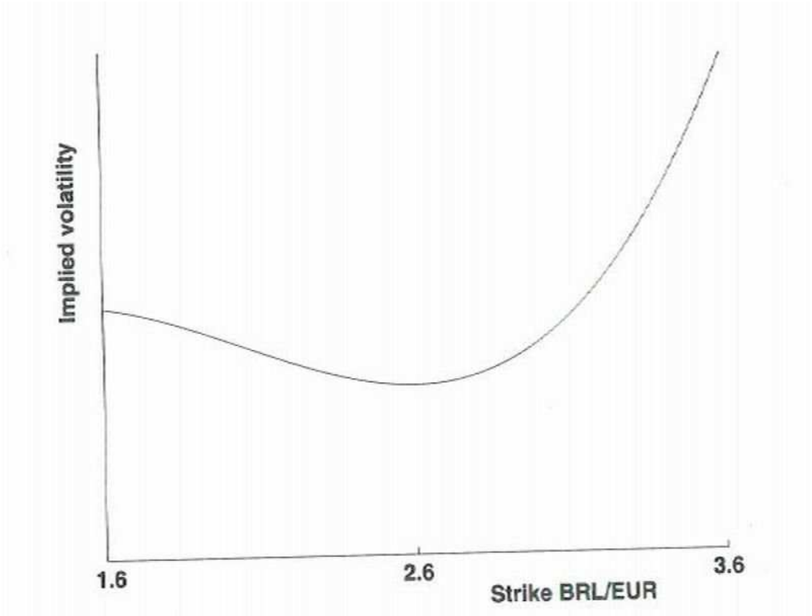

Figure 1.5: The smile for the BRLEUR which is not symmetrical.

strike price. The existence of the skew is apparently telling us that this increase in volatility has a bigger impact on lower strike options than on higher strike options.

The reason behind this becomes apparent when thinking in terms of realised gamma losses as a result of rebalancing the delta of an option in order to be delta hedged [DW 08]. In a downward spiraling market the gamma on lower strike option increases, which combined with a higher realised volatility causes the option seller to rebalance the portfolio more frequently, resulting in higher losses for the option seller. Naturally the option seller of lower strike options wants to get compensated for this and charges the option buyer by assigning a higher implied volatility to lower strikes. This principle applies regardless of the of the in- or out-of-the-moneyness of an option. Whether it is a lower strike in-the-money call or a lower strike out-the-money put, makes no difference from a skew perspective. Indeed if there were a difference there would be an arbitrage opportunity.

1.11.5 The Term Structure of Volatility

Another aspect of volatility that is observed in the market is that at-the-money options with different expiries trade at different volatilities. The at-the-money volatilities for different expiry dates are usually decreasing in time meaning shorter dated options trade at higher volatilities to longer dated ones. However, since the 2008 financial crisis we see the term structure tend to slope upward more often. This provides another method for traders to gauge cheap or expensive options. A downward sloping term structure is natural in the market because short dated downside options need to be delta hedged more often resulting in higher losses — downside short dated options have higher Gammas than long dated ones.

The term structure of volatility arises partly because implied volatility in short options changes much faster than for longer options and partly due to the assumed mean reversion of volatility. The effect of changes in volatility on the option price is also less the shorter the option.

It is well-known that volatility is mean reverting; when volatility is high (low) the volatility term structure is downward (upward) sloping. We therefore postulate the following functional form for the ATM volatility term structure

$$\sigma_{atm}(\tau) = \frac{\theta}{\tau^{\lambda}}. \tag{1.41}$$

Here we have

- τ is the months to expiry,
- λ controls the overall slope of the ATM term structure; $\lambda > 0$ implies a downward sloping ATM volatility term structure, whilst a $\lambda < 0$ implies an upward sloping ATM volatility term structure, and
- θ controls the short term ATM curvature.

We show the current term structure for USDZAR in Fig. 1.11.5)

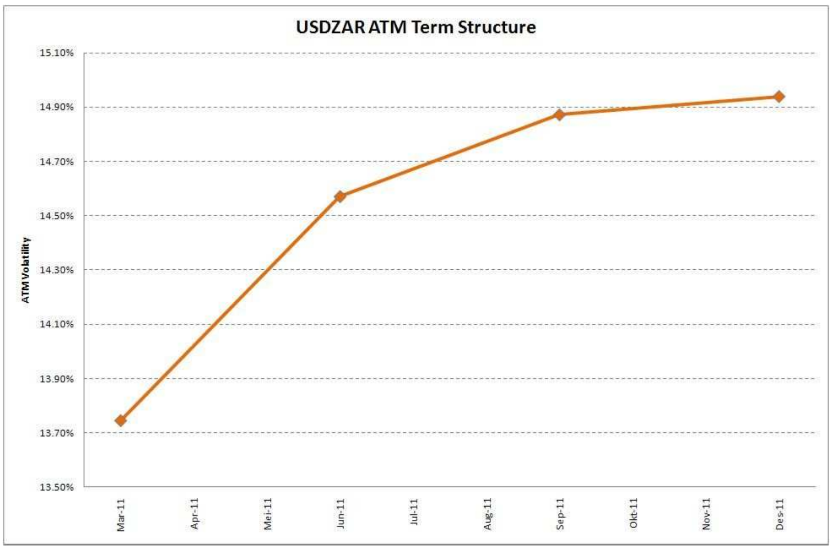

Figure 1.6: The USDZAR market fitted at-the-money volatility term structure during February 2011.

Moneyness	Mar-11	Jun-11	Sep-11	Des-11
70%	3.515%	8.263%	5.376%	3.556%
75%	4.873%	5.327%	3.542%	1.929%
80%	6.745%	2.634%	1.618%	0.535%
85%	7.460%	0.561%	0.062%	-0.439%
90%	5.655%	-0.561%	-0.687%	-0.815%
95%	2.431%	-0.704%	-0.628%	-0.618%
100%	0.000%	0.000%	0.000%	0.000%
105%	0.082%	1.401%	0.982%	0.900%
110%	2.404%	3.266%	2.215%	1.996%
115%	6.193%	5.340%	3.625%	3.215%
120%	10.678%	7.372%	5.137%	4.485%
125%	15.087%	9.106%	6.678%	5.731%
130%	18.660%	10.298%	8.173%	6.887%

Table 1.1: The official floating Yield-X volatility skew for USDZAR during February 2011

1.11.6 What is a Volatility Surface?

Combining the ATM term structure of volatility and the skew per expiry date, will render a 3 dimensional graph (time to expiry versus strike versus volatility). This is known as the volatility surface.

1.11.7 Skews in South Africa

In Table (1.1) we give the February 2011 volatility surface for the USDZAR contracts traded on Yield-X. We show the skew on a relative or floating basis where the ATM strikes are given as 100% and the ATM floating volatilities are given as 0%. The first column shows the percentage of moneyness[7] and the second column shows the relative volatility i.e., what number has to be added or subtracted from the ATM volatility to give the real volatility.

The current skews are plotted in Figure 1.7. Yield-X supplies the skews on a monthly basis. The currency option market in South Africa is new and the liquidity not great. Yield-X does not generate these itself. They are supplied by a London based company called 'Super Derivatives[8].' There are other companies who supply skew data to market players. Another one active in the South African market is called 'Markit[9].'

[7]Moneyness shows how far the strike is from the ATM strike.
[8]https://www.superderivatives.com/
[9]http://www.markit.com/en/?

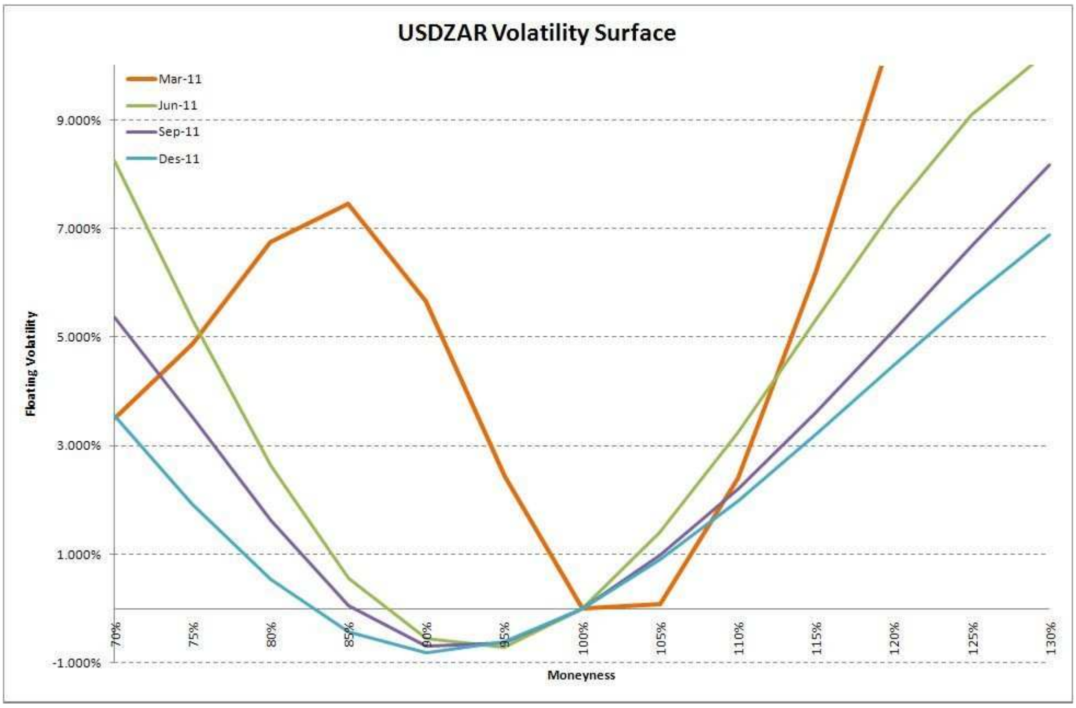

Figure 1.7: USDZAR volatility surface during February 2011.

So what happens if the strike price lies between two strikes given in the skew? Then we have to interpolate. We do straight line linear interpolation and this is explained next.

1.12 Sticky Volatilities

Sticky strike and sticky delta are trader phrases to describe the behaviour of volatility when the price moves. It can be shown that both the sticky Delta and sticky strike rules produce arbitrage opportunities should the volatility surface move as predicted by them. This is the reason why they are mainly regarded as quoting mechanisms and not expressions of actual behaviours of volatility surfaces [Ca 10].

1.12.1 Sticky Delta

In the sticky delta model (also known as a relative or floating skew), the implied volatility depends on the moneyness only (spot divided by strike - S/K). The ATM implied volatility does not change as the underlying spot changes. This entails that the smile floats with moneyness (or Delta) as spot is shifted such that the delta of options are preserved. This means that when the underlying asset's price moves and the Delta of an option changes accordingly, a different implied volatility has to be used in the *Black & Scholes* model. In this model, moneyness is plotted against

Delta	Mar-11	Jun-11	Sep-11	Des-11
Δ^p_{25}	1.4873%	-0.7198%	-0.0999%	-0.3640%
Δ^p_{30}	1.1899%	-0.5758%	-0.0799%	-0.2912%
Δ^p_{35}	0.8924%	-0.4319%	-0.0600%	-0.2184%
Δ^p_{40}	0.5949%	-0.2879%	-0.0400%	-0.1456%
Δ	0.0000%	0.0000%	0.0000%	0.0000%
Δ^c_{40}	0.0579%	0.7247%	1.3071%	3.0039%
Δ^c_{35}	0.0672%	1.0569%	1.6176%	3.0032%
Δ^c_{30}	0.0764%	1.3891%	1.9281%	3.0026%
Δ^c_{25}	0.0857%	1.7213%	2.2386%	3.0020%

Table 1.2: The official Yield-X volatility skew for USDZAR for February 2011 using Deltas instead of moneyness

relative volatility (difference in volatility from the ATM volatility). The sticky-delta rule quantifies the intuition that the current level of at-the-money volatility — the volatility of the most liquid options — should remain unchanged as spot changes. We listed such a skew in Table 1.1. We show the same skew in Table 1.2 where we map the volatilities against the put Deltas on the left of the ATM and call Deltas on the right of the ATM.

1.12.2 Sticky Strike

In a sticky strike model ("absolute skew"), the implied volatility of each option is constant as the spot changes or the volatility of a given strike is unaffected by a change in price. Another way to put this is that skew is kept fixed at strikes as the spot is shifted. This means the volatility is independent of the spot, it depends on the strike only. A sticky strike skew plots volatility against actual strikes. In Table 1.3 we give the February USDZAR skew for options traded on Yield-X. This is the same skew as that listed in Table 1.1. Intuitively, "sticky strike" is a poor man's attempt to preserve the Black-Scholes model. It allows each option an independent existence, and doesn't worry about whether the collective options market view of the spot is consistent.

1.12.3 Which is better: sticky strike or sticky delta?

There is no conclusion yet although market players seem to prefer the sticky strike model. *Rubinstein* and *Jackwerth* in 1997 compared several models and found that sticky-strike best predicts future smiles [JR 96]. However, *Derman* found in 1999 that market conditions should set the tone [De 99].

Strike	14-Mar-11	13-Jun-11	19-Sep-11	19-Des-11
6.7500	16.5626	13.8135	14.1262	14.1205
6.8000	16.1105	13.8181	14.1295	14.1259
6.8500	15.6748	13.8394	14.1455	14.1409
6.9000	15.2617	13.8770	14.1736	14.1653
6.9500	14.8770	13.9308	14.2133	14.1985
7.0000	14.5270	14.0002	14.2639	14.2403
7.0500	14.2177	14.0850	14.3248	14.2902
7.1000	13.9552	14.1849	14.3954	14.3480
7.1500	13.7455	14.2994	14.4751	14.4131
7.2000	13.5934	14.4282	14.5633	14.4854
7.2500	13.4980	14.5710	14.6593	14.5643
7.3000	13.4574	14.7274	14.7626	14.6495
7.3500	13.4693	14.8968	14.8725	14.7407
7.4000	13.5317	15.0785	14.9885	14.8375
7.4500	13.6424	15.2718	15.1104	14.9395
7.5000	13.7994	15.4761	15.2379	15.0464
7.5500	14.0005	15.6908	15.3709	15.1579
7.6000	14.2436	15.9151	15.5092	15.2739
7.6500	14.5266	16.1483	15.6526	15.3942
7.7000	14.8474	16.3899	15.8009	15.5187
7.7500	15.2040	16.6391	15.9539	15.6471
7.8000	15.5940	16.8954	16.1114	15.7793
7.8500	16.0156	17.1579	16.2733	15.9151
7.9000	16.4664	17.4261	16.4393	16.0544
7.9500	16.9446	17.6993	16.6094	16.1969
8.0000	17.4478	17.9769	16.7832	16.3424

Table 1.3: The official sticky strike Yield-X volatility skew for USDZAR during February 2011

Sticky Delta

He found that if the markets are trending, where the market is undergoing significant changes in levels without big changes in realized volatility, sticky delta rules [De 99]. Then, in the absence of a change in risk premium or an increased probability of jumps, the realised volatility will be the dominant input to the estimation of the implied volatility of (high-Gamma) at-the-money options. As the underlying moves to new levels, it is sensible to re-mark the current at-the-money implied volatility to the value of the previous at-the-money volatility. The at-the-money volatility "stick" to the ATM spot level.

Sticky Strike

Sticky strike rules if the market trades in a range (jumps are unlikely) without a significant change in the current realised volatility. As markets are range bound most of the time, sticky strike is the most common rule used.

Daglish, Hull and *Suo* concluded in 2006 that all versions of the sticky strike rule are inconsistent with any type of volatility smile or volatility skew [DHS 06]. They state that "If a trader prices options using different implied volatilities and the volatilities are independent of the asset price, there must be arbitrage opportunities." They further found that the relative sticky delta rule can be at least approximately consistent with the no-arbitrage condition.

1.13 The Binomial Tree

One of the assumptions *Black & Scholes* made was that the underlying asset is traded on a continuous basis and that delta hedging is done on a continuous basis — returns are normally distributed. The discrete version of the normal distribution is the binomial distribution. In 1976 *Cox, Ross* and *Rubinstein* realised this and constructed a tree based method to value derivatives [CR 85]. Stock returns are assumed to be governed by a discrete probability measure; in this case the binomial distribution. The binomial model breaks down the time to expiration into potentially a very large number of time intervals, or steps. Using probability theory, a tree of stock prices is produced working forward from the present to expiration. This is graphically shown in Fig. 1.8. The consequence of this methodology is that, at expiry of an option, we have a discrete set of possible stock prices. We can now use these prices to value any derivative security. We are doing what Louis Bachelier stated more than 100 years ago. In his famous dissertation he mentioned that "you have to 'model' spot price movements' before you 'model' option values [We 06]. He realised that the unknown spot price in the future is just a scaled version of the current spot price.

We start by making a very restrictive assumption. Let's start with a stock price S and we assume that after a small time interval Δt the price can take two values

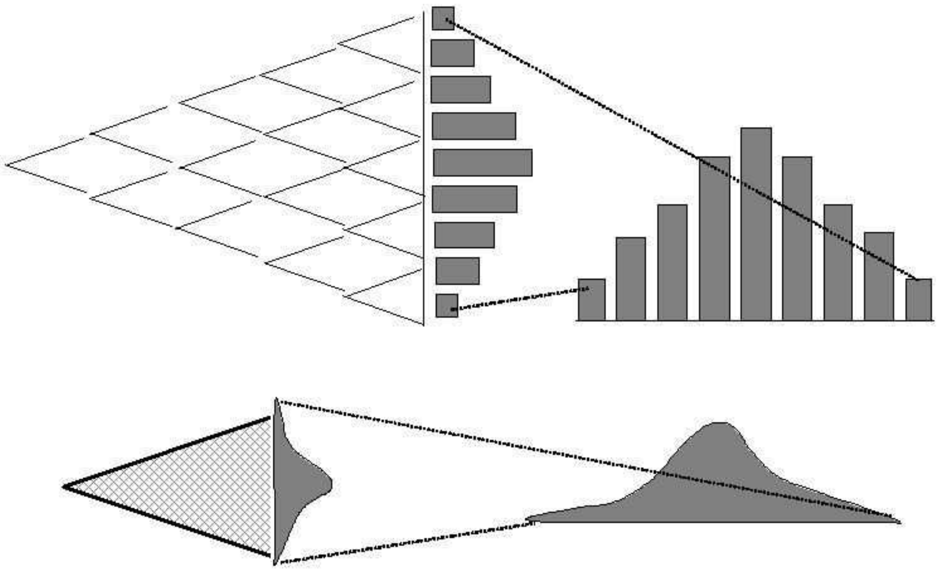

Figure 1.8: The binomial distribution is the discrete version of the normal distribution.

only: it can either move up to a new price of uS or down to dS (where u and d are numbers such that $u > d$). We also assume that the transition probability of moving to uS is p and to dS, $(1-p)$. Further, after another small time interval Δt all these prices can only take two values each. This leads to a tree of prices as shown in Fig. 1.9.

The probabilities are given by

$$\begin{aligned} p &= \frac{a-d}{u-d} \\ u &= \frac{1}{2a}\left[a^2+b^2+1+\sqrt{(a^2+b^2+1)^2-4a^2}\right] \simeq e^{\sigma\sqrt{\Delta t}} \\ d &= \frac{1}{u} \end{aligned} \tag{1.42}$$

where

$$a = e^{(r_d-r_f)\Delta t}. \tag{1.43}$$

$$b^2 = a^2(e^{\sigma^2\Delta t}-1).$$

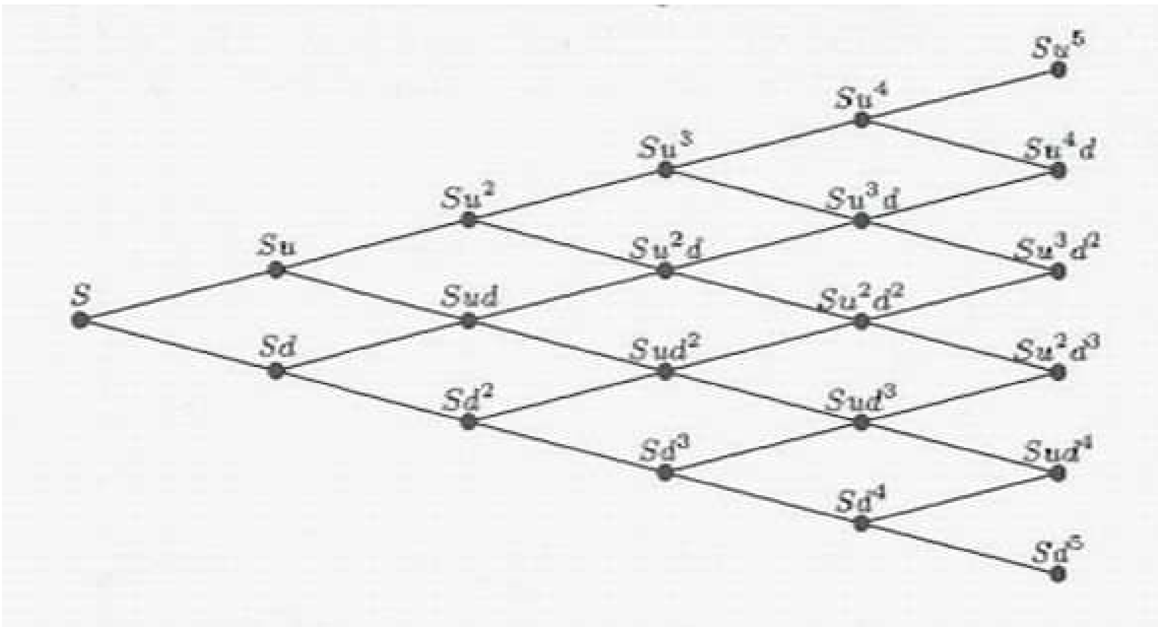

Figure 1.9: Five step tree [Ha 07].

Chapter 2

Practical Use of Option Models

2.1 Hedging Options

If a trader sells an option to a client, he is faced with problems in managing its risk. If a similar option trades on an exchange the trader can lay off his risk efficiently. But, if the option has been tailored to the needs of the client, hedging his exposure is far more difficult.

A trader who writes an option and who takes the risk onto his book, is given the ability to trade that risk. It is in the efficient trading of this risk where huge profit potential lies.

An option's price depends on six parameters. They are

- the current FX rate;
- the strike price;
- the time to expiration;
- the local riskfree interest rate;
- the foreign riskfree interest rate;
- the volatility.

In each case the relationship is nonlinear and the trader needs to manage these risks. There are two main approaches (not mutually exclusive) to managing market risk. One involves quantifying and controlling each of the above risks separately and the other involves scenario analysis and stress testing. We will describe both approaches in this section.

Most option traders use sophisticated hedging schemes. As a first step they attempt to make their portfolio immune to small changes in the price of the underlying asset in the next small interval of time. This is known as *delta hedging*. Then they also look at the *Gamma* and *Vega*. By keeping *Gamma* close to zero, a portfolio can be made relatively insensitive to fairly large changes in the price of the asset; by keeping *Vega* close to zero it can be made insensitive to changes in the volatility. In the next few sections we will describe these approaches in more detail.

2.1.1 The Delta

The Delta is an important parameter in the pricing and hedging of options. But, what is the Delta? The word delta comes from the fourth letter of the Greek alphabet and is universally abbreviated as a triangle. The Delta is the amount that an option will theoretically change in price for a one-point move in the stock. Analytically we define the Delta as

$$\Delta = \frac{\partial V}{\partial S} = \phi N(x)e^{-d\tau} \tag{2.1}$$

with V the price of an option defined in Eq. (3.3). From Eq. (2.1) we see the Δ is the rate of change of the option's price with respect to the price of the underlying stock. We can approximate the Delta with

$$\frac{\Delta V}{\Delta S} \tag{2.2}$$

where ΔS is a small change in stock price and ΔV is the corresponding change in the option's price[1]. For example, if an option price was R2.00 and that option had a Delta of 50 and the underlying stock were to move from R50 to R51 (and everything else remained constant), the option should move from 2.00 (where it was) to 2.50 - 50% of the stock move of R1.00.

From (2.1) we can now derive the following relationships: if the option is far out-of-the-money ($S \ll K$) we have

$$\Delta \simeq 0.$$

If the option is deep in-the-money we have

$$\Delta \simeq \phi.$$

Δ is the probability that the option will end in-the-money.

Remember, Deltas are not constant. They change as the stock moves. In Figure (2.1) we show the variation of the Δ of a call and put. In Figure (2.2) we show the typical patterns for the variation of delta with time to expiry for ATM, ITM and OTM options.

Delta measures the slope of the option's price curve.

[1]the Δ of a stock or future is just 1 or -1.

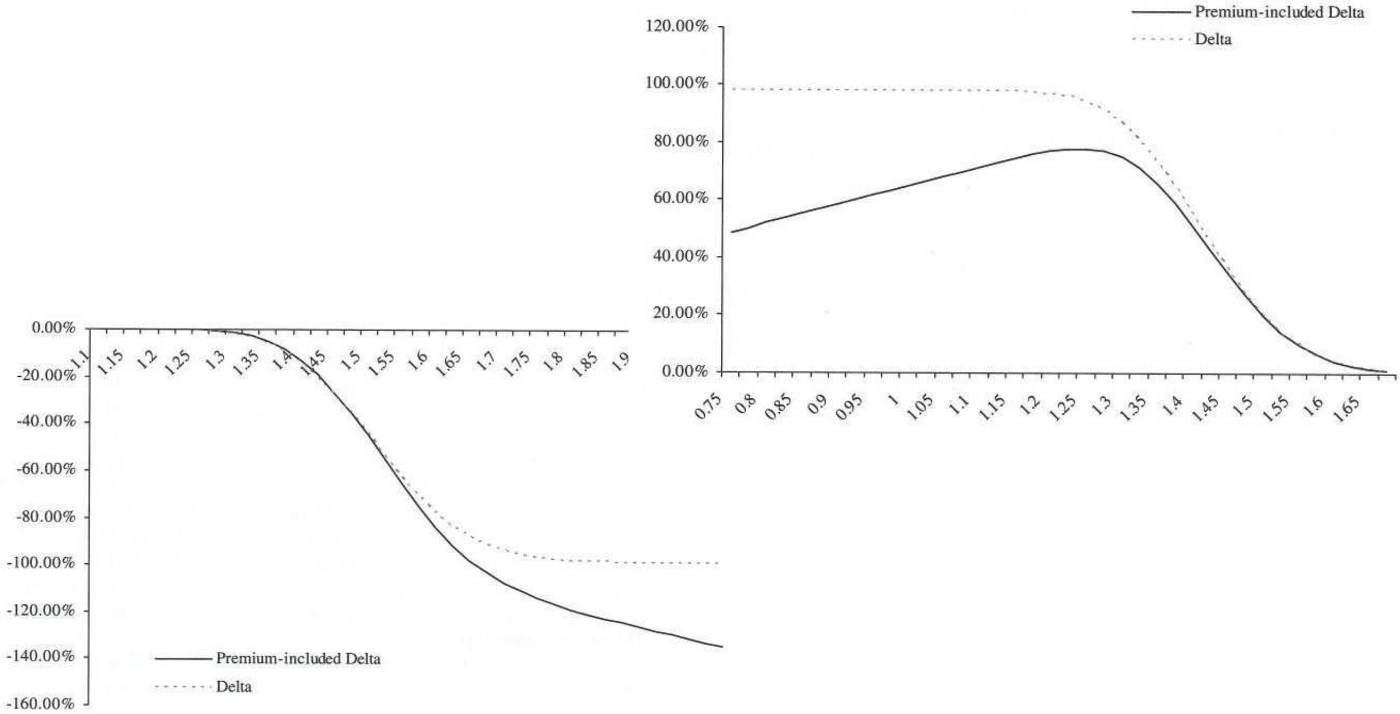

Figure 2.1: The Delta of a call.

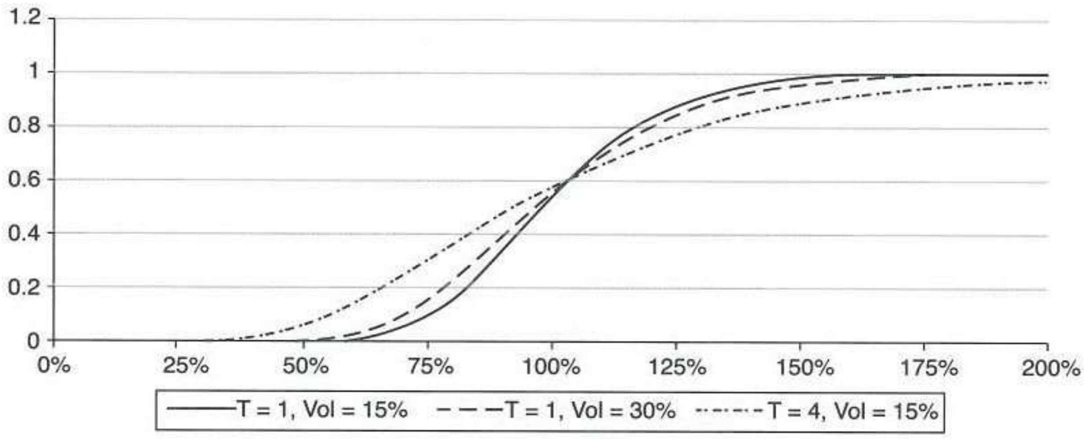

Figure 2.2: The Delta as a function of the time to expiry.

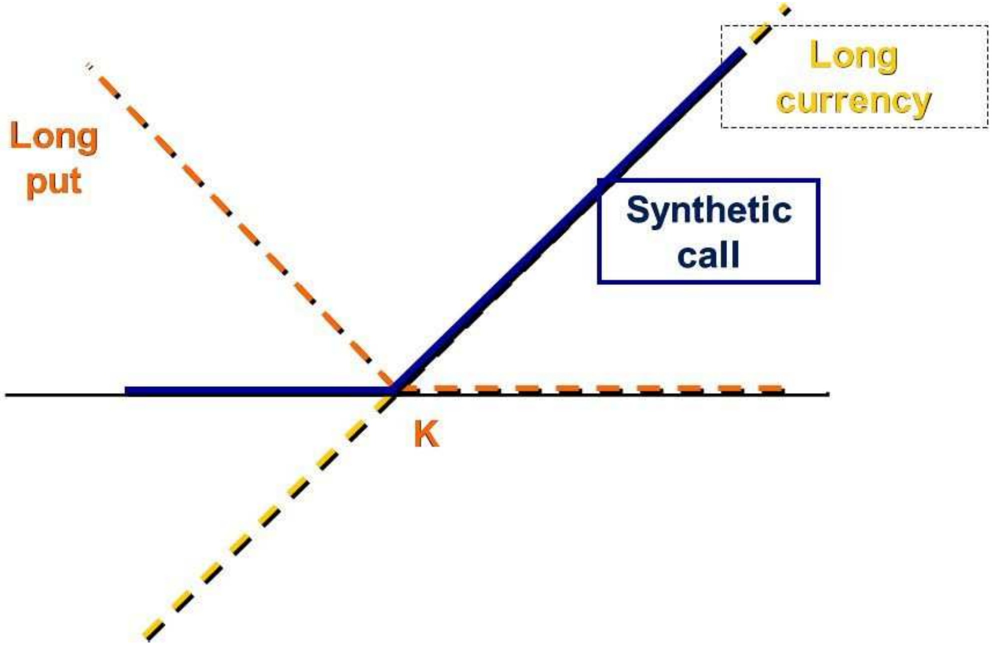

Figure 2.3: Delta hedging a put.

Delta as the Hedge Ratio

Black & Scholes showed that the Δ is a measure of the ratio of option contracts to the underlying asset in order to establish a neutral hedge - we can now say that *Black & Scholes* valued options by setting up a delta-neutral position and arguing that the return on the position in a short period of time equals the riskfree interest rate.

Stated differently we can say that the Δ is a measure of how fast an option's value changes with respect to changes in the price of the underlying asset and allows you to equate your options with shares of stock. For example, if you are long a call with a delta of 37, you can expect the call to move up 37 cents for a one-point move in the stock. This means that owning one call is equivalent in price movement to owning 37 shares of that stock. This results in the name *hedge ratio*. This is very helpful when you consider how many calls or puts to use to hedge a portfolio of stocks or vice versa.

Delta hedging aims to keep the total wealth of the financial institution as close to unchanged as possible [Hu 06]. Figure (2.3) shows the total position if we fully hedge a put option.

If we purchase an option we can trade in the cash instrument (called "trading spot" or "trading the cash"). By buying the option we pay the premium upfront. The Delta will then show us that we have to buy the cash at the lows and sell at the highs - we thus make money by delta hedging. In theory, if we hedge continuously we should make exactly what we paid for the option. In practise traders do not hedge continuously and they usually have a view. If they think the option was underpriced (volatility lower than it actually is) they then hope to realise more profit from trading the cash than what they paid for the option.

On the other hand, if we sell the option, we earn the premium. The Delta will then show us to buy the cash at the highs and sell it at the lows - we thus loose money by delta hedging. Again, if we sold the option at a higher than actual volatility, we can hope to loose less money than the premium we received.

The Delta as a Probability

In practise Deltas are quoted like integers but are actually percentages. A 50 delta means 50 percent — that is, a call with a delta of 50 would be expected to rise (or fall) one-half point if the FX rate rose (or fell) one point. A call with a delta near 100 (100 percent) could be expected to move almost point-for-point as the underlying FX rate moves. The same applies to puts, but with one difference: puts have negative deltas. Thus, a put with a delta of 50 would fall (or rise) 0.50 points if the FX rate rose (or fell) one point.

Of what use is this information? Delta is used to set up expectations. Imagine that you are considering purchasing a call on USDKES trading at 82. Should you buy the 82-strike call or the 86-strike call? The delta is an important piece of information that may help you with this decision. Let's assume the 86-strike call is offered in the market for KSh0.75 and has a delta of 25. Since you expect the rate to rise to 84, with all else being equal, the delta tells you that the option might rise around KSh0.50 (.25 for each one point move in the FX rate). This may or may not be the return you were expecting. Using delta, you can compare how much the option is expected to move for each of the calls in which you are interested.

Using Futures

In practise hedging is often carried out using a position in futures rather than one in the underlying stock. If you are due to sell an asset in the future, then it is possible to hedge prices by taking a short futures position (short hedging). If you are due to buy an asset in the future, then it is possible to hedge prices by taking a long futures position (long hedging).

When (and only when) the exact contract details can be replicated with futures contracts does short or long hedging produce a perfect hedge. In most situations it doesn't and the futures hedge either needs to be closed out before or after the options expiry date. This introduces basis risk. Basis risk = spot price of asset to be hedged - futures price of contract used. If perfectly hedged: basis risk = 0 at expiry (or + or - otherwise).

Because one uses a future to hedge in the spot market, the delta for the spot market should be mapped to the equivalent Δ for a futures contract. The correct number of futures to be sold or bought are then given by

$$\Delta_F = e^{-(r-d)(T^\star - t)}\Delta \tag{2.3}$$

Figure 2.4: The basis for a currency futures contract.

where $T^\star$ is the maturity of the futures contract. This shows that $e^{-(r-d)(T^\star - t)}$ futures contracts has the same sensitivity to stock price movements as one stock.

In hedging exchange traded currency futures, we can use the underlying currency pair and not the future (future contracts illiquid). In doing that we thus need to map the futures Δ to the spot Δ. This is done by inverting Eq. (2.3).

In using the future to hedge a spot position one should take the following risks into account:

Basis Risk If the OTC option expires on a different date than a close-out date there can be basis risk. The basis is defined as follows:

$$\text{basis} = \text{spot price of asset to be hedged} - \text{futures price of contract used}.$$

Figure (2.4) shows the basis for the Mar 02 Alsi futures contract. If the asset to be hedged and the assets underlying the future contract are the same, the basis should be zero at expiration of the futures contract (this is called "pull to par"). In general, basis risk increases as the time difference between the hedge expiration and OTC's expiration increases.

There are some sources of basis risk:

- Changes in the convergence of the futures price to the spot price.
- Changes in factors that affect the cost of carry: storage and insurance costs, opportunity cost.

- Different natures of mismatched assets.
- Maturity mismatch.
- Liquidity difference.
- Credit risk difference.
- Random Deviation from the Cost-of-Carry Relation.

Due to basis risk, the equivalent hedge given in Eq. (2.3) might not seem perfect. A trader can get profit and loss swings due to the fact that the spread between the spot and futures contract changes as the market moves. What one needs to realize is that the delta for the spot market is actually mapped to the equivalent Δ for the forward contract that trades at fair value when 2.3) is used. This is not always the value where the equivalent future is trading on the exchange. Basket arbitrageurs are then active in the market. Such a hedge is still perfect though, because, over the life of the futures contract, due to the "pull to par" effect, the total profit and loss scenario is equivalent to a hedge with the actual underlying.

Rollover Risk Sometimes, the expiration date of the OTC option is later than the maturity dates of all the futures contracts that can be used. The hedger must then roll the hedge forward. Hedges can be rolled forward many times.

When rolling a contract forward, there is uncertainty about the difference between the futures price for the contract being closed out and the futures price for the new contract. Hedgers reduce the rollover risks by switching contracts at favourable times. The hedger hopes that there will be times when the basis between different futures contracts are favourable for a switch.

2.1.2 Gamma

The Γ is the rate at which an option gains or loses deltas as the underlying asset's price move up or down. It is a measure of how fast an option is changing its market characteristics and is thus a useful indication of the risk associated with a position. We have

$$\Gamma = \frac{\partial^2 V}{\partial S^2} = \frac{N'(x)e^{-d\tau}}{S\sigma\sqrt{\tau}} \tag{2.4}$$

where $N'(x)$ is the derivative of the cumulative normal distribution function $N(x)$ defined as

$$N'(x) = \frac{1}{2\sqrt{\pi}}e^{-\frac{x^2}{2}}. \tag{2.5}$$

If Γ is small, Δ changes slowly and adjustments to keep the portfolio Δ-neutral need only be made relatively infrequently. If Γ is large, however, changes should be

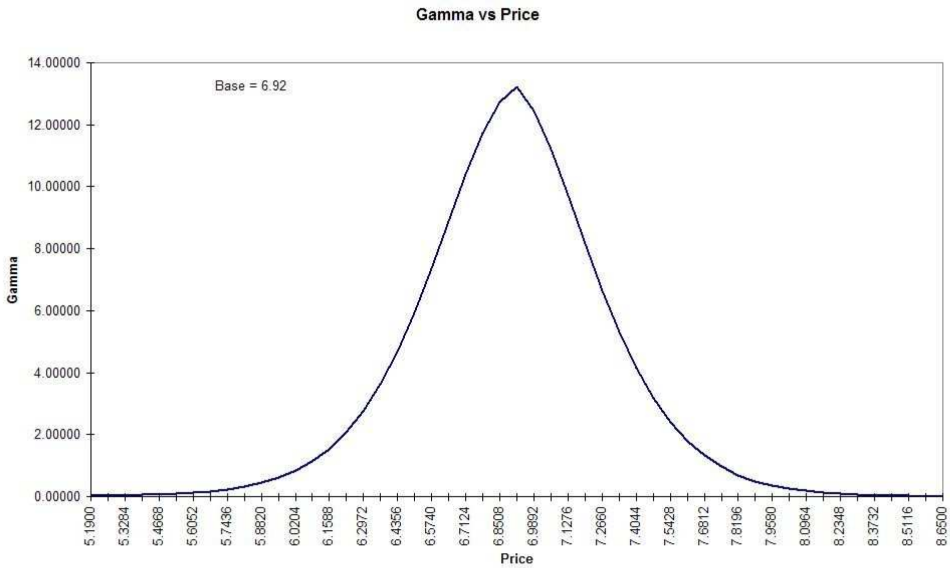

Figure 2.5: The gamma of an option.

made frequently because the Δ is then highly sensitive to changes in the price of the underlying asset. In Figure (2.5) we show the variation of the Gamma with the stock price.

Gamma measures the curvature of the option's price curve. Intuitively gamma jointly measures how close the current market is to the strike and how close the option is to expiration. The closer the market is to the strike price and the closer the maturity of the option is to the expiration date, the higher the gamma. This is shown in Figure (2.5).

On the date of expiry if the option is very close to ATM, the Gamma is infinite and the Delta jumps around violently. Consider the situation where you are a few minutes away from expiry and the spot price jumps from ITM to OTM. The Delta wil jump from 1 to 0 all the time. What do you do? Pray!

All buyers of options gain from the movements in the price of the underlying asset. Holders of options are thus "long Gamma" and writers of options are "short Gamma". We can also say that "Gamma is a measure of the exposure the position has to a change in the actual volatility of the underlying market" [To 94].

If one wants to ascertain whether a total position is positive or negative Gamma, one looks at the profit/loss profile prior to expiry. If the curved relationship to price of the underlying asset is curved convex[2], then the position is Gamma positive and you make money if the market moves. If the curve turns downward like a frown the position is Gamma negative and you loose money if the market moves.

The Gamma of a future or spot FX contract is zero.

[2]Curve turns upwards like a smile.

2.1.3 Theta

Remember, options lose value as expiry approaches. The Θ is the "time decay factor" and measures the rate at which an option loses its value as time passes. In practise we use two Thetas, the theoretical one and an approximate time decay value.

Theoretical Theta

$$\Theta = \frac{\partial V}{\partial \tau} = \phi r K e^{-r\tau} N(\phi y) - dSN(\phi x)e^{-d\tau} + \frac{SN'(x)\sigma e^{-d\tau}}{2\sqrt{\tau}} \tag{2.6}$$

Time Decay

$$\Theta_{td} = V(t+x) - V(t). \tag{2.7}$$

Here x is a day count parameter. Usually $x = 1$ meaning that one calculates the option's value tomorrow ($V(t+1)$) and one subtracts from that the option's value for today ($V(t)$). This give the amount of cash one looses or gains every day. For a weekend set $x = 3$.

Theta versus Gamma

We can show that if a portfolio is delta-neutral ($\Delta = 0$) that

$$\Theta + \frac{1}{2}\sigma^2 S^2 \Gamma = rf \tag{2.8}$$

with f the value of our portfolio. From this we can conclude that the size of the Γ correlates to the size of the Θ position where a large positive Γ goes hand in hand with a large negative Θ. A large negative Γ correlates with a large positive Θ. This means that every option position is a trade-off between market movement and time decay. Thus if Γ is large, market movement will help the trader but time decay will hurt him and vice versa.

2.1.4 Rho

The Rho of a portfolio of derivatives is the rate of change of the value of the portfolio with respect to the interest rate.

$$\rho = \frac{\partial V}{\partial r} = \phi K \tau e^{-r\tau} N(\phi y). \tag{2.9}$$

An increase in interest rates will decrease the value of an option by increasing carry costs. This effect is, however, outweighed by considerations of volatility, time to expiration and the price of the underlying asset.

2.1.5 Vega

The Vega measures the change in the option's price as volatility changes. In practise traders uses two Vega measurements: the theoretical Vega and the approximate one. "At-the-money" options are the most sensitive to volatility changes. If your Vega is positive and volatility increases, you make money; if volatility decreases, you loose money.

Theoretical Vega

$$\mathcal{V} = \frac{\partial V}{\partial \sigma} = S\sqrt{\tau}N'(x)e^{-d\tau}. \tag{2.10}$$

$N'(x)$ was defined in Eq. (2.5).

Volatility Decay

As with the approximate time decay we now say that

$$\mathcal{V} = V(\sigma \pm x) - V(\sigma). \tag{2.11}$$

Here x is a percentage parameter. Usually $x = 1\%$ meaning that one calculates the option's value with the volatility plus 1% and one subtracts from that the option's value. This give the amount of cash one gains or looses if volatility increases or decreases with 1%.

2.1.6 Other Risk Parameters

The following three risk parameters are also used by practitioners [Ha 07]

$$\begin{aligned} \text{Speed} &= \frac{\partial^3 V}{\partial S^3} \\ \text{Charm} &= \frac{\partial^2 V}{\partial S \partial t} \\ \text{Colour} &= \frac{\partial^3 V}{\partial S^2 \partial t} \end{aligned}$$

2.2 Hedging in Practise

In practise traders do not rebalance their portfolios continually to maintain delta-neutrality, gamma-neutrality, vega-neutrality and so on. Problems might be transaction costs that make frequent rebalancing very expensive. Liquidity of the hedge stock also plays a role. Rather than trying to eliminate all risks, option traders therefore usually concentrate on assessing risks and deciding whether they are acceptable.

If the downside risk is acceptable, no adjustment is made; if it is unacceptable, they take an appropriate position in either the underlying or another derivative.

If liquidity of a specific stock is a problem, traders might look at other similar stocks which is highly correlated to the original stock. This carries risk but this risk might be acceptable. For instance, Rand Merchant Bank Holding (RMH) and Firstrand (FSR) are highly correlated due to RMH's holding in FSR. Investec SA and Investec Plc are similarly highly correlated. Currently Harmony (HAR) and Gold Fields (GFI) are highly correlated due to the gold price. Also, Implats (IMP) and Amplats (AMS) are also highly correlated.

We have mentioned before that traders can short stock if they can borrow scrip. This carries a cost and sometimes scrip is just not available which can make hedging problematic.

Traders also carry out scenario analysis. This involves calculating the gain/loss on their portfolio over a specified period under a variety of different scenarios. This is done by choosing a time period like one day, one week or one month. Then one chooses say two variables upon which the portfolio depends like volatility and interest rates. One then calculates a matrix of profit/loss experienced under these different scenarios.

Traders also stress test their portfolios. This involves testing the effect on a portfolio of extreme movements in the underlying variables. One would for instance test the potential losses if there is a market crash of say 20%. This would not just mean that the prices would fall by 20% but the volatility would increase as well.

2.3 More Realistic Greeks

Black & Scholes assumed that the underlying market is continuous, that is, the stock price moves continuously as time progresses. The delta calculated in Eq. (2.1) and the trading thereof as described above, thus assumes that a trader can and should re-balance his hedge continuously as time progresses. In practise this is not possible and traders need risk parameters that reflect the real dynamics of the market more closely.

2.3.1 Impact Delta

We have stated before that Delta hedging aims to keep the total wealth of the financial institution as close to unchanged as possible. So, is this achievable in practise? Can we calculate a Delta that will keep a trader's P/L constant over a range of FX rates AND where we keep the volatility skew in mind? The answer is yes. Another point to keep in mind is that the ordinary *Black & Scholes* Delta actually only hedges the cash flow and not the P/L.

We can replicate the theoretical delta and gamma numerically by changing the underlying's price by a small amount and noting the change in the option values. We then have

$$\begin{aligned}\Delta_n &= \frac{\delta V}{\delta S} \\ &= \frac{V(S+\delta S)-V(S)}{\delta S}. \qquad (2.12)\end{aligned}$$

This means we obtain the Black-Scholes (BS) value with an underying's price at $S$and we also determine the BS value if we move the underlying's price up by a small amount to $S+\delta S$. We call this the "up delta" because we shift the underlying's price up. δS is usually chosen to be 0.5% or 1% of the value of S or it can be chosen to be one i.e., $\delta S = 1$. Here, when we obtain the BS values for the different underlying prices, we keep all the input parameters the same e.g., if the volatility is 30%, we use 30% in calculating $V(S)$ and we use 30% in calculating $V(S+\delta S)$.

However, moving the underlying in one direction (up), is not necessarily a great approximation of the behaviour of the function on the way down. A more powerful tool is to move the underlying's price down as well. We can then calculate the "down delta". The correct numerical delta is then the average of the "up delta" and "down delta" given by

$$\begin{aligned}\Delta_n &= \frac{1}{2}\left[\frac{V(S+\delta S)-V(S)}{\delta S}+\frac{V(S)-V(S-\delta S)}{\delta S}\right] \\ &= \frac{1}{2\delta S}\left[V(S+\delta S)-V(S-\delta S)\right]. \qquad (2.13)\end{aligned}$$

From equations (2.12) and (2.13) we deduce that the delta is dependent on the magnitude of the change in the price of the underlying security i.e., δS. The increment is chosen at the discretion of the trader. It could be a function of either his utility curve or his estimation of future volatility. The numerical delta as defined in equation (2.13) has advantages in that it incorporates a little of the second and third derivatives that should complete the mathematical delta in any form of analysis[3]

The numerical gamma is defined as the change in the numerical delta for a small change in the underlying's price. We thus obtain the "up delta" and the "down delta" and take the difference between the two – we normalise this by dividing by δS. The numerical gamma is given by

$$\begin{aligned}\Gamma_n &= \frac{1}{\delta S}\left[\frac{V(S+\delta S)-V(S)}{\delta S}-\frac{V(S)-V(S-\delta S)}{\delta S}\right] \\ &= \frac{1}{\delta S^2}\left[V(S+\delta S)+V(S-\delta S)-2V(S)\right]. \qquad (2.14)\end{aligned}$$

This will be the change in the delta if the market changes by a very small amount.

[3] N. Taleb, *Dynamic Hedging*, Wiley (1997)

We mentioned that using the numerical delta as defined in (2.13) is usually "better" than the theoretical BS delta because it takes some of the nonlinear effects of the BS equation into account. It is also a more appropriate delta because it is unbiased between up and down market moves.

However, we can refine our estimation of the delta and gamma even further. We do this by realising that, as markets change, volatilities and dividend yields may change as well. We cannot incorporate these market realities into the theoretical BS delta and gamma but we can in the numerical delta and gamma. From (2.13) we then have

$$\Delta_I = \frac{1}{2\delta S}\left[V(S+\delta S,\sigma^+,d^+) - V(S-\delta S,\sigma^-,d^-)\right]. \tag{2.15}$$

Here, σ^+ is the volatility obtained from the relevant volatility surface when the underlying's price is $S+\delta S$ (remember, the moneyness between the strike K and S is different to the moneyness between K and $S+\delta S$ hence the volatilities are different as well). σ^- is the skew volatility when the underlying's price is $S-\delta S$.

If we use cash dividends we'll have $d^+ = d^-$. However, if we calculate a dividend yield from a cash dividend and use the dividend yield in the BS equation, $d^+ \neq d^-$ and one will have to obtain the correct yields for the up and down moves. The delta in equation (2.15) is called the impact or modified delta.

As before, δS is at the discretion of the trader. However, the size of the up and down move is usually a percentage of the price S. We now define an up move and a down move as follows

$$\begin{aligned} S^+ &= S\,(1+i) \\ S^- &= S\,(1-i)\,. \end{aligned} \tag{2.16}$$

i represents a basis point or percentage change in the underlying. i is typically 0.5%, 1%, 1.5%, 2%, 3%, 4% or 5%. This means the trader will calculate his delta for different size moves reflecting his view on volatility. Equation (2.15) now becomes

$$\begin{aligned} \Delta_I &= \frac{1}{2}\,S\imath\left[V(S^+,\sigma^+,d^+) - V(S^-,\sigma^-,d^-)\right] \\ &= \frac{V(S^+,\sigma^+,d^+) - V(S^-,\sigma^-,d^-)}{S^+ - S^-}. \end{aligned} \tag{2.17}$$

We thus calculate the option value at S^+ and we also calculate the option value at S^-. This is graphically illustrated in Fig. 2.6

Some VBA pseudo code for calculating the impact delta is given in Fig. 2.7.

2.3.2 Impact Gamma

The impact gamma will give the change in the impact delta if the underlying's level changes by the chosen percentage i. To calculate the gamma we now "bump" the

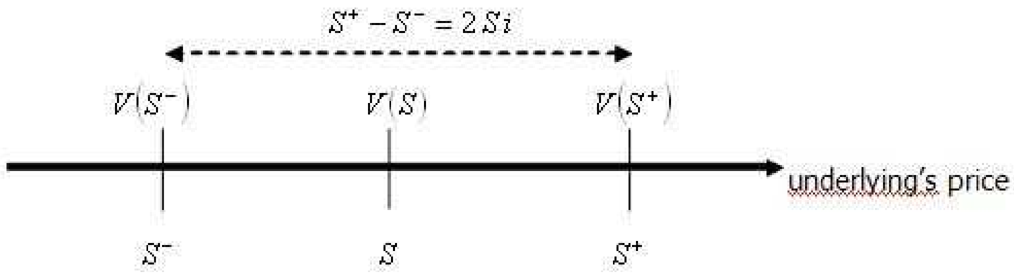

Figure 2.6: Calculating the Impact Delta

```
Function Delta_Impact(s, Strike, ValueDate, ExpiryDate, r_d, _
                      r_f, ATMVOL, phi, i)
‘ This function calculates the impact delta for a given option
‘ ATMVOL = the at-the-money volatility for this option
‘ r_d = domestic the risk-free interest rate
’ r_f = foreign interest rate
’ i = impact points
T = (ExpiryDate -- ValueDate)/365
‘ First, bump the spot price down by the impact points
Sdown = s * (1 - i)
’ sigma = volatility. Find it from volatility surface for the
’ strike and current spot price Sdown
sigma = FindSkewVol(Sdown, Strike, ExpiryDate, ValueDate, ATMVOL, _
                    r_d, r_f)
‘ Calculate the BS option value using Sdown
VDown = valuation(Sdown, Strike, r_d, r_f, sigma, T, phi)
‘ Phi = 1 for a call and -1 for a put
‘ Secondly, bump the spot price up by the impact points
Sup = s * (1 + i)
’ Find volatility from volatility surface for the strike
’ and current spot price SUp
sigma = FindSkewVol(Sup, Strike, ExpiryDate, ValueDate, _
                    ATMVOL, r_d, r_f)
VUp = valuation(Sup, Strike, r_d, r_f, sigma, T, phi)
‘ Impact delta is now
Delta_Impact = (VUp - VDown) / (Sup - Sdown)
End Function
```

Figure 2.7: VBA pseudo-code for calculating Impact Delta.

underlying's level up by i and we calculate the impact delta at that level. We then "bump" the underlying's level down by i and we again calculate the impact delta. Similar as before we then have

$$\begin{aligned} \Gamma_I &= \frac{1}{\delta S}\left[\Delta_I^{up} - \Delta_I^{down}\right] \\ &= \frac{1}{2\,S\,i}\left[\Delta_I^{up} - \Delta_I^{down}\right] \end{aligned} \tag{2.18}$$

where we have

$$\begin{aligned} \Delta_I^{up} &= \frac{1}{S^{++} - S^{+-}}\left[V\left(S^{++}, \sigma^{++}, d^{++}\right) - V\left(S^{+-}, \sigma^{+-}, d^{+-}\right)\right] \\ \Delta_I^{down} &= \frac{1}{S^{-+} - S^{--}}\left[V\left(S^{-+}, \sigma^{-+}, d^{-+}\right) - V\left(S^{--}, \sigma^{--}, d^{--}\right)\right]. \end{aligned} \tag{2.19}$$

Here we define (using equation (2.16))

$$\begin{aligned} S^{++} &= S^{+}(1+i) = S(1+i)^2 \\ S^{+-} &= S^{+}(1-i) = S(1+i)(1-i) = S(1-i^2) \\ S^{-+} &= S^{-}(1+i) = S(1-i)(1+i) = S(1-i^2) \\ S^{--} &= S^{-}(1-i) = S(1-i)^2 \end{aligned}$$

and σ^{++} is the volatility obtained from the volatility surface if the underlying's price is at S^{++} and the same holds for the dividend yield. From (2.19) and (2.20) we now find that

$$\Delta_I^{up} = \frac{1}{2Si\,(1+i)}\left[V\left(S^{++}, \sigma^{++}, d^{++}\right) - V\left(S^{+-}, \sigma^{+-}, d^{+-}\right)\right] \tag{2.20}$$

$$\Delta_I^{down} = \frac{1}{2Si(1-i)}\left[V\left(S^{-+}, \sigma^{-+}, d^{-+}\right) - V\left(S^{--}, \sigma^{--}, d^{--}\right)\right]. \tag{2.21}$$

We can depict this graphically as shown in Fig. 2.8.

Note the following:

$$S^{++} - S^{+} = S^{+} - S^{+-} \neq S^{-+} - S^{-} = S^{-} - S^{--}.$$

The impact gamma as defined in (2.18) will give a gamma similar in size to the theoretical gamma. This means it is normalised to reflect the change in delta for a very small change in the level of the spot price. We, however, want the gamma if the spot level changes by i. The correct impact gamma for a spot change of i is then given by

$$\begin{aligned} \Gamma_I &= \frac{1}{2\,S\,i}\left[\Delta_I^{up} - \Delta_I^{down}\right]\,S\,i \\ &= \frac{1}{2}\left[\Delta_I^{up} - \Delta_I^{down}\right]. \end{aligned} \tag{2.22}$$

Some VBA pseudo code for calculating the impact gamma is given in Fig. 2.9.

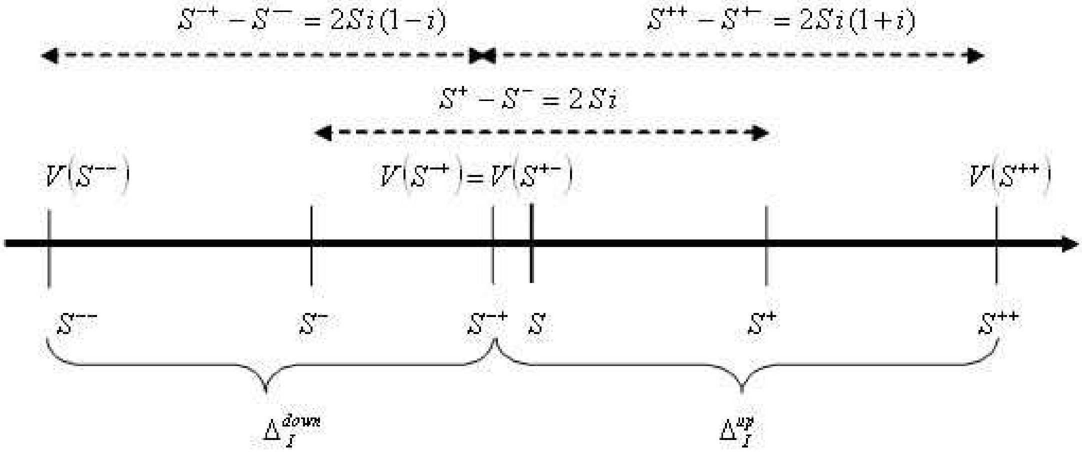

Figure 2.8: Calculating the Impact Gamma

```
Function Gamma_Impact(s, Strike, ValueDate, ExpiryDate, r_d, _
                      r_f, ATMVOL, phi, i)
‘ This function calculates the impact gamma for a given option
‘ ATMVOL = the at-the-money volatility for this option
‘ r_d = the risk-free interest rate
’ r_f = foreign interest
’ i = impact points
‘ First, bump the spot price down by the impact points
Sdown = s * (1 - i)
DeltaDown = Delta_Impact(Sdown, Strike, ValueDate, ExpiryDate, _
                         r_d, r_f, ATMVOL, phi, i)
‘ Phi = 1 for a call and -1 for a put
‘ Now bump the spot level up by the impact points
Sup = s * (1 + i)
DeltaUp = Delta_Impact(Sup, Strike, ValueDate, ExpiryDate, _
                       r_d, r_f, ATMVOL, phi, i)
‘ Impact gamma is now
Gamma_Impact = (DeltaUp - DeltaDown) * 0.5
End Function
```

Figure 2.9: VBA pseudo-code for calculating Impact Gamma.

2.4 Formalising Hedging Schemes

Consider a portfolio of derivatives written on some underlying asset with price S [JT 00]. These derivatives could include forwards, futures, options and the underlying asset itself. Although we are not restricting the type of instruments, we are assuming that the derivatives are written on the same underlying.

We know that derivatives depend on certain variables and the value of the portfolio can be written as

$$V = V(S, K, \sigma, t, d). \tag{2.23}$$

Let us now study how the value of the portfolio changes as time, asset price, volatility and interest rates changes. From calculus we can use a Taylor's series expansion to answer this question:

$$\Delta V = \frac{\partial V}{\partial t}\Delta t + \frac{\partial V}{\partial S}\Delta S + \frac{\partial V}{\partial \sigma}\Delta\sigma + \frac{\partial V}{\partial r}\Delta r + \frac{\partial V}{\partial d}\Delta d + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}(\Delta S)^2 + \cdots \tag{2.24}$$

Expression (2.24) is a statement about the changes in the value of the portfolio caused by changes in the underlying variables. Remembering how the Greeks were defined we have

$$\begin{aligned} \Delta V &= Theta \times \Delta t + Delta \times \Delta S + Vega \times \Delta\sigma + Rho \times \Delta r + Rhod \times \Delta d \\ &\quad + \frac{1}{2} Gamma \times (\Delta S)^2 + \cdots \end{aligned} \tag{2.25}$$

2.4.1 Delta Hedging

Let's assume we have the following portfolio

$$V(t = 0) = n_1 c + n_2 S + B \tag{2.26}$$

where $V(t = 0)$ is the current value of our portfolio, c represents the value of an option and we have n_1 options. S is the value of the underlying share and we have n_2 shares. Finally, B is an amount invested in a riskless asset. If $n_1 > 0$ it implies that we have bought the option and if $n_1 < 0$ it implies we have sold the option.

By taking the partial derivative of both sides with respect to S we obtain

$$\frac{\partial V}{\partial S} = n_1 \frac{\partial c}{\partial S} + n_2 \tag{2.27}$$

or

$$\text{Changeinvalueofportfolio} = n_1 Delta_c + n_2.$$

If the portfolio needs to be self-financing at date $t = 0$ we also add

$$V(t = 0) = 0 = n_1 c + n_2 S + B. \tag{2.28}$$

Now, using the above-mentioned philosophy, we can make a portfolio insensitive to small changes in the value of the underlying. In Sec. 2.1.1 we have seen that we can delta hedge an option. The same principles are used when we are dealing with a portfolio of derivatives.

To make a portfolio delta-neutral we have to set Eq. (2.27) equal to zero such that

$$0 = n_1 Delta_c + n_2. \tag{2.29}$$

Usually, the number of options, n_1 is known, Eq. (2.29) can thus be solved for the optimal number of shares n_2. If we further want our portfolio to be self-financing we also use Eq. (2.28). We then have a system of two equations in two unknowns, n_2 and B. This can be solved to obtain the right values.

Let's look at an example. Assume we are short 1000 options with a value of 2.734 and Delta of 0.562 per option. Also assume the underlying asset's price is 50. We then have

$$\begin{aligned} 0 &= -1000 \times 2.734 + n_2 \times 50 + B \\ 0 &= -1000 \times 0.562 + n_2. \end{aligned}$$

This can be solved to give $n_2 = 562$ and $B = -25,726$.

The same approach is used when one wants to hedge only the gamma, or only he vega or only the theta or only the rho. However, these Greeks can only be hedged with another option and not the underlying. The portfolio should thus be at least

$$V(t = 0) = n_1 c_1 + n_2 c_2 + B. \tag{2.30}$$

2.4.2 Delta-Gamma Hedging

If a portfolio is delta-neutral but the Gamma is negative, large changes in the underlying will still cause the portfolio to loose money. We now want to make the portfolio gamma-neutral as well.

If a self-financing portfolio needs to be delta-neutral, we need two assets. If we add the requirement that the hedged portfolio must also be gamma-neutral, we must add a third asset to the portfolio. The value of the portfolio is then

$$V(t = 0) = 0 = n_1 V_1 + n_2 V_2 + n_3 V_3 + n_4 V_4. \tag{2.31}$$

and for the portfolio to be delta-neutral we have

$$0 = n_1 Delta_1 + n_2 Delta_2 + n_3 Delta_3 + n_4 Delta_4 \tag{2.32}$$

and for the portfolio to be gamma-neutral we have

$$0 = n_1 Gamma_1 + n_2 Gamma_2 + n_3 Gamma_3 + n_4 Gamma_4. \tag{2.33}$$

If n_1 is known, we have 3 equations in 3 unknowns. This can be solved for n_2, n_3 and n_4.

Let's revisit our example from the previous section. We now add another option with a value of 1.1466 and a delta of 0.2965. We also have the gamma of option one as 0.0747 and that of option 2 as 0.0529. Eqs. (2.31-2.33) then become

$$\begin{aligned} 0 &= -1000 \times 2.734 + n_2 \times 50 + n_3 \times 1.1466 + B \\ 0 &= -1000 \times 0.562 + n_2 + n_3 \times 0.2965 \\ 0 &= -1000 \times 0.0747 + n_2 \times 0 + n_3 \times 0.0529. \end{aligned}$$

Solving these equations give $n_2 = 143.31$, $n_3 = 1,412.1$ and $B = -6,410.75$.

2.4.3 Theta Neutral

If our portfolio is delta-neutral, gamma-neutral and self-financing, it is also theta-neutral. Why? Because we have the identity

$$\Theta = \frac{1}{2}\sigma^2 S^2 \Gamma + rS\Delta - rV. \tag{2.34}$$

For this neutral portfolio we have Theta=Gamma=Delta=$V(t = 0)$=0 and thus Theta=0.

2.4.4 Vega Neutral

One can make a portfolio Delta-Gamma-Vega neutral or just Delta-Vega neutral by implementing the principles described above.

2.5 Imperfections of the Black-Sholes Model

We know that the *Black & Scholes* model is all but realistic compared to the real market. However, the model is used by everyone working in derivatives, whether they are salesmen, traders or quants. It is used confidently in situations for which it was not designed for, usually successfully. The ideas of delta hedging and risk-neutral pricing have taken formidable grip on the minds of academics and practitioners alike. In many ways, especially with regards to commercial success, the *Black & Scholes* model is remarkably robust [Wi 98].

Let's now look at these imperfections and how traders handle these in practise.

2.5.1 Lognormality

Real markets are not described by lognormality very well — this is per se reflected by the volatility skew. If we look at the lognormal return distribution of the USDZAR

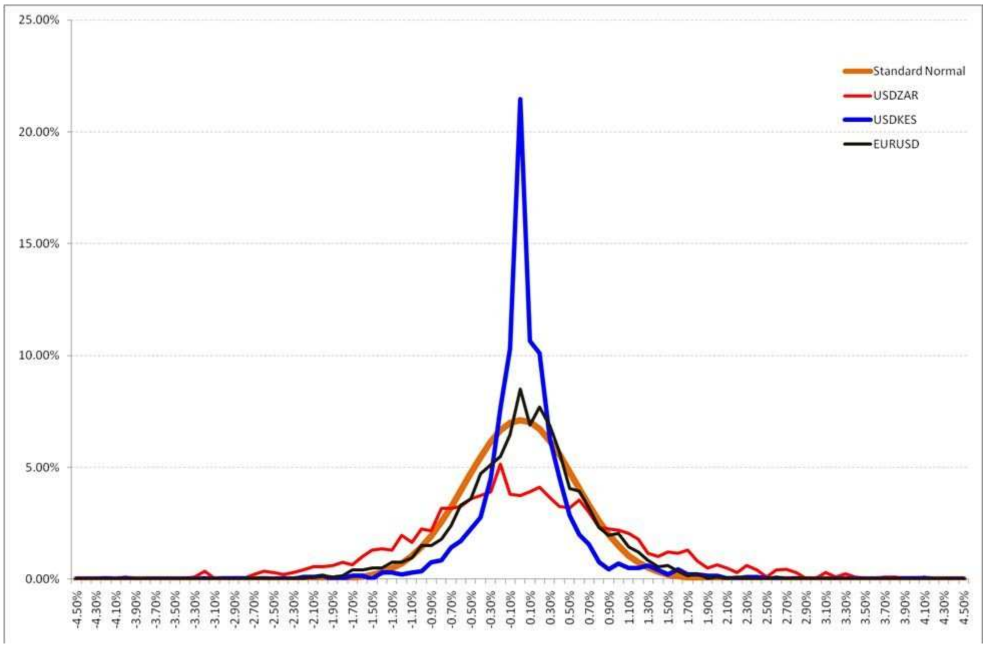

Figure 2.10: The distribution of USDZAR and EURUSD. The fat tails and skewness is clearly visible.

or USDKES we see some skewness and kurtosis — fat tails. This is shown in Figure 2.10. This shows there is a higher probability than normal that jumps will occur. This suggest that we should use another process to describe stock price behaviour. Many people have proposed different distributions like the parabolic or Levy distributions. These have not really catched on due to the complex pricing formulas these distributions leads to.

How do traders account for this? They play with the volatility (an accurate market related skew is imperative), and they look if they can lay off the risk to an exchange like Safex. To eliminate possible loss scenarios further, we do scenario analysis and stress testing. Risk management is thus crucial — this implies that a trading desk needs stable and user friendly risk management software or spreadsheets.

2.5.2 Delta Hedging

Black & Scholes assumed continuous hedging. In practise this is impossible. Traders calculate the number of FX contracts to buy or sell for a certain market move. Hedging is then done dynamically at these increments.

Monte Carlo simulation can be employed to calculate deltas discretely in time.

2.5.3 Transaction Costs

Black & Scholes assumed there are no costs in delta hedging. This is not the case.

We might have to pay brokerage and some taxes. Taxes and brokerage are added to the option premium.

2.5.4 Volatility

Volatility has to be estimated. Use implied or historical volatility as a base.

2.5.5 Interest Rates

Most options traded are ATM. This translates into a delta of roughly 50%. We can fix the interest rate for the initial money we borrow or deposit. Then we have the risk of our dynamic hedging strategy. Like a margining account at an exchange, we'll every day either deposit more money or borrow more.

If the yield curve is relatively flat, this should not be a big problem; especially for short dated options. The risk is, however, if interest rates move violently as they did during the Asian crises in 1998/99. For long dated options (3 to 5 years) interest rate risk must be considered. Part of the risk can be hedged by buying a coupon bond with a similar maturity or fixing the rate by buying a fixed for floating interest rate swap.

2.5.6 Price Gaps

What happens if there is a crash and liquidity dries up. Also, say the Dollar strengthens against the Euro. The Rand or KES can weaken over night and open at very different levels to the previous day's close. This is risky. Can use a jump-diffusion model to calculate more reliable hedge ratios.

2.5.7 Liquidity

Black & Scholes assumed a perfect market with infinite liquidity. What if one cannot hedge?

2.6 Tricks of the Trade

2.6.1 Data

Reliable data sources are very important. Data should be clean and the dates should be checked. You must be able to load data smoothly into Excel.

2.6.2 Risk Holes

Beware of risk holes, concentration of risks — localisation of risk that is difficult to hedge.

2.6.3 Yield Curve

It is imperative to have an up to date domestic and foreign yield curve at hand when one prices an option. Deals can be lost due an incorrect interest rates in your pricing. Also, ensure that you know what type of rate it is that you are inputting into your model. Imperative is the knowledge of the day count conventions in the different markets.

2.6.4 Technicals, Economic Information and Company Analysis

A trader should keep his ear to the ground for any information regarding currencies. Central bank reports and economic analysis are essential reading.

Every trader should also look at technical analysis and have a basic understanding thereof. Technical analysis is one of the tools used to obtain information on the market.

2.6.5 Put-Call Parity

In illiquid markets put-call-parity can help us in creating a synthetic hedge or exposure that might be beneficial. Look at the following example: let's assume we are long stock and we want to buy puts to protect the stock. The puts are however, very expensive. We then execute the following strategy: sell the stock buy a call (put + long stock is just a call). If, at expiry the price is above the strike, exercise and buy stock, if not, do not exercise but buy stock back at lower price.

2.7 Using Volatility

2.7.1 Volatility Spread

The fat is in the spread. When buying and selling options, there is usually a volatility spread involved. This means that a trader will buy an option from a client at a different volatility than what he will sell it to a client. Remember to take skew into consideration as well as at what volatilities the risk might be laid off in the market.

Some players buy market share by selling loss-leaders. This mean they sell options below the current market value. This is done just to get the deal on their books or to open potential lucrative lines with a client they never dealt with before.

2.7.2 Volatility Based Option Strategies

Note that implied volatility and historical volatilities are not good predictors of the actual realised volatilities. We also mentioned that this creates trading opportunities. How?

Professional option traders, market makers and institutions trade volatility by running "delta-hedged" positions. This means they buy or sell options and maintain a hedge against the option position in the underlying stock. This removes any net exposure to a small move in the stock. They continuously adjust this hedge as the market moves.

Because the hedge is in the underlying stock, these traders effectively capture historical volatility on the hedges while capturing implied volatility on the option price. That is, if they sell options at a higher implied volatility than the historical volatility of their hedges, they make money. Similarly, if they buy options at a lower implied volatility than the historical volatility of the hedges, they make money.

This strategy has a relatively low risk profile, but it involves a significant number of transactions. It also requires proper portfolio risk management systems.

While this type of delta-hedging volatility trading is difficult to implement and not very appropriate for the individual investor or non-institutional trader, it illustrates how volatility analysis can be translated into a practical trading strategy.

2.7.3 Option Volume and Volatility Changes

Option volume and volatility changes can be important indicators. Sudden jumps in call or put volume, combined with jumps in implied volatility, signal extreme market activity and possible market bias.

Combining implied volatility changes with technical analysis can be a powerful tool as well. It is not uncommon to see a rise in put volume and implied volatility as a stock is hitting technical levels on a rally. This can signal the market is worried about a downside correction and traders are buying puts as protection.

2.7.4 The Volatility Term Structure

This is plot of the variation of implied volatility with time to expiry of the option. In theory volatility for long dated contracts should be lower than that for short dated options. In the South African markets this is not always the case.

2.7.5 Volatility Matrices

Traders set up a matrix of implied volatilities. One dimension is the strike price (or moneyness) and the other time to expiry. One can then interpolate between the times and strikes.

2.7.6 Volatility as an Trading Indicator

One trap traders using volatility analysis tend to fall into is interpreting volatility itself as a directional indicator. High or low volatility by itself does not imply a certain direction or expected direction of the stock.

However, careful analysis of volatility patterns, combined with other indicators and stock movements, can lead to some interesting direction-based trading strategies. Different stocks behave differently, but in many cases, implied volatility tends to be a leading indicator of stock direction. This is why the volatility index (see Sec. 7.2.6) is such an important indicator of possible market moves in the USA.

When a stock is falling, every trader is looking for an indication of whether the stock will continue in that direction or whether it will stabilize and present a possible buying opportunity.

When a stock is declining and the implied volatility does not change (or falls), it suggests the market is not too nervous about the stock. On the other hand, if the implied volatility rises, it means the market continues to be nervous about the stock's downside potential.

2.8 The Skew and its Uses

The volatility skew is determined by the market. The skew spread expresses the consensual estimate of option traders in the derivatives market, and there must be some information you can extract from that [Ko 99]. We can thus ask: is the skew a true reflection of the market? If so, what does that tells us and how can we use the skew to our benefit? We will now touch on these issues.

2.8.1 Trading the Skew

A volatility skew can be handled as you handle any other asset. If you think an asset is overpriced, you sell it. If you believe a current option skew is higher than its average, and you think it should be coming down, then you should be selling that skew.

One way to execute the strategy is by selling a put bear spread - sell a put with a low strike price (high volatility) and buy a put with a high strike price (lower volatility). Another is to sell the out-of-the-money volatility through straddles[4], although this would be a speculative bet on the direction of the high-priced option volatility. If one had a market-directional view based on the option skew, the market index could be bought or sold short.

[4]Buy a put and call at the same strike with the same expiry date.

2.8.2 What about the Future?

Traders, and especially fund managers should also consider the implications of these skews for the future direction of the market or future profits? If today's skews are quite high, can that tell us something about the future expectations of the market? Yes it can. The skew tells us that option traders are implicitly assuming and pricing into their options a fear of future market decline.

Incidentally, the skew has historically been extremely wrong. When the S&P 500 was around 400 in 1992, the skew predicted that volatility should be 5% should the market rally to 1,300 within the next seven years. What happened is that volatility went to 22% instead. This tells you something about the quality of the prediction delivered by the market.

2.8.3 Counterintuitive Thinking

Consider the following counterintuitive thinking: when the market skew is quite steep, the market is quite likely to go up. Why?

Big jumps or big crashes happen as a surprise. When you have an extremely steep skew, it means that a crash has been discounted already and people are already paying this risk premium and expecting a crash to happen. And that's precisely why it doesn't happen. Because there is no surprise.

In the past strong moves of the market up or down were concomitant with a low level of volatility. Whenever volatilities are quite high, the risk premium people are willing to pay is too expensive, and the market doesn't move as much as the option market tells you it will. The implied volatilities are pricing in a risk premium that is too high. Using the designation "bearish" or "bullish" in this context is entirely futile.

When the skew is positive and is followed by higher returns, it thus doesn't necessarily mean that it's rational to be long. If the skew is positively sloped, it tells traders that the market is more likely to go down. If it's negatively sloped, it's telling traders that the delivered implied expected returns are going to be small positive returns, and that there will be plenty of them.

2.8.4 Supply and Demand

The skew is principally a result of the restriction of the supply of some category of options, rather than some consensus estimate by option traders. Since 1987, equity option traders have not been able to sell too many lower-strike puts, because the clearinghouse[5] would look at what would happens if there were a repeat of the crash and prevents them from having a large loss in such a situation. You have a similar phenomenon in the currencies.

[5] In the USA.

You can also infer valuable information about the skew by observing client OTC deal flows. For example, if institutional or corporate clients are hedging by either buying puts or executing collars, then the effect on skew will be quickly reflected in the market and can be dramatic. Similarly, if large retail products are being issued with capped calls, skew may move considerably as the banks involved in providing the assets to back these products hedge themselves. So traders need to keep a close eye on what clients are up to, because it will have quite a significant effect.

Another aspect of supply and demand is the number of players in the skew market is limited. In the long-term skew there is a big imbalance between what clients want and what professionals can provide. As a result, we have a huge increase in long-term skew. There's no theoretical explanation for it. As long as clients are willing to pay the level they are paying, and as long as the supply remains limited, the trend will remain the same.

2.8.5 Other Influences

There are other indirect pressures that influence the skew. Regulatory aspects can, for instance have an effect. If regulations change that, for instance, forces insurance companies to up their solvency measures, it can impact on the market. That kind of size can have an extremely dramatic effect on the skew, and yet tells you absolutely nothing about investor sentiment — since the hedge has not been put in place because of a bearish view, but rather to meet regulatory requirements. Skew is heavily influenced by supply and demand factors, more than by academic or technical arguments or fundamental views on the direction of a market.

2.8.6 The Skew in Other Markets

The skew is quite pronounced in the emerging-market currencies. In the Russian ruble, the peso and the eastern European currencies, we see pronounced skews over time. But it takes the shape of a volatility smile, as opposed to the one-sided skew we're seeing the equity markets.

Before 1987, there was no significant skew in the equities in USA, but there was a skew in the bonds. The calls on the bonds were cheap, because people would do covered writes on the bonds and sell puts. It was common wisdom among market-makers that if you sold out-of-the-money calls on a bond, you could always buy them back in a rally, since volatility did not increase. Then in the crash of 1987, when bonds snapped back up close to 10%, you could not buy them back. Since then, the skew has become symmetric in the bonds. In fixed-income markets, the skew is symmetric, but it's more of a smile then a skew.

Skew appears periodically in markets, depending on the mood of risk managers, typically after the fact, because they fear a repeat of the events. Usually this should tell you that these events will not repeat themselves, because, people are prepared and these events usually happen when people are not prepared. In U.S. dollar/yen, the skew fluctuates. The only market that is relatively flat is the euro. Practically everything else in the financial markets has a skew.

2.9 Buying and Selling Volatility

Options are like a 3D chess game. The three dimensions are price (of the underlying), time, and volatility [Ya 01]. The most misunderstood and neglected dimension is volatility. We have mentioned before that option prices are very sensitive to changes in volatility.

Every options trader needs to understand volatility and appreciate its effects. When a trader buys an option, he goes long volatility and when he sells an option he goes short volatility. As we all try to do, traders try to buy options when they are cheap and sell them when they are dear. The reason it is called volatility based trading comes from the way we measure cheapness or dearness — volatility. We have seen before that a high volatility is synonymous with expensive options and low volatility with cheap options.

There are two ways of judging the cheapness or dearness of options. The first is simply by comparing current volatility with past levels of volatility on the same underlying asset. The second is by comparing current traded or implied volatility with the volatility of the underlying itself. Both approaches are important. Opportunities arise when options are cheap or dear by both measures.

The volatility trader typically uses puts and calls in combination, selecting the most appropriate strikes, durations, and quantities, to construct a position that is said to be *delta neutral*. A delta neutral position has nearly zero exposure to small price changes in the underlying. Once a position is set up, it is held, and adjusted at times when necessary to re-establish the appropriate delta. These adjustments can be costly, in terms of transaction costs, and should be minimized, but not to the point where you expose yourself to too much delta risk. A general rule is: "If you give the market a chance to take money away from you (through delta), it will".

Once option prices return to a more normal, average level, then the position can be closed. If not many adjustments were required in the meantime, the trader should see a profit.

The investor can always count on volatility returning to normal levels after going to an extreme. This principle is called *the mean reversion tendency of volatility*, and it is the foundation of volatility based trading. That volatilities mean revert is a well established fact and you can see it for yourself just by looking at a few historical volatility charts. In Fig. (2.9) we plot the historical volatility for USDZAR and USDKES together with their long term menans. Volatility always comes back to "normal". Sometimes it does not happen right away. It may take anywhere from days to months, but sooner or later it always comes back.

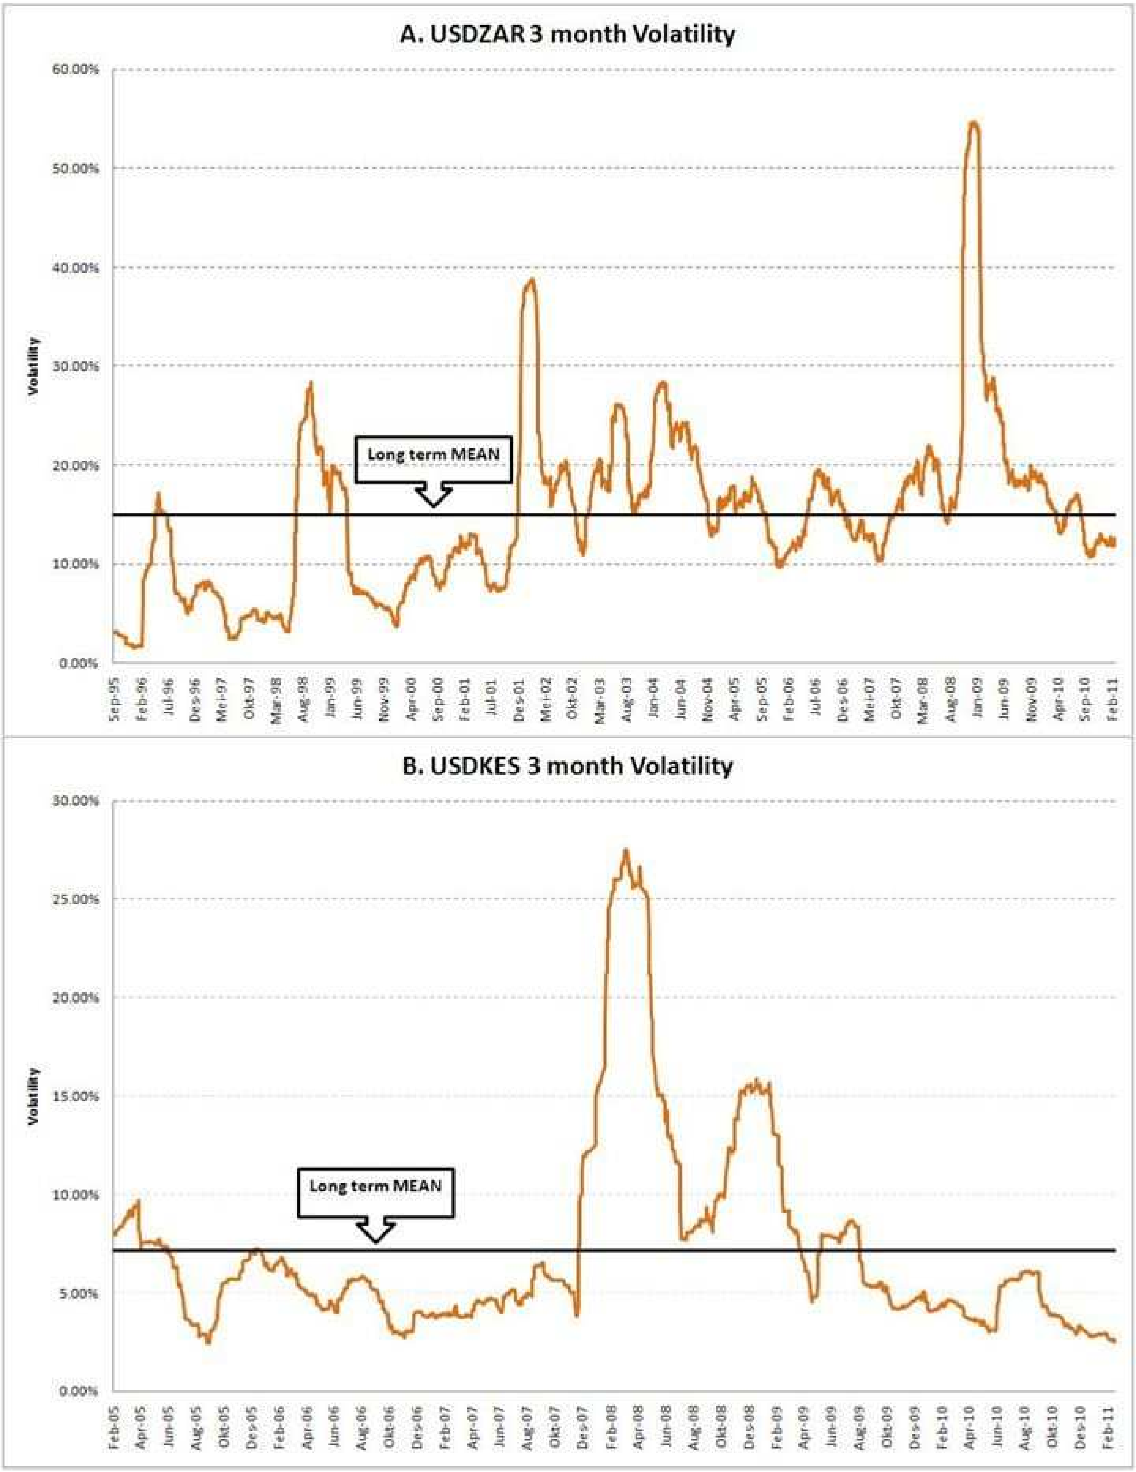

Figure 2.11: Historical volatilities for USDZAR and USDKES

2.9.1 Shorting Volatility

When the options of a particular asset are more expensive than usual, sometimes that additional expense is justified by unusually high volatility in the underlying. While this may be a decent opportunity to sell options, it is even more advantageous to sell options when the extra implied volatility is **not** accompanied by a high historical volatility. This means there is an anomaly in the market and thus a good trading opportunity.

Generally, any position in which you are short more options than you are long is short volatility. One prefers out-the-money options when one shorts volatility because it gives the underlying some room to wander, and increases the likelihood of realizing a profit. Generally, the farther out-the-money you go, the lower your returns, but the greater the probability of achieving those returns. By giving the underlying room to move, the trader minimizes his chances of having to make costly adjustments.

Longer term options work best, provided they have decent liquidity. Longer term options have higher Vega, and will therefore respond best when implied volatility comes down. Longer term options have the additional advantage of having lower "gamma". Gamma measures how fast delta changes with price changes in the underlying. By using lower gamma options, it takes a bigger price change in the underlying to imbalance your position.

2.9.2 Buying Volatility

Low volatility situations can be just as lucrative as high volatility situations. We have mentioned time decay that is against an option buyer but time decay, at best, is a funny concept. It says that if the underlying asset's price holds perfectly still, the option will decay at a certain rate. But what underlying asset price holds still? None, obviously. In fact, time is what gives the asset its freedom to move!

Let's assume I have a short volatility position, and let's say it has a theta of 100. This means I'm making R100 per day from time decay. Should I feel gratified to see this? Not really. It is a false gratification because today's movement in the underlying could take away R100, or perhaps many times that.

There is nothing wrong with buying options. When an option is fairly valued, by definition there is no advantage to the buyer nor the seller. If you buy a fairly valued option, you have not taken on a latent disadvantage in the guise of "time decay". Why? Because the underlying is in constant motion.

When buying options, it makes more sense to buy near-the-money, although it doesn't have to be a pure straddle (call and put at the same strike). That way a sharp move in the underlying has a better chance of helping the position. When that happens, not only does implied volatility normally get a boost, but the move may drive one of the sides deep in-the-money and give you a gain just from price movement.

It is interesting that long volatility positions have a completely different "feel" than short volatility positions. Short volatility positions often gratify the holder with steady, almost daily, gains, but can suddenly lose money if the underlying makes a sharp move. Long volatility positions often seem to dribble away value day by day for many weeks, and suddenly gain very quickly.

Deciding when to close a long volatility position is usually more difficult, since the position has blossomed into a larger position with a sharp move in the underlying, and has probably become imbalanced. Often there is the potential to make (or lose) more money with each additional day that you hold the position. What can help you make a decision is to identify whether volatility has returned to normal levels. If it has, you should consider closing the position. If it has not, you might consider continuing with an adjusted (re-balanced) position.

When buying volatility, just as when selling volatility, use the longest dated options you can find that give you decent liquidity. The reason is the same as when selling: high Vega. The long dated options, with their higher Vega, will respond the best when implied volatility increases.

Chapter 3

Exchange Traded FX Derivatives

3.1 Advantages of a Futures Market

- Very low trading costs
- Easy to short sell
- Gearing
- It is very easy and inexpensive to buy or sell a future on an index whilst it is more difficult, and expensive, to transact in a portfolio of shares that spans an index — and this has to be done all at once.

3.2 Making a Market in Futures/Forwards

Let's look at the example of making a market in a FX futures. If the trader sells the future, his resultant position will be short the future. To hedge he should buy the spot currency — he will borrow money to do this — and he will earn the foreign interest rate. Buying the spot means there are trading costs to consider. The futures price is then given by

$$\begin{aligned} F &= S(e^{(r_d - r_f)\tau} + B) \qquad (3.1) \\ &= S(1 + (r'_d - r'_f)\tau + B) \\ &= S\left(\frac{(1+r''_d)^\tau}{(1+r''_f)^\tau} + B\right) \end{aligned}$$

where B is the brokerage paid in percentage. Quantities denoted with a $'$ refer to simple rates and quantities denoted with a $''$ refer to NACA rates.

If there are any other fees we can lump it all together into a parameter l. The futures price is then given by

$$\begin{aligned} F &= S(e^{(r_d - l - r_f)\tau} - B) \\ &= S(1 + (r_d' - r_f' - l')\tau - B) \\ &= S\left(\frac{(1+r_d'')^\tau}{(1+r_f'')^\tau}\frac{1}{(1+l'')^\tau} - B\right). \end{aligned} \tag{3.2}$$

3.3 The Cost of Carry

The relationship between futures prices and spot prices can be summarised in terms of what is known as the *cost of carry*. This measures the storage cost plus the interest that is paid to finance the asset less the income earned on the asset. For an investment asset like a share we define

$$c = r_d - r_f - l. \tag{3.3}$$

3.4 Currency Futures Dispensation in South Africa

South Africa has a system of exchange control prohibiting certain foreign exchange transactions. These controls are implemented and overseen by the reserve bank. Currency futures were launched in 2007 predominately as a retail product. The initial dispensation granted by the Minister of Finance in 2007 allows individuals to trade over and above their foreign allocation allowance stipulated by the South African Reserve Bank. Individuals, in other words, have no limits to the value traded in the currency futures market.

The Minister of Finance in his 2008 budget speech extended the currency futures qualifying audience to include all South African corporate entities. Corporate entities, including limited or unlimited companies, private and public companies, close corporations, partnerships, trusts, hedge funds and banks are authorised to trade currency futures with no restrictions on the value traded. Corporate entities do not need to apply to Reserve Bank for approval to trade the currency futures nor do they have to report their trades.

Unfortunately, pension funds and long term insurance companies are subject to their 15% foreign allocation limits while asset managers and registered collective investment schemes are subject to their 25% foreign allocation limits.

3.5 Justification for a Futures Market

From an economic point of view, the function of a futures market is to allow for the transfer of risk. These markets have the special function of allowing those who do not wish to take the risks to nevertheless run their business enterprises. Take a farmer who has acquired considerable skills in agriculture but is totally put off by the prospect of volatile prices in the grain market. Futures allow him to exercise his skills - growing the normal crop risks, which he bears anyway. In short, the futures markets enable many productive entrepreneurs and businessman to operate without exposing themselves to risks greater than they are willing to bear. This holds true for importers and exporters as well. Hedging their currency exposures allows them to focus on their core businesses and not on the by-products of currency risks which can have unexpected consequences

The futures market is also valuable to the economy in that it facilitates “price discovery” and the rapid dissemination of prices. In a traditional forward market contracts are not standardised and are entered into “over the counter” between buyers and sellers. The prices at which forward contracts are fixed are not relayed to the market because they are ”private” deals. Price determination in the overall market is therefore not as efficient as it could be and buyers and sellers cannot be sure that they are getting the best possible price. In the futures market, by contrast, the competitive nature of the market ensures that commodities trade at or very close to what the market thinks they are worth, and the smallest market user has as much knowledge as the largest user as to the current value attached to the commodity.

3.6 Futures versus Forwards

A forward contract is one where the buyer and the seller agree on a price, but the actual transfer of payment for property is deferred until a later time. Forward contracts are arranged between two principals with complete flexibility as to exactly what property is being transferred and when the transfer will occur.

In contrast, futures contracts are transacted in the arena of a futures exchange. Transactions must be made in prescribed increments (i.e., whole numbers of futures contracts covering a designated “size” per contract), where the price-setting capability applies to a limited number of prospective settlement dates [Ko 02]. Transactions take place at the best bids and offers provided by the exchange members who trade through an electronic trading system. Using internet trading systems, clients of exchange members, trade directly onto the exchange via the exchange’s “direct market access” (DMA) platform.

Cash flow obligations are very different for forward contracts and futures contracts. With a forward contract, a price is established on the trade date; but cash changes hands only on the value (or settlement) date, when, as agreed, the buyer pays the seller and takes possession of the property. With a futures contract, the change in value of the futures is passed between the two parties to the trade following movements of the futures price each day, making use of the clearinghouse[1] as an intermediary.

[1] Clearing houses in all countries use a common set of safeguards to limit the likelihood of defaults

When the futures price rises, the buyer (who holds the long position) “earns” the change in value of the contract, and the seller (the short-position holder) loses. Opposite adjustments are made when the futures price declines. This daily cash adjustment thus collects from the loser and pays to the winner each day, with no extension of credit whatsoever. The daily Rand value that changes hands is called the “variation margin.”

This cash-flow aspect of the futures contact is perhaps the most difficult conceptual hurdle, as well as the hardest operational feature, for a potential futures market user. Maintaining a futures position requires that the position taker, both the buyer and the seller, be ready and able to pay funds into the clearinghouse (via a broker) each day that the futures position generates losses.

Alternatively, efficient participation in the futures market requires that the trader or hedger be ready and able to employ funds that may be generated from profitable futures positions. Naturally, the former situation is the one that would cause potential problems. Due to the high leverage nature of the futures contract, the cash-flow requirements of a losing futures position may be quite onerous. The futures participant must either have the cash readily available or have the prearranged capability of financing this cash flow requirement. The "silver lining" to this process is that the cash requirement fosters a discipline that focuses attention on a market situation as it is happening-not months after the fact when it is too late to take corrective action.

Parties to forward contracts may require some form of collateral security in the form of compensating balances or a performance letter of credit. With futures contracts, customers must provide their brokers with initial margin. Initial margin is a Rand value per contract and is set by the exchange. These amounts are determined through a statistical analysis and are estimates about what losses are possible in the future – usually 1 trading day. Participants are required to lodge margins with the exchange which are sufficient to cover these possible future losses. Should the losses eventuate and the participant be unable to bear them, the margin is available to the exchange to meet the shortfall. The Rand value of initial margin requirements vary depending on the particular futures contract traded; and this amount is adjusted as volatility conditions change [Ko 05].

Currency futures transactions tend to be used primarily as price-setting mechanisms rather than as a means of transferring property. That is, when using futures contracts, buyers and sellers typically offset their original positions prior to the delivery date specified by the contract, and then they secure the desired currency via a spot market transaction. This offset of the futures hedge is accomplished simply by taking a position opposite from the initial trade. For example, if one were to enter a long futures position, the offset would require selling the futures contracts. Conversely, if one started with a short position, offset would be arranged by buying the

by clearing members and to ensure that if defaults do occur, the clearing house has adequate resources to cover any losses and to meet its own payment obligations without delay.

contracts. The complete buy/sell (or sell/buy) is referred to as a "round turn" and, with the completion of a round turn, commissions are charged on a "per contract" basis.

Please note that, for Rand futures specifically, no physical exchange of currencies ever takes place even on the expiration date - the contract is cash settled; that is, following the final trading session (on the third Monday of the expiry month), one last mark-to-market and cash adjustment takes place.

The size of the commission is negotiated, reflecting the amount of support and assistance that the broker provides, as well as the volume of trade generated by the customer. On the forward side, commissions may or may not be charged, depending on whether the trade is arranged directly with the dealer or if a broker serves as an agent. Importantly, it is not safe to assume that direct dealing necessarily reduces transaction costs. Often, the use of a broker-whether a futures broker or an interbank currency broker - allows customers to access more competitive market prices than they can otherwise. The factor most likely to determine whether futures or forwards provide the better prices is the size of the required transaction.

3.7 Economics of Hedging with Currency Futures

The difference between hedging and speculating relates to risk existing before entry into the futures/forward market. The speculator starts with no risk and then enters into a transaction that takes on risk in order-one hopes-to make profits. The hedger, on the other hand, starts with a pre-existing risk generated from the normal course of his or her traditional business. Futures (forwards) are then used to reduce or eliminate this pre-existing exposure. These contracts may be used to hedge some or all of such risk, essentially by fixing the price or exchange rate associated with the relevant exposure. Once so hedged, the manager is insulated from the effects of subsequent changes in the exchange rate, either beneficial or adverse.

As of October 2010, 8 different currency futures contracts are listed and actively traded at the JSE. They are:

- USA Dollar Futures
- British Pound Futures
- Euro Futures
- Australian Dollar Futures
- Japanese Yen Futures
- Swiss Frank Futures
- Canadian Dollar Futures

- Chinese Yuan Futures

EXHIBIT 1: *Perfect Long Futures Hedge*
Exposed to the risk of strengthening Rand – weakening US Dollar. The EXPORTER
Size: $10,000
Hedge Instrument: 10 long futures contract

Exchange Rate and Interest Rate Data

	Initiation of Hedge	Liquidation of Hedge
Transaction Date	June 2 2010	Sep 13 2010
Spot Value Date	June 4 2010	Sep 15 2010
Futures Delivery Date	Sep 15 2010	Sep 15 2010
Spot Price (ZAR/USD)	7.6405	7.1398
Futures Price	7.8313	7.1398

Results

Rands obtained for $10,000 on Sep 15: $10,000 x R7.1398/$ = R71,398.00
Hedge result: $10,000 x (R7.8313/$ -R7.1398/$) = R6,915
Effective exchange rate = (R71,398 + R6,915)/$10,000 = R7.8313/$

Strictly speaking, each futures contract locks in an exchange rate for a specific value date or delivery date. This result is demonstrated above in Exhibit 1, which shows the case of the hedger (exporter) who initiates a long hedge of 10 futures contracts on June 2 to protect against a weakening US Dollar. The size of the exposure is $10,000 (equal to 10 futures contracts), and the desired value date is precisely the same as the futures delivery date (Sep 15 2010).

Following a 10% rise in the strength of the Rand, the Rands are purchased at the new, lower FX rate; 'but profits on the hedge foster an effective exchange rate equal to the original futures price'. At the time the hedge is initiated, highest quality bank customers would likely find the price of the forward contract for the same futures value date to be virtually identical to the futures contract, so an analogous trade with a forward contract with the same settlement date in September would foster the same economic result. Lesser quality (i.e., smaller) customers, however, might find discriminatory pricing in forward markets, resulting in a slightly disadvantaged outcome.

Of course, the assumption that the currency requirement coincides with the futures value date schedule is overly restrictive. A more likely scenario would be one in which the hedge value date differs from the available futures delivery (value) dates. In such cases, it may seem that forward contracts have an advantage over futures, given the flexibility to select a value date that coincides precisely with the exposure being hedged. This judgment typically turns out to be overstated, however, and thus this preference may not be justified. Even when using forwards, the date for which the currency exchange is expected to take place may need to be altered, so additional transactions might be required, adding to the cost of the currency hedge. Also, many users of forwards have to "bundle" their exposures, thus having individual forward contract hedges cover the exposures of several planned cash transactions. The capacity to select a specific value date therefore involves somewhat of a compromise.

EXHIBIT 2: *Long Futures Hedge: Early Liquidation, Strengthening Rand*
Exposed to the risk of strengthening Rand – weakening US Dollar. The EXPORTER.
Size: $10,000
Hedge Instrument: 10 long futures contract

Exchange Rate and Interest Rate Data

	Initiation of Hedge	Liquidation of Hedge
Transaction Date	June 2 2010	Sep 13 2010
Spot Value Date	June 4 2010	Aug 31 2010
Futures Delivery Date	Sep 15 2010	Sep 15 2010
Spot Price (ZAR/USD)	7.6405	7.3645
Futures Price	7.8313	7.3922

Results

Rands obtained for $10,000 on Aug 31: $10,000 x R7.3645/$ = R73,645.00
Hedge result: $10,000 x (R7.8313/$ -R7.3922/$) = R4,391
Effective exchange rate = (R73,645 + R4,391)/$10,000 = R7.8036/$

When the hedge value date differs from one of the available futures delivery dates, the hedger simply initiates a futures hedge with the contract that expires as soon as possible after the desired currency exchange date. The hedge would then simply be liquidated before expiration.

Mechanically, when the need for the currency is at hand, the hedger would secure the desired currency using the spot market and simultaneously offset the futures hedge. An example is shown above in Exhibit 2. Here, as before, the hedge is initiated on June 2; but now the hedge must take possession of the Rands on 31 August - approximately three weeks prior to the expiration of the September futures contract. On August 31 the hedger simultaneously sells the required $10,000 with a spot market trade at a price of R7.3645/$ and offsets the futures hedge at a price of R7.3922/$. At the time of the hedge liquidation or offset, the difference between futures and spot prices (the basis) thus equals R0.0277. The consequence of this non-convergence is that the effective exchange rate realized from hedging the futures is R7.8036/$ - a difference of 0.0277 from the original futures price.

The outcome shown is predicated on the assumption that the differential between U.S. interest rates and South African interest rates present in the market on June 2, when the futures value date was 105 days away, remains in effect on August 31, when the futures have 15 days to go before expiration. Relatively higher South African (versus U.S. interest rates) on August 31 would have fostered a higher effective exchange rate, and vice versa. Clearly the futures hedge necessarily has some small degree of uncertainty in terms of the ultimate exchange rate realized; but this incremental effect can be either beneficial or adverse.

Again, the hedger might have chosen to operate with a forward contract rather than with the futures. When the need for the currency arises before the futures value date, however, the relevant forward price would not be the same as the futures price. Typically, interbank market forward prices are quoted as spot prices plus some premium (or less some discount), where premiums and discounts are expressed as "forward swap points," or "swap prices." In this example where the desired currency exchange is scheduled for September 2, the swap points would likely be roughly proportional to the basis, where the constant of proportionality would reflect the ratio of time to the desired forward date divided by the time to the futures delivery date. In this case, that ratio is 90/105. The forward pricing, therefore, could be estimated as follows[2]:

Future basis = 7.8313-7.6405=-0.1908 (for 105 days)

Approximate swap price = -0.0050 x (90/105) = -0.1635 (for 90 days)

Approximate forward price = 7.6405 +0.1635 = 7.8040 (for 90 days)

Thus, the hedger should be comparing a forward price of R7.8040/$ for a September 2 settlement with a September futures contract, traded at R7.8313/$ but expected to realize an effective exchange rate of R7.8036/$ as a consequence of early liquidation. It should be clear, then, that the effective rate realized from a futures hedge will likely be quite close to the outcome of a forward hedge (i.e., within a few basis points – 4 basis points in our example) irrespective of whether the timing of the risk coincides with the futures value date schedule.

For completeness, Exhibit 3 starts with the same problem as that shown in Exhibit 2. In this case, however, we hypothesise that the Rand weakens against the Dollar. Regardless, comparing Exhibits 2 and 3 shows the same effective exchange rate whether the Rand appreciates or depreciates. This example thus demonstrates the robust outcome of a futures hedge. That is, once hedged, the hedger is indifferent about the prospective direction of exchange rates in the future, as the effective rate (R7.8036/$ in this case) is unaffected by subsequent spot market moves[3].

---

[2] Actual forward prices quoted may differ somewhat from this estimate; but the closer the hedge value date is to the futures value date, the greater the confidence one should have for this approach to estimation.

[3] This conclusion requires that the hedge is implemented with no rounding error, and it assumes consistent basis conditions upon hedge liquidation regardless of the level of spot exchange rates.

Exhibit 3: *Long Futures Hedge: Early Liquidation, Weakening Rand*
Exposed to the risk of strengthening US Dollar
Size: \$10,000
Hedge Instrument: 1 long futures contract

Exchange Rate and Interest Rate Data

	Initiation of Hedge	Liquidation of Hedge
Transaction Date	June 2 2010	Sep 13 2010
Spot Value Date	June 4 2010	Aug 31 2010
Futures Delivery Date	Sep 15 2010	Sep 15 2010
Spot Price (ZAR/USD)	7.6405	8.1562
Futures Price	7.8313	8.1839

Results

Rands obtained for \$10,000 on Aug 31: \$10,000 x R8.1562/\$ = R81,562.00
Hedge result: \$10,000 x (R7.8313/\$ -R8.1839/\$) = -R3,526
Effective exchange rate = (R81,562 – R3,526)/\$10,000 = R7.8036/\$

Note: The general rule for choosing the "correct" futures contract month is to pick the contract expiration concurrent with or immediately following the desired date of the actual currency conversion. For example, if you plan to make an actual conversion on November 1, the closest futures contract expiration following November 1 is available with the December contract.

Liquidity conditions, however, may justify a departure from this practice when the planning horizon extends beyond the date for which futures contracts are actively traded. In these cases, hedges temporarily rely on nearby futures positions. After deferred contracts (i.e., later expirations) develop greater liquidity, the original hedge contract is offset and a new position is established in the more distant contract month. This process is called "rolling the hedge." It necessarily introduces a certain amount of uncertainty in that the price differentials between successive futures expirations (i.e., "spread prices") cannot be known with certainty before the roll.

3.8 Choosing between Futures and Forwards

Choosing between futures contracts and forward contracts for managing currency exchange rate risk involves consideration of a number of trade-offs. Perhaps most important is the fact that forwards lock in a prospective exchange rate with virtual certainty. Futures contracts, on the other hand, will foster approximately that same exchange rate. The source of risk for the futures contract pertains to the uncertainty associated with the size of the basis at the time the futures hedge needs to be liquidated. Depending on prevailing interest rate differentials in the market at that time, this uncertainty may prove to be beneficial or adverse.

Beyond this consideration, a further issue deals with hedge management practices. Forwards tend to be maintained consistently until the value date arrives when currencies are then exchanged even when the forwards are generating losses. The mark-to-market aspect of futures and the required daily cash settlements tend to foster a re-examination of the desirability of hedging when hedges generate losses, thus allowing for the curtailment of these losses. Put another way, futures provide greater flexibility in that they are more easily offset than forwards if the need for hedging is obviated. And finally, futures have the ancillary benefit that they do not introduce any added credit risk for the hedger as a consequence of the rigorously practiced marking-to-market requirement, while forwards do.

3.9 The Role of the Stock Exchange

3.9.1 A Brief History

Some stories suggest that the origins of the term "bourse" come from the latin bursa meaning a bag because, in 13th century Bruges, the sign of a purse (or perhaps three purses), hung on the front of the house where merchants met[4].

However, it is more likely that in the late 13th century commodity traders in Bruges gathered inside the house of a man called Van der Burse, and in 1309 they institutionalized this until now informal meeting and became the "Bruges Bourse". The idea spread quickly around Flanders and neighbouring counties and "Bourses" soon opened in Ghent and Amsterdam.

The Dutch later started joint stock companies, which let shareholders invest in business ventures and get a share of their profits - or losses. In 1602, the Dutch East India Company issued the first shares on the Amsterdam Stock Exchange. It was the first company to issue stocks and bonds.

3.9.2 What is a Stock Exchange

A stock exchange, share market or bourse is a corporation or mutual organization which provides facilities for stock brokers and traders, to trade company stocks and other securities. Stock exchanges also provide facilities for the issue and redemption of securities, as well as, other financial instruments and capital events including the payment of income and dividends. The securities traded on a stock exchange include: shares issued by companies, unit trusts and other pooled investment products, bonds and derivative instruments like futures and options.

To be able to trade a security on a certain stock exchange, it has to be listed there. Usually there is a local & central location at least for recordkeeping, but trade is less and less linked to such a physical place, as modern markets are electronic networks, which gives them advantages of speed and cost of transactions. Trade on an exchange is by members only.

Advantages of bourses

[4]See http://en.wikipedia.org/wiki/Stock_exchange

- Raising capital for businesses
- Mobilizing savings for investment
- Enlarges the investment horizon
- Higher liquidity will lead to more efficient pricing
- Facilitate company growth
- Redistribution of wealth
- Corporate governance
- Creates investment opportunities for small investors
- Government raises capital for development projects
- Barometer of the economy

3.9.3 Objectives for Using Financial Instruments

The question is always asked: why would an institution trade in a financial instrument? We list a few answers to this question

- outperform the benchmark rate;
- participate in favourable currency / interest rate movements;
- dynamic hedging;
- effective use of credit lines;
- structural change of the underlying asset;
- diversification of investments and cash flows;
- proprietary trading.

3.10 The Role of the Clearing House

The dictionary definition of a clearing house is: "an office where banks exchange checks and drafts and settle accounts" – in today's terms, this is the bank's back office.

A futures exchange also has a back office. It is, however, still known as a clearing house. A clearing house is an agency or separate corporation of a futures exchange responsible for settling trading accounts, clearing trades, confirming trades, collecting and maintaining margin monies by calculating gains and losses, regulating delivery and reporting trading data

Some clearinghouses interpose between buyers and sellers as a legal counter party, i.e., the clearinghouse becomes the principal buyer to every seller and vice versa. This obviates the need for ascertaining credit-worthiness of each counter party and the only credit risk that the participants face is the risk of clearing house committing a default. This actually means the following: if SBSA makes a price in a currency future, and a client trades on that price, the client trades with the exchange (clearinghouse) and the clearing house trades with SBSA. The clearing house will thus buy the future from SBSA and on sell it to the client. The advantage of such a structure is that every trade is guaranteed – there is virtually no credit risk. This process is referred to as novation, whereby the clearing house guarantees the performance on each trade.

The clearinghouse puts in place a sound risk-management system to be able to discharge its role as a counter party to all participants. In South Africa, the performance of the contracts registered by the exchange is guaranteed by the clearing house SAFCOM. The exchange also maintains its own fidelity fund[5] and insurance which can be used in case of a default. The exchange also puts in place membership criteria and some of the new exchanges have also prescribed certain minimum capital adequacy norms.

What happens in practise? Each trade concluded is matched daily by YieldX i.e. the exchange ensures that there is a buyer and a seller to each contract. YieldX's clearing house (SAFCOM[6]) then becomes the counterparty to each trade once each transaction has been matched and confirmed. To protect itself against non-performance, SAFCOM employs a process known as margining.

3.11 Member Brokers

The rules of most exchanges state that any person who wishes to trade on an exchange needs to be either an exchange member or a client of a member. An exchange usually has strict membership criteria and membership is limited. Members must also, at all times, comply with the rules of the exchange. The rules sometimes prescribe capital adequacy requirements and necessary administrative systems. Such rules are set up as an additional risk measure to ensure that all business activities on the exchange are done with the highest integrity.

Exchanges have members because they cannot handle the administrative burden of individual clients who trade through the exchange – this is the task of the members. The broker acts on behalf of the investor. Every investor must have an account with a broker if he wishes to trade through the exchange. With internet trading, some brokers become administrative agents only. Clients trade "directly" through

[5]The rules of the JSE Fidelity Fund were approved by the Financial Services Board (FSB)

[6]SAFCOM is a separate legal entity to the exchange itself.

the exchange with the supplied software and internet. However, the rules are clear: the broker is still liable for his clients' position if disaster might strike.

There are usually different classes of membership. Private clients will usually trade through a broking member. A broking member may trade for and on behalf of clients and enter into client agreements with clients. Members with clients must administer all aspects of all their clients e.g., their FICA requirements and margin accounts.

The exchange can be seen as the factory of goods. The members are the wholesale agents who distribute the factory's products to the retail market. Such a structure enables the exchange to concentrate on being an exchange: listing investment products and ensuring that trade happens orderly.

3.12 Margining

Margining is a risk management procedure. Risk management may be defined as identifying the risk of loss in a portfolio and ensuring that the losses can be borne.

This mechanism of margining is two-fold: *marking-to-market* and *margining.* Marking-to-market ensures that all losses up to the present are absorbed. Participants with losses are required to make cash payments to the exchange equal to their losses. Margining then estimates what losses are possible in the future. Participants are required to lodge margins with the exchange which are sufficient to cover these possible future losses. Should the losses eventuate and the participant be unable to bear them, the margin is available to the exchange to meet the shortfall.

Practically this means: Firstly, when a position is opened (either long or short), the investor is called on to pay an *initial margin* in cash to the clearing house. This cash is deposited into what is termed a *margin account.* This amount remains on deposit as long as the investor has an open position. It attracts a market related interest rate (RODI less 25 basis points), which is refunded to the investor once the position is closed out, or the contract expires. The RODI index is calculated by reference to an average of the overnight call deposit rates paid by the banks where SAFEX deposits margin. It is important to remember that this margin is purely a deposit with the clearinghouse.

Secondly, the Exchange re-values each position against the market price at the close of trade daily. This process is referred to as *Marking-to-Market* (MTM) with the market price being termed as the MTM price. Any difference from the previous day's market price is either paid to the investor, or paid by the investor to the clearinghouse (all flows go through the member), in cash - the profit or loss is thus realised on a daily basis. This is possible, because the clearinghouse is the central counterparty to all contracts, and the number of long contracts equals the number of short contracts. This payment is called *variation margin (also maintenance margin).*

The MTM levels are distributed at the end of trading by the exchange. This is either done by email or the levels can be downloaded from their website. The trading system also has a report with all the levels.

All trades done on the JSE is reported through the BDA system. This is a flat file that is distributed to members at around 5:30 each day. This report holds all relevant information: all trades for each client and margin movements. The exchange usually sends a margin report to each member stating the total margin either due to a member or that has to be paid by the member. This will be the total margin taking all of that member's clients into account. If the member receives a margin call, he has till 12:00 the next day to pay. The member must then ensure that he receives the appropriate margins from all of his clients.

If the balance in the margin account falls below the initial margin level, the client receives a margin call from his member broker and he is requested to top up the margin account to the total initial margin level. If the client fails to do so, the broker has the right to liquidate any position up to a point where the margin account is above the initial margin level. Any balance above the initial margin level can be withdrawn by the investor.

3.13 Spread Margining

YieldX's initial margin requirement is determined by spread margining which is similar to the Standard Portfolio Analysis of Risk (SPAN) method. Many different variations of this method are used by derivatives exchanges world-wide. SPAN is a risk-based, portfolio-approach, for calculating margin requirements on futures, options on futures, and other derivative and non-derivative instruments. Contracts are grouped into classes of similar underlying instruments but expiry dates can also be factor. These classes are examined over a range of price and volatility changes to determine the potential gains and losses.

The basis of Span is that the whole of the portfolio on the exchange is valued ("scanned") at a number of points over a wide range of market moves. The range is selected to cover almost all conceivable market moves within the next day. The lowest of the portfolio values is identified and from this is found the greatest loss which the participant could suffer on the next day. The initial margin, due in cash the next morning, is then set equal to this greatest loss.

Statistically YieldX determines the possible loss in one day with a confidence level of 99.5% - this means only 0.5% of daily losses is further than 3.5 standard deviations from the mean or 99.95% of all possible daily changes in the market will be covered by the IMR. We show this in Fig. 3.13. This is in essence a Value at Risk (VAR) calculation. The result is an initial margin requirement that will, 99.5% of the time, cover any loss during a one day period. The precise margin required thus varies from one exchange-traded product to another. Note, a confidence level of 99.95% theoretically means there should not be more than one breach per 2000 days.

Figure 3.1: Lines show 3.5 standard deviations from the mean.

In practise YieldX only adjusts the initial margin monthly but adverse market movements can result in more frequent adjustments. The margins used by YieldX during February 2011 is given in Table 3.13. We show the calculation of the IMR in Fig. 3.13. Of interest to risk managers is the number of breaches per annum shown at the bottom.

The VSR (volatility scanning range) is used to calculate the margins for options. Remember, initial margin is required with every new trade — longs and shorts. That means that if a client buys a December Alsi future and he buys a March Alsi future, he will have to post margin on both trades.

However, the exchange recognises that some contract have very similar risk profiles. Spread margins are used when a client buys or sells a spread e.g., he goes long a March USDZAR future and sells a June USDZAR future (contracts must be in the same class of instruments). This is calculated by looking at the interest rate spread between the dates — YieldX uses the JIBAR rates. The series spread margin is required when spreads are traded between different classes e.g. a client buys a March USDZAR future and sells a June EURZAR future. Series spread margins are calculated by taking the correlations between the different contracts into account.

3.14 Offsetting Margins

The exchange can specify certain contract whose margin requirements can be offset against one another – the margins are netted off. An example explains: let's assume we are long a February futures contract and short a May option contract. If we have

Contract Code	Expiry Date	Fixed Initial Margin Requirement	Spread Margin Requirement	Series Spread Margin Requirement	VSR
Dollar/Rand ($/R)	14 March 2011	R300.00	R15.00	R40.00	2.5
Dollar/Rand ($/R)	13 June 2011	R315.00	R15.00	R40.00	2.5
Dollar/Rand ($/R)	19 Sept 2011	R320.00	R15.00	R40.00	2.5
Dollar/Rand ($/R)	19 Dec 2011	R325.00	R15.00	R40.00	2.5
Dollar/Rand ($/R)	16 March 2012	R330.00	R20.00	R40.00	2.5

Contract Code	Expiry Date	Fixed Initial Margin Requirement	Spread Margin Requirement	Series Spread Margin Requirement	VSR
Euro/Rand (€/R)	14 March 2011	R355.00	R20.00	R30.00	2.5
Euro/Rand (€/R)	13 June 2011	R385.00	R20.00	R30.00	2.5
Euro/Rand (€/R)	19 Sept 2011	R390.00	R20.00	R30.00	2.5
Euro/Rand (€/R)	19 Dec 2011	R395.00	R20.00	R30.00	2.5
Euro/Rand (€/R)	16 March 2012	R400.00	R20.00	R30.00	2.5

Table 3.1: Margins for USDZAR and EURZAR futures during February 2011.

	B	C	D	E	F	G	H
3	7.47		**No St Devs**	3.5			
4	7.43	-0.0053691	Total Standard Deviations	4.352%			
5	7.419	-0.0014816	Maximum Market Movement	**4.448%**			
6	7.48	0.00818851	Volatility - Historical	19.74%			
7	7.495	0.00200334	Expected level at upper limit	7.22			
8	7.4888	-0.0008276	Expected level at lower limit	6.62			
9	7.6001	0.01475283	Differential (Rand)	0.31			
10	7.5576	-0.0056077					**3 Mo JIBAR**
11	7.56	0.00031751	Initial Margin Requirement - USD/RAND spc	308.00			5.28%
12	7.53	-0.0039761	Futures Expiry Date	Futures Level	**Initial Margin**	**IM Old**	**Spread Margin**
13	7.59	0.00793655	14-Mar-11	6.9671	**310**		**15**
14	7.65	0.00787406	13-Jun-11	7.0617	**315**		**15**
15	7.7276	0.01009269	19-Sep-11	7.1650	**320**		**15**
16	7.71	-0.0022801	19-Des-11	7.26520	**325**		**15**
17	7.6276	-0.0107449	16-Mar-12	7.37460	**330**		**15**
18	7.58	-0.00626	14-Des-15	9.28040	**415**		**20**
19	7.5476	-0.0042836	**Statistics**				
20	7.48	-0.0089968	**Mean**	(0.00003870948865)			
21	7.49	0.00133601	**St Deviation**	0.01243663219812			
22	7.3951	-0.0127512	**Non-Centered St Dev**	0.01243358290862	=SQRT(SUMSQ(C4:C2003)/COUNT(C4:C2003))		
23	7.327	-0.0092515	**Data Counts**				
24	7.3176	-0.0012837	Frequency BELOW lower limit	4			
25	7.33	0.00169311	Frequency ABOVE upper limit	6			
26	7.3466	0.00226211	Total between U and L limits	1,990.00			
27	7.45	0.01397641	Percentage of Data INSIDE limits	99.500%			
28	7.42	-0.004035	# Datapoints Outside	10.00			
29	7.3851	-0.0047146	% Outside	0.500%			
30	7.34	-0.0061256	Breach per annum (# days)	1.3			

	A	B	C	D	E	F	G	H
1	Zar/Gbp	GBPZAR=R		Current Spot Level	**11.2865**			
2	Date (GMT)		2000	Date of last Data Point	**28-Feb-11**			
3	30/06/2003	12.3569		**No St Devs**	3.5			
4	01/07/2003	12.3516	-0.00043	Total Standard Deviations	3.951%			
5	02/07/2003	12.3593	0.000623	Maximum Market Movement	**4.030%**			
6	03/07/2003	12.4774	0.00951	Volatility - Historical	17.92%			
7	04/07/2003	12.4964	0.001522	Expected level at upper limit	11.74			
8	07/07/2003	12.3445	-0.01223	Expected level at lower limit	10.85			
9	08/07/2003	12.4262	0.006597	Differential (Rand)	0.45			
10	09/07/2003	12.3718	-0.00439					**3 Mo JIBAR**
11	10/07/2003	12.3538	-0.00146	Initial Margin Requirement - GBP/RAND Spot	455.00			5.28%
12	11/07/2003	12.2784	-0.00612	Futures Expiry Date	Futures Level	**Initial Margin**	**IM Old**	**Spread Margin**
13	14/07/2003	12.2449	-0.00273	14-Mar-11	11.3140	**455**		**25**
14	15/07/2003	12.1704	-0.0061	13-Jun-11	11.4551	**460**		**25**
15	16/07/2003	12.3332	0.013288	19-Sep-11	11.6015	**470**		**25**
16	17/07/2003	12.3052	-0.00227	19-Des-11	11.7388	**475**		**25**
17	18/07/2003	12.1363	-0.01382	16-March_2012	11.8876	**480**		**25**
18	21/07/2003	12.1288	-0.00062	**Statistics**				
19	22/07/2003	12.0467	-0.00679	**Mean**	(0.00004530364513)			
20	23/07/2003	12.0181	-0.00238	**St Deviation**	0.01129139528757			
21	24/07/2003	12.1038	0.007106	**Non-Centered St Dev**	0.01128866299243			
22	25/07/2003	11.9764	-0.01058	**Data Counts**				
23	28/07/2003	11.9042	-0.00605	Frequency BELOW lower limit	5			
24	29/07/2003	11.8831	-0.00177	Frequency ABOVE upper limit	6			
25	30/07/2003	11.8467	-0.00307	Total between U and L limits	1,989.00			
26	31/07/2003	11.8288	-0.00151	Percentage of Data INSIDE limits	99.450%			
27	01/08/2003	11.9997	0.014344	# Datapoints Outside	11.00			
28	04/08/2003	11.9536	-0.00385	% Outside	0.550%			
29	05/08/2003	11.9284	-0.00211	Breach per annum (# days)	1.43			

Figure 3.2: Initial margin calculation for the USDZAR and GBPZAR futures contracts during February 2011.

a loss on the February contract, we might get a margin call. However, if we have a profit on the option contract, this profit can subsidise the loss and, we might end up with no margin call at all. Even initial margin can be offset against a profit in our total portfolio.

The exchange thus takes all open positions of a client into account when calculating the gross margin requirements.

3.15 Credit Risk

This was mentioned before. The big advantage of trading a futures contract on an exchange is that the playing field is leveled for everyone. The individual investor can now trade on a similar footing than a big bank or fund manager. The only credit risk is that of the clearinghouse because the clearinghouse is counterparty to every trade.

SAFCOM underwrites settlement of all positions of all participants on the exchange. This means that it has a resultant exposure to these positions. If settlement is likely to fail for any reason for spot bonds, the JSE Settlement Authority steps in and takes the necessary action to ensure that settlement is effected.

A liquidity provider like SBSA should thus not care whether they trade with me or with another bank like Absa. This makes trading on an exchange especially advantageous to the small investor. Bid/offer spreads are determined by liquidity (not credit risk). The small investor is now able to get exactly the same price as those a bank will offer to a big corporate client. However, initially, when the liquidity is not high, corporates might feel they will get better prices by doing OTC forwards.

3.15.1 Who Trades Futures?

We have already stated that there are only two types of trades: speculators and hedgers. Speculators love futures due to the leverage they can get. It is sometimes also cheaper to trade futures than to trade the underlying — this is the case for the single stock futures. The single stock futures (SSF) market on Safex is currently the second biggest SSF market in the world. It is growing rapidly.

However, hedging plays a big role. The hedger starts with a preexisting risk generated from the normal course of his or her traditional business. Futures (or forwards) are then used to reduce or eliminate this pre-existing exposure. The white maize future listed on the agricultural division of Safex is a good example. It is an extremely liquid contract. A lot of farmers and millers use it to hedge their physical maize positions. The Alsi future is another example. This contract (and its options) is used by the bigger banks to hedge OTC derivative structures on their books. Such structures are done by asset managers and pension funds who want to lay off risk to the risk takers (usually banks). The banks then lay off their risk to the market through Safex.

Yield-X is still in its infancy and the JIBAR futures are only now starting to take off. Currency futures has attracted many speculators, but it is also a means for the small investor to hedge his own currency risk. Small importing or exporting companies might also find a currency future useful. Liquidity in the currency futures contracts has picked up substantially, especially in the USDZAR contracts. Options on these futures are picking up nicely.

3.15.2 Risk Measurement

The forward value of an exchange rate is given by

$$F_T = S_0 \frac{1 + r_d T_d}{1 + r_f T_f}. \tag{3.4}$$

To understand the risk associated with a position in a futures contract, we need to understand what this equation really means and what the holder of a futures contract exposure is.

If a client buys a future, he buys a contract or agreement. He does not buy the underlying physical currency; he only gets exposure to that currency. As a matter of fact, he does not pay anything for that exposure! Remember that the initial margin is paid back to him when he exits the position. The client thus gets exposure to the currency for free. Exposure to any security has an element of risk to it. How did the client obtain exposure "for free?" The equation above holds the true tail.

On the other side of this long position is either another client or a liquidity provider who assumes a short position. Let's assume the liquidity provider is a bank and the client bought 1 USDZAR contract where the spot rate is 7.25 R/\$. What the bank has done is the following

- Lend the client \$1000
- Charge the client the Dollar interest rate
- Convert the Dollars into Rand at the spot rate
- Put the R7,250 on deposit at the domestic Rand interest rate.

The equation above has all these elements. If we break the trade down like this we see that the client is long a deposit in Rand and long a loan in Dollars. His true risk is at expiry (or when he unwinds) when the opposite will be done and the R7,250 + interest is converted back into Dollars at the then prevailing spot exchange rate. If the prevailing spot rate is below the quoted futures rate, he looses money and he makes money if the prevailing spot rate is above the futures price.

There are thus 3 quantities that underpin the risk: the spot exchange rate, S_0, the Dollar interest, r_f and the domestic interest r_d. The client has exposure to all of these but only paid a small initial margin amount. From the liquidity provider's perspective: all three risks need to be measured and managed.

3.15.3 Risk Parameters

The Delta

Let's first look at the spot exchange rate risk also called the market risk. This is the risk due to the change in value of the underlying asset. From Equation (3.4) we deduce there is a direct or linear relationship between the spot price S_0 and the value of the future. The risk parameter to help us in hedging this risk is called the delta.

The delta is the ratio of the change in the value of the futures contract to the change in the value of the underlying spot FX rate. The Delta is the amount that the derivative will theoretically change in price for a one-point move in the underlying. This ratio will give us the number of units of the spot currency to hold or sell in order to create a riskless hedge. The delta relates the derivative's price to that of the underlying asset. This is mathematically given by

$$\Delta = \frac{\Delta F_T}{\Delta S_0} \approx \frac{\partial F_T}{\partial S_0} = \frac{1 + r_d T_d}{1 + r_f T_f}. \tag{3.5}$$

This delta is similar to an option's delta risk measure — the gamma of a future is zero. Note that the following holds

$$\Delta \quad is \quad \begin{cases} > 1 & if \quad r_d > r_f \\ = 1 & if \quad r_d = r_f \\ < 1 & if \quad r_d < r_f \end{cases}$$

The delta shows us how many of the underlying spot contracts must be bought or sold in order to hedge the position in the future.

Delta hedging is done continuously because the Delta is not constant — for the return on the hedge portfolio to remain riskless, the portfolio must continuously be adjusted as the asset price changes over time.

Rho Risk

Rho risk is risk due to interest rate exposure. From Equation (3.4) we ascertain that interest rates are a big risk factor in any forward or futures contract. There are, however, two interest rate risks: risk due to exposure to the domestic interest rate and risk due to the exposure to foreign interest rates.

The Rho of a portfolio of derivatives is the rate of change of the value of the portfolio with respect to the interest rate. We now define the mathematical Rho as follows

$$\rho_d = \frac{\Delta F_T}{\Delta r_d} \approx \frac{\partial F_T}{\partial r_d} = \frac{S_0 T_d}{1 + r_f T_f} \tag{3.6}$$

and

$$\rho_f = \frac{\Delta F_T}{\Delta r_f} \approx \frac{\partial F_T}{\partial r_f} = \frac{-S_0 T_f \ (1 + r_d T_d)}{(1 + r_f T_f)^2} \tag{3.7}$$

Another way to express the interest rate risk is to define the Present Value of a Basis Point (the so-called PV01). PV01 measures the change in value of interest rate sensitive exposures resulting from a 0.01% increase in interest rates. We see this as the numerical Rho measure. It is defined as follows where we now write$F_T = F(r_d, r_f)$, i.e. F_t is a function of both r_d and r_f

$$\rho_d = F(r_d + 0.0001) - F(r_d) \tag{3.8}$$

$$\rho_f = F(r_f + 0.0001) - F(r_f) \tag{3.9}$$

If we have a portfolio of contracts we "bump" the yield curve by 1 basis point and determine the difference between our original value and the new value.

The PV01 shows us how much money we will make or loose if the yield curve moves up by 1 basis point. This is similar to the "rand per point" measured used in the bond market.

Theta Risk

Theta is the risk due to time. We see from Equation (3.4) that F_T is also a function of time. This is called time value. The time value for an option is extremely important; it is less important for a future. Let's define the numerical Theta as follows: it is the change in the value of the futures contract from one day to the next. We can express it as follows

$$\Theta_d = F\left(T_d - {}^{1}\!/_{365}\right) - F(T_d) \tag{3.10}$$

$$\Theta_f = F\left(T_f - {}^{1}\!/_{\alpha}\right) - F(T_f) \tag{3.11}$$

where $\alpha = 360$ for the USD and Euro and $\alpha = 365$ for the GBP.

3.15.4 Hedging

Convergence of Futures Prices to Spot Prices

As the expiry date of the futures contract approaches, the futures price will eventually converge to the spot price of the underlying asset[7]. Two alternative ways of doing so are presented in Fig. 3.3. The future/forward can be higher than the spot or vice versa.

On the close-out date, the futures price must be equal to the spot price. If this was not the case, there would be clear arbitrage opportunities. For instance, if the futures were above the spot price, one could exploit the following arbitrage opportunity

1. Short a futures

[7]See http://www.theponytail.net/DOL/DOLnode1.htm

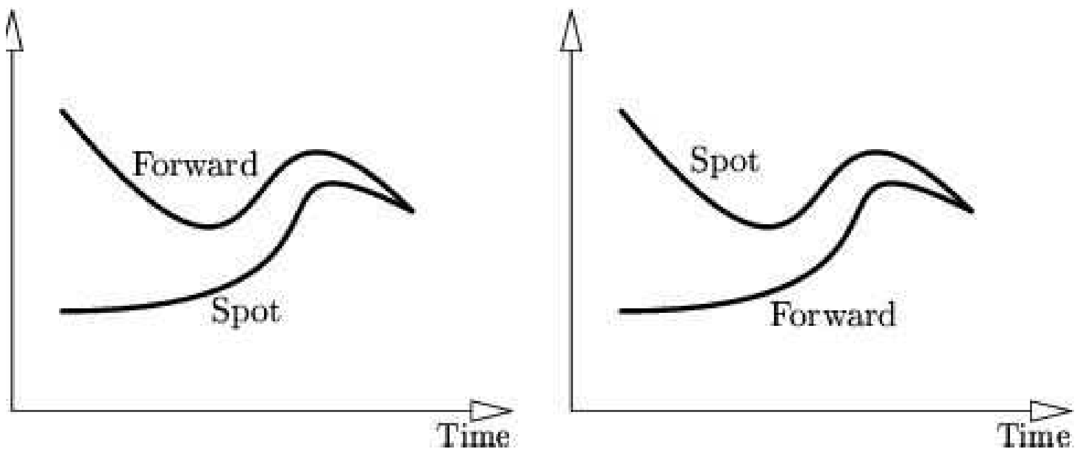

Figure 3.3: Future and spot convergence.

2. Buy the asset
3. On close-out, sell the asset.

As arbitrageurs exploit this opportunity, the price of the futures will decline and the price of the underlying asset will rise. This pattern will continue until the futures and the spot price become equal.

Basis Risk

Basis risk in finance is the risk associated with imperfect hedging using futures. It could arise because of the difference between the asset whose price is to be hedged and the asset underlying the derivative, or because of a mismatch between the expiration date of the future and the actual selling date of the asset.

Under these conditions, the spot price of the asset, and the price of the future do not converge on the expiration date of the future. The amount by which the two quantities differ measures the value of the basis risk. That is,

$$\boxed{\text{Basis} = \text{Spot price of hedged asset - Futures price of contract used}}$$

If the asset to be hedged and the assets underlying the future contract are the same, the basis should be zero at expiration of the futures contract (this is called "pull to par"). This is shown in Fig. 3.13. Prior to expiration, the basis may be positive or negative. In general, basis risk increases as the time difference between the hedge expiration and OTC's expiration increases.

There are some sources of basis risk

- Changes in the convergence of the futures price to the spot price.
- Changes in factors that affect the cost of carry: storage and insurance costs, opportunity cost.
- Different natures of mismatched assets.
- Maturity mismatch.
- Liquidity difference.
- Credit risk difference.
- Random Deviation from the Cost-of-Carry Relation.

Due to basis risk, the equivalent hedge given in Eq. (3.5) might not be optimal. A trader can get profit and loss swings due to the fact that the spread between the spot and futures contract changes as the market moves. We can define a “better” hedge ratio as follows

$$h = \Delta\,\rho\frac{\sigma_S}{\sigma_F} = \frac{1+r_dT_d}{1+r_fT_f}\rho\frac{\sigma_S}{\sigma_F}. \tag{3.12}$$

In this equation, ρ is the correlation coefficient between the spot price and the futures price, σ_S is the volatility of the spot price and σ_F is the volatility of the futures price. If we hedge with the underlying to the future, we should have $\rho = 1$ and $\sigma_S = \sigma_F$ giving $h = \Delta$ with Δ given in (3.5).

Rolling the Hedge

Let’s assume a trader has exposure to an OTC forward. He wants to hedge using futures. However, the longest future’s expiry date is before the expiry date of the forward. On the expiry date of the future, the hedger must then roll the hedge forward – this means he has to open a position in another futures contract to stay hedged. Hedges can be rolled forward many times.

When rolling a contract forward, there is uncertainty about the difference between the futures price for the contract being closed out and the futures price for the new contract. Hedgers reduce the rollover risks by switching contracts at favourable times. The hedger hopes that there will be times when the basis between different futures contracts is favourable for a switch. Such switches usually happen the week before the futures close-out date. This is evident in the Alsi futures contract. Volumes in Alsi futures traded usually rise during the week or two before a close-out date.

3.15.5 Arbitrage Opportunities

In the previous section we mentioned that the futures price must converge to the spot price on expiry. However, before expiry and in a liquid market, the futures price and the equivalent theoretical forward price might not always be the same.

We might have a situation where

$$F > S_0 \frac{1 + r_d T_d}{1 + r_f T_f}.$$

We say that the future is trading at a premium to the forward. Such a mismatch can be turned into riskless profits by buying the underlying asset and selling the future. Such trades will decrease the difference between the future and forward. The opposite might also be true where

$$F < S_0 \frac{1 + r_d T_d}{1 + r_f T_f}.$$

We then say the future trades at a discount. Profits can be made by buying the future and by selling (or shorting) the underlying asset. On expiry, due to the convergence of the future and spot prices, both strategies would have made money. However, such strategies are seldom held till expiry. Traders roll out of them when it is opportune to do so.

3.16 What is Margin, Novation and Safcom?

From the introduction we deduce that margining is an important part of the risk management process utilised by an exchange. Let's define what we mean by "margin". To minimize credit and market risk to the exchange, derivative traders must post margin.

Margin helps derivative exchanges to avoid credit and market risk, i.e., the chance of one or more counterparties to a trade, defaulting on their obligations. They accomplish this in two ways. Firstly, all trades on an exchange are settled or "cleared" through a clearinghouse which may be a separate legal entity to the exchange itself. The JSE's clearing house is SAFCOM. The clearinghouse acts as the principal counterparty to all trades through an exchange. Thus, it interposes itself as the 'buyer to every seller' and the 'seller to every buyer' – known as novation. Through novation Safcom guarantees to its members the financial performance of all contracts traded. SAFCOM becomes the guarantor of all futures transactions allowing members participants to deal freely with each other without counterparty credit risk constraints. This process is graphically shown in Fig. 3.16. Secondly, exchanges employ a system of margining. Accordingly, a counterparty to a transaction on an exchange is required to pay a sum over to it at the inception of the derivative transaction to cover any potential losses arising from a default.

There are 5 different types of margin

Figure 3.4: The clearing house (Safcom) becomes guarantor to each trade — the process of novation.

1. Initial Margin: is the amount of money determined by the clearing house on the basis specified by the risk management committee (RMCO) and held in respect of the aggregate position of a member or a client – this is paid by both buyers and sellers. Initial margin shall be paid to, or by, a member or client whenever the risk of loss changes with respect to the aggregate position (it is also called a good faith deposit). This margin is reinvested at a competitive rate and at close out of the positions of the client/member the initial margin is paid back plus the interest earned for the period. The initial margin may be reduced or increased based on changes in the margin parameters.

2. Variation Margin: is paid by the members or clients on a daily basis as the result of the mark-to-market process of the clients'/members' position. Mark-to-market refers to the present loss/profits of the position.

3. Additional Margin: clearing members may require additional margin from his members and members may require additional margin from their clients.

4. Retained Margin: Member may require a client to deposit retained margin with him which may be used to furnish initial and additional margin requirements.

5. Maintenance Margin: The client may have to top up his account with the member with maintenance margin. The client has to pay an amount of money to restore additional margin when the additional margin has been used to meet payments of variation margin.

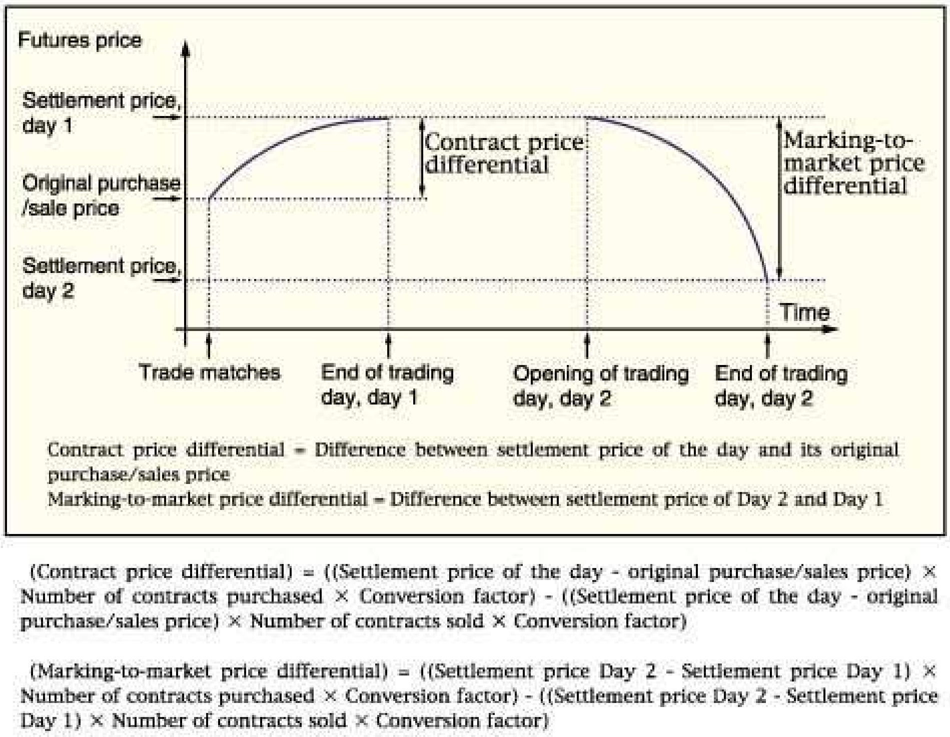

Figure 3.5: Marking-to-market and variation margin.

3.17 Initial and Variation Margin

Risk management may be defined as identifying the risks of loss in a portfolio and ensuring that the losses can be borne. In the case of a futures exchange, market risk management is performed in two steps: marking-to-market and margining.

Marking-to-market ensures that all losses up to the present are absorbed. Participants with losses are required to make cash payments to the exchange equal to their losses – this is called variation margin. Safcom operates under the T+1 method of "pays and collects", meaning that all profits/losses (change in value) in all accounts is received or paid by clearing participants by noon of the business day following the day the change occurred. This entire process goes a long way to insure market integrity and is graphically depicted in Fig. 3.17.

The exchange then also estimates what losses are possible in the future — usually 1 trading day. Participants are required to lodge margins with the exchange which are sufficient to cover these possible future losses – this is called initial margin. Should the losses eventuate and the participant be unable to bear them, the margin is available to the exchange to meet the shortfall.

There are two stages to estimating possible future losses and the initial margin requirements

- The exchange does a statistical analysis of historical market moves and subjective assessments of the state of the market. They express the maximum anticipated price and volatility moves between the present and the next mark-to-market day.

- Secondly, the exchange re-values each position at this maximum anticipated price and volatility at the next mark-to-market day. The margin covers this maximum conceivable mark-to-market loss that the position (entire portfolio) could suffer.

This is called the “Spread Margining” methodology and is in use at most derivative exchanges across the globe. This methodology is similar to a Value-at-Risk (VAR) analysis[8] [Al 01].

An exchange controls the risk exposure (manages the risk) by changing the margin requirements. In times of uncertainty or high volatility, margins should be adjusted higher whilst in quite times it can be adjusted lower. At all times, the exchange will want to be confident that it has allowed for the sudden unanticipated shocks which characterise the markets.

3.18 Safex Can-Do Structures

Can-Do’s, the JSE’s latest derivative offering, gives investors the advantages of listed derivatives with the flexibility of OTC contracts. Can-Do’s allow two counter-parties to negotiate the terms of a contract, specifically the expiry date and underlying asset. The expiry date will not be required to be the close-out dates specified for the single stock futures and index futures and options. Can-Dos will soon also be available on the currency derivatives and commodity derivatives markets.

Now investors will be able to select any business day as the expiry date of the contract, which will allow funds to refine and co-ordinate their hedging strategies. In addition, the reference asset may be a basket of assets, the constituents of which will be negotiated by the long and short-position holder (previously investors had been limited to specific stocks and indices). This basket will be valued by the JSE on a daily basis in accordance with the accepted practice for listed derivatives.

The JSE will then margin each contract holder and in this way nullify any risk of default. Can-Do’s therefore provide the investment community with a derivative that is a hybrid of OTC and traditional listed derivatives: a product that is flexible, transparent and is independently valued.

In Fig. 3.18 we show a screen shot of the current Can-Dos on Yield-X. These futures have broken expiry dates.

[8] Value at Risk is defined as the potential market loss in a portfolio over a specified period of time – usually 1 or 10 days. The analysis is based on volatility and correlation.

Currency_Cando (SAFMABA)

Instrument	QtyB	Bid	/	Offer	QtyO	Change	Last	Time	High	Low	Volume	OpenInt
18 MAR11 ANY DAY EXPIRY CAAB $	0	0.000	/	0.000	0	0.000	6.879	00:00	0.000	0.000	0	1000
18 MAR11 ANY DAY EXPIRY CAAC €	0	0.000	/	0.000	0	0.000	9.558	00:00	0.000	0.000	0	1000

Figure 3.6: Yield-X Can-Dos.

3.18.1 Advantages of Listed Derivatives

Listed derivatives are standardised contracts. This means that the contract is available to all market participants. Listed derivatives have several advantages over OTC contracts.

- The exchange acts as the central counter-party to all trades, and any counterparty risk is off-set using a method known as margining. This means that if you enter a listed contract, there is virtually no risk of default.
- Listed contracts can be bought and sold. If a party enters a contract and later decides to exit it, he is able to do so by 'closing out' the contract.
- The JSE values all listed contracts on a daily basis and provides a 'market value'. This generally gives auditors a great deal of comfort as holders of long and short positions must value their positions equal and opposite. This gives the market a third part independent valuation that ensures adherence to accounting law AC133 whereby certain instruments have to be valued by an independent 3rd party.
- Can-Do's provide a great deal of flexibility whereby a client can decide on the type of product to trade, the underlying of the product, expiry dates and time, expiry valuation method and the nominal underlying size.

In Fig. 3.18.1 we show a screen shot of the currency futures traded on Yield-X in South Africa.

3.18.2 Disadvantages of OTC Derivatives

- Prior to entering the contract, the parties must assess each others creditworthiness. This adds an element of cost to the transaction and also implies a certain amount of risk. Institutions with small balance sheets are often considered to be too high a credit-risk for the major banks and are therefore excluded from the market, or receive less favourable pricing than bigger companies.

Currency_Contracts (SAFMABA)

Instrument	QtyB	Bid	/	Offer	QtyO	Change	Last	Time	High	Low	Volume	OpenInt
14 MA 11 AU$ / R	200	6.9070	/	6.9510	200	0.0000	6.9424	00:00	0.0000	0.0000	0	7,349
13 JU 11 AU$ / R	200	6.9120	/	6.9720	200	0.0000	6.9574	00:00	0.0000	0.0000	0	10,408
19 SEP11 AU$ / R	200	6.9140	/	6.9950	200	-0.0072	6.9650	09:54	6.9650	6.9650	100	6,240
14 MAR11 CAD/ R	100	7.1030	/	7.1320	100	0.0000	7.0911	00:00	0.0000	0.0000	0	1,641
13 JUN11 CAD/ R	100	7.1630	/	7.2340	100	0.0000	7.1733	00:00	0.0000	0.0000	0	0
14 MAR11 € / R	200	9.5170	/	9.5630	200	-0.0285	9.5300	10:15	9.5562	9.5300	16	17,193
13 JUN11 € / R	200	9.6260	/	9.6860	200	0.0000	9.6749	00:00	0.0000	0.0000	0	8,958
19 SEP11 € / R	200	9.7380	/	9.8130	200	0.0000	9.7942	00:00	0.0000	0.0000	0	3,261
19 DEC11 € / R	0	0.0000	/	0.0000	0	0.0000	9.9048	00:00	0.0000	0.0000	0	0
14 MAR11 CHF / R	100	7.3810	/	7.4190	100	0.0000	7.4089	00:00	0.0000	0.0000	0	0
13 JUN11 CHF / R	100	7.4700	/	7.5400	100	0.0000	7.5151	00:00	0.0000	0.0000	0	0
14 MAR11 £ / R	200	11.1110	/	11.1600	200	0.0000	11.1214	00:00	0.0000	0.0000	0	29,472
13 JUN11 £ / R	200	11.2430	/	11.3060	200	0.0320	11.2930	10:36	11.2930	11.2865	110	4,187
19 SEP11 £ / R	200	11.3820	/	11.4600	200	0.0000	11.4064	00:00	0.0000	0.0000	0	2,669
19 DEC11 £ / R	0	0.0000	/	0.0000	0	0.0000	11.5422	00:00	0.0000	0.0000	0	87
14 MAR11 ¥ / R	100	0.082940	/	0.083950	100	0.000000	0.083100	00:00	0.000000	0.000000	0	400
13 JUN11 ¥ / R	100	0.083040	/	0.085550	100	0.000000	0.084300	00:00	0.000000	0.000000	0	573
14 MAR11 $ / R MAXI	10	6.8961	/	6.9038	10	0.0000	6.8746	00:00	0.0000	0.0000	0	70
13 JUN11 $ / R MAXI	10	6.9901	/	6.9991	10	0.0000	6.9685	00:00	0.0000	0.0000	0	330
14 MA 11 $ / R	302	6.8965	/	6.9038	500	0.0254	6.9000	10:44	6.9191	6.8965	4,294	205,749
13 JU 11 $ / R	500	6.9901	/	6.9993	500	0.0250	6.9935	10:36	7.0160	6.9935	886	83,733
19 SE 11 $ / R	500	7.0911	/	7.1040	500	0.0295	7.1010	10:23	7.1020	7.1010	16	13,993
19 DE 11 $ / R	200	7.1909	/	7.2068	500	0.0000	7.1712	00:00	0.0000	0.0000	0	890
19 MAR12 $ / R	0	0.0000	/	0.0000	0	0.0000	7.2811	00:00	0.0000	0.0000	0	277

Figure 3.7: Trading front-end on Yield-X.

- Parties to a contract are not able to sell their contractual obligation to a third party. Once a contract has been entered into, the only way that a party can get out of its obligation is by way of early settlement of the contract (if this is catered for) or to default.

- The contract is not valued by an independent valuation agent. This means that each party to the contract may attach its own value to the derivative position. It is not uncommon for two parties to assign vastly different value to the same OTC contract - a nightmare for auditors and a practice that has resulted in several cases of fraud overseas.

3.18.3 Exotic Derivatives

The JSE and SAFEX is the first exchange in the world to list, trade and clear exotic options. Safex envisaged that the appetite would only be for vanilla exotics. On 8 January 2007 the first exotic was traded; a discrete look-back put spread. Since then the types of exotics traded grew in leaps and bounds with most of the traded exotics being complex in nature. Most exotics have the Alsi or JSE/FTSE Top 40 index as underlying instrument but trades have been done on the DTOP and some single names as well. To date the following exotics have been listed and traded:

- Fixed and floating strike lookback options

  - put spreads with discrete monitoring
  - partial time lookbacks
  - including Asian option features

- Barrier Options (Knock-in and Knock-outs)
  - discrete and continuous monitoring of barrier
  - partial time barrier monitoring
  - ladders or timer put spreads or strike resetting options
  - double barriers (KIKO options)

- Forward start Options

- Asian Options

- Digital/Binary Options
  - cash or nothing and asset or nothing
  - including barriers

- Cliquets

Can-Do Exotics are generally structured products where the details are agreed upon by two counterparties. The counterparties will approach the JSE who will inspect all aspects of the deal. These include the adherence to the rules and regulations of the JSE and the daily valuation and margining of the instrument. If the need exists, the JSE will develop a pricing model independent from the two counterparties. The model is parallel tested with the counterparties.

The test phase does not just concentrate on the pricing but also on how interest rates, dividends and volatilities should be meaningfully incorporated into the model.

Types of models developed thus far include closed form, Monte Carlo and tree based approaches. A new instrument will only be listed once all parties involved approve the JSE's valuation methodology. To date we have not yet come across an option that we were not able to value. Our motto has always been (subject to the JSE rules and regulations): "If we can value and clear it, we will list it".

3.18.4 Exotics: the way Forward

Currently, exotics are only written with equity as the underlying product. However, in the near future, traders and investors will be able to trade exotic option structures on a multiple of underlying instruments including interest rate products, currencies and commodities.

Equity linked swaps were just introduced in the very near future whereby an investor/trader will be able to trade one contract with multiple positions. For example, in the simplest form, a basket could be created with a long position in stock ABC and a short position in stock XYZ. Offset margins will apply.

Chapter 4

Shariah Compliant Derivatives

There are various prohibitions in Islam regarding banking that must be abided by and in this regard the Shari'ah prohibits uncertainty (gharar) and gambling (qimar). As a result, many of the structures that have been created to provide the characteristics of conventional derivatives while still maintaining Shari'ah compliance are proprietary and are often not generally openly available. But, 'there is the concept of ibaha which means that if something is not banned then it is permitted'. Under this principle, 'because something appears to be similar to something that is banned, then don't assume that it too is banned'.

Looking looking at some of the financial products on the market, it seems that everything is possible using murabaha. For instance, using 'murabaha an investor can "invest" in an "arm's length Special Purpose Vehicle" (a specially formed company) that in turn could create "trades" in anything —- from options to futures to warrants.

4.1 Introduction

Risk-Management refers to the process/techniques of reducing the risks faced in an investment. It generally involves three broad steps

- Identifying the source and type of risk.
- Measuring the extent of the risk.
- Determining the appropriate response (either on Balance Sheet or Off Balance Sheet) methods.

What makes risk management challenging is the fact that risks and returns are generally positively correlated. Thus, the risk-return tradeoff. The challenge of risk-management is to protect the expected returns while simultaneously reducing or laying-off the risks. Note that all risk management techniques involving derivatives are Off Balance Sheet. What this means is that, the hedging mechanism/method is "detached" from the underlying transaction.

The advantage: No need to change the way one does business. No loss of competitiveness, customer convenience etc. An On Balance-Sheet technique is one where a transaction is structured in such a way as to manage the inherent risk.

Example: Malaysian Exporter; Foreign Customer. On Balance Sheet Technique

- Quote only in Ringgit (domestic currency)
- Increase the foreign currency price equivalent to cover risk (pricing strategy)
- Currency Risk Sharing Agreement.

Off Balance Sheet

- Forwards; Short foreign currency (FC) forward contracts.
- Futures; Short FC futures contracts.
- Options; Long FC Put Options.
- Swaps; FC payer, domestic currency receiver

Off Balance Sheet techniques have become tremendously popular due to

- Cheap and flexible
- No inconvenience to customer
- Can enhance competitiveness.

Despite the popularity of derivatives based off balance sheet techniques, Islamic Jurists have generally not been in favor.

With help of Bahrain-based International Islamic Financial Market and New York-based *International Swaps and Derivatives Association* (ISDA), global standards for Islamic derivatives were set in 2010. The "Hedging Master Agreement" provides a structure under which institutions can trade derivatives such as profit-rate and currency swaps.

4.2 Shari'ah Compliant Derivatives

Shariah Compliant derivative instruments should adhere to the following principles to be considered halal (acceptable)

- At a primary level all financial instruments and transactions must be free of at least the following five items
  1. riba (usury)

2. rishwah (corruption)
3. maysir (gambling)
4. gharar (unnecessary risk) and
5. jahl (ignorance).

- Riba can be in different forms and is prohibited in all its forms. For example, Riba can also occur when one gets a positive return without taking any risk.
- As for gharar, there appears to be no consensus on what gharar means. It has been taken to mean, unnecessary risk, deception or intentionally induced uncertainty.
- In the context of financial transactions, gharar could be thought of as looseness of the underlying contract such that one or both parties are uncertain about possible outcomes.
- Masyir from a financial instrument viewpoint would be one where the outcome is purely dependent on chance alone - as in gambling.
- Finally, jahl refers to ignorance. From a financial transaction viewpoint, it would be unacceptable if one party to the transaction gains because of the other party's ignorance.

In addition to the above mentioned requirements for financial instruments, the Shari'ah has some basic conditions with regards to the sale of an asset (in this case a real asset as opposed to financial assets). According to the Shari'ah for a sale to be valid,

1. the commodity or underlying asset must currently exist in its physical sellable form and
2. the seller should have legal ownership of the asset in its final form.

If something is forbidden by Islamic law, it is called haram.

From all of these conditions and principles, we see that trading of derivatives would be very difficult if not impossible. But all is not lost because the Shari'ah provides exceptions to these general principles to enable deferred sale where needed.

4.3 Futures Contracts and Islamic Finance

The futures contract, although standardized and supervised by law, is not permissible under Shari'ah as it contains elements of both gharar and maisir/qimar. Moreover, there is no balance between profit and risk sharing between the parties. In respect of gharar, the result of the contract is in the future, hence is unknown for the parties at the time it is entered into. The parties bet on the future result of the contract. The risk is borne solely by the person who pays the margin, as a financial institution is secured thanks to the margin and in fact never loses [Pa 09].

Another interesting aspect in case of derivatives is arbitrage. Arbitrage is speculating on differences in prices of commodities and assets on different markets. The persons involved in arbitrage called arbitrageurs earn on price differences of the assets. The practice of arbitrage would also be considered haram as the profit is based on speculating.

Another aspect to consider is leverage. Under Shari'ah, using leverage would be deemed haram and would invalidate the whole transaction. The leverage transaction involves payment of interest (loans, facilities). The aim of the transaction is to gain maximum profit. There are no intentions of the transaction being beneficial to society or for charitable purposes. Also, the parties to the transaction make a bet on the prospected profit or loss of the investment.

A number of instruments/contracts exist in Islamic finance that could be considered a basis for forward/futures contracts within an Islamic framework. We will examine three such contracts. These are

1. the Salam Contract,
2. the Istisna Contract,
3. Joa'la Contract and
4. Wa'ad Contract.

Each of these contracts concern deferred transactions, and would be applicable for different situations. The first and probably the most relevant of these to modern day forward/futures contracts would be the Salam Contract or Ba'i Salam.

4.3.1 Ba'i Salam

Salam is essentially a transaction where two parties agree to carry out a sale/purchase of an underlying asset at a predetermined future date but at a price determined and fully paid for today. This is similar to a conventional forward contract however, the big difference is that in a Salam sale, the buyer pays the *entire amount in full* at the time the contract is initiated. The contract also stipulates that the payment must be in cash form.

The idea behind such a 'prepayment' requirement has to do with the fact that the objective in a Ba'i Salam contract is to help needy farmers and small businesses with working capital financing. Since there is full prepayment, a Salam sale is clearly beneficial to the seller. As such, the predetermined price is normally lower than the prevailing spot price.

This price behavior is certainly different from that of conventional futures contracts where the futures price is typically higher than the spot price by the amount of the carrying cost. The lower Salam price compared to spot is the "compensation" by the seller to the buyer for the privilege given him. Despite allowing Salam sale, Salam is still an exception within the Islamic financial system which generally discourages forward sales, particularly of foodstuff. Thus, Ba'i Salam is subject to several conditions

1. Full payment by buyer at the time of effecting sale.
2. The underlying asset must be standardizable, easily quantifiable and of determinate quality.
3. Cannot be based on an uniquely identified underlying.
4. Quantity, quality, maturity date and place of delivery must be clearly enumerated.

It should be clear that current exchange traded futures would conform to these conditions with the exception of the first, which requires full advance payment by the buyer. Given the customized nature of Ba'i Salam, it would more closely resemble forwards rather than futures. Thus, some of the problems of forwards like "Double-coincidence", negotiated price and counterparty risk can exist in the Salam sale.

Counterparty risk however would be one sided. Since the buyer has paid in full, it is the buyer who faces the seller's default risk and not both ways as in forwards/futures. In order to overcome the potential for default on the part of the seller, the Shariah allows for the buyer to require security which may be in the form of a guarantee or mortgage.

4.3.2 The Salam Contract & Islamic Financial Institutions

Since the Salam Contract involves transacting in the underlying asset and financial institutions may not want to be transacting in the underlying asset, there are a number of alternatives available. These are in the form of parallel Salam Contracts. Jurists however are not all in agreement of the permissibility. We have

- Parallel with Seller
  - After entering into the original Salam Contract, the bank can get into a parallel Salam sale to sell the underlying commodity after a time lapse for the same maturity date.
  - The resale price would be higher and considered justifiable since there has been a time lapse. The difference between the 2 prices would constitute the bank's profit. The shorter the time left to maturity, the higher would be the price.

  - Both transactions should be independent of each other. The original transaction should not have been priced with the intention to do a subsequent parallel Salam.

- Offsetting Transaction with Third Party (Istisna and Joala contracts).
  - Here, the bank which had gone into an original Salam Contract enters into a contract promising to sell the commodity to a third party on the delivery date.
  - Since this is not a Salam Contract the bank does not receive advance payment.
  - It would be a transaction carried out on maturity date based on a predetermined price.

Note: This is very much like modern day forward/futures. The difference here being that the Islamic bank is offsetting an obligation - not speculating.

4.3.3 Istisna and Joala Contracts

In addition to Ba'i Salam, there are two other contracts where a transaction is made on a "yet to" exist underlying assets. These are the Istisna and Joala contracts. The Istisna Contract has as its underlying, a product to be manufactured. Essentially, in an Istisna, a buyer contracts with a manufacturer to manufacture a needed product to his specifications. The price for the product is agreed upon and fixed. While the agreement may be cancelled by either party before production begins, it cannot be cancelled unilaterally once the manufacturer begins production.

Unlike the Salam Contract, the payment here is not made in advance. The time of delivery too is not fixed. Like Ba'i Salam, a parallel contract is often allowed for in Istisna.

The Joala Contract is essentially a Istisna but applicable for services as opposed to a manufactured product.

4.3.4 The Bai'bil-wafa and Bai 'bil Istighlal Contracts

The Bail bil-wafa is a composite of bai (sale) and rahnu (pledge). Under this contract, one party sells an asset to a buyer who pledges to sell back the asset to the original owner at a predetermined future date.

The rahnu (pledge) being to sell back to the owner and not to a third party. Looks like a REPO? Except that the resale price must be the same as the original purchase price. But like a REPO, the buyer has rights to benefits from ownership of the asset.

The Bai bil-Istighlal is really a combination of the Bai wafa and Ijarah. Under this contract, the buyer not only promises to resell at a predetermined future price but to also lease the asset to the seller in the interim period. The Bai bil-Istighlal can therefore be a convenient means by which an investment bank can provide short/medium term financing. The investment bank first purchases the asset, leases it to the customer before finally reselling it to the customer.

4.3.5 Wa'ad

A third common Shari'ah compliant hedging mechanism that has been developed over the last few years has been based on the concept of wa'ad. The Wa'ad is used in managing FX risk. Essentially, party A, who is looking for a hedge, will provide an undertaking (a wa'ad) to purchase a specific currency at a future date. The promise cannot be conditional on any event, and will have details of the amount of the currency to be purchased along with the future date of purchase [Hu 09].

4.4 Options in Islamic Finance

From the Islamic banking perspective, options contracts are considered haram and thus are not permissible under Shari'ah. The reason is that the option contracts generally trade in possibility, which means that no commodity is actually transferred between the parties. Both in call and put option contracts, the buyer/the seller intend to buy/sell the asset, but at the time the contract is entered into, there is no transfer of commodity. Furthermore, option contracts are entered into for future purposes, hence the buyer/seller may not execute the contract and in the end, may not sell/buy the asset. The option contracts trade in possibility of buying or selling asset in the future.

The whole concept of option contracts is based on bet and chance. If we were to use Islamic terms, we would say that an option contract is based on riba as interest is involved in the form of a premium paid by the client. And just this one element renders the contract haram and a heavy sin by Islam standards.

Furthermore, the contract contains gharar as the result is uncertain and the parties do not possess the knowledge what the outcome of the option will be. The gharar may also refer to the lack of sufficient experience and know-how regarding the consequences of entering into an option contract. One may not earn anything, but might suffer serious financial losses, one did not expect upon entering into the contract. The gharar may also cover the unclear and ambiguous provisions of an option contract, which are not comprehensible to the client. The financial institution can hire a team of both lawyers and economists to draft such a contract, that its terms and conditions will be hard to comprehend for the customer. Such occurrence of gharar would also invalidate the contract [Pa 09].

Recall our earlier argument that to be acceptable an instrument/investment must be free of gharar and not have zero risk in order to provide some positive return. The

Istijrar Contract is a recently introduced Islamic financing instrument. The contract has embedded options that could be triggered if an underlying asset's price exceeds certain bounds. The contract is complex in that it constitutes a combination of options, average prices and Murabaha or cost plus financing.

4.4.1 Overview of Istijar

The Istijrar involves two parties, a buyer which could be a company seeking financing to purchase the underlying asset and a financial institution. A typical Istijrar transaction could be as follows; a company seeking short term working capital to finance the purchase of a commodity like a needed raw material approaches a bank. The bank purchases the commodity at the current price (P_0), and resells it to the company for payment to be made at a mutually agreed upon date in the future — for example in 3 months. The price at which settlement occurs on maturity is contingent on the underlying asset's price movement from t_0 to $t_{90} = T$.

Unlike a Murabaha contract where the settlement price would simply be a predetermined price; $P(T)$ where

$$P(T) = P_0(1 + r)$$

with r being the bank's required return/earning, the price at which the Istijrar is settled on maturity date could either be $P(T)$ or an average price $\bar{P}$ of the commodity between the period t_0 and T.

As to which of the two prices will be used for settlement will depend on how prices have behaved and which party chooses to "fix" the settlement price. The embedded option is the right to choose to fix the price at which settlement will occur at anytime before contract maturity.

At the initiation of the contract both parties agree on the following two items

1. in the predetermined Murabaha price; $P(T)$ and
2. an upper and lower bound around the P_0 (bank's purchase price at t_0).

We show the possible settlement prices in Fig. 4.4.1. In Fig. 4.4.1 we have

$$\begin{aligned} P_0 &= \text{the price the bank pays to purchase the underlying} \\ P^* &= P(T) = P_0(1 + r) = \text{Murabaha price} \\ P_{LB} &= \text{the lower bound price} \\ P_{UB} &= \text{the upper bound price} \end{aligned}$$

The basic idea behind such a contract is to spread out the benefits of favourable price movement to both parties, i.e. it is not a zero sum game. Such a contract fulfills the need to avoid a fixed return on a riskless asset which would be considered "riba" and also avoids gharar in that both parties know up front what $P(T)$ is and also the range of other possible prices (by definition between the upper and lower bounds).

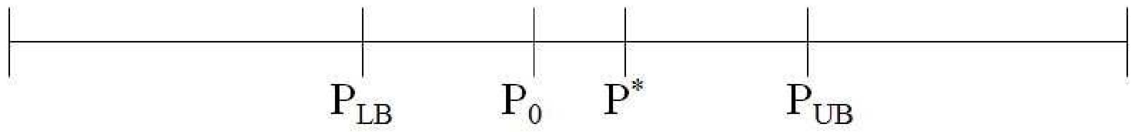

Figure 4.1: Settlement prices.

Given our description of the Istijrar Contract, the contract comes across as something that is the result of modern day financial engineering. Many of the products of financial engineering tend to have the complexities, bounds, trigger points etc. similar to that of the Istijrar.

4.4.2 Concept of Urban

When analysing the option contract, one cannot ignore urbun, known also as a down-payment sale. The buyer who intends to buy a certain commodity in the future pays a certain amount to the seller as a down-payment. If the buyer purchases the commodity, the down- payment is counted towards the total price for the commodity. If the buyer decides not to buy the commodity, the down- payment is forfeited and is treated as a gift from the buyer to the seller. The use of a down-payment sale (urbun) was declared permissible after long discussions in Shari'ah scholarly circles. At first glance, the terms and conditions of urbun seem similar to an option contract; nevertheless, there are serious differences between them.

Firstly, in urbun, the down payment serves as a kind of collateral or guarantee to the seller that the potential buyer has a serious intention to purchase the commodity. The down payment also serves as compensation to the seller in case the buyer withdraws from the purchase. Secondly, the down- payment upon purchase of the commodity is counted towards the total price of the commodity, whereas in an option contract, premium is paid by the client as consideration for the possibility/entitlement to purchase underlying for an agreed price on an agreed term in the future. Thirdly, in urbun, both parties bear the risk of the transaction.

The seller, as he may not sell the item and the buyer as he may forfeit his down-payment. In an option contract, the client bears always the risk of his unlucky bet; the financial institution is secured thanks to the premium and in fact never loses. The premium is paid at the moment the contract is entered into and the financial institution does not have to worry about the result of the client's forecast. One of the principles of Islamic banking is the profit and risk sharing principle. It implies that any transaction between two parties has to create a healthy balance in sharing the risks and the profits. A contract where the burden of risks lays mostly on one party is not permissible. Fourthly, the purpose of urbun is to secure future purchase of the commodity, the total price of which is known to the parties. If the contract is successful, there will be a transfer of commodity which is regarded as the core issue of any Islamic sale contract. The client does not bet what the future price of the commodity might be; he knows it at the time the contract is entered into.

Options have generally been examined under the fiqh doctrine of al-khiyarat (contractual stipulations) or under the bai-al-urbun concept. Urbun being a transaction in which a buyer places an initial good faith deposit.

4.4.3 Arboon

The arboon sale contract, in which a sale contract is affected with a condition of revocation by the purchaser, works exactly like a call option. The down payment can be retained by the seller thus working as a fee for the option. We highlight two features

- the option should be seen as a sale contract for the underlying asset and there should absolutely not be growth of derivatives beyond actual need of the real transaction;
- no arboon sale contract is to be entered into unless the seller actually owns the underlying asset and continues ownership for the whole duration of the option.

4.5 Fuqaha (jurists) Viewpoints on Conventional Derivative Instruments

4.5.1 Futures

Fatwa of Omam Al-Haramaini Al-Jauwaini

Futures Trading is Halal if the practice is based on Darurah and the Needs or Hajaat of the Ummah

Syariah Advisory Council (SAC) of Securities Commission

- Futures trading of commodities is approved as long as underlying asset is halal.
- Crude Palm Oil Futures Contracts are approved for trading.
- For Stock Index Futures contract, the concept is approved. However since the current KLCI SE based stock index future (SIF) has non halal stocks, it is not approved.

- Thus is implies that a SIF contract contract of a halal index would be acceptable.

Ustaz Ahmad Allam; Islamic Fiqh Academy

Stock index futures trading is HARAM, since some of the underlying stocks are not halal. Until and unless the underlying asset or basket of securities in the SIF is all Halal, SIF trading is not approved.

Mufti Taqi Usmani

Futures transactions not permissible due to

1. According to Syariah, sale or purchase cannot be affected for a future date.
2. In most futures transactions delivery or possession is not intended.

4.5.2 Options

When viewed solely as a promise to buy or sell an asset at a predetermined price within a stipulated period, Shariah scholars find nothing objectionable with options. It is in the trading of these promises and the charging of premiums when objections are raised. The gain in derivatives depends on chance; it is affected by Maisir and Qimar.

Ahmad Muhayyuddin Hassan

Objects to option trading for 2 reasons

1. Maturity beyond three days as in al-khiyarat is not acceptable.
2. The buyer gets more benefits than the seller - injustice.

Abu Sulayman (1992) (Fiqh Academy - Jeddah)

Acceptable when viewed in the light of bai-al-urbun but considers options to have been detached and independent of the underlying asset - thereforethey are unacceptable

Mufti Taqi Usmani (Fiqh Academy - Jeddah)

Promises as part of a contract is acceptable in Shariah, however the trading and charging of a premium for the promise is not acceptable. Yet others have argued against options by invoking “maisir” or unearned gains. That is, the profits from options may be unearned.

Hashim Kamali (1998)

Finds options acceptable because

- Invokes the Hanbali tradition
- Cites Hadiths of Barira (RA) and Habban Ibn Munqidh (RA).
- Also draws parallels with the al-urbun in arguing that premiums are acceptable.
- Also cites that contemporary scholars such as Yusuf al-Qaradawi and Mustafa al-Zarqa have authenticated al-urbun. (similar stand by Iranian scholars)

Shariah Advisory Council; Securities Commission

Though no formal opinion on stock or Index Options, the SAC has allowed other option-like instruments

- Warrants
- TSRs
- Call Warrants

Each of these are really option like instruments. Call Warrants for example, are simply long dated Call Options. Have similar risk/payoff profile.

4.5.3 Conclusion

The overall stance of Fuqaha, of conventional derivative instruments appears to be one of apprehension even suspicion. That these instruments could easily be used for speculation appears to be the key reason for objection. The fact that derivatives form the basis of risk-management appears to have been lost.

Key Problem: Evaluation has always been from a purely juridical viewpoint. And like most juristic evaluation, have relied on precedence? But there isn't a precedence nor equivalence for the kind of risk-management problems faced today.

When extrapolating/inferring: template may be wrong. The object of juridical analysis appears to be a micro examination of each and every feature of a derivative instrument to see if it passes, a often subjective religious filter.

The overall intended use of the instrument nor the societal benefits that could accrue do not seem to have been given due consideration. Aside from individual interpretation, the differing opinions among mazhabs/imams complicates the situation further. Thus, an options contract may be found objectionable for exactly opposite reasons.

While some mazhabs like the Hanbalis have been broader in their acceptance, the Shafi' and Hanafis have been less so. The Hanbalis for example are somewhat liberal when it comes to Option of Stipulation (Khiyar-al-Shart). The Hanbalis hold that stipulations that remove a hardship, fulfills a legitimate need, provide a benefit or convenience, or facilitate the smooth flow of commercial transactions are generally valid as a matter of principle.

4.6 Back to Basics

From the previous section it seems nearly impossible to trade derivatives under Islamic finance. However, we know that as almost anything is permitted in conventional banking, Islamic banking should be no different, merely a special case. All conventional banking products are built from four pillars: deposits, exchange, forwards and options. One should be able to built Islamic products using Islamic equivalents to these four pillars. To achieve this, without getting close to the edge, go back to first principles, to when derivatives were first created.

Back in the 1970s, US and UK companies made back-to-back loans to hedge foreign currency exposures — a forerunner of currency swaps. Now, instead of using back-to-back loans, Islamic products can be created using Islamic equivalents of back-to-back murabaha, back-to-back ijara, or back-to-back sukuk.

There is a major problem facing IFIs in terms of 'the lack of tools available for risk management and risk profile alteration'. The need to address asset liability management and the yield curve management in IFIs should be met and catered for.

The question that comes to mind is: Should the Islamic finance industry be scurrying around trying to replicate each and every complex derivative rather than focus on what is actually needed? Maybe the way ahead is not to talk about structuring Islamic derivatives with all of the connotations of gambling and uncertainty; let's focus on financial takaful, is a suggestion. Going back to basics and addressing the problems that derivatives were originally designed to address may be the way forward. Islamic banks would then have Shari'ah-compliant tools that would help solve the real asset/liability and risk management issues faced by Islamic institutions, both financial and commercial

4.7 Islamic Business

The evaluation of derivative instruments obvious needs a more coordinated approach; needs based rather than purely juristic/precedent driven.

But remember, Muslim businesses operate in the same environment and so face the same risks as any other. Yet, in the current state of affairs, Shariah compliance can impede risk management needs. Unless there is a convergence between Shariah compliance and risk management needs, Muslim business can be seriously handicapped.

Chapter 5

The Volatility Surface

5.1 Introduction

The exchange traded index option market in South Africa has seen tremendous growth during the last couple of years. The biggest liquidity is in options on the near and middle Alsi future contracts. Alsi futures are listed future contracts on the FTSE/JSE Top40 index, the most important and tradable equity index in South Africa. OTC and listed options trade on a skew and most market makers have implemented their own proprietary skew generators.

In this Chapter we show how to generate the implied volatility surface by fitting a quadratic deterministic function to implied volatility data from Alsi index options traded on Safex[1]. This market is mostly driven by structured spread trades, and very few at-the-money options ever trade. It is thus difficult to obtain the correct at-the-money volatilities needed by the exchange for their mark-to-market and risk management processes. We further investigate the term structure of at-the-money volatilities and show how the at-the-money implied volatilities can be obtained from the same deterministic model. This methodology leads to a no-spread arbitrage and robust market related volatility surface that can be used by option traders and brokers in pricing structured option trades.

5.2 Stochastic and Nonparametric Volatility Models

The idea that the price of a financial instrument might be arrived at using a complex mathematical formula is relatively new. This idea can be traced back to the seminal paper by Myron Scholes and Fischer Black [BS 73]. We now live in a world where it

[1]The South African derivatives exchange based in Johannesburg — http://www.safex.co.za/ed/

is accepted that the value of certain illiquid derivative securities can be arrived at on the basis of a model (this is the practise of marking-to-model) [Re 06].

In order to implement these models, practitioners paid more and more attention to, and began to collect, direct empirical market data often at a transactional level. The availability of this data created new opportunities. The reasonableness of a model's assumptions could be assessed and the data guided many practitioners in the development of new models. Such market data led researchers to models whereby the dynamics of volatility could be studied and modeled.

Stochastic, empirical, nonparametric and deterministic models have been studied extensively. In this section we will give a brief overview of the first 3 models before we delf into the deterministic model in greater detail in §5.3.

5.2.1 Stochastic Models

Black & Scholes defined volatility as the standard deviation because it measures the variability in the returns of the underlying asset. They determined the historical volatility and used that as a proxy for the expected or implied volatility in the future. Since then the study of implied volatility has become a central preoccupation for both academics and practitioners [Ga 06]. Volatility changes over time and seem to be driven by a stochastic process. This can be seen in Figure 5.1 where we plot the 3 month rolling volatility for the FTSE/JSE Top40 index. Also evident from this plot is the mean reversion phenomenon. The 3 month long term mean volatility of this index is 21.38%.

The stochastic nature of volatility led researchers to model the volatility surface in a stochastic framework. These models are useful because they explain in a self-consistent way, why options with different strikes and expirations have different *Black & Scholes* implied volatilities. Another feature is that they assume realistic dynamics for the underlying [Le 00]. The stochastic volatility models that we briefly appraise are the popular Heston model [He 93] and the well-known SABR model by Hagan et al. [HK 02].

In the Heston model, volatility is modeled as a long-term mean reverting process. The Heston model fits the long-term skews well, but fails at the shorter expirations [Ga 09]. Heston models the volatility surface as a joint dynamic in time and strike space. A viable alternative to the Heston model is the well known SABR model [We 05]. Here, volatility is modelled as a short term process by assuming that the underlying is some normally distributed variable. The SABR model assumes that strike and time to expiration dynamics are disjoint, i.e. the skews and term structure of the skews are calibrated separately. This model works better for shorter expirations but because volatilities do not mean revert in the SABR model, it is only good for short expirations. Another problem with the SABR model[2] is that its parameters are

[2]The SABR stochastic volatility model has the at-the-money volatility as input, which means

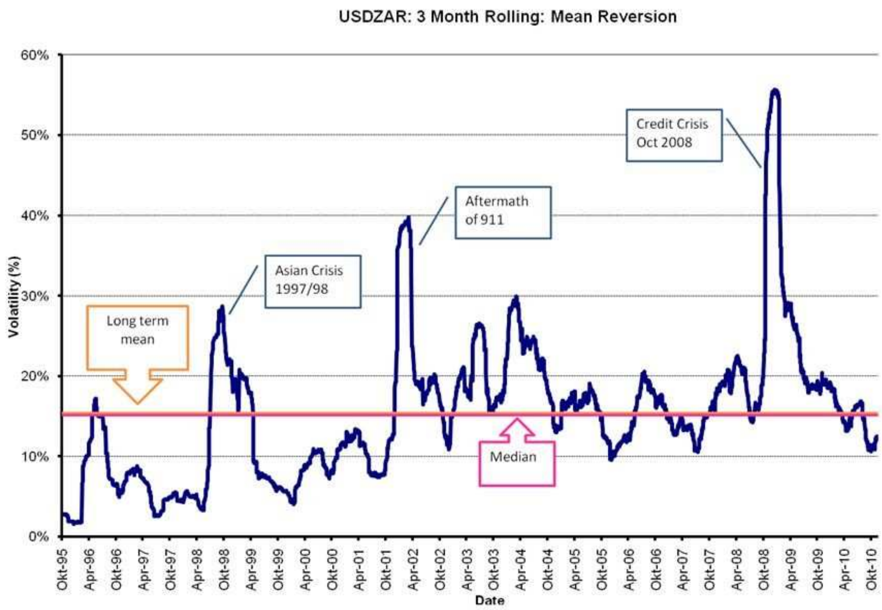

Figure 5.1: FTSE/JSE Top40 3 month historical volatility during the period, June 1995 through October 2009. The plot shows that volatility is not constant and seems to be stochastic in nature. Also evident is the phenomenon of mean reversion.

time-homogenous. This means that the model implies that future volatility surfaces will look like today's surface. West has shown how to calibrate the SABR model using South African index option data [We 05]. Bosman et al. also showed how to obtain a representative South African volatility surface by implementing the SABR model using Alsi option data [BJM 08]. Issues with the calibration of the SABR model has been tackled by Tourrucô [To 08] while the original SABR formula has been rectified by Oblój [Ob 08].

A general problem associated with stochastic volatility models is that they fail to model the dynamics of the short term volatility skews. This arises from the fact that the at-the-money volatility term structure can be so intricate in the short-end[3] that these models just fail at accurately modelling the short end volatility dynamics. In fact, Jim Gatheral states [Ga 06]

> *"So, sometimes it's possible to fit the term structure of the at-the-money volatility with a stochastic volatility model, but it's never possible to fit the term structure of the volatility skew for short expirations... a stochastic volatility model with time homogenous parameters cannot fit market*

that in an illiquid options market, the SABR generated market surface accuracy will have additional errors arising from the illiquidity of the at-the-money options trade data.

[3]Intricate in the sense that the short term volatility skews can include volatility jumps that are correlated with the equity index level, that are usually not assumed when stochastic volatility models are derived to model the equity index volatility surface.

prices!"

The possibility of using an extended stochastic volatility model with correlated jumps in the index level and volatility might fit the short-term market volatility skews better, but in practice it is difficult (if not impossible) to calibrate such a complex model [Ga 09].

The above-mentioned problems associated with stochastic volatility models led us to look beyond these models.

5.2.2 Empirical Approaches

The empirical approach known as the Vanna-Volga method has been studied extensively. This approach was introduced by Lipton and McGhee [LM 02]. The Vanna-Volga method is also known as the traders' rule of thumb. It is an empirical procedure that can be used to infer an implied-volatility smile from three available quotes for a given maturity; it is thus useful in illiquid markets. It is based on the construction of locally replicating portfolios whose associated hedging costs are added to corresponding Black-Scholes prices to produce smile-consistent values. Besides being intuitive and easy to implement, this procedure has a clear financial interpretation, which further supports its use in practice [Wy 08].

The Vanna-Volga approach considers an option price as a Black-Scholes price corrected by hedging costs caused by stochasticity in price-forming factors (volatility, interest rate, etc.) observed from real markets. While accounting for stochasticity in volatility, it differs from stochastic volatility frameworks (Heston, SABR, and modern Levy process-based generalizations) in the following way: rather than constructing a parallel (possibly correlated) process for the instantaneous volatility and defining the price as the risk-neutral expectation, the Vanna-Volga approach puts much more weight on the self-financing argument, considering an option price as a value of the replicating Black-Scholes portfolio plus additional corrections offsetting stochasticity in volatility [Sh 08].

This approach is very popular in the foreign exchange market but has been applied to equities as well.

5.2.3 Nonparametric Estimation of the Skew

In illiquid markets it might not be possible to calibrate models (stochastic or deterministic) due to a lack of data. In such circumstances nonparametric option pricing techniques might be feasible.

Nonparametric option pricing techniques utilise spot market observed security prices in order to determine the probability distribution of the underlying asset. Derman and Kani [DK 94] and Rubinstein [Ru 94] utilise implied binomial tree (IBT) approaches to find risk-neutral distributions which result in estimated option prices that match observed option prices. These nonparametric approaches impose no assumptions about the nature of the probability distribution of the asset. This leaves little room for misestimation of option pricing resulting from an incorrect choice of the probability distribution of the underlying asset. Thus nonparametric methods do not suffer from the smile bias that is evidenced in parametric models. Using nonparametric methods to estimate the price of an option and reversing this price through the Black-Scholes formula will yield the conventional implied volatility. Following this procedure for a range of possible strikes will result in the so-called nonparametric volatility skew. The nonparametric volatility skew is free of any model-specific assumptions and, as it is based only on observed asset data, it is a minimally subjective estimate of the volatility skew [AM 06].

An IBT is a generalization of the Cox, Ross, and Rubinstein binomial tree (CRR) for option pricing [Hu 06]. IBT techniques, like the CRR technique, build a binomial tree to describe the evolution of the values of an underlying asset. An IBT differs from CRR because the probabilities attached to outcomes in the tree are inferred from a collection of actual option prices, rather than simply deduced from the behavior of the underlying asset. These option implied risk-neutral probabilities (or alternatively, the closely related risk-neutral state-contingent claim prices) are then available to be used to price other options. Arnold, Crack and Schwartz showed how an IBT can be implemented in Excel [ACS 04].

De Araújo and Maré compared nonparametrically derived volatility skews to market observed implied volatility skews in the South African market [AM 06]. Their method is based only on the price data of the underlying asset and does not require observed option price data to be calibrated. They characterised the risk-neutral distribution and density function directly from the FTSE/JSE Top40 index data. Using Monte Carlo simulation they generated option prices and found the implied volatility surface by inverting the Black-Scholes equation. This surface was found to closely resemble the Safex traded implied volatility surface.

5.3 The Deterministic Volatility Approach

Implementing stochastic volatility models and implied binomial trees can be very difficult. With stochastic volatility, option valuation generally requires a market price of risk parameter, which is difficult to estimate [DFW 98]. However, if the volatility is a deterministic function of the asset price and/or time, the estimation becomes a lot simpler. In this case it remains possible to value options based on the Black-Scholes partial differential equation although not by means of the Black-Scholes formula itself.

5.3.1 Deterministic Models

Deterministic volatility functions (DVF) are volatility models requiring no assumptions about the dynamics of the underlying index that generated the volatility. These models can only be implemented if one assumes that the observed prices of options reflect in an informationally efficient way everything that can be known about the true process driving volatility — the efficient market hypothesis [Re 06].

Deterministic volatility functions were introduced in the 1990s. During 1993, Shimko studied the risk neutral densities of option prices and chose a simple parabolic function as a possible parametric specification for the implied volatility [Sh 93]. Dumas et al. studied the applicability of such simple functions. They started by rewriting the general Black-Scholes differential equation as a forward partial differential equation (applicable to forward or futures contracts)

$$\frac{1}{2}\sigma^2\left(F,T\right)K^2\frac{\partial^2 C}{\partial K^2}=\frac{\partial C}{\partial t}$$

with the associated initial condition, $C(K,0)\equiv\max(S-K,0)$. Also, S is the spot price, F the forward or futures price, K the absolute strike price, T is the time to expiration and $C(K,T)$ is the call price. The advantage of the forward equation approach is that all option series with a common time to expiration can be valued simultaneously. They further mentioned that $\sigma(K,T)$ is an arbitrary function. Due to this they posit a number of different structural forms (deterministic functions) for the implied volatility as a function of the strike and time to expiry [DFW 98]

$$\begin{aligned}
\text{Model 0}:\ \sigma &= a_0\\
\text{Model 1}:\ \sigma(K) &= a_0+a_1K+a_2K^2\\
\text{Model 2}:\ \sigma(K,T) &= a_0+a_1K+a_2K^2+a_3T+a_5KT\\
\text{Model 3}:\ \sigma(K,T) &= a_0+a_1K+a_2K^2+a_3T+a_4T^2+a_5KT
\end{aligned}\tag{5.1}$$

Here, the variables a_0 through a_5 are determined by fitting the functional forms to traded option data. Model 0 is the Black- Scholes model with a constant volatility where Model 1 attempts to capture variation with the asset price. Models 2 and 3 capture additional variation with respect to time. They estimated the volatility function $\sigma(K,T)$ by fitting the functional forms in Eq. (5.1) to observed option prices at time t (today). They used S&P 500 index option data captured between June 1988 and December 1993. Estimation was done once a week. The parameters and thus skews were estimated by minimising the sum of squared errors of the observed option prices from the options' theoretical deterministic option values given by the functional forms[4]. Models 1 through 3 all reflected volatility skews or smirks. They found that the Root Mean Squared Valuation Error of Model 1 is half of that of the Black-Scholes constant volatility model. Another important result from their study is that

[4]They used an algorithm based on the downhill simplex method of Nelder and Mead.

the functional forms loose their predictive ability rather quickly. This means that these functional forms need to be refitted to traded data on a regular basis.

Beber followed the Dumas et al. study but used options written on the MIB30 — the Italian stock market index [Be 01]. He optimised two models

$$\begin{aligned} \text{Model 1}: \ \sigma &= \beta_0 + \beta_1 K + \epsilon \\ \text{Model 2}: \ \sigma &= \beta_0 + \beta_1 K + \beta_2 K^2 + \epsilon \end{aligned} \tag{5.2}$$

One difference between his and the Dumas study is that he defined K as the moneyness and not absolute strike. Model 1 is linear in the moneyness and Model 2 quadratic i.e. iit is a parabola. ϵ is just an error estimate. The simplicity of the two models is determined by the endeavour to avoid overparametrization in order to gain better estimates' stability over time. Beber also decided to give the same weight to each observation, regardless the moneyness, as the strategy to assign less weight to the deep out of the money options owing to the higher volatility has not proven to be satisfactory. He fitted the traded option data by ordinary least squares. Beber found that the average implied volatility function is fitted rather well by a quadratic model with a negative coefficient of asymmetry. Hence the average risk neutral probability density function on the Italian stock market is fat tailed and negatively skewed. The interpretation of the parameters in general are as follows

- β_0 represents a general level of volatility which localizes the implied volatility function (it is also the constant of regression),
- β_1 characterizes the negative profile which is responsible for the asymmetry in the risk neutral probability density function; it is the coefficient that controls the displacement of the origin of the parabola with respect to the ATM options, and
- β_2 provides a certain degree of curvature in the implied volatility function or it controls the wideness of the smile.

One of the most comprehensive studies was done by Tompkins in 2001 [To 01]. He looked at 16 different options markets on financial futures comprising four asset classes: equities, foreign exchange, bonds and forward rate agreements. He compared the relative smile patterns or shapes across markets for options with the same time to expiration. He also used a data set comprising more than 10 years of option prices spanning 1986 to 1996. The individual equities examined were: S&P500 futures, FTSE futures, Nikkei Dow futures and DAX futures. He fitted a quadratic volatility function to the data and found his graphs of the implied volatility to be similar to that shown by Shimko in 1993 [Sh 93] and Dumas et al [DFW 98]. Tompkins then states:

"*If the sole objective was to fit a curved line, this has been achieved*".

He concluded that regularities in implied volatility surfaces exist and are similar for the same asset classes even for different exchanges. A further result is that the shapes of the implied volatility surfaces are fairly stable over time.

Many studies followed the Dumas and Tompkins papers, most using different data sets. All of these studied the models listed in Eqs. (5.1) and (5.2). Sehgal and Vijayakumar studied S&P CNX Nifty index option data [SV 08]. These options trade on the derivatives segment of the National Stock Exchange of India[5]. Badshah used out-the-money options on the FTSE 100 index and found the quadratic model to be a good fit [Ba 09]. Zhang and Xiang also studied S&P500 index options and found the quadratic function fits the market implied volatility smirk very well. They used the trade data on 4 November 2003 for SPX options expiring on 21 November 2003 - thus very short dated options [ZX 05]. They used all out-the-money puts and calls and fitted the quadratic function by minimising the volume weighted mean squared error and found the quadratic function to work very well.

Another comprehensive study was done by Panayiotis et al. [PCS 08]. They tackled the deterministic methodology from a different angle. They considered 52 different functional forms to identify the best DVF estimation approach for modelling the implied volatility in order to price S&P500 index options. They started with functions as given by Dumas et al. and listed in Eq. (5.1) where K is the absolute strike. Next they changed the strike to $\ln K$, S/K (moneyness) and lastly to $\ln(S/K)$. All in all they considered 16 functions similar to Eqs. (5.1). They also considered a number of asymmetric DVF specifications. Their dataset covered the period January 1998 to August 2004 — 1675 trading days. They recalibrated all 52 functions on a daily basis. Their main result is that the deterministic specification with strike used as moneyness (S/K) works best in-sample while the model with strike used as $\ln K$ works best out-of-sample.

Modeling the volatility skew as a deterministic process has other benefits too [Bu 01]

- They allow one to model volatility separately in expiry time and strike. This means that each expiry's skew can be independently calibrated minimising compounding errors across expiries. This property is useful in modeling volatility surfaces in illiquid markets where data is sparse across strikes.

- Pricing using deterministic volatility preserves spread arbitrage market conditions because no assumptions about the underlying process is made.

- The whole surface can be calibrated with minimal model error.

[5] The S&P CNX Nifty is an index comprising the 50 largest and most liquid companies in India with about 60% of the total market capitalisation of the Indian stock market.

5.3.2 Principle Component Analysis

Using a 3 parameter quadratic function was further motivated by the research findings due to Carol Alexander [Al 01]. She did a principle component analysis (PCA) on FTSE 100 index options and found that 90% of the dynamics of the volatility skew are driven by three factors

- parallel shifts (trends),
- tilts (slopes), and
- curvature (convexity).

Badshah also did a PCA on FTSE 100 options. He used the implied volatility surfaces for the March and October months for the years 2004, 2005, 2006 and 2007. His results are in line with that of Alexander although he found that on average 79% of the dynamics of the skews are driven by the first 3 components — this is similar to the findings by Alexander [Ba 09]. Le Roux studied options on the S&P500 index with strikes ranging from 50% to 150% [Le 07]. He found that 75.2% of the variation of the implied volatility surface can be described by the first principle component and another 15.6% by the second! His first component reflects the slope or tilt of the skew. The difference between his study and that of Alexander's is that he used moneyness instead of absolute strikes.

Bonney, Shannon and Uys followed Alexander's methodology and studied the principle components of the JSE/FTSE Top 40 index [BSU 08]. They found that the trend affect explains 42% of the variability in the skew changes, the slope 19% and the convexity an additional 14%. Their results show that the first three components explain 76.24% of the variability in skew changes. They attribute the difference between Alexander's 90% and their 76% to differences between a liquid market and less liquid emerging market.

5.3.3 The SVI Model

Another interesting research finding was presented by Jim Gatheral [Ga 04]. He derived the Stochastic Volatility Inspired (SVI) model. This is a 5 parameter quadratic model (in moneyness) based on the fact that many conventional parameterisations of the volatility surface are quadratic as discussed in §5.3.1. This parametrisation has a number of appealing properties, one of which is that it is relatively easy to eliminate calendar spread arbitrage. This model is "inspired" by the stochastic volatility models due to the fact that implied variance is linear in moneyness as $K \to \pm\infty$ for stochastic volatility models. Any parametrisation of the implied variance surface that is consistent with stochastic volatility, needs to be linear in the wings and curved in the middle. The SVI and quadratic models exhibit such properties. Gatheral also states that:

> *"if the wings are linear in strike (moneyness), we need 5 and only 5 parameters to cover all reasonable transformations of the volatility smile."*

This model was extensively tested using S&P500 (SPX) index option data with excellent results.

5.3.4 The Quadratic Function for the ALSI Implied Volatility Surface

The results from all the studies on deterministic volatility functions and the PCA studies mentioned above form the basis in modelling the South African ALSI index volatility surface. In scrutinising Eqs. (5.1) and (5.2) and taking illiquidity into account, we postulate that the following three parameter quadratic function should be a good model of fit for the ALSI implied volatility data (following the Beber notation in Eq. (5.2)

$$\sigma_{model}(\beta_0,\, \beta_1,\, \beta_2) = \beta_0 + \beta_1\, K + \beta_2\, K^2. \tag{5.3}$$

In this equation we have

- K is the strike price in moneyness format (Spot/Strike),
- β_0 is the constant volatility (shift or trend) parameter, $\beta_0 > 0$. Note that
$$\sigma \underset{K\to 0}{\longrightarrow} \beta_0,$$
- β_1 is the correlation (slope) term. This parameter accounts for the negative correlation between the underlying index and volatility. The no-spread-arbitrage condition requires that $-1 < \beta_1 < 0$ and,
- β_2 is the volatility of volatility ('vol of vol' or curvature/convexity) parameter. The no-calendar-spread arbitrage convexity condition requires that $\beta_2 > 0$.

Note that Eq. (5.3) is also linear in the wings as $K \to \pm\infty$. In Fig. 5.3.4 we plot the volatility skew for the near Alsi as obtained by implementing Eq. (5.3).

5.3.5 Volatility Term Structure

The functional form for the skew in Eq. (5.3) is given in terms of moneyness or in floating format (sticky delta format). This DVF does also not depend on time. The optimisation is done separately for each expiry date. In order to generate a whole implied volatility surface we also need a specification or functional form for the *at-the-money* (ATM) volatility term structure. It is, however, important to remember that the ATM volatility is intricately part of the skew. This means that the two

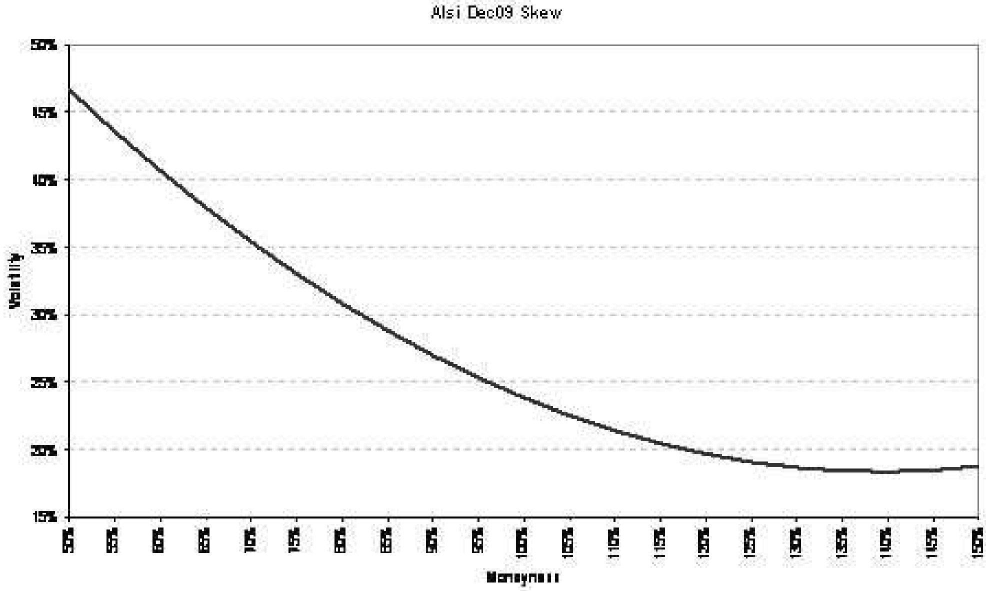

Figure 5.2: The ALSI volatility skew for the December futures contract at the beginning of October 2009.

optimisations (one for the skews and the other for the ATM volatilities) can not be done strickly separate from one another. Taking the ATM term structure together with each skew will give us the 3D implied volatility surface.

It is well-known that volatility is mean reverting; when volatility is high (low) the volatility term structure is downward (upward) sloping [Al 01, Ga 04]. This was shown for the JSE/FTSE Top 40 index in Fig. 5.1. We therefore postulate the following functional form for the ATM volatility term structure

$$\sigma_{atm}(\tau) = \frac{\theta}{\tau^{\lambda}}. \tag{5.4}$$

Here we have

- τ is the months to expiry,
- λ controls the overall slope of the ATM term structure; $\lambda > 0$ implies a downward sloping ATM volatility term structure (this is plotted in Fig. 5.3.5), whilst a $\lambda < 0$ implies an upward sloping ATM volatility term structure, and
- θ controls the short term ATM curvature.

Please note that τ is not the annual time to expiry. It actually is the "months to expiry". It is calculated by

$$\frac{date_{expiry} - date_0}{365} * 12.$$

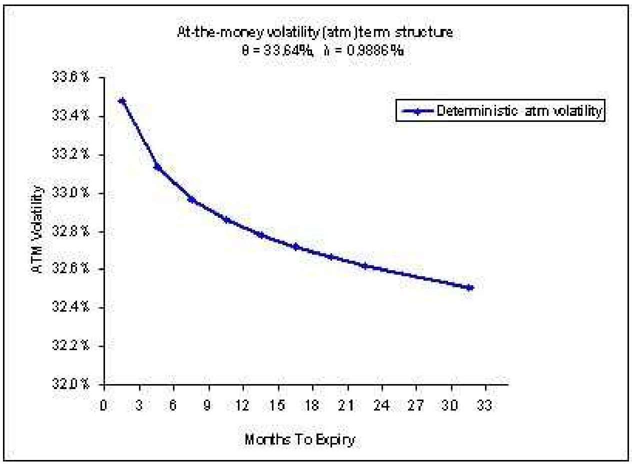

Figure 5.3: The market fitted at-the-money volatility term structure for the ALSI at the beginning of April 2009.

The reason for this notation is that if $\tau = 1$, the ATM volatility is given by θ; the 1 month volatility is thus just θ. This makes a comparison between the South African 1 month volatility, and the 1 month volatilities offshore, like the VIX, easier.

A better understanding of these parameters is obtained if we consider the deterministic term of the Heston stochastic differential equation [He 93]

$$d\sigma(\tau) = \frac{\lambda}{\tau}(\omega - \sigma_I(\tau))\,dt \tag{5.5}$$

where ω is a long term mean volatility and λ/τ is the mean reversion speed.

The solution to the ordinary differential equation in (5.5) is given by

$$\sigma(\tau) = \omega + \frac{\sigma_0 - \omega}{\tau^\lambda} \tag{5.6}$$

Comparing Eqs. (5.4) and (5.6) let us deduce that

- θ is a term that represents the difference between the current at-the-money volatility σ_0 and the long term at-the-money volatility ω, and
- λ is a parameter defined such that λ/τ is the mean reversion speed useful for ATM calendar spreads.

In using the volatility skew function given in Eq. (5.3) and the volatility term structure function shown in Eq. (5.4), we can generate the market equity index volatility surface in the deterministic framework. We show the Alsi surface as determined on 8 March 2011 in Fig. 5.3.5.

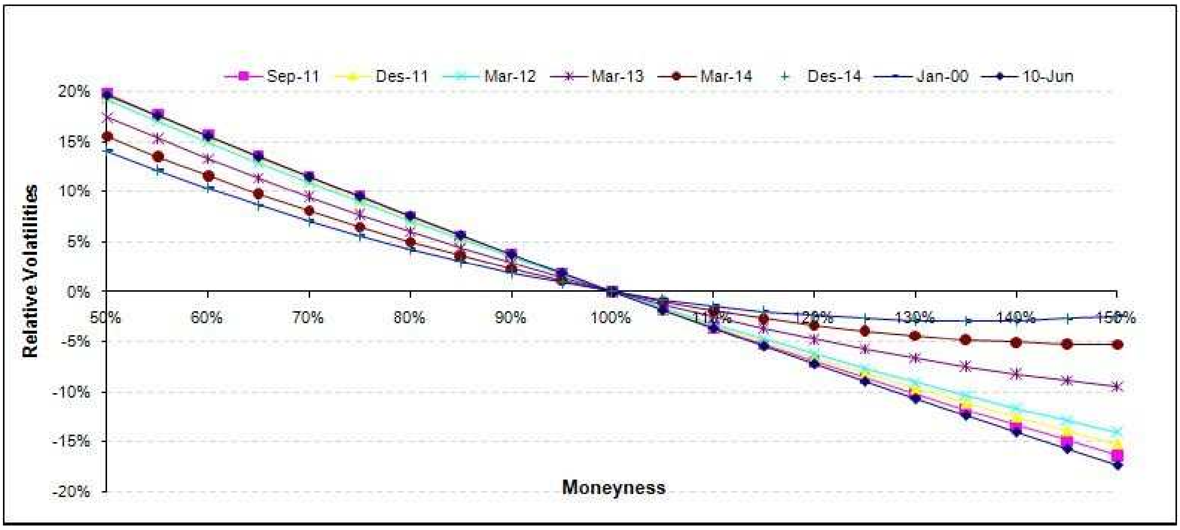

Figure 5.4: The ALSI volatility surface on 8 March 2011.

5.3.6 FX Delta and Strike Relationship

There are four types of Deltas in the FX markets

- Spot Delta
- Forward Delta
- Spot Delta Premium-Adjusted
- Forward Delta Premium-Adjusted

We will denote the deltas by Δ_S, Δ_f, Δ_{Spa} and Δ_{fpa}.

The problem arises in the calculation of the strike K given the Delta Δ and volatility σ. This is a straightforward procedure for the first two delta types, as there are closed form solutions available

$$\begin{aligned} K &= fe^{-\phi N^{-1}(\phi\Delta_f)\sigma\sqrt{\tau}+\frac{1}{2}\sigma^2\tau} \\ K &= fe^{-\phi N^{-1}\left(\phi e^{r_f\tau}\Delta_S\right)\sigma\sqrt{\tau}+\frac{1}{2}\sigma^2\tau} \end{aligned} \tag{5.7}$$

with N^{-1} being the inverse normal cumulative density function. However, we also have

$$\Delta_{Spa} = \phi e^{r_f\tau}\frac{K}{f}N(\phi x). \tag{5.8}$$

This can only be solved numerically.

5.4 Vanna Volga and Implied Skews

The Vanna-Volga (VV) method is an empirical procedure that can be used to infer an implied-volatility smile from three available quotes for a given maturity. It is based on the construction of locally replicating portfolios whose associated hedging costs are added to corresponding *Black & Scholes* prices to produce smile-consistent values. Besides being intuitive and easy to implement, this procedure has a clear financial interpretation, which further supports its use in practice.

The VV method is commonly used in foreign-exchange options markets, where three main volatility quotes are typically available for a given market maturity

- the delta-neutral straddle, referred to as at-the-money (ATM). The premium of a straddle yields information about the expected volatility of the underlying - higher volatility means higher profit, and as a result a higher premium.
- the risk reversal for 25Δ call and put. The RR is a long out-of-the-money call and a short out-of-the-money put with a symmetric skew.
- the (Vega-weighted) butterfly with 25Δ wings. A buyer of a Vega-weighted butterfly profits under a stable underlying.

The application of VV allows us to derive implied volatilities for any option's delta, in particular for those outside the basic range set by the 25Δ put and call quotes.

A straddle together with a strangle give simple techniques to trade volatility. The AtM volatility σ_{AtM} is then derived as the volatility of the delta-neutral straddle and the volatilities of the risk-reversal (RR) and Vega-weighted butterfly (VWB) where we have the following relationships

$$\begin{aligned} \sigma_{RR} &= \sigma_{25\Delta Call} - \sigma_{25\Delta Put} \\ \sigma_{VWB} &= \frac{1}{2}(\sigma_{25\Delta Call} + \sigma_{25\Delta Put}) - \sigma_{AtM} \end{aligned} \tag{5.9}$$

Implied volatility of a risk-reversal incorporates information on the skew of the implied volatility curve, whereas that of a strangle on the kurtosis.

With the volatilities received in that way, Vanna-Volga allows us to reconstruct the whole smile for a given maturity from three given market quotes only. We assume a market where, for a given maturity T, three basic options are quoted: without loss of generality we assume that the options are all call. We denote the corresponding strike by $K_i,\ i = 1, 2, 3$ and $K_1 < K_2 < K_3$. The market implied volatility associated with K_i is denoted by $\sigma_i,\ i = 1, 2, 3$. This setting is consistent with the FX market environment, where three main strikes are dealt.

In most practical applications we set $\sigma_2 = \sigma_{AtM}$ and this $K_2 = K_{AtM}$ although it is not strictly required. *Castagna* showed that the implied volatility $\sigma(K)$ can be approximated by [Ca 10]

$$\sigma(K) = \sigma_2 + \frac{-\sigma_2 + \sqrt{\sigma_2^2 + d_1(K)d_2(K)(2\sigma_2 D_1(K) + D_2(K))}}{d_1(K)d_2(K)} \tag{5.10}$$

where we have

$$\begin{aligned}
D_1(K) &= \eta_1(K) - \sigma_2 \\
D_2(K) &= \frac{\ln\frac{K_2}{K}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_1}} d_1(K_1)d_2(K_1)(\sigma_1-\sigma_2)^2 \\
&\quad + \frac{\ln\frac{K}{K_1}\ln\frac{K}{K_2}}{\ln\frac{K_3}{K_1}\ln\frac{K_3}{K_2}} d_1(K_3)d_2(K_3)(\sigma_3-\sigma_2)^2 \\
d_1(K) &= \frac{1}{\sigma\sqrt{T}}\left[\ln\frac{S_0}{K} + \left(r_d - r_f + \frac{1}{2}\sigma^2\right)T\right] \\
d_2(K) &= d_1(K) - \sigma\sqrt{T}
\end{aligned}$$

and further

$$\eta_1(K) = \frac{\ln\frac{K_2}{K}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_1}}\sigma_1 + \frac{\ln\frac{K}{K_1}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_2}}\sigma_2 + \frac{\ln\frac{K}{K_1}\ln\frac{K}{K_2}}{\ln\frac{K_3}{K_1}\ln\frac{K_3}{K_2}}\sigma_3.$$

5.4.1 Vanna-Volga: From Theory to Market Practice

See article by *Bossens*, *Rayeé*, *Skantzos* and *Deelstra*

Chapter 6

Vanilla Currency Exotic Options

What is an “exotic option?” It is defined as being “anything but vanilla.” By this we mean any option that has any associated feature that differs from a standard European option [We 06].

A vanilla European call option has essentially two features: firstly, you have the right to decide whether you want to exercise at expiry, Secondly, if you exercise you have the right to purchase a specific amount of an underlying asset at the agreed upon strike price. Exotic options can be quite different. While the payoff from a vanilla option is a linear function of the underlying, the payoff from an exotic option can be nonlinear like a digital that has a constant payoff. Note, due to the more complex of exotic options, they are mostly European in nature.

6.1 Introduction

We now know people use options primarily ”to hedge” and to “implement a view on the market.” Why would someone use an exotic option? In short, to either get more tailored financial insurance or to position a more precise view on the market. This means there are more than two features linked to exotic options.

We know the payoff from a standard vanilla option is

$$V = \phi max[0, S_T - K]. \tag{6.1}$$

This means the only thing that matters, is where the underlying is trading at expiry. A vanilla option is said to “not be path-dependent” on the underlying. The downside of vanilla call options, for instance, is that they can be trading in-the-money most of the time, but then a day or so before expiry there is a shock in the market and the underlying falls in value and the option expire worthless. The question is, can such a situation be overcome?

Let’s say we change the above payoff to be the following

$$V = \phi \max[0, S_{average} - K] \tag{6.2}$$

where $S_{average}$ is the average of the underlying FX rate. We can even specify how the average is calculated. It might be over the life of the option or it might be the average over the last ten days before expiry. This means that if there is sudden drop in the underlying the day before expiry, this call option might still pay out because the average is still above the strike price. Such an option is called an "Asian option" or "average rate option". This option is *path-dependent* because the payout is dependent on what happens with the underlying over the life of the option.

"Timing the markets" has been a topic of investigation for many decades. Every investor dreams of the perfect system that will tell him to buy right at the bottom or sell right at the top. The same holds for investors who want to hedge or gear their portfolio by using options. If an investor buys a call and the market turns down he always thinks that he should have waited to get the call at a lower strike. The same holds if the investor bought a put as a hedge. If the market rallies and then retraces he could have hedged at a better level if he waited. Now let's define another option to have a payoff

$$C = \max[0, S_{max} - K] \tag{6.3}$$

where we define S_{max} to be the maximum level that the underlying had over the life of the option. This is called a "lookback" option and is sometimes described as the "lazy" fund manager's instrument because you will always get the best payout without once trading in the market.

We will now describe some exotic options that have become standard across the world. These options have been trading for such a long time already that people started to call them "vanilla exotics" or "first generation exotics."

6.2 Digital Options

Digital options (synonyms: binary options, all-or-nothing options, cash-or-nothing options, asset-or-nothing options) pays a set amount if the underlying is above or below the strike price, or nothing at all — hence, the names for these options, because they pay 1 of 2 values[1]. A digital call pays off if the underlying asset is above a certain value at expiration - the strike price; a digital put pays off if the underlying asset is below the strike price at expiration. We show the discontinuous payoff functions in Fig. 6.2.

6.2.1 Where Binary Options are Used

Binaries are typically bought and sold in the Over the Counter (OTC) markets between sophisticated financial institutions, hedge funds, corporate treasuries, and large

[1] In computer or mathematics jargon, a binary number is one which is given a value of either 0 or 1 and nothing else

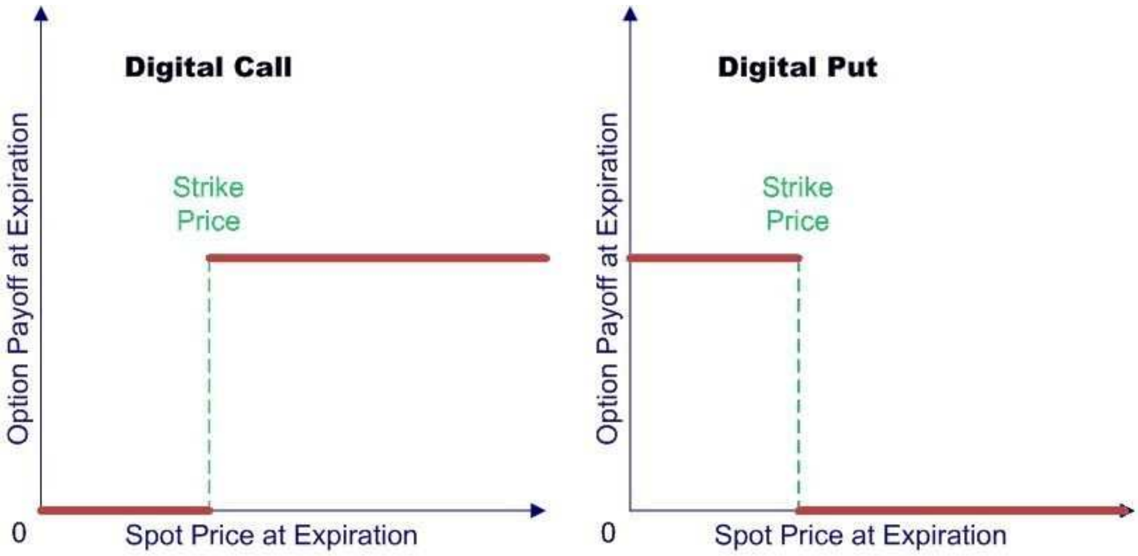

Figure 6.1: Payoff for a digital call and put option.

trading partners. They are widely used where the underlying instrument is a commodity, currency, rate, event, or index. For example:

Binary call and put options are popular in the platinum market, struck on the mid-market price of the metal of a certain quality, quoted by several dealers over a stated time period. Platinum trades in large varying quantities among major producers and manufacturers, as well as between speculators and dealers. Prices are determined between disparate parties, with varying frequency, and are not centrally reported or confined to a centralized exchange. A third party calculation agent is often agreed upon as part of the deal, to guarantee an uninterested price estimate obtained by sampling various dealers on the expiration date.

Binary options are used widely to hedge weather events, such as hurricanes, temperature, rainfall, etc. Major agricultural and transportation companies can be severely affected by adverse weather conditions. Weather is highly unpredictable and difficult to measure (e.g. what is a hurricane? How fast do the winds have to be? How long does it need to last? Does it need to touch ground or can it remain over water? What must the temperature be? Where is the exact location of the measurement to take place?). This makes a binary option a perfect tool for hedging weather events, as it allows the option seller (option writer) to assume a fixed amount of risk tied to the occurrence of a future event whose magnitude is impossible to predict. An uninvolved and highly reliable third party such as a government weather bureau is typically used to determine whether the weather event has occurred.

Binary options are also traded on inflation figures such as the Consumer Price Index (CPI) or Producer Price Index (PPI) in the U.S. These figures are reported fairly infrequently based on independent sampling methods, and are often revised after they are released once input values are further verified. There is no continual stream of prices because inflation is not an actual traded instrument (aside from recent developments in exchange-traded inflation futures). Without continual input prices, it is very difficult to mark-to-market vanilla American or European options, whose value is highly dependent on dense volatility and price data. A binary option allows the buyer to obtain inflation protection, while providing the option seller with limited risk in the event that inflation jumps or drops unexpectedly.

Finally, binary options are popular in the foreign currency markets, especially on illiquid and volatile currencies such as the Turkish Lira and Thai Bhat. Emerging market currencies are often subject to rapid "jump risk" caused by political or economic instability, or simply illiquidity due to the relatively small volume of foreign trade. Sophisticated currency speculators borrow low-rate developed economy currencies such as USD or EUR and invest in high-rate emerging market currencies, then purchase binary options as protection against currency risk in the high rate leg. This allows the speculator to earn "carry" while protecting against "jump risk."

6.2.2 Pricing Cash-or-Nothing

In the FX market, the cash can either be domestic or foreign. If we denote by S the FOR/DOM exchange rate (i.e. 1 unit of foreign currency is worth S units of domestic currency e.g., \$1 = KSh 82) we the following two types of cash-or-nothing options:

In case of a digital paying out R units of the domestic currency we have in a *Black & Scholes* world

$$V_{con} = Re^{-r_d T} N(\phi y) \tag{6.4}$$

where we use the same notation as in the *Garman-Kolhagen* formula given in Eq. 1.3. While in case of a digital paying out R units of the foreign currency we have,

$$V_{con} = RSe^{-r_f T} N(\phi x). \tag{6.5}$$

Note that the $R \times S$ term must reflect the way the FX rate is quoted — is it USDKES or KESUSD?

In general, an out-the-money binary option will be cheaper to purchase than an equivalent out-the-money vanilla option, assuming the same underlying, strike, and time to expiration. This is because the binary option has a fixed payout in the event it expires in the money. The vanilla option, on the other hand, can theoretically pay an infinite amount, limited only by the potential underlying price and credit of the option seller. Out-the-money vanilla options typically have greater "time value" than binary options.

This valuation difference between an out-the-money binary and vanilla option has two benefits. First, it enables the option seller to assume a known limited risk.

Second, from the perspective of the buyer, a binary option can offer significantly greater leverage since the up front premium investment is lower.

When a binary option moves from being out-the-money to in-the-money, its theoretical value profile is much different than a vanilla option. The binary option moves up in value very rapidly as it crosses the strike threshold. Conversely, when a binary option moves from in-the-money to out-the-money, its value changes very quickly, dropping towards zero in a steep fall, then leveling off.

Let's look at an example: On a slow day in the market, a trader is looking for trading opportunities and considers the "Spot EURUSD ¿ 1.2500 (3:00 p.m.) Binary."

This contract allows a trader to take a position on where the spot Euro/US$ (€/$) will be at 3 p.m. ET. A trader who believes that the rate will be above 1.2500 at 3 p.m. ET would buy the contract. A trader who believes that the rate will be at or below that level would sell it.

The price of the Binary at any point prior to settlement reflects the market's assessment of the probability of the specified outcome occurring.

- At 2 p.m. ET, spot €/$ stands at 1.2490. The market for this Binary is bid at 33 and offered at 37 (reflecting the market's assessment of an approximately 35% probability that the spot €/$ rate will be above 1.2500 at 3 p.m.). A trader believes that €/$ will drift upward in the next hour, so he places an order to buy 10 contracts at 37. His order is executed on the Nadex exchange opposite the existing offers at a cost of 10 contracts*$37 = $370.

- The trader's risk is strictly capped. The trader can lose a maximum of 37*10 = $370. Similarly, the trader's potential profit is limited by the contract size of $100, so if the contract expires in the money, the trader will make $100-$37 = $63 per contract or $630 in profit.

- At 2:55 p.m. ET, €/$ has risen 15 points, to 1.2505. The market for this contract is now 95-98. The trader decides to try to take a profit rather than waiting for expiration and places an order to sell 10 contracts at 95. This order is executed on the exchange and the trader takes a profit of (95-37)*10 = $580.

This example illustrates the way that Binaries can multiply trading opportunities in a quiet market — in this case a movement of just 0.12% in the underlying market has resulted in the Binary's value changing by 157%.

6.2.3 Pricing with a Skew

In the standard *Black & Scholes* model, one can interpret the premium of the binary option in the risk-neutral world as the expected value = probability of being in-the-money * unit, discounted to the present value.

To take volatility skew into account, a more sophisticated analysis based on call spreads can be used. As a matter of fact, from the perspective of a trader, he actually

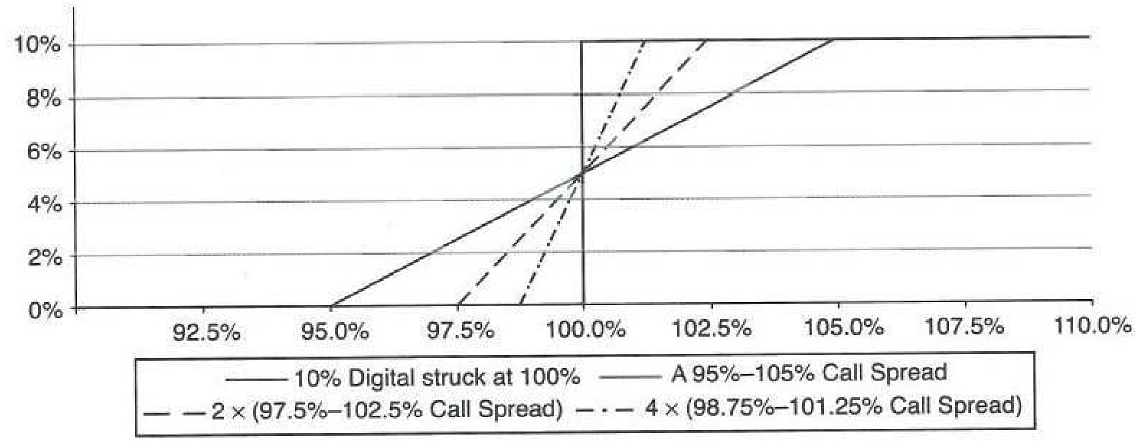

Figure 6.2: A 10% digital with three call spreads with different strikes. - BO10 bl 169 F11.3

trades a call spread [DW 08]. When a trader sells a digital he books a call spread in his risk management system instead of the exact terms of a digital. In this way he can accurately risk manage his digital as call spread.

A binary call option is, at long expirations, similar to a tight call spread using two vanilla options. One can model the value of a binary cash-or-nothing option, V_{con}, at strike K, as an infinitessimally tight spread, that is consider a call spread with strikes $K-\epsilon$ and $K+\epsilon$. Let C be the value of a vanilla call option, we then have

$$V_{con} = \lim_{\epsilon \to 0} \frac{C(S, K-\epsilon) - C(S, K+\epsilon)}{2\epsilon} = -\frac{\partial C(S,K)}{\partial K}. \tag{6.6}$$

This is intuitively explained in Fig. 6.2.3 where we have a 10% digital and three call spreads [BO 10]. Note that to obtain the payoffs in the diagram we need one 95%-105%, but we need two of the 97.5%-102.5% call spreads and four of the 98.75%-101.25% call spreads. This is a demonstration of the limit shown in Eq. 6.6 in the sense that, as the distance between the call strikes and digital strike, ϵ gets smaller, we need $1/\epsilon$ call spreads of width 2ϵ to replicate the digital.

When one takes volatility skew into account, σ is a function of K such that

$$V_{con} = -\frac{dC(S, K, \sigma(K))}{dK} = -\frac{\partial C(S,K)}{\partial K} - \frac{\partial C(S,K)}{\partial \sigma}\frac{\partial \sigma}{\partial K}. \tag{6.7}$$

The first term is equal to the premium of the binary option ignoring skew because

$$-\frac{\partial C}{\partial K} = -\frac{\partial \left(Se^{-r_f \tau} N(\phi x) - Ke^{-r_d \tau} N(\phi y)\right)}{\partial K} = e^{-r_d \tau} N(y) = V_{con}^{noskew}. \tag{6.8}$$

Note that $\partial C/\partial\sigma$ is the Vega of the vanilla call and $\partial\sigma/\partial K$ is sometimes called the "skew slope" or just "skew". Skew is typically negative, so the value of a binary call is higher when taking skew into account. In short

$$V_{con} = C(noskew) - Vega \times Skew.$$

6.3 Barrier Options

Barrier options are standard calls and put except that they either disappear (the option is knocked out) or appear (the option is knocked in) if the underlying asset price breach a predetermined level (the barrier) [BB 98]. Barrier options are thus conditional options, dependent on whether the barriers have been crossed within the lives of the options. These options are also part of a class of options called path-dependent options.

Apart from distinguishing between knock-in and knock-out options there is a second distinction to make. If the options knocks in or out when the underlying price ends up above the barrier level, we speak of an up-barrier. Likewise, if the price ends up below the barrier we speak of a down- or lower-barrier [CS 97].

Barrier options are probably the oldest of all exotic options and have been traded sporadically in the US market since 1967 [Zh 97]. These options were developed to fill certain needs of hedge fund managers. Barrier options provided hedge funds with two features they could not obtain otherwise: the first is that most "down-and-out" options were written on more volatile stocks and these options are significantly cheaper than the corresponding vanilla options. The second feature is the increased convenience during a time when the trading volume of stock options was rather low. In other words, barrier options were created to provide risk managers with cheaper means to hedge their exposures without paying for price ranges that they believe unlikely to occur. Barrier options are also used by investors to gain exposure to (or enhance returns from) future market scenarios more complex than the simple bullish or bearish expectations embodied in standard options. The features just mentioned have helped to make barrier options the most popular path-dependent options being traded world wide. They are also the most commonly traded kind of exotic options in the FX markets.

6.3.1 Types of Barrier Options

Having mentioned that barriers come in all shapes and sizes, we consider the most basic type of barrier option — the single barrier. This option comes in 8 flavours, each with its own characteristics [Ko 03, Ca 10]

1. Up & In
2. Up & Out

3. Down & In

4. Down & Out

Where each type can take the form of a call or a put, giving us a total of 8 single barrier types. An "In" barrier means that a barrier becomes active once crossing a particular barrier level; for example, an Up&In barrier becomes active when the underlying price hits a barrier from below. In the currency market, the 'out' options are known as 'knockouts'.

Barrier options can also have cash rebates associated with them. This is a consolation prize paid to the holder of the option when an out barrier is knocked out or when an in barrier is never knocked in. The rebate can be nothing or it could be some fraction of the premium. Rebates are usually paid immediately when an option is knocked out, however, payments can be deferred to the maturity of the option.

We talk about vanilla barrier options because variations on the basic barrier come in many types [CS97]. First, the barrier need not be active during the whole life of the option. In this case we talk of a partial barrier instead of a full barrier. A second variation concerns the monitoring frequency. It is not always necessary or desirable to check for a barrier hit continuously. Monitoring can be limited to once a day, a week or month. In that case we speak of a discrete barrier and not a continuous barrier. Thirdly, the barrier might not necessarily be linked to the underlying price. It may be linked to another variable like another interest rate or another exchange rate. This is referred to an outside barrier as opposed to an inside barrier.

6.3.2 Monitoring the Barrier

In practise one has to define precisely what it means for the barrier of an option to be crossed. The issue is how the spot price of the underlying is tracked. Is the barrier breached the moment the spot price crosses it intra-day? Further, does one use the last trade, the bid, the offer or the middle of the double? One can also use the official end-of-day closing prices meaning the barrier is only deemed breached if the closing price crossed the barrier. One can also specify that the price of the underlying should have breached the barrier level by at least a certain time period. These options are called Parisian options [Hu 06].

The key to barrier event monitoring is transparency [Hs 97]. The option writers need to be transparent as to what method is used to monitor whether a barrier has been breached or not. This process needs to be impartial, objective and consistent. For instance, when the last trade is used as monitor, the minimum size of the transaction needed to trigger a barrier event becomes crucial. This is to prevent dealers from trying to push the spot price through the barrier level, at their own benefit.

Option writing warehouses need to put policies in place to prevent dealers from deliberately triggering or defending barriers.

6.3.3 Pricing Barrier Options

Robert Merton was first at deriving a closed-form solution for a barrier option where he showed that a European barrier option can be valued in a *Black & Scholes* environment [Me 73]. Thereafter *Rubinstein* and *Reiner* generalised barrier option-pricing theory [RR 91]. *Rich* gave an excellent summary of barrier options [Ri 94]. With a rebate, continuous dividend yield and continuous monitoring of the barrier, the following general equations are obtained

$$
\begin{aligned}
A &= \phi S\, e^{-r_f\tau}\left(\frac{H}{S}\right)^{2\lambda} N\left(\eta y\right) - \phi K\, e^{-r_d\tau}\left(\frac{H}{S}\right)^{2\lambda-2} N\left(\eta y - \eta\sigma\sqrt{\tau}\right) \\
B &= R\, e^{-r_d\tau}\left[N\left(\eta x_1 - \eta\sigma\sqrt{\tau}\right) - \left(\frac{H}{S}\right)^{2\lambda-2} N\left(\eta y_1 - \eta\sigma\sqrt{\tau}\right)\right] \\
C &= \phi S\, e^{-r_f\tau} N\left(\phi x\right) - \phi K\, e^{-r_d\tau} N\left(\phi x - \phi\sigma\sqrt{\iota}\right) \\
D &= \phi S\, e^{-r_f\tau} N\left(\phi x_1\right) - \phi K\, e^{-r_d\tau} N\left(\phi x_1 - \phi\sigma\sqrt{\tau}\right) \\
E &= \phi S\, e^{-r_f\tau}\left(\frac{H}{S}\right)^{2\lambda} N\left(\eta y_1\right) - \phi K\, e^{-r_d\tau}\left(\frac{H}{S}\right)^{2\lambda-2} N\left(\eta y_1 - \eta\sigma\sqrt{\tau}\right) \\
F &= R\left[\left(\frac{H}{S}\right)^{a+b} N\left(\eta z\right) + \left(\frac{H}{S}\right)^{a-b} N\left(\eta z - 2\eta b\sigma\sqrt{\tau}\right)\right]
\end{aligned}
$$

where S is the spot FX rate, K is the strike price, H is the barrier (in the same units as S and K), R is the rebate (in currency units), τ is the annualised time till expiration, r_d is the domestic and r_f the foreign risk-free short term interest rates in continuous format, σ is the volatility and ϕ and $\eta$are binary variables set out in Table 6.1 below.

All the other variables are defined as follows (with lnthe natural logarithm)

$$
\begin{aligned}
x &= \frac{1}{\sigma\sqrt{\tau}}\left\{\ln\left(\frac{S}{K}\right)+\left(r_d-r_f+\frac{\sigma^2}{2}\right)\tau\right\}\\
x_1 &= \frac{1}{\sigma\sqrt{\tau}}\left\{\ln\left(\frac{S}{H}\right)+\left(r_d-r_f+\frac{\sigma^2}{2}\right)\tau\right\}\\
y &= \frac{1}{\sigma\sqrt{\tau}}\left\{\ln\left(\frac{H^2}{SK}\right)+\left(r_d-r_f+\frac{\sigma^2}{2}\right)\tau\right\}\\
y_1 &= \frac{1}{\sigma\sqrt{\tau}}\left\{\ln\left(\frac{H}{S}\right)+\left(r_d-r_f+\frac{\sigma^2}{2}\right)\tau\right\}\\
z &= \frac{1}{\sigma\sqrt{\tau}}\left\{\ln\left(\frac{H}{S}\right)+b\sigma^2\tau\right\} \qquad (6.9)\\
\lambda &= 1+\frac{\mu}{\sigma^2}\\
a &= \frac{\mu}{\sigma^2}\\
b &= \frac{1}{\sigma^2}\left[\sqrt{\mu^2+2r\sigma^2}\right] \qquad (6.10)
\end{aligned}
$$

$$\mu = r_d - r_f - \frac{\sigma^2}{2}.$$

$N(\bullet)$ is the cumulative of the normal distribution function.

The valuation formulas for the eight barrier options can be written as combinations of the quantities A to F given in Eq. 6.9. The value of each barrier is also dependent on whether the barrier H is above or below the strike price K. The pricing formulae of all barrier options are set out in Table 6.1.

The abbreviations used are: $DIC_{K<H}$ is short for "down and in call barrier option" where the strike value K is less than the barrier value H. If the payment of the rebate is deferred to maturity for the *knock-out* options, we put $F = B$ in the equations above. The following should be considered when implementing the formulas in Table 6.1

- The formulas given for the "in"-barriers are used to value the option before the barrier is hit. After the barrier has been hit (the option was knocked in), the buyer has an ordinary vanilla call or put. Use the Black-Scholes equations for vanilla options to value the option.

- The *knock-out's* value is zero after the barrier has been hit.

6.3.4 Reverse Knockout

In the currency market, a reverse knockout is an option that extinguishes an in-the-money option. For example, **a reverse knockout put has a barrier level lower**

Call	Put
Down and In Barriers	
$\phi = \eta = 1$ $DIC_{K \geq H} = A + B$ $DIC_{K<H} = C - D + E + B$	$\phi = -1, \eta = 1$ $DIP_{K \geq H} = D - A + E + B$ $DIP_{K<H} = C + B$
Up and In Barriers	
$\phi = 1, \eta = -1$ $UIC_{K \geq H} = C + B$ $UIC_{K<H} = D - A + E + B$	$\phi = -1, \eta = -1$ $UIP_{K \geq H} = C - D + E + B$ $UIP_{K<H} = A + B$
Down and Out Barriers	
$\phi = \eta = 1$ $DOC_{K \geq H} = C - A + F$ $DOC_{K<H} = D - E + F$	$\phi = -1, \eta = 1$ $DOP_{K \geq H} = C - D + A - E + F$ $DOP_{K<H} = F$
Up and Out Barriers	
$\phi = 1, \eta = -1$ $UOC_{K \geq H} = F$ $UOC_{K<H} = C - D + A - E + F$	$\phi = -1, \eta = -1$ $UOP_{K \geq H} = D - E + F$ $UOP_{K<H} = C - A + F$

Table 6.1: Pricing Formulas for European barrier options. The variables are defined in Eq. 6.9.

than the strike [Ta 10]. As an example, suppose you have a 5-month EURUSD put with strike $K = 1.4$ and barrier $H = 1.35$. The put thus initially benefits from falling EURUSD. However, if the EURUSD trades below 1.35, the put goes from in-the-money to worthless. This is shown in Fig. 6.3. We also show the Delta and Vega profiles. The discontinuity is reflected in the dramatic change in the Delta as we approach the barrier level. The Vega drops to zero close to the barrier.

6.3.5 Parity Relationship

In-out parity is the barrier option's answer to put-call parity: if we combine an in- and an out- barrier option, both with the same strike and expiration date, we get the price of a vanilla option. Here is how it works.

Consider a portfolio consisting of an in-option and its corresponding out-option without any rebate. If the barrier is never reached, the in-option will end up worthless, yet the out-option will have its corresponding vanilla option as payoff. On the other hand, if the barrier is reached any time during the life of the option, the in-option will have the corresponding vanilla option as its payoff whereas the out-option will end up worthless. Since the portfolio and the corresponding vanilla option have exactly the same payoff at maturity, arbitrage arguments guarantee that the portfolio and

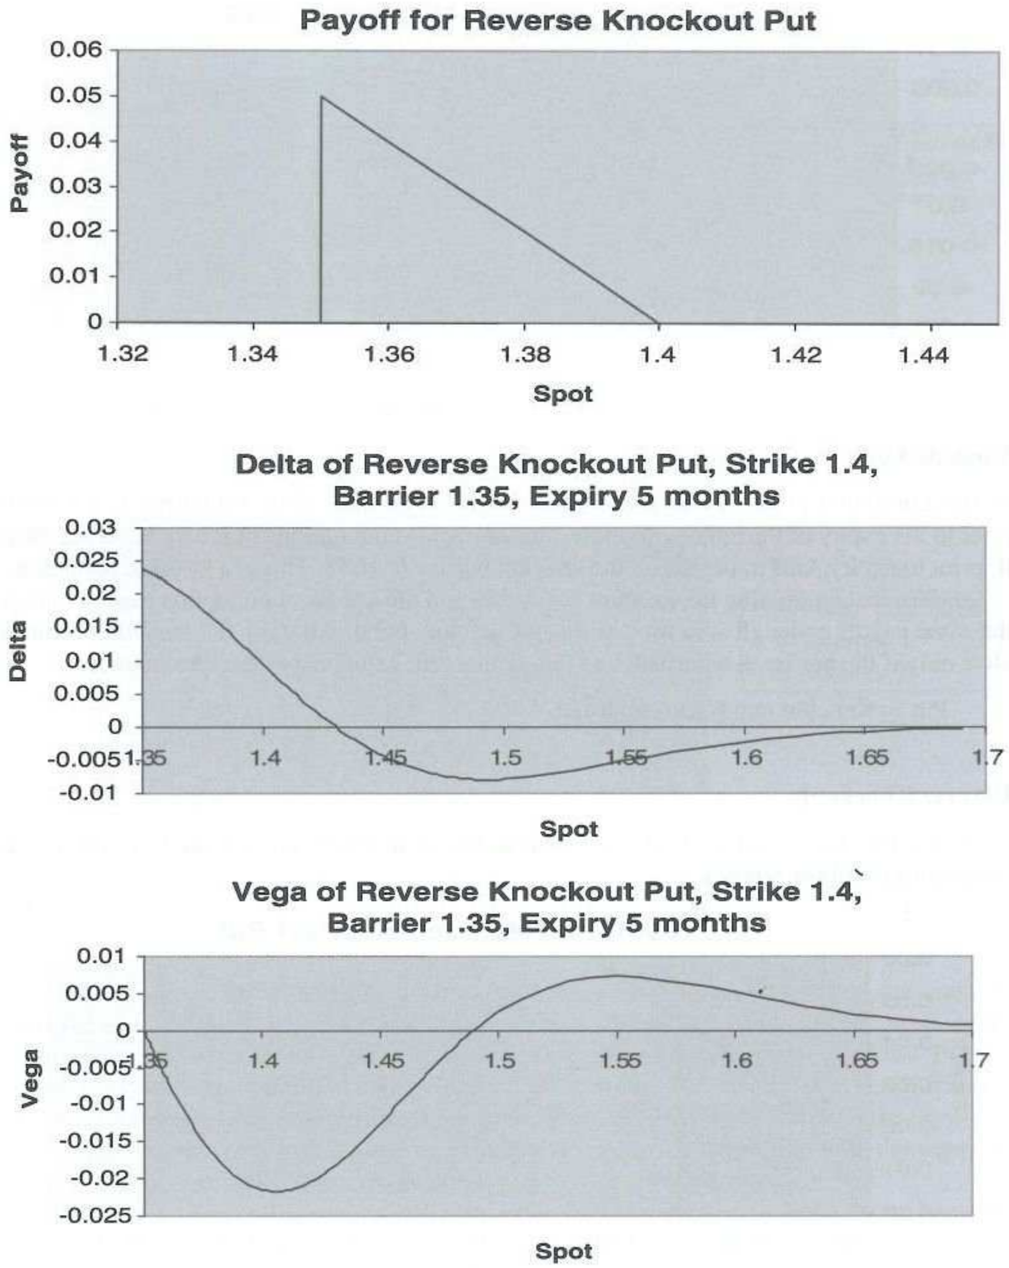

Figure 6.3: Payoff, Delta and Vega for Reverse Knockout Put - Tan bl 51.

the vanilla option must be the same. Algebraically this identity is

$$\text{Black} - \text{ScholesvanillaCall(Put)} = \text{OutCall(Put)} + \text{InCall(Put)}.$$

This is also known as the 'kick-in-kick-out' parity rule.

6.3.6 Behaviour of Barrier Options

Barrier options are less expensive than standard options but provide similar potential investment returns. We list some points that need to be considered when trading barrier options

- Knock-outs:
  - The closer the barrier level to the current spot, the lower the barrier option's premium;
  - The higher the volatility, the lower the barrier option's premium;
  - The longer the time to expiration of the option, the lower the barrier option's premium.

  For example, an investor that is long an up-and-out call, is forfeiting some of the upside potential of an ordinary call but the payoff can be the same if the barrier is never hit. If the barrier is hit, the investor looses his exposure and the barrier must thus be less expensive than a standard option.

- Knock-in options' behaviour is similar to standard options but the premiums are also less
  - The closer the barrier level to the current spot, the higher the barrier option's premium;
  - The higher the volatility, the higher the barrier option's premium;
  - The longer the time to expiration of the option, the higher the barrier option's premium.

6.3.7 Continuity Correction

The aforementioned analytic formulas present a method to price barrier options in continuous time, but often in industry, the asset price is sampled at discrete times, where periodic measurements rather than a continuous lognormal distribution of the asset prices is assumed. Most barriers are monitored at the end of day only meaning that the official close of the day is the level to monitor and see whether it has breached

Continuous	$\Delta t = 0$
Daily	$\Delta t = 1/365$
Weekly	$\Delta t = 1/52$
Monthly	$\Delta t = 1/12$
Hourly	$\Delta t = 1/(365 \times 24$

Table 6.2: Monitoring barriers dicretely in time.

the barrier or not. *Broadie, Glasserman & Kou* calculated an adjustment to the barrier value to account for discrete sampling as follows [BGK 97]

$$H = He^{\alpha\sigma\sqrt{\Delta t}} \tag{6.11}$$

For "up" options which hit the barrier from underneath, the value of α is 0.582597, whereas for "down" options where the barrier is hit from the top, the value of α is -0.582597, where Δt is the time interval. The time intervals used most often is set out in Table 6.2.

6.3.8 The Delta

The risk parameters can be obtained by calculating the relevant partial derivatives of the equations given in Eq. 6.9. To obtain the Δ we calculate the partial derivatives of the functions A to F with respect to S and by substituting these in the equations given in Table 6.1 e.g. substitute A with $\partial A/\partial S$. By taking the partial derivatives, with respect to S, we obtain

$$
\begin{aligned}
\frac{\partial A}{\partial S} &= \frac{2}{S}(1-\lambda)A - \phi\ e^{-d\tau}\left(\frac{H}{S}\right)^{2\lambda} N\left(\eta y\right) \\
\frac{\partial B}{\partial S} &= \frac{2}{S}(1-\lambda)R\ e^{-r\tau}\left(\frac{H}{S}\right)^{2\lambda-2} N\left(\eta y_1 - \eta\sigma\sqrt{\tau}\right) \\
&\quad + \frac{\eta R}{\sigma\sqrt{\tau H}}\ e^{-d\tau}\left[N'\left(\eta x_1\right) + \left(\frac{H}{S}\right)^{2\lambda} N'\left(\eta y_1\right)\right] \\
\frac{\partial C}{\partial S} &= \phi\ e^{-d\tau} N\left(\phi x\right) \qquad (6.12) \\
\frac{\partial D}{\partial S} &= \phi\ e^{-d\tau} N\left(\phi x_1\right) - \phi K\ e^{-d\tau}\frac{N'(\phi x_1)}{\sigma\sqrt{\tau}}\left[1 - \frac{K}{H}\right] \\
\frac{\partial E}{\partial S} &= \frac{2}{S}(1-\lambda)E - \phi\ e^{-d\tau}\left(\frac{H}{S}\right)^{2\lambda} N\left(\eta y_1\right) - \phi\eta\ e^{-d\tau}\left(\frac{H}{S}\right)^{2\lambda}\frac{N'(\eta y_1)}{\sigma\sqrt{\tau}}\left[\frac{K}{H} - 1\right] \\
\frac{\partial F}{\partial S} &= -\frac{R}{S}\left\{(a+b)\left(\frac{H}{S}\right)^{a+b} N\left(\eta z\right) + (a-b)\left(\frac{H}{S}\right)^{a-b} N\left(\eta z - 2\eta b\sigma\sqrt{\tau}\right)\right. \\
&\quad \left. + \frac{\eta}{\sigma\sqrt{\tau}}\left[\left(\frac{H}{S}\right)^{a+b} N'\left(\eta z\right) + \left(\frac{H}{S}\right)^{a-b} N'\left(\eta z - 2\eta b\sigma\sqrt{\tau}\right)\right]\right\}.
\end{aligned}
$$

6.3.9 Static Hedging

Some barriers are very difficult to hedge statically . With put-call-symmetry, however, we are able to create perfect static replications for some barriers (if $R = 0$) with zero cost strategies — this works best where the underlying instrument is a futures contract. These hedging techniques are easily implemented and are described in the Table 6.3 where we have assumed we are the writers of the barrier options and we want to hedge our exposure.

The strategies in Table 6.3 are zero cost if everything stays the same i.e., interest rates and volatilities. This hardly ever happens which means there are some risks involved in doing this. In illiquid or underdeveloped markets, one might also not be able to pick up the vanilla options at the strikes proposed.

Statically hedging barrier options not mentioned in Table 6.3 is impossible; those can only be hedged dynamically i.e., re-value and manage risk on a daily basis. Another hedging strategy often used is, if the number of barrier options is small, all can be dumped into a large portfolio of standard vanilla options and the risk is then managed all together.

Delta hedging becomes very difficult if the time to expiry is short and the spot price is near the barrier level. For a knock out, the delta can become very negative near the knock out boundary. The hedger is in an unstable situation. Because the

Up and In Call ($K \geq H$) Do nothing until barrier is hit. The moment the barrier is hit, buy a vanilla call with the same strike as that of the up and in call originally sold.
Down and In Put ($K \leq H$) Do nothing until the barrier is hit. The moment the barrier is hit buy a put with the same strike as that of the down and in put originally sold.
Down and In Call ($K \geq H$) With the sale of the barrier, go long K/H puts with strike H^2/K. The moment the barrier is touched, sell the puts and buy a call at the same strike as that of the down and in call. If the barrier is never touched, all options expire worthless.
Down and Out Call ($K \geq H$) With the sale of the barrier, go short K/H puts with strike H^2/K and buy a call with strike K. The moment the barrier is touched, sell the calls and buy back the puts.
Up and In Put ($K \leq H$) With the sale of the barrier, go long K/H calls with strike H^2/K. The moment the barrier is touched, sell the calls and buy a put at the same strike as that of the up and in put. If the barrier is never touched, all options expire worthless.
Up and Out Put ($K \leq H$) With the sale of the barrier, go short K/H calls with strike H^2/K and buy a put with strike K. The moment the barrier is touched, sell the puts and buy back the calls at the same strike as that of the up and out put.

Table 6.3: Hedging strategies for barrier options.

delta is so negative he should take a very large short position in the underlying stock and invest these proceeds in the money market. If the stock does not cross the barrier he covers his short position with the money market funds, pays off the option and is left with zero funds – the option would be alive and would expire in-the-money. If the stock crosses the barrier (the option is knocked out) the delta becomes zero. He should now cover his short position with the money market. This is more expensive than before because the stock price has risen and consequently he is left with no money. But, the option is not alive anymore so no money is needed to pay it off.

Because a large short position is being taken, a small error in hedging can create a significant effect. To circumvent this, do the following: rather than using the boundary condition $v(t, H) = 0$, $0 \leq t \leq T$, (i.e. the value of the barrier option should be zero at the barrier level), solve the barrier partial differential equation with the new boundary condition

$$v(t, H) + \alpha H \Delta(t, H) = 0, \qquad 0 \leq t \leq T \tag{6.13}$$

where α is a "tolerance parameter", say 1%. At the boundary, $H \Delta(t, H)$ is the currency value of the short position. The new boundary condition guarantees

- $H \Delta(t, H)$ remains bounded;
- the value of the portfolio is always sufficient to cover a hedging error of α times the currency value of the short position.

6.3.10 Pricing with the Binomial

Like most other path-dependent options, barrier options can be priced via lattice tree such as binomial, trinomial or adaptive mesh models by solving the PDE using a generalised finite difference method. The binomial tree method was discussed in §1.13. A tree with a barrier os depicted in Fig. 6.4. An important aspect to note is that with the existence of barriers, the 'zig-zag' movements of a binomial model will undoubtedly create problems, as the true barrier of the barrier option in question is often not the same barrier implied by the tree. We show the 'zig-zag' convergence in Fig. 6.5. There are actually two barriers: the specified barrier and the 'effective barrier' due to its position relative to the nearest nodes. This is shown in Fig. 6.6 [DK 95].

Boyle & Lau suggested a way to determine which nodes will give good approximate pricing of a barrier option (where the inner and outer nodes are closest together). The number of time steps which will give the most accurate prices when using the binomial lattice is given as [RS 95, BL 94]

$$Node(i) = \frac{i^2 \sigma^2 \tau}{\left(\ln \frac{S}{H}\right)^2} \tag{6.14}$$

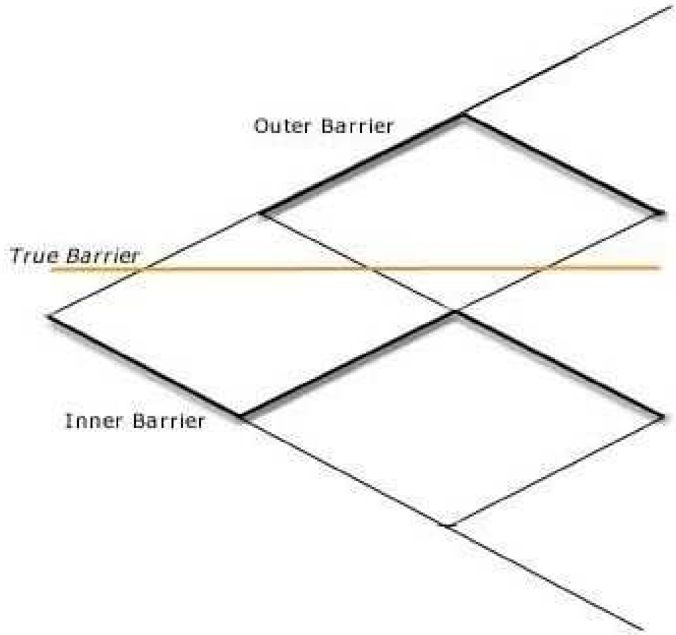

Figure 6.4: A three step tree with a barrier.

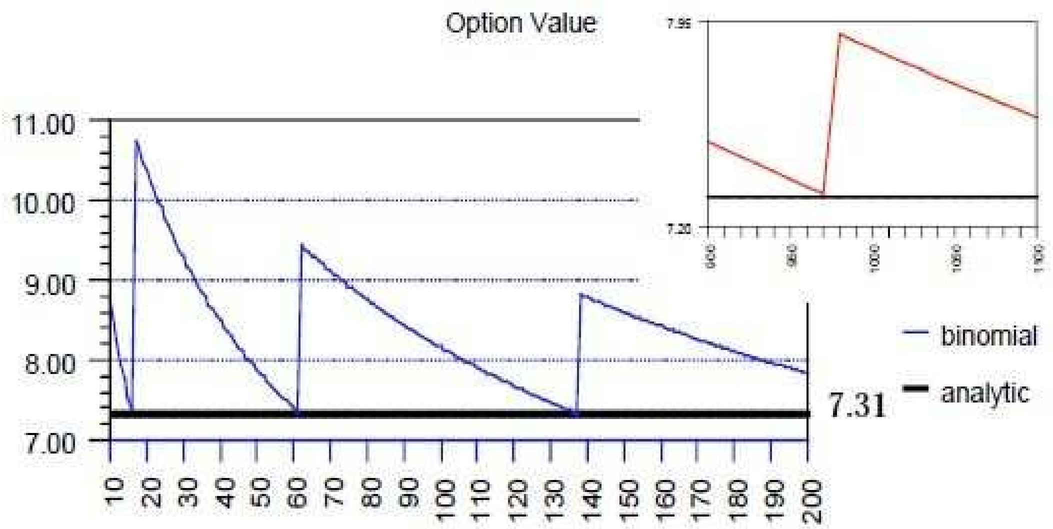

Figure 6.5: Convergence to analytic value of a binomially-valued one year European down-and-out call option as the number of binomial levels increases [DK 95].

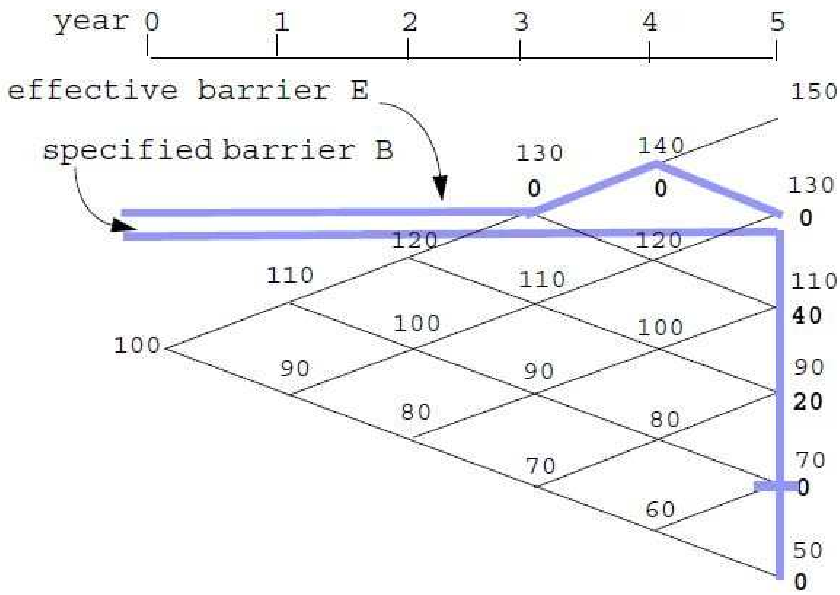

Figure 6.6: Specified and effective barrier. The specified barrier is at 125 but the effective barrier is at 130.

where $i = 1, 2, 3, \ldots, N$ nodes.

Nonetheless, although these optimal nodes will provide a more accurate value for the barrier, a large number of time steps should be used to determine a reasonable value. The convergence of the binomial model is shown in Fig. 6.7.

6.3.11 Partial Time Barrier Options

Options where the barrier H is only considered for some fraction of the option's lifetime are referred to as partial-time or window barrier options. The most popular is the early-ending barrier option, i.e. where the monitoring starts at the deal date but it ends some time t_1 before expiry where $t_1 \leq T$ [Ha 07]. We show this in Fig. 6.8.

Partial-time barrier options give more flexibility in comparison to vanilla barrier options, they also in general give lower premiums compared to the respective standard vanilla option.

6.4 One-Touch Digitals

A type of exotic option that gives an investor a fixed payout R once the price of the underlying asset reaches or surpasses a predetermined barrier H. This type of option allows the investor to set the position of the barrier, the time to expiration and the payout to be received once the barrier is broken.

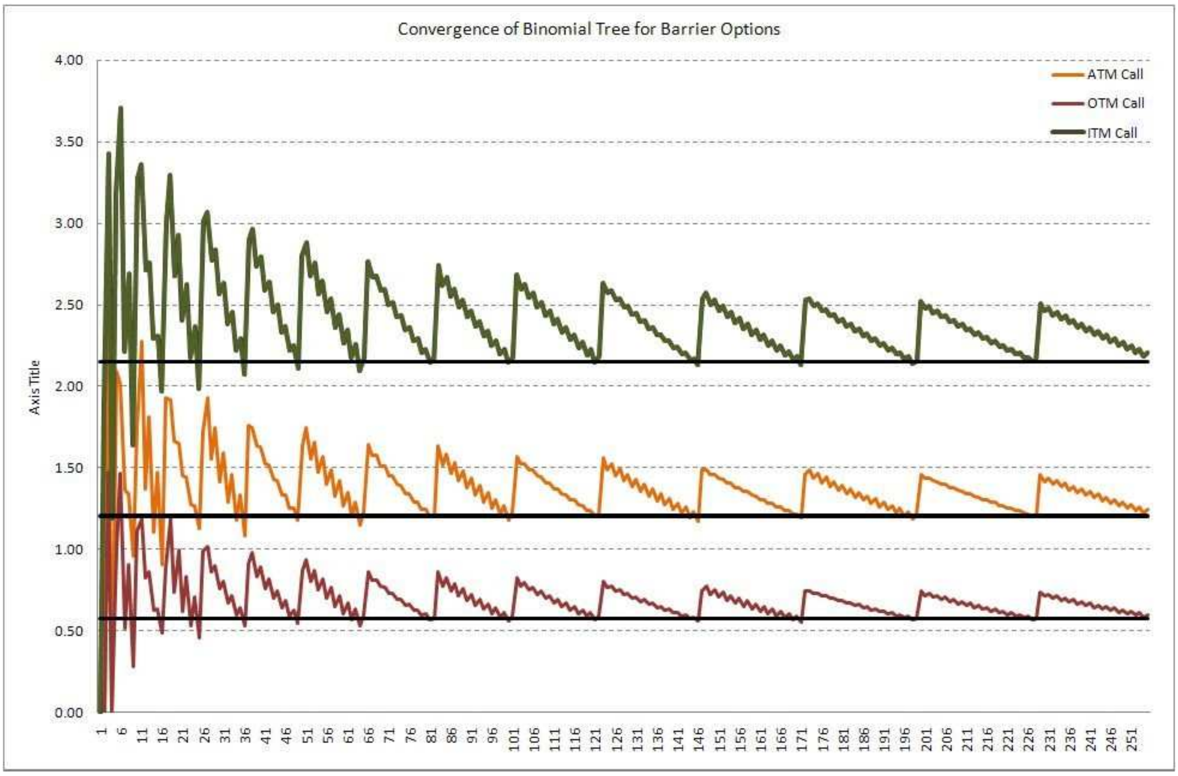

Figure 6.7: Wave-like pattern of the binomial model. We show the patterns for at-, in-, and out-the-money options.

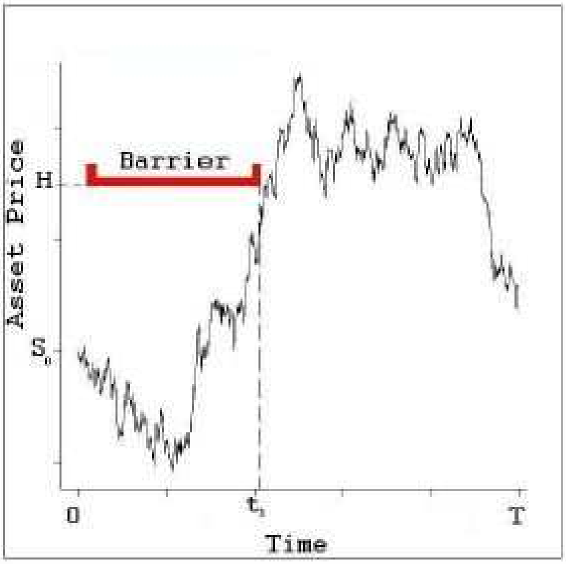

Figure 6.8: Early ending partial time barrier option.

Only two outcomes are possible with this type of option

1. the barrier is breached and the trader collects the full payout agreed upon at the outset of the contract, or

2. the barrier is not breached and the trader loses the full premium paid to the broker.

This type of option is useful for traders who believe that the price of an underlying asset will exceed a certain level in the future, but who are not sure that the higher price level is sustainable. Because a one-touch option only has one barrier level, it is generally slightly less expensive than a double one-touch option.

Speculative market participants like to use one-touch options as bets on a rising or falling exchange rate. Clients, who prefer to hedge, trade one-touch options as a rebate in order to secure themselves a compensation in case their strategy doesn't work out. One-touch options are also often integrated into structured products to increase returns on forward and interest rates.

Similar as for vanilla barrier options we obtain the following factors (see Eq. 6.9)

$$
\begin{aligned}
A_1 &= Se^{-r_f\tau}N(\phi x) \\
B_1 &= Re^{-r_d\tau}N(\phi x - \phi\sigma\sqrt{\tau}) \\
A_2 &= Se^{-r_f\tau}N(\phi x_1) \\
B_2 &= Re^{-r_d\tau}N(\phi x_1 - \phi\sigma\sqrt{\tau}) \\
A_3 &= Se^{-r_f\tau}\left(\frac{H}{S}\right)^{2\lambda} N(\eta y) \\
B_3 &= Re^{-r_d\tau}\left(\frac{H}{S}\right)^{2\lambda-2} N(\eta y - \eta\sigma\sqrt{\tau}) \\
A_4 &= Se^{-r_f\tau}\left(\frac{H}{S}\right)^{2\lambda} N(\eta y_1) \\
B_4 &= Re^{-r_d\tau}\left(\frac{H}{S}\right)^{2\lambda-2} N(\eta y_1 - \eta\sigma\sqrt{\tau}) \\
A_5 &= R\left[\left(\frac{H}{S}\right)^{a+b} N(\eta z) + \left(\frac{H}{S}\right)^{a-b} N(\eta z - 2\eta b\sigma\sqrt{\tau})\right]
\end{aligned}
\tag{6.15}
$$

Down-and-in cash-(at hit)-or-nothing $S > H$	Payoff: Value:	R (at hit)if for some $\tau \leq T$ $S(\tau) \leq H$ 0 if for all $\tau \leq T$, $S(\tau) > H$ A_5; $\eta = 1$
Up-and-in cash-(at hit)-or-nothing $S < H$	Payoff: Value:	R (at hit)if for some $\tau \leq T$ $S(\tau) \geq H$ 0 if for all $\tau \leq T$, $S(\tau) < H$ A_5; $\eta = -1$
Down-and-in cash-(at expiry)-or-nothing $S > H$	Payoff: Value:	R (at hit)if for some $\tau \leq T$ $S(\tau) \leq H$ 0 if for all $\tau \leq T$, $S(\tau) > H$ $B_2 + B_4$; $\eta = 1$, $\phi = -1$
Up-and-in cash-(at expiry)-or-nothing $S < H$	Payoff: Value:	R (at hit)if for some $\tau \leq T$ $S(\tau) \geq H$ 0 if for all $\tau \leq T$, $S(\tau) < H$ $B_2 + B_4$; $\eta = -1$, $\phi = 1$

Table 6.4: One touch options

where

$$
\begin{aligned}
x &= \frac{1}{\sigma\sqrt{\tau}} \ln\left(\frac{S}{K}\right) + \lambda\sigma\sqrt{\tau} \\
x_1 &= \frac{1}{\sigma\sqrt{\tau}} \ln\left(\frac{S}{H}\right) + \lambda\sigma\sqrt{\tau} \\
y &= \frac{1}{\sigma\sqrt{\tau}} \ln\left(\frac{H^2}{SK}\right) + \lambda\sigma\sqrt{\tau} \\
y_1 &= \frac{1}{\sigma\sqrt{\tau}} \ln\left(\frac{H}{S}\right) + \lambda\sigma\sqrt{\tau} \\
z &= \frac{1}{\sigma\sqrt{\tau}} \ln\left(\frac{H}{S}\right) + b\sigma\sqrt{\tau} \\
\lambda &= 1 + \frac{\mu}{\sigma^2} \\
a &= \frac{\mu}{\sigma^2} \\
b &= \frac{1}{\sigma^2}\left[\sqrt{\mu^2 + 2r\sigma^2}\right] \\
\mu &= r_d - r_f - \frac{\sigma^2}{2}.
\end{aligned}
\tag{6.16}
$$

Taking combinations of the factors in Eq. 6.15 leads to 28 different type of binary barrier options. One touchers are described by four of these. They are given in Table 6.4.

Digital options are simple, easy and inexpensive to trade. If you think, the EURUSD rate is going to be above 1.3500 after 2 months but you are not sure about the timing of this move taking place within the next two months, buy a digital option. If after 2 months, the EURUSD rate is indeed above 1.3500, you get your profit/rebate/cash payout. If not, your digital option will expire. You with lose only a small premium that you had to pay while purchasing the digital option.

One Touch Options are perfect for those currency traders who believe that there will be a retracement and the price of a given currency pair will test a support or resistance level with a false breakout. The one touch options will pay a profit if the market touches the predetermined barrier level.

One touchers are often part of packages. Recall our example in §6.3.4. There we had a reverse knockout. It is not uncommon for an investor to want his in-the-money put option to be valueless if the underlying drops below the barrier H, but he may well accept a certain fixed rebate should H be breached (since this would make the package cheaper than a normal put option). Hence, a downside one touch added to the reverse knockout put will satisfy this investor [Ta 10]. For this example we can consider adding a one toucher with barrier $H = 1.35$ and notional 1.4-1.35=0.05 to fully compensate the investor if the barrier is breached. Notice that this structure is more valuable than a capped put or put spread (i.e. where you have bought a put with strike at 1.4 and sold a put with strike 1.35). This must be case because if the barrier is breached the payout will be $0.05 regardless of the value of EURUSD at expiry. For this reason the notional of the rebate might be less than 0.05 to cheapen the structure.

It is worth mentioning that a one touch option cost about twice that of a normal digital (with stike = barrier) option for the following reason: you can sell a one touch option and buy two digital options to hedge. If the one toucher never breaches the barrier, then neither will the digitals and all are worthless. If the one touches does touch the barrier at a time prior to expiry, the digitals will roughly be at-the-money since spot will be at the barrier level which is also the strike. Now, at-the-money digitals has a value of roughly 0.5 whereas a one toucher pays off one and hence has a value of one.

6.5 No-Touch Digitals

A type of exotic option that gives an investor a fixed payout R if the price of the underlying asset does not reach or surpasses a predetermined barrier H. This type of option allows the investor to set the position of the barrier, the time to expiration and the payout to be received.

This type of option is useful for a trader who believes that the price of an underlying asset will remain range-bound over a certain period of time. A No-Touch Option is a great way to profit from a trending market. The no-touch option pays a fixed amount if the market never touches the barrier level that you choose. All you need to do is to determine the desired payoff, the currency pair, the barrier price and the expiration date.

Down-and-out cash-or-nothing $S > H$	Payoff: Value:	R (at hit)if for some $\tau \leq T$ $S(\tau) > H$ 0 if for all $\tau \leq T$, $S(\tau) \leq H$ $B_2 - B_4$; $\eta = 1$, $\phi = 1$
Up-and-out cash-or-nothing $S < H$	Payoff: Value:	R (at hit)if for some $\tau \leq T$ $S(\tau) < H$ 0 if for all $\tau \leq T$, $S(\tau) \geq H$ $B_2 - B4$; $\eta = -1$, $\phi = -1$
Down-and-in cash-(at expiry)-or-nothing	Payoff:	R (at hit)if for some $\tau \leq T$ $S(\tau) \leq H$

Table 6.5: No-touch options

No-touchers are the opposite to the one touchers. There are two of these that we list in Table 6.5.

6.6 Double Digital Options

Let's now extend the analysis up to this point and add another barrier to the option. We talk about a *'double barrier'* if there is a barrier below the spot price (called the *lower barrier*) and another barrier above the spot price (called the *upper barrier*).

A Double No-Touch option pays at expiry the notional amount contingent on the event that neither upper or lower barrier has been breached during the life of the option. This type of option is useful for a trader who believes that the price of an underlying asset will remain range-bound over a certain period of time. Double no-touch options are growing in popularity among traders in the forex markets.

For example, assume that the current EURUSD rate is 1.3000 and the trader believes that this rate will not change dramatically over the next 14 days. The trader could use a double no-touch option with barriers at 1.29 and 1.31 to capitalize on this outlook. In this case, the trader stands to make a profit if the rate fails to move beyond either of the two barriers.

A one-touch double barrier option is the opposite to the double no-touch. This means this options pays off a cash amount if one of the two barriers are touche before expiry. This options pays off zero if neither barrier is touche before expiry.

A semi-closed form formula for the valuation of both of these options has been published by *Hui* [Ha 07]. The formula is quite involved but do not give the market price in any case. This is due to the volatility skew that needs to be taken into account. Tree methods works quite well.

6.7 Forward Start Options

Vanilla options become effective immediately after settlement. Forward-start options, however, are options that are paid for now but will start at some pre-specified time in the future with the strike price set to be the underlying asset price at the time when it starts. Alternatively, forward-start options are at-the-money options when they actually start, yet, the strike price is not known at present.

Forward-start options are less expensive than vanilla options and are used by cost sensitive investors. These options also allow an investor or risk manager to lock in the level of volatility at a time when the market's volatility expectations are low and when they expect volatility to increase in the future.

6.7.1 Advantages

- Protection against spot market movement and against increasing volatility [Wy 06],
- Buyer can lock in current volatility,
- Spot risk easy to hedge.

The key reason for trading a forward start option is trading the forward volatility without any spot exposure. In quiet market phases with low volatility, buying a forward start is cheap. Keeping a long position will allow participation in rising volatility, independent of spot prices.

6.7.2 Valuation

A forward-start option with a time to maturity T starts at-the-money or proportionally in- or out-of -the-money after a known elapsed time ζ in the future. The strike is set equal to a positive constant αtimes the asset price S after the time ζ. If $\alpha < 1$, the call (put) will start $1-\alpha$percent in-the-money (out-of-the-money); if $\alpha = 1$, the option will start at-the-money; and if $\alpha > 1$, the call (put) will start $\alpha - 1$ percentage out-of-the-money (in-the-money) [Ha 07].

Rubinstein showed that a European forward-start option can be valued in a *Black & Scholes* environment [Ru 91]. The payoff of a European style forward-start option can be expressed as

$$V_{FSO} = \max\left[\phi\left\{S(T) - \alpha S(\zeta)\right\}, 0\right]$$

with $\phi = 1$ for a call and $\phi = -1$ for a put and ζ is a time such that $\zeta < T$.

It seems difficult to price a forward-start option because the strike price$K = \alpha S(\zeta)$ is not known at present. However, suppose for the time being that the underlying asset price at time ζ is known. The value of a forward start option, at time ζ, is then obtained by using the generalised Black-Scholes equation where we substitute $K = \alpha S(\zeta)$into Eq. 1.3 such that

$$V_{FSO}(\zeta) = \phi S\, e^{-r_d \zeta}\left[e^{-r_f \tau} N(\phi x) - \alpha\, e^{-r_d \tau} N(\phi y)\right] \tag{6.17}$$

and from (1.4) we have

$$\begin{aligned} x &= \frac{1}{\sigma\sqrt{\tau}}\left\{\ln\left(\frac{1}{\alpha}\right) + \left(r_d - r_f + \frac{\sigma^2}{2}\right)\tau\right\} \\ y &= x - \sigma\sqrt{\tau}. \end{aligned} \tag{6.18}$$

From Eq. 6.18 we see that a forward-start option is a vanilla option with strike αS and time to maturity $\tau = T - \zeta$, that is adjusted for any dividends paid over the period from time t to the grant time ζ. The value of the option at any time $t \geq \zeta$ (i.e., after time ζ has elapsed) is just given by the generalised *Black & Scholes* equation in Eq. 1.3 with strike K fixed at time $t = \zeta$.

In valuing a FSO, use the forward volatility that is calculated as follows: let σ_ζ be the volatility of a vanilla option contract that expires at time ζ and let σ_T be the volatility of a vanilla option contract that expires at time T. The forward volatility from ζ to T is then given by

$$\sigma = \frac{\sqrt{\sigma_T^2 T - \sigma_\zeta^2 \zeta}}{T - \zeta} \tag{6.19}$$

6.7.3 Peculiarities of Forward-Start Options

- FSO options are less expensive than standard options.
- All of the risk parameters exhibit discontinuities at the grant time.
- The theta and gamma is zero before the grant time.
- After the grant time, a FSO is exactly the same as a vanilla European option.

6.7.4 Risk Parameters and Hedging FSO

The risk parameters can be obtained by calculating the relevant partial derivatives of Equation (6.17). This gives

$$\Delta_{FSO} = \frac{\partial V_{FSO}}{\partial S} = \phi\, e^{-r_f\zeta}\left[e^{-r_f\tau}N(\phi\, x) - \alpha\, e^{-r_d\tau}N(\phi\, y)\right] \tag{6.20}$$

which leads to the fact that the Gamma is always zero. Similarly we can obtain the Vega where

$$Vega_{FSO} = \frac{\partial V_{FSO}}{\partial \sigma} = S e^{-r_f\tau}\sqrt{\tau}N'(x) \tag{6.21}$$

and the Theta is given by

$$\Theta_{FSO} = \frac{\partial V_{FSO}}{\partial \tau} = \phi\, S e^{-r_f\zeta}\left[-r_f e^{-r_f\tau}N(\phi\, x) - \alpha\, r_d e^{-r_d\tau}N(\phi\, y)\right] + S e^{-r_f T}\frac{\sigma N'(x)}{2\sqrt{\tau}} \tag{6.22}$$

Here, $N'(\bullet)$ is the cumulative normal probability function.

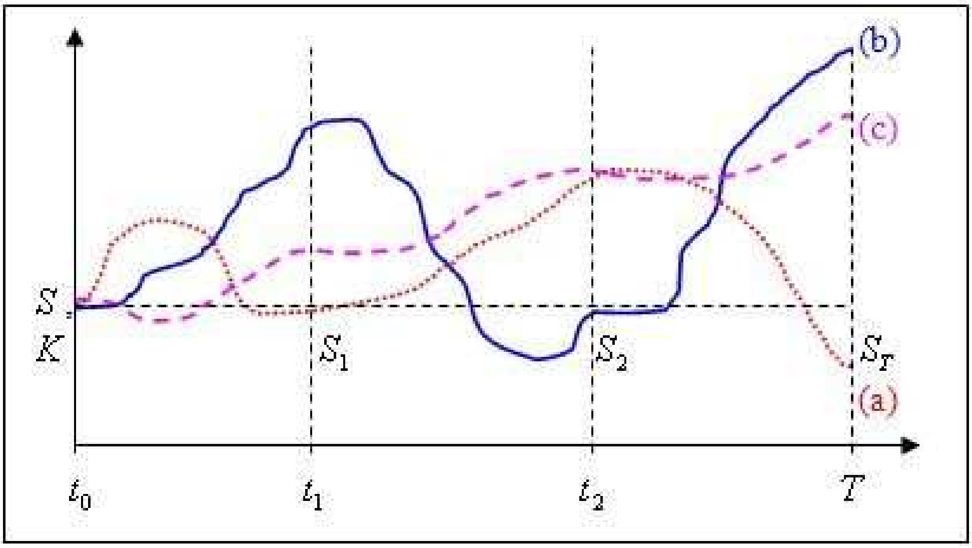

Figure 6.9: Possible price paths for some underlying asset.

6.8 Cliquet/Ratchet Options

These exotic options were first developed in France and were based upon the CAC 40 stock index. The cliquet starts out as an ordinary vanilla option with a fixed strike price, but the strike is reset to be a positive constant, times the asset price on a set of dates that has been predetermined [To 94]. When the strike price is reset, *any positive intrinsic value is locked in.* The profit can be accumulated until final maturity, or paid out at each reset date. If the underlying asset price, at the next reset date, is below the previous level, nothing happens except that the strike price has been reset at a lower level.

A Cliquet option thus consists of a vanilla option that expires at the first reset date together with a series of forward starting options where the start dates are the next reset dates and the expiry dates are the reset dates just after those dates. The locked in values can be paid out at the reset dates (the so-called "pay-as-you-go" Cliquet) or at expiry (the so-called compound or "pay-at-end" Cliquet). The latter Cliquet is slightly cheaper than the first. A Cliquet option is more expensive than a vanilla option due to the chances of an extremely favourable payoff that can be obtained. By buying a Cliquet the investor also locks in the current volatility and interest rate for every forward starting option in the series of options.

The major advantage of the Cliquet is that the probability of some payout is high. There is a high probability that even if the market closes lower/higher at expiry, that it will have closed higher/lower at, at least one of the reset dates. This is illustrated in Fig. 6.8 Figure 6.9 shows three possible price paths for some FX rate starting at time t_0 with a starting value of S. There are three reset dates at t_1, t_2 and T. The prices on these dates are S_1, S_2 and S_T respectively and the strike price is K. The payoffs are summarised in Table 6.8.

Path	**Vanilla Call Payoff**	**Cliquet Payoff**
(a)	0	$(S_2 - S_1)$
(b)	$S_T - K$	$(S_1 - K) + (S_T - S_2)$
(c)	$S_T - K$	$S_T - K$

The Cliquet is suitable for investors with a medium term investment horizon. It is less risky than ordinary medium term options, as there is less specific risk i.e. the reset facility gives the buyer a "second" and "third" chance. This increases the chance of a payout, but must be balanced with the higher premium cost. As a series of forward-start options, the Cliquet is attractive to passive investors as it requires no intermediate management. Investors use Cliquets to take advantage of future assumptions about volatility.

An age old complaint of investors is that when a fund has performed well in the past it usually loses all the profits in one particular bad patch. If this happens the investor will ask the question: why didn't you take profits when you had them? Fund managers who must report at the end of every quarter will find a Cliquet very handy.

The price of a ratchet option is

$$V_{cliquet} = \phi S \sum_{i=1}^{n} e^{-dt_i} \left[e^{-d\tau_i} N(\phi d_{1i}) - \alpha e^{-r\tau_i} N(\phi d_{2i}) \right] \tag{6.23}$$

where t_i are the times to the strike resets, $\tau_i = T_i - t_i$ where T_i are the expiry times of the forward-start options and x_i and y_i are defined in Eq. 6.18 with τ substituted with τ_i. Usually $\alpha = 1$ (options are at-the-money) and the first option is a vanilla European starting today ($t_1 = 0$) and it expires at T_1; the second option's strike is fixed at time t_2 and it expires at time T_2; etc.

To make the pricing of cliquets more rigid one should use the correct interest rate term structures to calculate the forward interest rates and forward volatilities.

6.9 Lookback Options

The dream contract has to be one that pays the difference between the highest and the lowest asset prices realized by an asset over some period. All speculators are trying to achieve such a trade. This can be achieved by a lookback option. There are two types of lookback options

1. Fixed Strike Lookback – The strike is fixed at the start. At maturity, the buyer can "lookback" over the life of the option and choose the expiry value of the underlying at the level that maximizes his payoff.

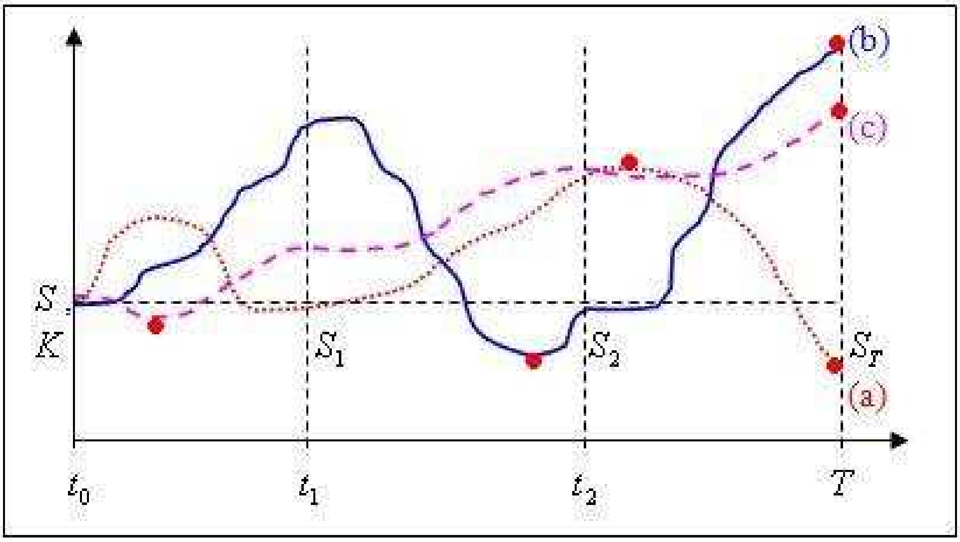

Figure 6.10: You always get the best payout from a lookback. The dots show the highs and lows of each price path.

2. Floating Strike Lookback – The strike is fixed at maturity. At maturity the buyer can "lookback" and get the most favorable strike that maximizes his payoff between the strike and expiry value of the underlying.

Lookbacks are useful if an investor is uncertain when to enter or exit a market. The floating strike lookback will always give an investor the best entry point into a market and the fixed strike lookback the best exit point out of a market — the investor always gets the "better deal". This is shown in Fig. 6.10 where we show three possible price paths. The dots show the highs and lows of each path.

Due to this, lookbacks are expensive. A rule of thumb is that an at-the-money lookback option when issued will be priced at about two times a standard option. Pricing formulas for continuous lookbacks are available (see [Ha 07]). However, due to their high prices, they are made cheaper by monitoring the lookback discretely in time. This means the highs and lows are only monitored on certain dates. This can be monthly of weekly or even quarterly. Discrete lookbacks can only be valued by numerical methods. One can also shorten the lookback period. The option is then called a partial lookback and these options can also be valued by numerical methods only.

Lookbacks are not very common in the FX markets.

6.10 Asian Options

Asian options are options based on some average of the underlying asset price. Generally, an Asian option is an option whose payoff depends on the average price of the underlying asset during a pre-specified period within the option's lifetime, and a pre-specified observation frequency.

They are very popular hedging instruments for corporates.

6.10.1 Uses in the FX Markets

- Protection against rapid price movements or manipulation in thinly traded underlyings at maturity. They lower the volatility risk.
- Reduction of hedging cost due to the lower option premium compared to vanilla options.
- Adjustment of option payoff to payment structure of the firm. Asians can be used to hedge a stream of (received) payments from offshore.

6.10.2 Fixed Strike Arithmetic Average Options

A fixed strike arithmetic average option is an option where the strike is set on the deal date and the average is taken of the underlying asset to determine the payoff. These options are also known as *Asian Out* options. The payoff is given by

$$V = \max\left[0, \phi\left(S_A - K\right)\right]$$

with K the strike and S_A the discretely sampled average of the asset price defined by

$$S_A = \frac{1}{n}\sum_{i=1}^{n} S_i.$$

Here we take the average over n intervals where S_i is the asset's spot price at every interval i. Also, ϕ is a binary variable where $\phi = 1$ for a call and $\phi = -1$ for a put.

There is currently no known closed form solution to the arithmetic average option problem. The problem is that investors usually measure the average at discrete dates. Vorst proposed an approximation where he adjusts the strike price of the option [KV 90]. His method is quite general and take the actual dates as inputs. This means the measuring of the average can start immediately or only a couple of days before expiry.

We define the option as follows: let t_0 be the deal date (annualised time, usually zero) and we have n averaging dates over which the average will be taken. We also let m be the number of averaging dates already passed if we have entered an averaging period. We can then define the times on which the averaging is done as t_i $(i = m+1, m+2, m+3, \ldots, n)$. This is explained by looking at the time line below. When the deal is done at first, the averaging starts at some point after the deal date. Let t_1 be the first averaging date. We show this in Fig. 6.11. However, as

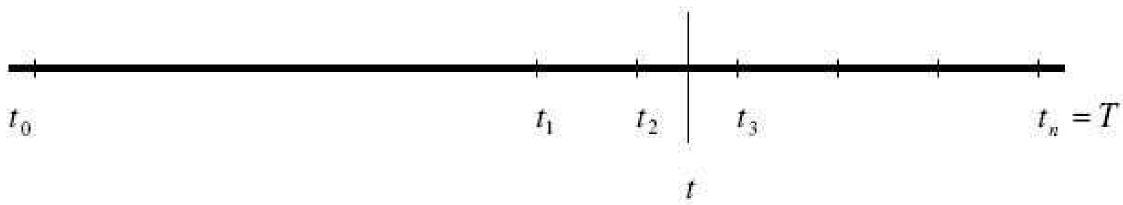

Figure 6.11: Time line for an Asian option.

time passes by the valuation date t can be after the averaging has started such that $t > t_1$. We then count how many averaging dates have we passed and put that equal to m. There are thus $n - m$ averaging dates left.

Vorst's solution holds in the Black-Sholes environment and he found the value of the option to be at any time t

$$V_A(t) = \frac{n-m}{n} e^{-r_d \tau} \phi \left(e^{M+V/2} N(\phi x) - K' N(\phi y) \right). \tag{6.24}$$

Now we can define the variables used in Eq. (6.24)

$$\begin{aligned}
x &= \frac{1}{\sqrt{V}} (M - \ln(K') + V) \\
y &= x - \sqrt{V} \\
\tau &= T - t \\
M &= \ln(S(t)) + \frac{1}{n-m} \sum_{i=1+m}^{n} \left(r_d - r_f - \frac{\sigma^2}{2} \right) \tau_i \qquad (6.25) \\
V &= \frac{\sigma^2}{(n-m)^2} \sum_{i=1+m}^{n} \sum_{j=1+m}^{n} \min(\tau_i, \tau_j) \\
\tau_i &= t_i - t_0 \\
K^* &= \frac{n}{n-m} \left(K - \frac{m}{n} S_A(m) \right) \\
K' &= K^* - \left(S(t) A - e^{M+V/2} \right) \\
S_A &= \frac{1}{m} \sum_{i=1}^{m} S(t_i) \\
A &= \frac{1}{n-m} \sum_{i=1+m}^{n} e^{(r_d - r_f)\tau_i}
\end{aligned}$$

Here $S(t)$is the spot price on the valuation date t and $S(t_i)$ is the spot price on the averaging date t_i. We then define $S_A(m)$ as the average of the spot price thus far if we have entered an averaging period and we have passed m averaging dates. As an example, look at the time line. We see $m = 2$ and $S_A(m) = \frac{1}{2}(S(t_1) + S(t_2))$.

Also, $N(\bullet)$ is the cumulative of the normal distribution function. Note the following, if we are outside of the averaging period $m = 0$ and thus $K^* = K$ and $S_A(0) = 0$.

Curran also gave an efficient routine for valuing Asians options [Cu 92].

6.10.3 The Greeks

Using Vorst's approximation, we can determine the Greeks. The Delta is given by

$$\Delta_A(t) = \frac{n-m}{n}\, e^{-r_d\tau}\left(\frac{1}{S(t)} e^{M+V/2}\ \left[N(\phi\, x) -\ N(\phi\, y)\right] + AN(\phi\, y)\right)$$

and the Gamma is given by

$$\begin{aligned}\Gamma_A(t) \quad &= \quad \frac{-(n-m)}{n}\, e^{-r_d\tau} e^{M+V/2}\frac{N'(x)}{\sqrt{V}}\,\phi\left(\frac{1}{S(t)}\left(1 - e^{M+V/2}\right) + \frac{A}{K'}\right) \\ &\times\left(\frac{1}{S(t)} + \frac{1}{K'}\left[A - \frac{1}{S}e^{M+V/2}\right]\right).\end{aligned}$$

Here, $N'(\bullet)$ is the cumulative normal probability function and all other quantities are defined in Eq. 6.26. We calculate the Theta by obtaining the 1 day time decay such that

$$\Theta(t) = V_A(t + 1 \text{ day}) - V_A(t)$$

and the Vega is obtained similarly by adding 1% to the current volatility such that

$$\Lambda(t, \sigma) = V_A(t, \sigma + 1\%\ ) - V_A(t, \sigma)$$

and the Rho is obtained similarly by adding 1% to the current domestic interest rate such that

$$\rho(t, r_d) = V_A(t, r_d + 1\%\ ) - V_A(t, r_d)$$

6.10.4 Example

Let's look at the USDZAR FX rate. In Fig. 6.12 we plot the exchange rate since the beginning of the year. On 25 January you want to buy a call option. The Rand is busy appreciating against the Dollar. On this date the Rand was trading at 7.0428. Let's assume this option expires on Friday 4 March, the volatility is 15%, $r_d = 6\%$ and r_f=2%. A vanilla call cost R0.14993. On 4 March the ZAR rate settled at 6.866. A vanilla call would have been out-the-money and expired worthless.

However, if you considered an *Asian Out* with an average taken from 5 February daily. The cost of this Asian was R0.1066 - cheaper than the vanilla. At expiry, the average was R7.12078 - meaning the Asian expired in-the-money.

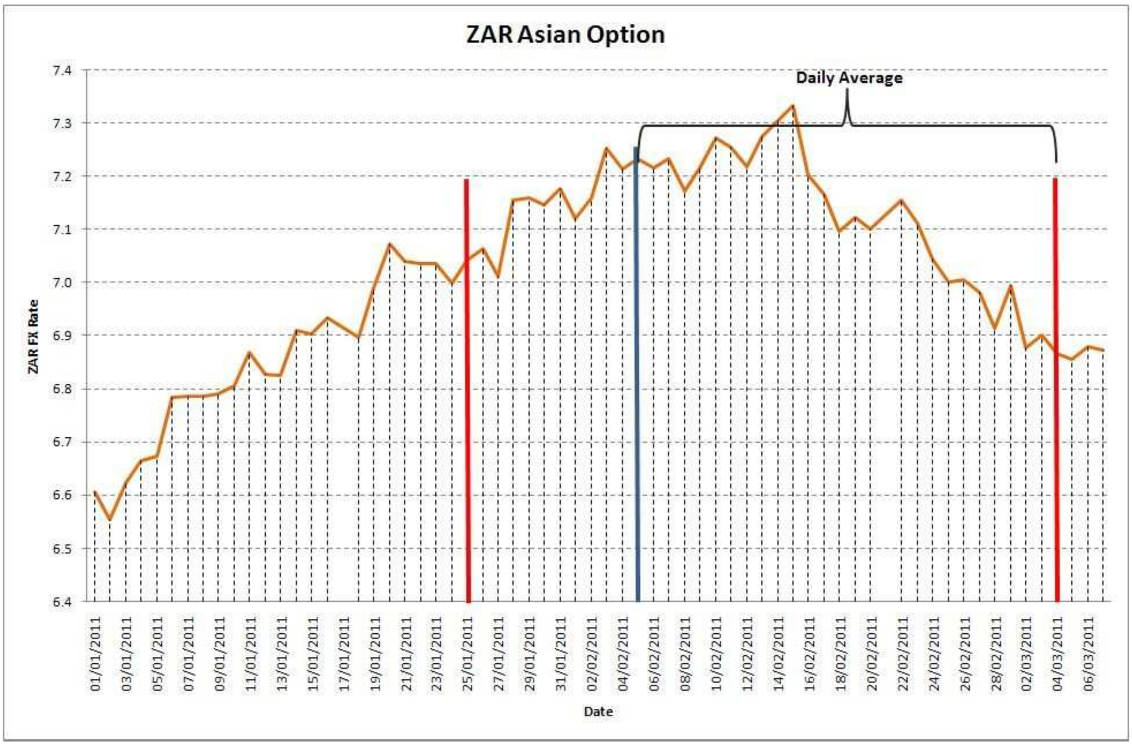

Figure 6.12: Asian option on USDZAR.

6.10.5 Floating Strike Arithmetic Average Options

These options are also known as *Asian In* options. Here, the strike price is determined by an average of the underlying on a set of predetermined observation dates. We then have by

$$V = \max\left[0, \phi\left(S - K_A\right)\right]$$

with S the spot price and K_A the discretely sampled average of the asset price defined by

$$K_A = \frac{1}{n}\sum_{i=1}^{n} S_i.$$

This option is useful in illiquid markets where one needs to hedge a structure like a spread by either buying or selling the Delta. If the volumes are such that the Delta will only be done over a few days, structure the product such that the strike is the average price over a few days.

These options can only be valued by numerical methods.

6.10.6 In-Out Asian Options

An option can be defined whereby the strike is set as the average of the underlying over the first couple of days of the options and the final spot to use is the average

Figure 6.13: In-Out Asian on USDNGN.

over the last couple of days. This option's payoff is given by

$$V = \max\left[0, \phi\left(S_A - K_A\right)\right].$$

Unfortunately, these options can only be priced by numerical methods.

We show an example of this type of option in Fig. 6.13 for an option on USDNGN. Figure 6.13 shows that USDNGN had a big drop from 1 October 2010 until the beginning of November. If you bought a vanilla call on 1 October 2010, the ATM strike would have been 154.5. Let's assume the option expired on 4 March 2011. On that date the FX rate was 154.2. Your vanilla call expired worthless. However, if you bought an In-Out Asian where the strike was the average over the first twenty trading days, and the payout rate the average over the last ten trading days you would have had $K_A = 151.814$ and $S_A = 153.394$ meaning the option would have expired in-the-money.

Chapter 7

Complex Currency Derivatives

These options are also called 'second generation exotic options.'

7.1 Roll Up Puts and Roll Down Calls

These options are also called '*timer*' options.

These structures appeal to investors who feel that they may be early in implementing a bullish or bearish position i.e., they may feel the market will go down or up with sudden reversals later on. An example of such a view is given in Fig. 6.12.

They are combinations of barrier down-and-out calls and up-and-out puts (see §6.3. Roll-ups and -downs are synthetics hedging structures that are put in place on top of existing long or short positions in the underlying — these are overlay structures.

The value of the roll down call can be calculated as follows: let $H_1, H_2, \ldots, H_n$ be a decreasing sequence of positive barrier levels. Similarly, let $K_0, K_1, \ldots, K_n$ be a decreasing sequence of strikes, with $K_i \geq H_i$, $i = 1, 2, \ldots n$. The roll down call is decomposed into down-and-out calls such that[CEG 97]

$$RDC(K_i, H_i) = DOC(K_0, H_1) + \sum_{i=1}^{n} [DOC(K_i, H_{i+1}) - DOC(K_i, H_i)] \qquad (7.1)$$

where $DOC(K_i, H_i)$ is a down-and-out call with strike K_i and barrier level at H_i.

For a roll up put we have an increasing set of barrier levels where $K_i \leq H_i$, $i = 1, 2, \ldots n$ such that

$$RUP(K_i, H_i) = UOP(K_0, H_1) + \sum_{i=1}^{n} [UOP(K_i, H_{i+1}) - UOP(K_i, H_i)]. \qquad (7.2)$$

These structures are depicted in Figs. 7.1.

To explain the mechanics of a roll-up put lets look at an example where there are two roll-up points. We thus have the current market level at K_0 and two levels H_1

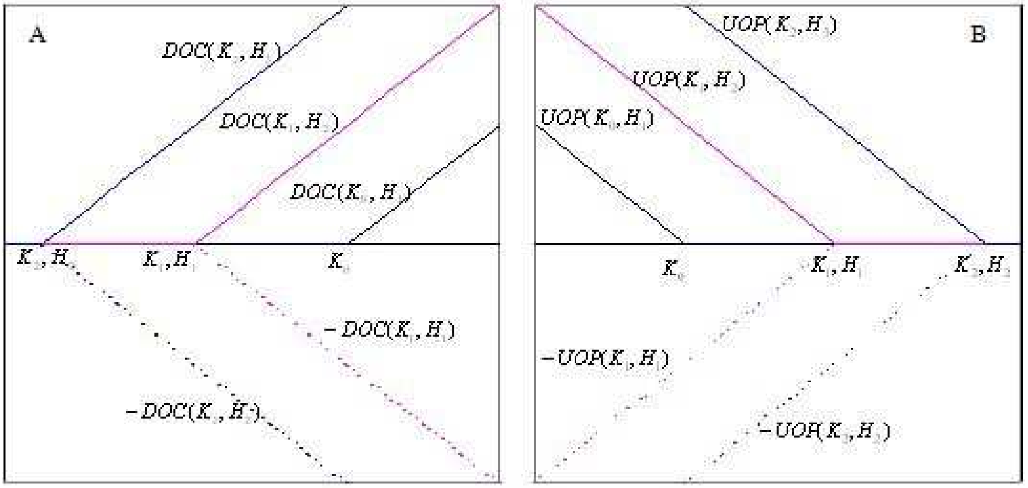

Figure 7.1: A: Roll down call. B: Roll up put.

and H_2 where H_1, $H_2 > K_0$. The value of the roll-up put given by

$$RUP(K_0) = UOP(K_0, H_1) + UOP(K_1, H_2) - UOP(K_1, H_1)$$

where $K_1 = H_1$.

The buyer of this roll-up put option buys an up-and-out-put struck at K_0 which vanishes as soon as the barrier at H_1 is breached. This put is active and behaves just like a vanilla put which means the investor's long position in the underlying is immediately hedged against the market turning down. He/she also buys an up-and-out-put struck at $K_1 = H_1$ which vanishes as soon the barrier at H_2 is breached. He/she then shorts an up-and-out-put struck at K_1 that will vanish once the level H_1 has been breached. Graphically this is be depicted in Fig. 7.2. The thick line is the original long position and the thin lines depict the up-and-out-puts.

The dashed lines in Fig. 7.2 show the possible payoffs that are possible. If H_1 is never breached this structure behaves like an ordinary vanilla call (the put struck at K_0 together with the long position in the underlying gives a synthetic call struck at K_0). This means that the investor is immediately hedged against the market moving below K_0. Also, the long and short puts struck at K_1 cancel one another. The payoff is thus given by

$$S_T + \max(0, K_0 - S_T)$$

where S_T is the value of the underlying at expiry.

If H_1 is breached the following happens: the long put struck at K_0 and the short put struck at K_1 are knocked out. In essence the strike of the synthetic call has thus moved up to K_1 and the investor is hedged such that if the market moves below K_1, he/she will at least receive $K_1 - K_1$ is thus locked in/guaranteed to the holder of the

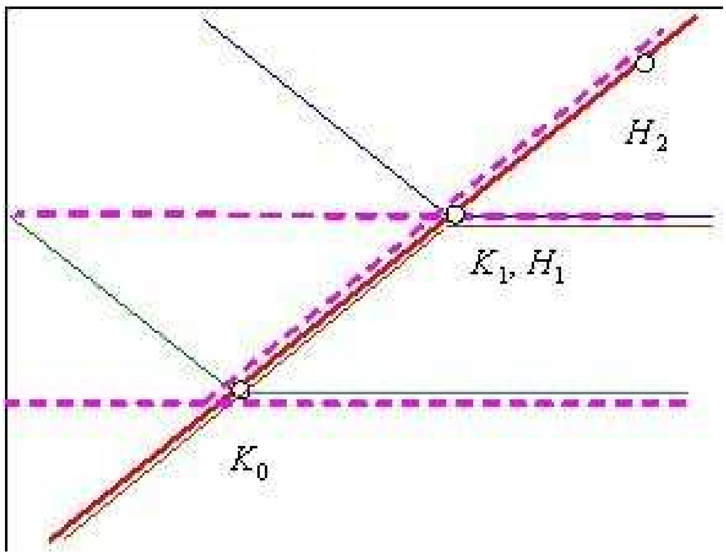

Figure 7.2: A roll up put structure.

underlying long position and the payoff is given by

$$S_T + \max(0, K_1 - S_T).$$

If H_2 is breached the whole synthetic hedging structure vanishes and the investor is left with the original long position with payoff just S_T. To circumvent this from happening it is advised that the last roll-up level be placed so far from the original level at K_0 that the chances of it being breached is nearly zero. On the other hand, the investor might decide that once the last level is breached, the original long position is so far in-the-money that he/she might risk market conditions and a subsequent reversal.

As an example we look at Fig. 6.12 and the example discussed in §6.10.4. An investor wants to buy a put on 25 January 2011. The cost of the vanilla put is R0.12184. He, however feels it is early in implementing this and add two barriers at 103% and 106% (reset levels). We then have $S = K = 7.0428$, $H_1 = 7.2541$ and $H_2 = 7.4654$. This *timer put* cost R0.15506. It should be more expensive than the vanilla due to the opportunity to get a better strike. Now, on 4 March, the ZAR level was at 6.866. The vanilla would have been in-the-money by $7.0428 - 6.866 = 0.1768$ — a hansom profit. However, on 10 Feb 2011 the ZAR rate moved to 7.2726 breaching the first barrier level. It means that the strike of the put is reset at this level. The second reset level is never breached and the payout would have been $7.2541 - 6.866 = 0.3881$ — an even more hansom profit!

7.1.1 Ladder Options

There has recently been a lot of interest in derivatives which guarantee the return of capital invested or allow the purchaser to periodically lock in gains. One way to achieve this is by using ladder options. Ladder options are similar to a lookback options.

The buyer of a floating strike lookback call buys an option at a certain strike, say K_0. Now, as soon as a more advantageous level is registered, the strike is adjusted to this level, say K_1 where $K_0 < K_1$. This means that the difference between the new strike and original strike level, $(K_1 - K_0)$, is locked in and guaranteed to the holder of the lookback.

It is, however, also possible to do this only if the difference with the original strike exceeds a certain minimum predetermined value. This is then a ladder option. This option is similar to the Cliquet. The Cliquet is linked to time where a Ladder is linked to levels of the underlying.

A ladder option has a set of Npredetermined levels (rungs) $L_i, i = 1, 2, \ldots, N$ where, every time the market breaches a certain ladder level, the intrinsic value at that level is locked in and will be paid out to the holder at expiry. A European ladder will have the following payoff

$$LO = \max\left[\phi\left(S_T - K\right), \max\left\{\phi\left(L_i - K\right), 0\right\}, 0\right]$$

where $LO$is the terminal value of the ladder option, S_T is the value of the underlying quantity at expiry time T, K is the initial strike price, $\phi = 1$ for a call and $\phi = -1$ for a put and $L_i$is the i^{th} ladder level reached in the life of a option. This function shows that the pay-off for a ladder option is the greater of a plain vanilla call option with strike K and the highest ladder level reached, or zero.

As an example, look at Fig. 7.3. This shows paths (a), (b), (c) and (d) that are different possibilities of the underlying over time. L_1 and L_2 are two rungs of a two-rung ladder call option. The payoff for path (a) at expiry is the terminal spot value minus the strike $S_T - K$ — the same as for a vanilla call; the payoff for path (b) at expiry will be $L_2 - K$; the payoff for path (c) at expiry will be $L_1 - K$; and the payoff for path (d) is zero. The payoff is thus similar to that of a lookback option with discrete lock-in levels. In the limit when the number of rungs,$N \to \infty$ we get the vanilla fixed strike lookback option.

A ladder is a total synthetic instrument that gives synthetic exposure to the market. The exposure is obtained through the vanilla call and the gains are locked in by using barrier options.

To synthesise a ladder call with initial strike K_0, we let $L_0, L_1, L_2, \ldots, L_N$ be a predetermined increasing set of barrier levels such that $K_0 = L_0$, then

$$LC(K_0, L_i) = C(K_0) + \sum_{i=1}^{N} [UIP(L_i, L_i) - UIP(L_{i-1}, L_i)];$$

and for a ladder put we have a predetermined decreasing set of barrier levels

$$L_0, L_1, L_2, \ldots, L_N \text{ where } K_0 = L_0$$

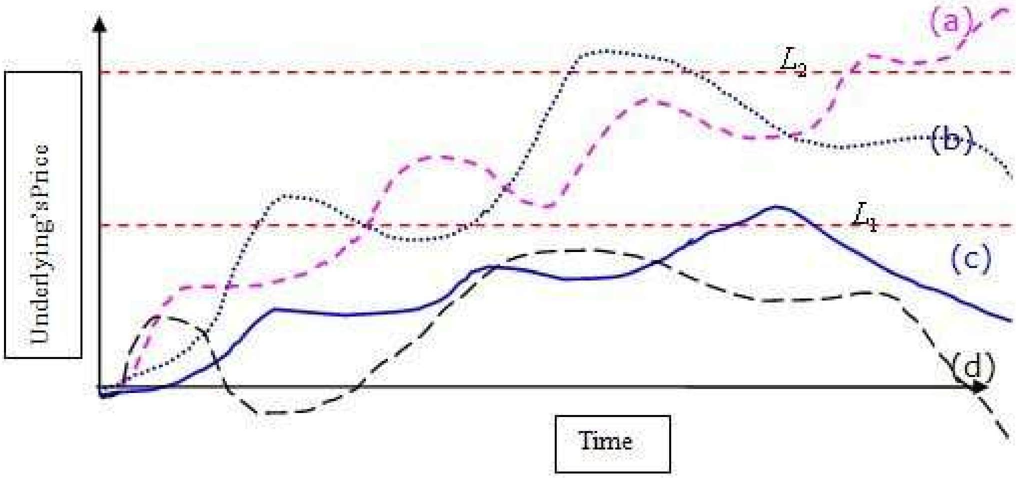

Figure 7.3: Possible price paths showing the dynamics of a Ladder option.

and thus

$$LP(K_0, L_i) = P(K_0) + \sum_{i=1}^{N} [DIC(L_i, L_i) - DIC(L_{i-1}, L_i)]$$

where $C(K_0)$ and $P(K_0)$ is a vanilla call and vanilla put struck at K_0 respectively and $UIP(x, y)$ is an up-and-in-put with strike x and barrier level y and $DIC(x, y)$ is a down-and-in-call with strike x and barrier level y.

This shows that ladder options are packaged from ordinary vanilla and barrier options. A ladder option is cheaper than a lookback but more expensive than the corresponding vanilla option. The reason for this is that a ladder consists of a vanilla option and two barrier options. Ladders are hedged by using the strategies given for barrier options.

To explain the mechanics of a ladder lets look at an example where there is only one rung at level L_1 such that $L_1 > K$and $L_0 = K$. The value of this ladder call is given by

$$LC(K_, L_1) = C(K_) + UIP(L_1, L_1) - UIP(L_0, L_1).$$

The buyer of a ladder call option buys a vanilla call, struck at K; he/she also buys an up-and-in-put (UIP), struck at L_1, that will only be activated once the underlying has breached the level at L_1. He/she then shorts an up-and-in-put struck at K that will be activated once the level L_1 has been breached. Graphically this is be depicted in Fig. 7.4.

The dashed line in Fig. 7.4 shows the payoffs that are possible. If L_1 is never breached this structure behaves like an ordinary vanilla call (the puts are not activated and do not exist) which means that if the expiry value of the underlying is below K,

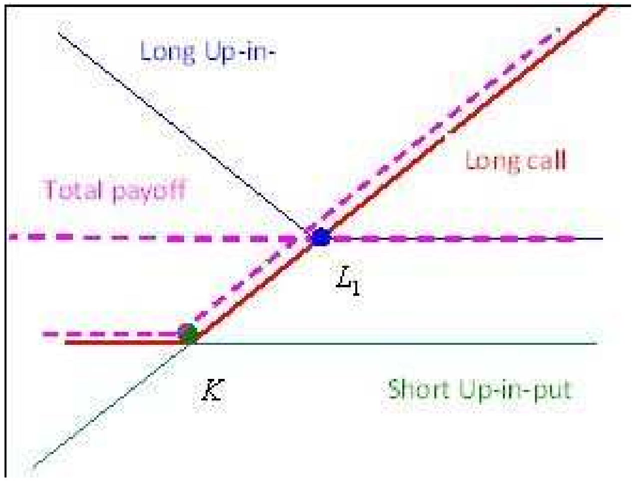

Figure 7.4: The mechanics of a Ladder call.

the payoff of the whole structure is zero. If L_1 is breached, $L_1 - K$ is guaranteed to the holder and, in essence, the vanilla call's strike level is increased to L_1. This is explained as follows:
the two up-and-in-puts are activated and are now vanilla put options. The short put and long call (both struck at K) together form a synthetic long position in a forward contract where $F = C(K) - P(K)$. The payoff for the ladder call now is just $LC = F + P(L_1)$. At expiry, if the underlying's price ends above the barrier, $P(L_1)$ is out-the-money and the payoff is just the payoff for F. Actually $P(K)$ is also out-the-money and the total payoff is just that of the vanilla call i.e., $S_T - K$. If the price ends below the barrier, the total payoff is just the payoff for LO. The total payoff is given by

$$C(K) - P(K) + P(L_1) = \max(0, S_T - K) - \max(0, K - S_T) + \max(0, L_1 - S_T).$$

This will always give a payoff of $L_1 - K$, for any $S_T \leq L_1$ and $S_T - K$ if $S_T \geq L_1$.

As an example let's look at the same one we discussed previously under *timers* and in §6.10.4. An investor wants to buy a call on 25 January 2011. We have the expiry as Friday 4 March, the volatility is 15%, $r_d = 6\%$ and r_f=2%. The cost of the vanilla call is R0.14993. He, however feels it is early in implementing this and add two rungs at 103% and 106%. We then have $S = K = 7.0428$, $H_1 = 7.2541$ and $H_2 = 7.4654$. This Ladder cost R0.20312. It should be more expensive than the vanilla due to the opportunity to lock in gains. Now, on 4 March, the ZAR level was at 6.866. The vanilla would have been out-the-money. However, on 10 Feb 2011 the ZAR rate moved to 7.2726 breaching the first rung and locking in an amount of $7.2541 - 7.0428 = 0.21128$. The second rung is never breached. However, even though the ZAR was at a level of 6.866 on 4 March 2011, the payout of the Ladder was 0.21128.

7.2 Variance Swaps

Volatility is a measure of the risk or uncertainty and it has an important role in the financial markets. Volatility is defined as the variation of an asset's returns — it indicates the range of a return's movement. Large values of volatility mean that returns fluctuate in a wide range.

Derivatives market professionals know that managing volatility is central to hedging the risk in an options portfolio. Although the Black-Scholes-Merton framework for hedging options is both well-established and well-understood, spectacular losses in volatility trading have been dealt to broker-dealers and hedge funds in the past.

Alas, the tools for managing volatility risk are few. But a relatively new product-the variance swap-offers investors a straightforward vehicle for achieving long or short exposure to market volatility. Although it's called a swap contract, it is fundamentally an option-based product with properties similar to those of options. The product consequently represents a significant addition to the overall landscape of volatility-driven instruments and can fill a useful role for investors seeking optionality in one form or another.

Black & Scholes defined volatility as the amount of variability in the returns of the underlying asset. They determined, what today is known as, the historical volatility and used that as a proxy for the expected or implied volatility in the future. Since then the study of implied volatility has become a central preoccupation for both academics and practitioners.

In Fig. 7.5 we show the 3 month volatilities for USDZAR and USDKES. It is clear that volatility is not constant. Volatility is, however, statistically persistent, i.e., volatility trends: if it is volatile today, it should continue to be volatile tomorrow. This is also known as volatility clustering and can be seen in Fig. 7.6. This is a plot of the logarithmic returns of USDKES since June 1995 using daily data.

7.2.1 How it Works

The variance swap is a contract in which two parties agree to exchange cash flows based on the measured variance of a specified underlying asset during a certain time period. On the trade date, the two parties agree on the strike price of the contract (the reference level against which cash flows are exchanged), as well as the number of units in the transaction.

Variance swaps or variance futures (also called variance contracts) are derivative instruments offering pure exposure to daily realised future variance. Variance is the square of volatility (usually denoted by the Greek symbol σ^2). At expiration, the swap buyer receives a payoff equal to the difference between the annualised variance of logarithmic stock returns and the swap rate fixed at the outset of the contract. The swap rate (or delivery price) can be seen as the fixed leg of the swap and is chosen such that the contract has zero present value.

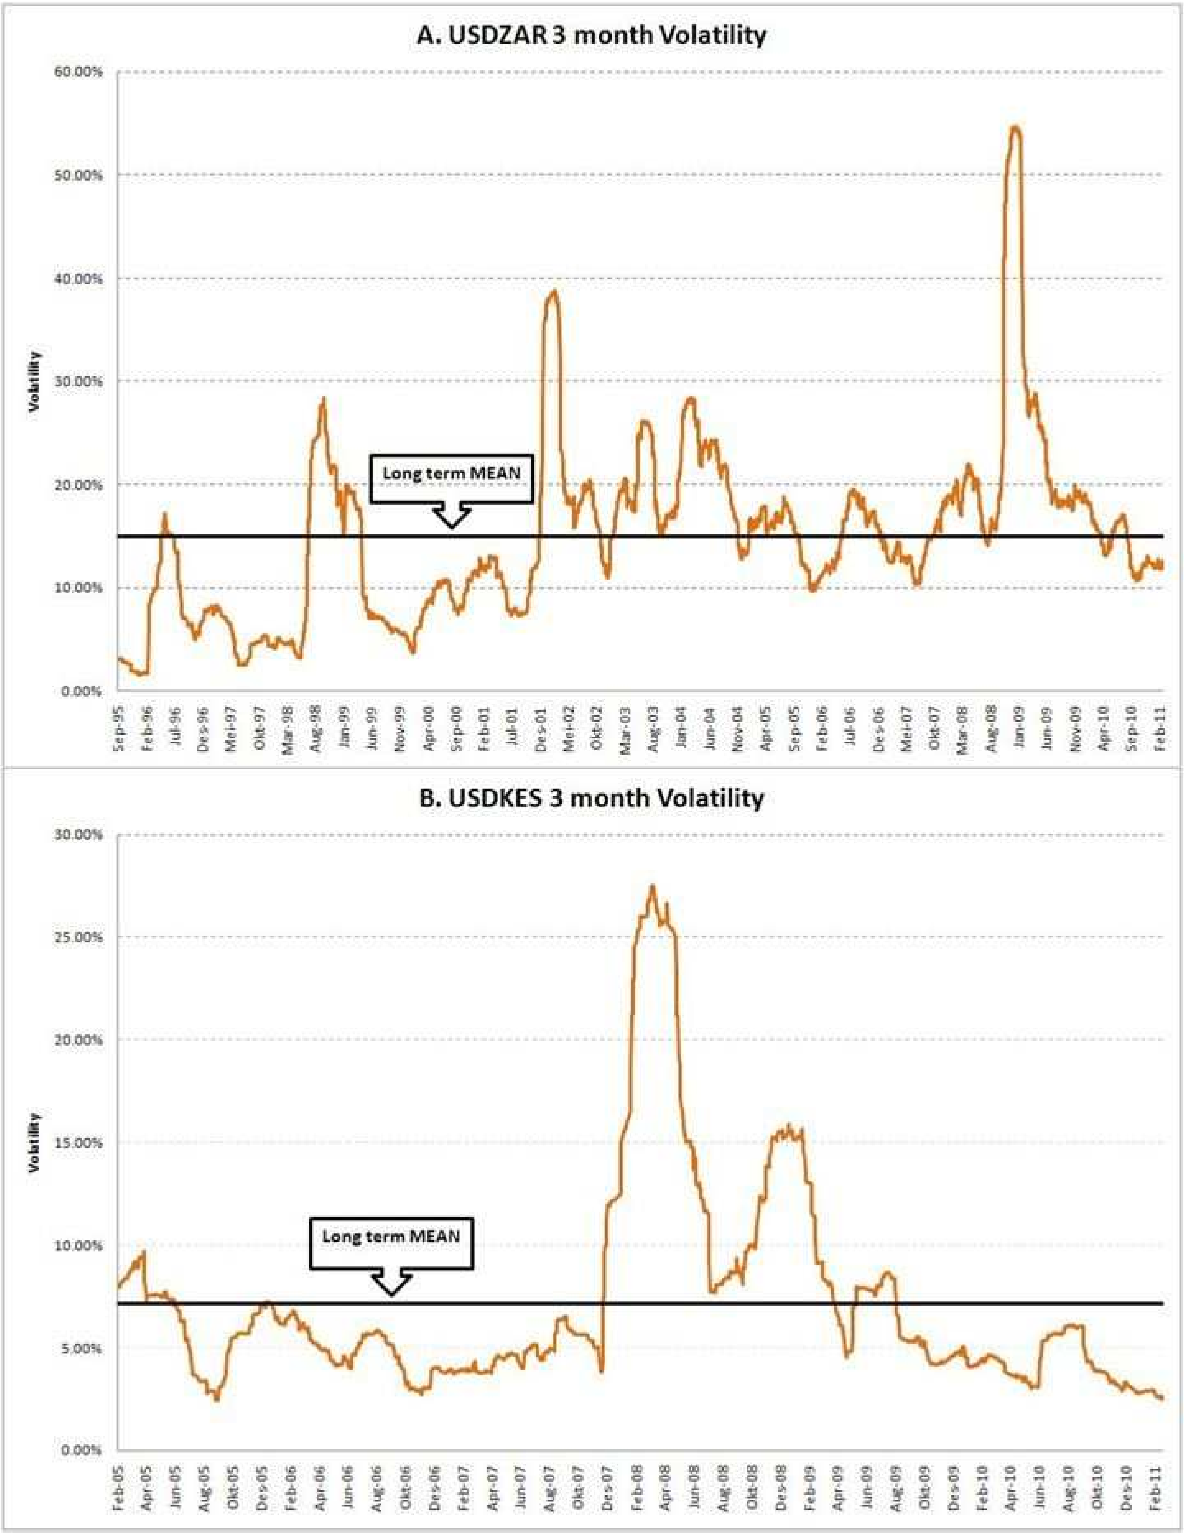

Figure 7.5: 3 Month Historical volatility for A: USDZAR and B: USDKES.

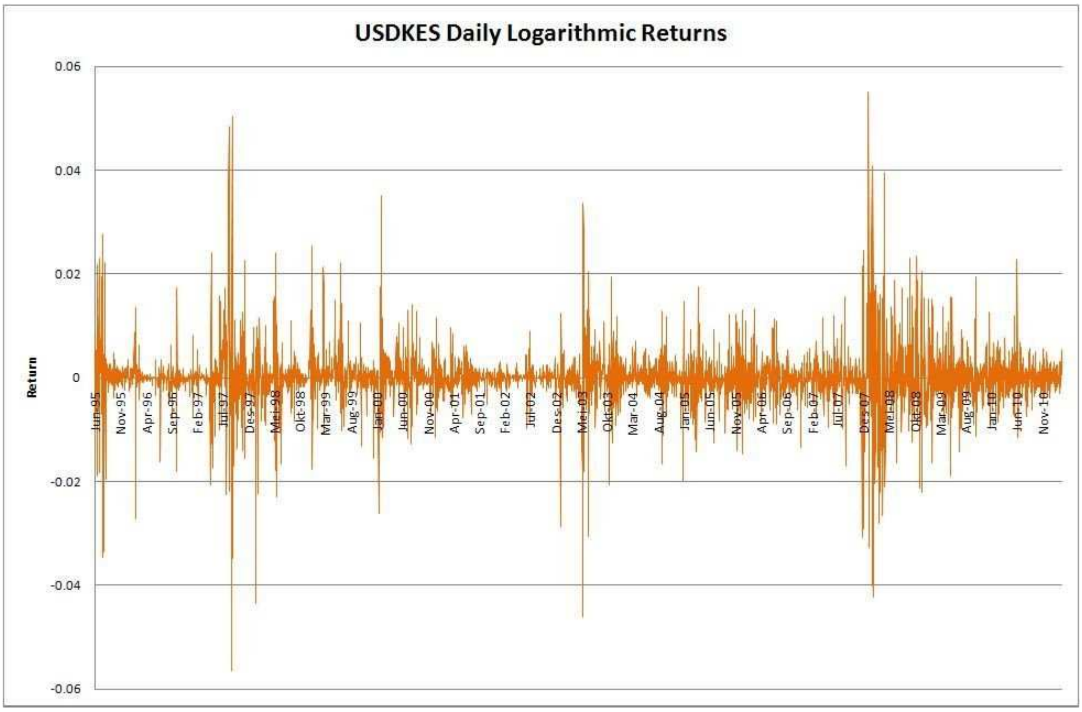

Figure 7.6: Logarithmic returns for USDKES.

In short, a variance swap is not really a swap at all but a forward contract on realised annualised variance. A long position's payoff at expiration is equal to [DD 99]

$$VNA\ \left[\sigma_R^2 - K_{var}\right] \tag{7.3}$$

where VNA is the Variance Notional Amount, $\sigma_R^2$is the annualised non-centered realised variance of the daily logarithmic returns on the index level and K_{var} is the delivery price. Note that VNA is the notional amount of the swap in Rand per annualised variance point. The holder of a variance swap at expiration receives VNA currency for every point by which the underlying's realised variance has exceeded the variance delivery price. The "swap" is diagrammatically depicted in Fig. 7.7.

A capped variance swap is one where the realised variance is capped at a predefined level. From Eq. (7.3) we then have

$$VNA\ \left[\min\left(cap, \sigma_R^2\right) - K_{var}\right].$$

The realised variance is defined by

$$\sigma_R^2 = \frac{252}{n}\sum_{i=1}^{n}\left[\ln\left(\frac{S_i}{S_{i-1}}\right)\right]^2 \tag{7.4}$$

with S_i the index level and n the number of data points used to calculate the variance].

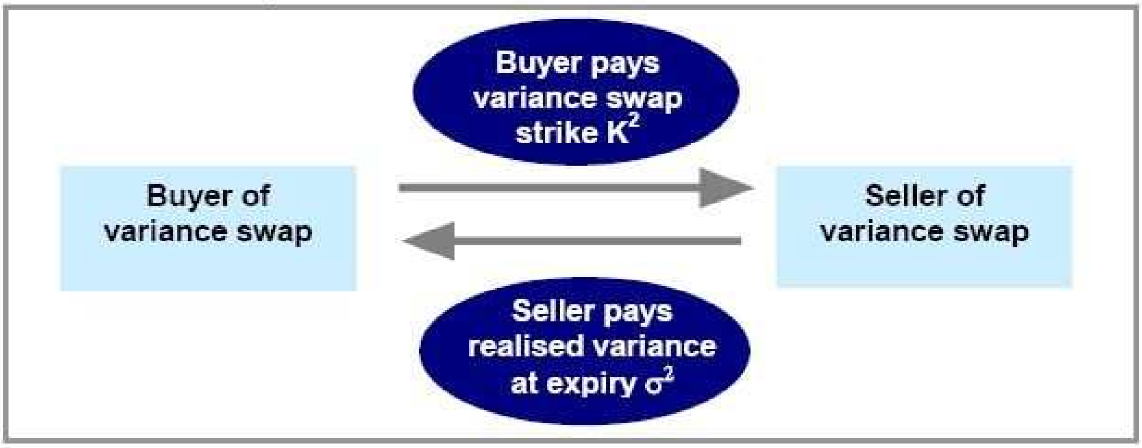

Figure 7.7: The variance swap.

If we scrutinised Eq. (7.4), we see that the mean logarithmic return is dropped if we compare this equation with the mathematically correct equation for variance. Why? Firstly, its impact on the realised variance is negligible. Secondly, this omission has the benefit of making the payoff perfectly additive[1]. Another reason for ditching the mean return is that it makes the estimation of variance closer to what would affect a trader's profit and loss. The fourth reason is that zero-mean volatilities/variances are better at forecasting future volatilities. Lastly *Figlewski* argues that, since volatility is measured in terms of deviations from the mean return, an inaccurate estimate of the mean will reduce the accuracy of the volatility calculation [Fi 94]. This is especially true for short time series like 1 to 3 months (which are the time frames used by most traders to estimate volatilities).

Note, historical volatility is usually taken as the standard deviation whilst above we talk about the variance. The question is: why is standard deviation rather than variance often a more useful measure of variability? While the variance (which is the square of the standard deviation) is mathematically the "more natural" measure of deviation, many people have a better "gut" feel for the standard deviation because it has the same dimensional units as the measurements being analyzed. Variance is interesting to scientists, because it has useful mathematical properties (not offered by standard deviation), such as perfect additivity (crucial in variance swap instrument development). However, volatility is directly proportional variance.

[1] Suppose we have the following return series: 1%, 1%, -1%, -1%. Using variance with the sample mean, the variance over the first two observations is 0, and the variance over the last two observations is zero. However the total variance with sample mean over all 4 periods is 0.01333%, which is clearly not zero. If we assumed the sample mean to be zero, we get perfect variance additivity.

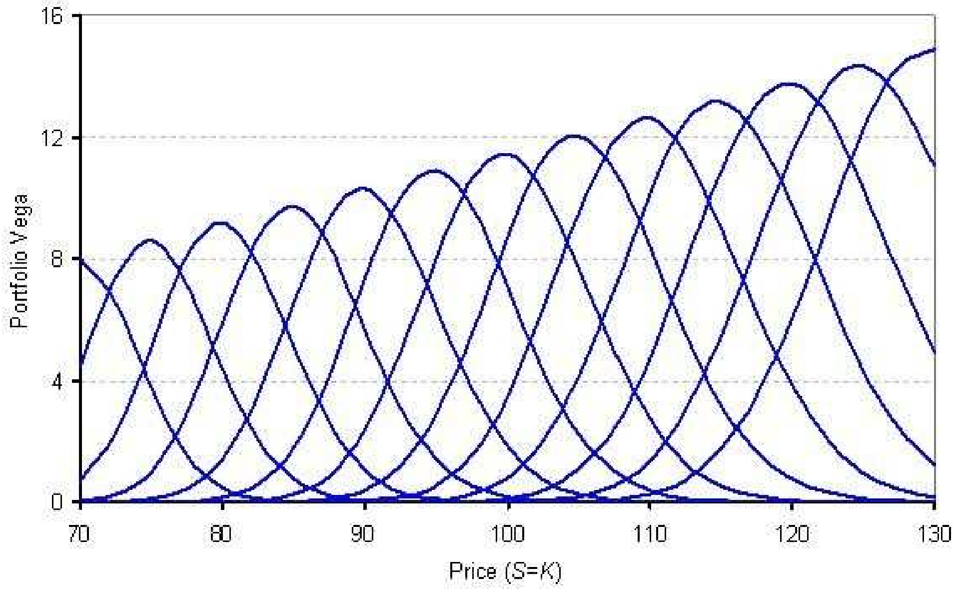

Figure 7.8: Option sensitivity to volatility (Vega) per at the money ($S = K$) index level for 13 options with strikes ranging from 70% to 130% in increments of 5%.

7.2.2 Variance Swap Pricing in Theory

The price of the variance swap per variance notional at the start of the contract is the delivery variance. So how do we find the delivery price such that the swap is immune to the underlying index level? We start by deriving a static hedge.

Consider the following ingenious argument. We know that the sensitivity of an option to volatility, Vega, is centered (like the Gaussian bell curve) around the strike price and will thus change daily according to changes in the underlying's level. We show the Vega for different strikes in Fig. 7.8. Also evident: the higher the strike, the larger the Vega. If we can create a portfolio of options with a constant Vega, we will be immune to changes in the stock or index level — a static hedge. From Fig. 7.8 we see that the contribution of low-strike options to the aggregate Vega is small compared to high-strike options. Therefore, a natural idea is to increase the weights of low-strike options and decrease the weights of high-strike options. A sensible first guess is to weight each option Vega with the corresponding inverse strike. By induction, it turns out that the Vega is constant for a portfolio of options inversely weighted by the square of their strikes — shown in Fig. 7.9. This hedge is independent of the stock level and time.

From the previous argument we deduce that the fair price of a variance swap[2]

[2]The fair price of a variance swap, given that it is the price of variance going forward, it is also referred to as the fair value of future variance. This is an intuitive rational for the forward factor, $e^{(r_d - r_f)T}$ in equation (7.5).

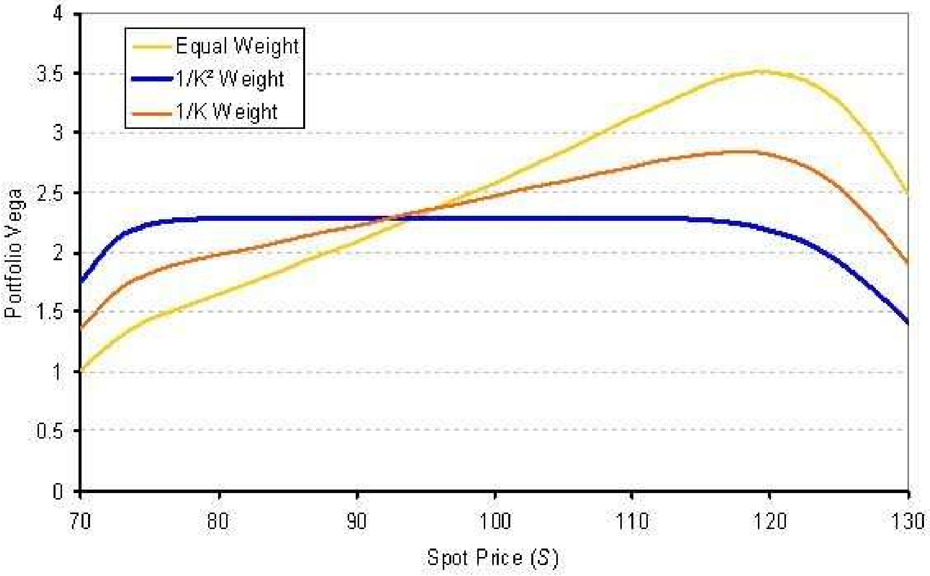

Figure 7.9: A portfolio of 13 options with Vega equally weighted, inversely weighted by their strikes and also inversely weighted by the square of their strikes.

is the value of a static option portfolio, including long positions in out-the-money (OTM) options, for all strikes from 0 to infinity. The weight of every option in this portfolio is the inverse square of its strike. The theoretical fair delivery value at time $t = 0$, for a variance swap maturing at time $t = T$, is then given by [DD 99, BS 05]

$$\begin{aligned} K_{var} &= \frac{2}{T}\Bigg[(r_d - r_f)\,T - \left(\frac{S}{F}\,e^{(r_d-r_f)T} - 1\right) - \ln\left(\frac{S}{F}\right) \\ &\quad + e^{(r_d-r_f)T}\left(\int_0^F \frac{P(K)}{K^2}dK + \int_F^\infty \frac{C(K)}{K^2}dK\right)\Bigg] \end{aligned} \tag{7.5}$$

Here F is the approximate at-the-money forward stock index level with risk-free interest rate r and dividend yield d. F marks the boundary between the liquid put options $P(K)$ and call options $C(K)$ of the same expiry time T with strikes K. K_{var} is the fair variance and we see from Eq. (7.5) it is defined in terms of an indefinite integral expression.

7.2.3 Pricing in Practice

In practice, integrals can be approximated by summations. *Demeterfi* and *Derman et. al.* showed that the discrete version of Eq. (7.5) is given by [DD 99]

$$K_{var} = \frac{2}{T}\left[A + e^{(r_d - r_f)T}\left(\int_0^F \frac{P(K)}{K^2}dK + \int_F^\infty \frac{C(K)}{K^2}dK\right)\right]$$

where we need to integrate over the whole strike range. The discrete version is given by

$$K_{var} = \frac{2e^{(r_d - r_f)T}}{T}\left(\sum_{i=1}^{n=F}\frac{P_i(K_i)}{K_i^2}\,\Delta K_i + \sum_{i=n+1}^{\infty}\frac{C_i(K_i)}{K_i^2}\,\Delta K_i\right) \tag{7.6}$$

with

$$A = (r_d - r_f)\,T - \left(\frac{S}{F}\,e^{(r_d - r_f)T} - 1\right) - \ln\left(\frac{S}{F}\right).$$

The fair delivery variance in Eq. (7.6) approaches the theoretical fair delivery variance, in (7.5), from below, and converges in the limit as the strike increments approach zero and the number of options tend to infinity.

As with any numerical estimation of an indefinite integral expression such as Eq. (7.5), the optimum strip width, and integrand range limits, for good accuracy must be established. This brings two issues to the fore

- The strikes trade in standard fix spacings[3] ΔK that are not infinitesimal small. This introduces a discretisation error[4] in the approximation of the fair variance.
- Only a limited number of option strikes are available. The strike range is finite. This introduces a truncation error in the approximation of the fair variance.

Derman et al addresses the discretisation problem by numerically approximating the fair delivery variance using the fact that the fair variance can be replicated by a log contract of the form

$$f(F_T) = \frac{2}{T}\left[\frac{F_T}{F_0} - \log\frac{F_T}{F_0} - 1\right] \tag{7.7}$$

Here $f(F_T)$ is the maturity T payoff from the long position in the forward price F, and short position in the log contract. At the inception of the swap, the log contract in (7.7), is equivalent to the fair delivery variance defined by (7.5). Since there is no log

[3]For example for index options, the strike spacing are 50 index points

[4]We implicitly use the term error, here and not uncertainty. This is done because an error implies that

the theory /true value is known, whereas uncertainty do not.

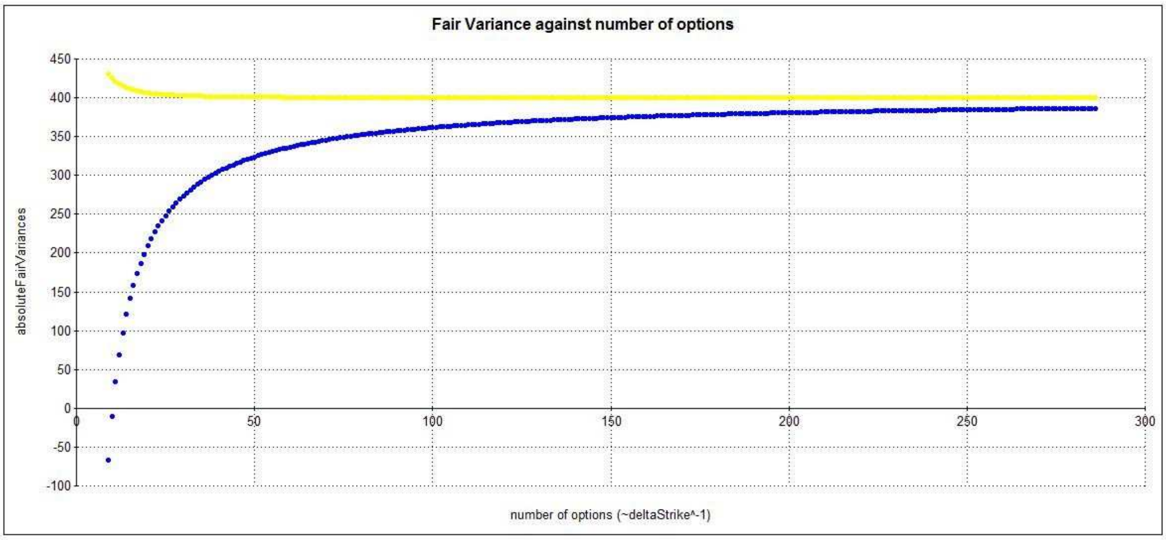

Figure 7.10: Fair delivery variances, Eq. (7.5) (blue) versus Eq. (7.7) (yellow) as a function of the number of options.

contract traded, the log contract is approximated using piecewise linearly weighted options.

It is the piecewise linear approximation (really super-replication) to equation (7.7) that gives the option weightings such that convergence to the accurate delivery variance is obtained quicker using an minimal amount of strike spacing, or number of options (See Fig. 7.10).

The delivery variance, using a limited number of options, according to *Derman et al*, is then

$$K_{var}(0,T) = \frac{2}{T}\left(\sum_{i=1}^{n=F} w_{iP} P_i(K_i) + \sum_{i=n}^{\infty} w_{iC} C_i(K_i)\right) \tag{7.8}$$

with piecewise linear recurring option weightings of

$$w_{iP} = \frac{f(K_{i+1}) - f(K_i)}{K_i - K_{i+1}} - \sum_{j=0}^{i-1} w_{jP} \quad ; \quad w_{iC} = \frac{f(K_{i+1}) - f(K_i)}{K_{i+1} - K_{i+1}} - \sum_{j=0}^{i-1} w_{jC}$$

The *Derman et al.* delivery variance in (7.8) '*super replicates*' the theoretical fair delivery variance given in (7.5), but only in the limited strike range from lowest to highest strike. Moreover, the Vega is constant only within the chosen strike limits of the underlying.

7.2.4 Limit Tests

The *Derman* weighting scheme for quick strike convergence with strike spacing, is used to look at the sensitivity of the option strike range on the fair variance accuracy. It is clear that the fair variance accuracy is more sensitive to the strike range, than the strike spacings. In other words, the accuracy of the fair variance, is more prone to the truncation errors. We also note that

- For a small number of symmetrical options (<20) and the lower the put bound strike, the more over-estimated the fair variance. This over-estimation is less pronounced the higher the number of symmetrical options.

- For a low put bound strike (<40%) and a higher number of symmetrical options, the more underestimated the fair variance. This under-estimation is considerably less pronounced the higher the left wing bound.

Moreover, the greater the strike range, the more the approximated fair variance approaches the theoretical fair variance from below. Truncation thus underestimates the fair variance.

Also, the higher the number of options, or the smaller the strike spacing, the more the approximate variance approaches the fair variance from above. Discretisation overestimates the fair variance. This is consistent with the findings of *Jiang and Tian*[5] [JT 07]

From all of the above we conclude that, the strike spacing, discretisation, and the strike range, truncation errors, are a trade-off between over and under estimation of the fair variance. Finding the strike spacing, and strike range optimum pair that minimises the fair variance accuracy, is a three dimensional problem.

7.2.5 Volatility Indices

7.2.6 VIX

VIX is the ticker symbol for the Chicago Board Options Exchange Market Volatility Index, a popular measure of the implied volatility of S&P 500 index options. Often referred to as the *fear index* or the *fear gauge*, it represents one measure of the market's expectation of stock market volatility over the next 30 day period. The VIX was introduced during 1993.

On September 22, 2003, the Chicago Board Options Exchange made some major changes to the way the implied volatility index (VIX) is constructed. It is now based on the more liquid S&P 500 index options, instead of the S&P 100. Even more

[5] In fact, *Jiang and Tian* found that: "When volatility is low (high), the underlying index level is less (more) likely to move beyond the range of option strikes, so that in this case, the discretisation (truncation) dominates the truncation error. In order to quantify the error involved with approximating the fair variance. It is important to know the direction of the uncertainty"

important is a change in the methodology for constructing the index. The old index was a weighted average of implied volatilities inverted from the Black-Scholes option pricing model. The new index eliminates this dependence on a specific model and uses a model-free approach similar to the one proposed in Britten-Jones and Neuberger (2000).

The Old VIX

The old VIX index is based on the Black-Scholes implied volatility of S&P 100 options. To construct the old VIX, two puts and two calls for strikes immediately above and below the current index are chosen. Near maturities (greater than eight days) and second nearby maturities are chosen to achieve a complete set of eight options. By inverting the Black-Scholes pricing formula using current market prices, an implied volatility is found for each of the eight options. These volatilities are then averaged, first the puts and the calls, then the high and low strikes. Finally, an interpolation between maturities is done to compute a 30 calendar day (22 trading day) implied volatility.

Because the Black-Scholes model assumes the index follows a geometric Brownian motion with constant volatility, when in fact it does not, the old VIX will only approximate the true risk-neutral implied volatility over the coming month. In reality the price process is likely more complicated than geometric Brownian motion. Limiting it to a very specific form and deducing an implied volatility from market prices may lead to substantial error in the estimation. Since the S&P 100 index options are American, an approximation is involved to compute the implied volatility.

The New VIX

The VIX is the square root of the par *variance swap* rate for a 30 day term initiated today. Note that the VIX is the volatility of a variance swap and not that of a volatility swap (volatility being the square root of variance). A variance swap can be perfectly statically replicated through vanilla puts and calls whereas a volatility swap requires dynamic hedging. The VIX is the square-root of the risk neutral expectation of the S&P 500 variance over the next 30 calendar days. The VIX is quoted as an annualized variance.

Note the following: The VIX is quoted in percentage points and translates, roughly, to the expected movement in the S&P 500 index over the next 30-day period, which is then annualized. For example, if the VIX is 15, this represents an expected annualized change of 15% over the next 30 days; thus one can infer that the index option markets expect the S&P 500 to move up or down $15\%/\sqrt{12} = 4.33\%$ over the next 30-day period. That is, index options are priced with the assumption of a 68% likelihood (one standard deviation) that the magnitude of the S&P 500's 30-day return will be less than 4.33% (up or down).

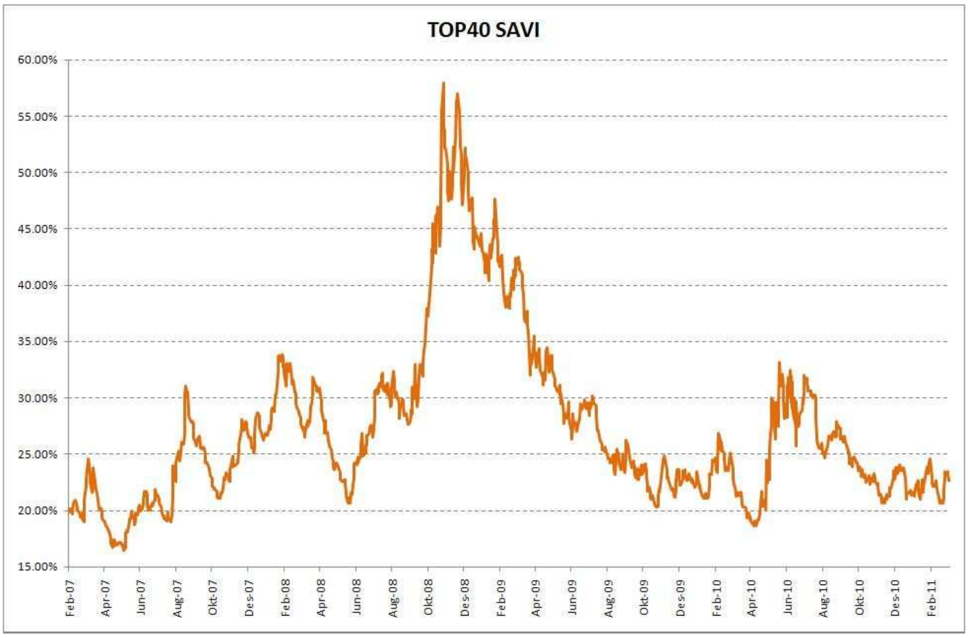

Figure 7.11: The SAVI fear guage.

One can trade futures on the VIX.

7.2.7 SAVI

There are currently volatility indices on the TOP40 index, the USDZAR exchange rate and white maize.

SAVI

The old SAVI was launched, in 2007, as an index designed to measure the market's expectation of the 3-month volatility. The SAVI is based on the FTSE/JSE Top40 index level and it is determined using the at-the-money volatilities. Since it is well documented that there exist a negative correlation between the underlying index level and its volatility, the SAVI can be though of as a "fear" gauge

The old SAVI was calculated on a daily basis, via polling the market. The polled at-the-money volatilities are then used to calculate the 3-month at-the-money volatility. The average 3-month at-the-money volatility as determined from the polled volatilities, were then published as the SAVI.

The SAVI was updated two years later, in 2009, to reflect a new way of measuring the expected 3-month volatility. The new SAVI is also based on the FTSE/JSE Top40 Index, but it is not only determined using the at-the-money volatilities but also using the volatility skew. Given that the volatility skew is the market's expectation of a crash, the new SAVI can be thought of as a more efficient "fear" gauge, since it incorporates a market crash protection volatility premium. The "new" SAVI is thus similar to the new VIX. The SAVI is plotted in Fig. 7.11

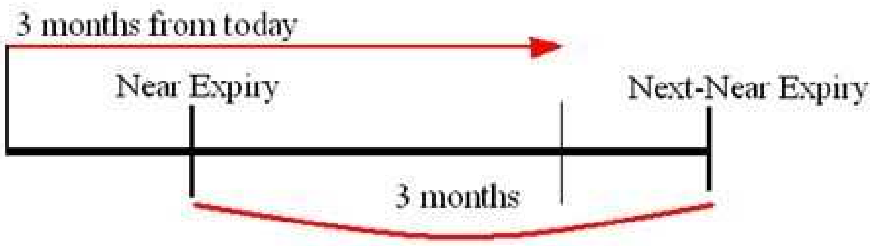

Figure 7.12: A volatility index is forward looking.

Figure 7.13: The SAVI-$.

SAVI-$

The SAVI-$ is merely an expansion on the work that was done with the SAVI-T40 that was launched at the start of 2007. With the SAVI-$ we are constructing a 3 month forward looking index. In essence we want to know what the market volatility is three month from today every day. On the JSE we don't have option expiring every day and thus can't reference the volatility for one of the contracts that is trading on our market. We have 4 expiry dates per underlying contract each year. This means that we have an expiry date every three months. That means that the time period that we are interested in will fall within the period of our next expiry and the one just after that (then next-near expiry), as can be seen in Fig. 7.12.

In essence this serves as a volatility indicator. Because it's a forward looking indicator and not based entirely on the historic values but rather more on people's opinions one will be able to notice that as people get more fearful the value of the indicator will start to rise. This can be clearly seen in the graph plotted in Fig. 7.13

The SAVI-$ is a very useful indicator to find out what the market sentiment is and how people see our stability compared the Dollar. Bellow we can see how the fear indicator rises as the rand weakens and the markets lose confidence in the value of the rand.

7.3 Range Accruels and Corridors

From [BO 10]

From [Ta 10]

.

.

.

.

From [Wy 06]

.

.

.

7.4 Quantos or Currency Translated Options

From [Ta 10]

.

From [Ha 07]

.

.

.

.

.

.

.

Chapter 8

Implied Binomial Trees

Bruno Dupire stated

> "Implied volatility is the wrong number to put into the wrong formula to obtain the correct price. Local volatility on the other hand has the distinct advantage of being locally consistent. It is a volatility function which produces, via the Black-Scholes equation, prices which agree with those of the exchange traded options".

Given the computational complexity of stochastic volatility models and the difficulty of fitting parameters to the current prices of vanilla options, practitioners sought a simpler way of pricing exotic options consistently with the volatility skew [Ga 06].

8.1 Introduction

Most derivative markets exhibit persistent patterns of volatilities varying by strike. In some markets, those patterns form a smile curve. In others, such as equity index options markets, they form more of a skewed curve. This has motivated the name "volatility skew". In practice, either the term "volatility smile" or "volatility skew" (or simply skew) may be used to refer to the general phenomena of volatilities varying by strike.

The aim with a stochastic volatility model is to incorporate the empirical observation that volatility appears not to be constant and indeed varies, at least in part, randomly. The idea is to make the volatility itself a stochastic process. In this document we present the exact solutions for a European option using both constant and *Heston* volatility.

In 1994, *Dupire* showed that if the spot price follows a risk-neutral random walk and if arbitrage-free market prices for European vanilla options are available for all strikes K and expiries T, then the local volatility can be extracted analytically from these option prices. Remember, it is unlikely that *Dupire* ever thought of local volatility as representing a model of how volatilities actually evolve. Rather, it is likely that they thought of local volatilities as representing some kind of average over all possible instantaneous volatilities in a stochastic volatility world. Local volatility model therefore do not represent a separate class of volatility models; the idea is more to make a simplifying assumption that allows practitioners to price exotic options consistently with the know prices of vanilla options.

8.2 Questions to be Considered

We can estimate the market view using the volatility which is implied by the market prices. Using this real information, we can simulate the future asset price path with its corresponding implied, local and stochastic volatility.

- How useful is it?
- Can we take advantage of this information to obtain a more accurate price for an exotic option and have a better understanding of hedging?

8.3 Local Volatility

Given the prices of call or put options across all strikes and maturities, we may deduce the volatility which produces those prices via the **full** Black-Scholes equation. Note that

- This function has come to be known as local volatility.
- Unlike the naive volatility produced by applying the Black-Scholes formulae to market prices, the local volatility is the volatility implied by the market prices and the one factor Black-Scholes.

In order to obtain a theoretical value for vanilla options, Black and Scholes assumed the following stochastic behaviour for the underlying stock price

$$dS = \mu S dt + \sigma S dW \tag{8.1}$$

where dW denotes a Wiener process or random walk or Brownian motion. *Black & Scholes* used this together with their other assumptions to derive the *Black & Scholes* formula.

8.3.1 Dupire's Formula

An extention to this was introduced by *Dupire* in 1994 [Du 94]. He allowed the volatility of the underlying to depend on both strike and time. The stochastic process used to model the behaviour of the underlying asset, is a simple generalisation of Eq.

8.1 where we extend the *Black & Scholes* model to make full use of its diffusion setting without increasing the dimension of uncertainty

$$dS = \mu S dt + \sigma(S,t) S dW. \tag{8.2}$$

The risk neutral process for the underlying asset is found by setting $\mu = r_d - r_f$ in Eq. 8.2, where r_d and r_f denote the domestic and foreign interest rates in continuous format respectively. The parameter $\sigma(S,t)$ is the volatility of the underlying, when its value is S at time t. It is known as the *instantaneous* or *local volatility* and is assumed to be a deterministic function. Equation 8.2 is a differential equation that holds for a small time increment dt. $\sigma(S,t)$ is thus the volatility that holds over this short time interval. It should not be confused with the implied volatility, defined earlier. This instantaneous volatility can be calibrated to observed market values of vanilla options.

If arbitrage-free market prices for European vanilla options are available for all strikes K and expiries T, then $\sigma_l(K,T)$ can be extracted analytically from these option prices. *Dupire* showed that

$$\frac{\partial C}{\partial T} = \sigma_l(K,T)\frac{K^2}{2}\frac{\partial^2 C}{\partial K^2} - (r_d - r_f)K\frac{\partial C}{\partial K} - r_f C \tag{8.3}$$

where C denotes a European call with strike K and expiry T and $\sigma_l(K,T)$ is the local volatility.

If we now rearrange Eq. 8.3 we get

$$\sigma_l^2(K,T) = \frac{\frac{\partial C}{\partial T} + (r_d - r_f)K\frac{\partial C}{\partial K} + r_f C}{\frac{K^2}{2}\frac{\partial^2 C}{\partial K^2}}. \tag{8.4}$$

We can view this formula as a definition of the local volatility function regardless of what kind of process (stochastic volatility for example) actually governs the evolution of volatility.

We can further show that if $C(S,K,\sigma_I,T)$ is the *Black & Scholes* value for a European call with strike K, expiry T and *implied volatility* σ_I, making this substitution in Eq. 8.4 gives us an alternative expression for local volatility in terms of the derivative of the implied volatility

$$\sigma_l^2(K,T) = \frac{\sigma_I^2 + 2\sigma_I\tau\left(\frac{\partial \sigma_I}{\partial T} + (r_d - r_f)K\frac{\partial \sigma_I}{\partial K}\right)}{\left(1 + Kx\frac{\partial \sigma_i}{\partial K}\sqrt{\tau}\right)^2 + \sigma_I^2 K^2 \tau\left(\frac{\partial^2 \sigma_I}{\partial K^2} - x\left(\frac{\partial \sigma_I}{\partial K}\right)^2\sqrt{\tau}\right)} \tag{8.5}$$

with x defined in Eq. 6.18.

To implement any one these two equations, we need a volatility skew that is deterministic, i.e. a formula that we can differentiate. If you do not have that you'll have to do the differentiation numerically.

One potential problem of using the Dupire formula given in Eq. 8.4 is that, for some financial instruments, the option prices of different strikes and maturities are not available or are not enough to calculate the right local volatility. Another problem is that, for strikes far in- or out-the-money, the numerator and denominator of this equation may become very small, which could lead to numerical inaccuracies. Note, the implied vol surface obviously has to be arbitrage free, which is equivalent to showing that the Dupire local volatility is a real number (not a complex number). Now, if you plug in the Dupire formula based on Call and Put prices then it may not work so well mainly because call/put Vega is very small as soon as you move away from the forward, resulting in numerical noise.

Gatheral has shown that we get a better formula if we describe everything in terms of variances. This leads to

$$v_L = \frac{\partial w}{\partial T}\left[1 - \frac{y}{w}\frac{\partial w}{\partial y} + \frac{1}{4}\left(-\frac{1}{4} - \frac{1}{w} + \frac{y^2}{w^2}\right)\left(\frac{\partial w}{\partial y}\right)^2 + \frac{1}{2}\frac{\partial^2 w}{\partial y^2}\right]^{-1} \tag{8.6}$$

where

$$\begin{aligned} w &= \sigma_I^2 \\ y &= \ln\left(\frac{K}{F_T}\right) \\ F_T &= \text{forward of } S \\ v_l &= \sigma_L^2. \end{aligned}$$

In summary, calculating the local volatility with the implied volatility gives us a more accurate and stable result. Furthermore we can save huge amounts of computation time (30 times less). The last results tell us that a flat implied volatility surface automatically yields a flat local volatility surface

8.4 Implied Binomial Trees

See article by Derman and *Kani* — "The Volatility Smile and Its Implied Tree."

Bibliography

[ACS 04] Tom Arnold, Timothy Falcon Crack and Adam Schwartz, *Implied Binomial Trees in Excel*, Working Paper (2004)

[Al 01] Carol Alexander, *Market Models, Guide to Financial Data Analysis*, John Wiley and Sons (2001)

[AM 06] Mark de Araújo and Eben Maré, *Examining the Volatility Skew in the South African Equity Market using Risk-Neutral Historical Distributions*, Investment Analysts Journal, **64**, pp. 15-20 (2006)

[Ba 09] Badshah, Ihsan Ullah, *Modeling the Dynamics of Implied Volatility Surfaces*, Working paper (February 23, 2009). Available at SSRN: http://ssrn.com/abstract=1347981

[BB 98] E. Briys, M Bellalah, H. M. Mai and F. de Varenne, *Options, Futures and Exotic Derivatives: Theory, Application and Practice*, John Wiley and Sons (1998)

[Be 01] Alessandro Beber, *Determinants of the implied volatility function on the Italian Stock Market*, Laboratory of Economics and Management Sant'Anna School of Advanced Studies Working Paper, (2001)

[BGK 97] M. Broadie, P. Glasserman and S.G. Kou, *A Continuity Correction for Discrete Barrier Options*, Journal of Mathematical Finance, **Oct** (1997)

[BJM 08] Petrus Bosman, Samantha Jones and Susan Melmed, *The Construction of an Alsi Implied Volatility Surface: Smiling at the skew*, Cadiz Quantitative Research, February (2008)

[Bl 76] F. Black, *The Pricing of Commodity Contracts*, J. Fin. Econ., **3**, 167 (1976)

[Bl 88] F. Black, *How to use the Holes in Black-Scholes*, Journal of Applied Corporate Finance **1**, Issue 4, p67-73 (1989)

[BL 94] P.P Boyle and Sok Hoon Lau, *Bumping Up Against the Barrier with the Binomial Method*, The Journal of Derivatives, **1**, No. 4, 6-14 (1994)

[BO 10] M. Bouzoubaa and A. Osseiran, *Exotic Options and Hybrids: A Guide to Structuring, Pricing and Trading*, John-Wiley and Sons (2010)

[BS 72] F. Black and M Sholes, *The Valuation of Option Contracts and a test of Market Efficiency*, Proceedings of the Thirtieth Annual Meeting of he American Finance Association, 27-29 Dec. 1971, J. of Finance, **27**, 399 (1972)

[BS 73] F. Black and M Sholes, *The Pricing of Options and Corporate Liabilities*, J. Pol. Econ., **81**, 637 (1973)

[BSU 08] L. Bonney, G. Shannon and N Uys, *Modelling the Top40 Volatility Skew: A Principle Component Analysis Approach*, Investment Analysts Journal, **68**, pp. 31-38 (2008)

[Bu 01] Burashchi A., *The Forward Valuation of Compound options*, The Journal of Derivatives, (Fall 2001)

[Ca 10] A. Castagna, *FX Options and Smile Risk*, John-Wiley & Sons (2010)

[CEG 97] P. C. Carr, K. Ellis and V. Gupta, *Static Hedging of Exotic Options*, Working Paper: Johnson Graduate School of Management, Cornell University (1997)

[Cl 11] I. J Clark, *Foreign Exchange Option Pricing: A Practitioner's Guide*, John-Wiley & Sons (2011)

[CR 85] J. C. Cox and M. Rubinstein, *Option Markets*, Prentice-Hall (1985)

[CS 97] L. Clewlow and C. Strickland, *Exotic Options: the state of the art*, Thomson Business Press (1997)

[Cu 92] M. Curran, *Beyond Average Intelligence*, Risk 5, 10, 60 (1992)

[DD 99] K. Demeterfi, E. Derman, M. Kamal and J. Zou, *More Than You Ever Wanted To Know About Volatility Swaps*, Quantitative Strategies Research Notes, Goldman Sachs, March, (1999)

[De 99] E. Derman, *Regimes of Volatility. Some Observations on the Variation of S&P 500 Implied Volatilities*, Goldman Sachs Quantitative Strategies Research Notes (1999)

[DFW 98] Bernard Dumas, Jeff Fleming, and Robert E. Whaley, *Implied Volatility Functions: Empirical Tests*, The Journal of Finance, Vol III, **No. 6**. (1998)

[DHS 06] T. Daglish, J. Hull and W. Suo, *Volatility Surfaces: Theory, Rules of Thumb, and Empirical Evidence*, Working Paper (2006)

[DK 94] E. Derman and I. Kani, *The Volatility Smile and Its Implied Tree*, Goldman Sachs Quantitative Strategies Research Notes (1994)

[DK 95] E. Derman, I. Kani, D. Ergener and I. Bardhan, *Enhanced Numerical Methods for Options with Barriers*, Goldman Sachs Quantitative Strategies Research Notes (1995)

[DW 08] F. de Weert, *Exotic Options Trading*, *John Wiley & Sons* (2008)

[Du 94] B. Dupire, *Pricing With A Smile*, RISK, **7** (January), 18-20 (1994)

[Fi 94] S. Figlewski, *Forecasting Volatility using Historical Data*, Working Paper, Stern School of Business, New York University, (1994)

[Ga 04] Jim Gatheral, *A Parsimonious Arbitrage-free Implied Volatility Parameterization with Application to the Valuation of Volatility Derivatives*, Working paper (2004); also see http://www.math.nyu.edu/fellows_fin_math/gatheral/madrid2004.pdf

[Ga 06] Jim Gatheral, *The Volatility Surface: a practitioner's guide*, Wiley Finance (2006)

[Ga 09] Jim Gatheral, *The Volatility Surface*, Lecture at the AIMS Summer School in Mathematical Finance, Muizenberg, Cape Town (Feb 2009)

[GK 83] M.B. Garman and S.W. Kohlhagen, *Foreign Currency Option Values*, Journal of International Money and Finance, **2**, 231-237 (1983)

[Ha 07] E. G. Haug, *The Complete Guide to Option Pricing Formulas*, *McGraw-Hill* (2007)

[He 93] S. Heston, *A closed-form solution for options with stochastic volatility, with application to bond and currency options,* Review of Financial Studies, **6**, pp. 327-343 (1993)

[HK 02] Patrick S. Hagan, Deep Kumar, Andrew S. Lesniewski , and Diana E. Woodward, *Managing Smile Risk*, Wilmott Magazine (2002)

[Hs 97] H. Hsu, *Surprised Parties*, RISK, **April** (1997)

[Hu 06] J. Hull, *Options, Futures, and other Derivatives*, Sixth Edition, *Prentice-Hall* (2006)

[Hu 09] K. Hussain, *Hedging Market Risk in Islamic Finance*, World Commerce Review, **2**, Issue 3, (2009) [JR 96] J. Jackwerth and M. Rubinstein, em Recovering Probability Distributions from Option Prices, Journal of Finance **51**, 1611-1631 (1996)

[JT 00] R. Jarrow and S. Turnbull, *Derivative Securities*, 2nd. Ed., *South-Western College Publishing* (2000)

[JT 07] G. J. Jiang and Y. S. Tian, *Extracting Model-Free Volatility from Option Prices: An Examination of the Vix Index*, Journal of Derivatives, **14**, No. 3 (2007)

[Ko 99] J. Kolman, *Roundtable: Learning from the Skew*, Derivatives Strategy Magazine, **November** (1999) - http://www.derivativesstrategy.com/magazine/archive/1999/1199fea4.asp.

[Ko 02] A. A. Kotzé, *Equity Derivatives: Effective and Practical Techniques for Mastering and Trading Equity Derivatives*, Working paper (2002)

[Ko 03] A. A. Kotzé, *Black-Scholes or Black Holes?*, The South African Financial Markets Journal, **2**, 8 (2003)

[KV 90] A. G. Kemna and A. C. F. Vorst, *A Pricing Method for Options Based on Average Asset Values*, Journal of Banking and Finance, **14**, 113-129, (1990)

[Le 00] Alan Lewis, *Option Valuation under Stochastic Volatility*, Finance Press (2000)

[Le 07] Martin le Roux, *A Long-Term Model of The Dynamics of the S&P500 Implied Volatility Surface*, North American Actuarial Journal, **11**, No. 4 (2007)

[LM 02] Lipton, A. and McGhee, W. *Universal Barriers*, Risk, **May** (2002)

[Ma 95] S. Mayhew, *Implied Volatility*, Financial Analyst Journal, **July-August**, 8 (1995)

[Me 73] R. C. Merton, *Theory of Rational Option Pricing*, Bell Journal of Economics, **4**, 141-183 (1973)

[MM 79] J. D. Macbeth and L. J. Mervile, *An Empirical Examination of the Black-Scholes Call Option Pricing Model*, Journal of Finance, **XXXIV**, 1173 (1979)

[MS 00] R.N. Mantegna and H.E. Stanley, *An Introduction to Econophysics: Correlations and Complexity in Finance*, Cambridge University Press (2000)

[Na 88] S. Natenberg, *Option Volatility and Pricing Strategies*, *Probus Publishing Company* (1988)

[Ob 08] Jan Oblój, *Fine-Tune your Smile: Correction to Hagan et al*, Imperial Collge of London Working Paper, (2008)

[Pa 09] B.M. Paxford, *Derivatives - An Islamic Finance Perspective*, Islamic Finance Today, 19-27 (2009)

[PCS 08] Andreou Panayiotis, Charalambous Chris and Martzoukos Spiros, *Assessing Implied Volatility Functions on the S&P500 Index Options*, University of Cyprus Working Paper, (2008)

[Pe 91] E. E. Peters, *Chaos and Order in the Capital Markets - a new view of cycles, prices, and market volatility*, John Wiley & Sons (1991)

[Re 06] Riccardo Rebonato, *Volatility and Correlation; the perfect hedger and the fox*, Second Edition, Wiley & Sons (2006)

[Ri 94] D. R. Rich, *The Mathematical Foundations of Barrier Option-Pricing Theory*, Advances in Futures and Options Research, **7**, 267 (1994)

[RR 91] M. Rubinstein and E. Reiner, *Breaking Down the Barriers*, RISK, **September** (1991)

[RS 95] M. Reimer and K. Sandmann, *A Discrete Time Approach for European and American Barrier Options*, working paper, Department of Statistics, Bonn University (1995)

[Ru 91] M. Rubinstein, *Forward-Start Options*, RISK, **April** (1991)

[Ru 94] M. Rubinstein, *Implied Binomial Trees*, Journal of Finance, **69**, pp. 771-818 (1994)

[Sh 93] David Shimko, *Bound of Probability*, Risk, **6**, pp. 33-37 (1993)

[Sh 08] Yuriy Shkolnikov, *Generalized Vanna-Volga Method and Its Applications*, NumeriX Quantitative Research, (July 2008)

[Sh 10] L. Shover, *Trading Options in Turbulent Markets*, John Wiley & Sons (2010)

[SV 08] Sanjay Sehgal and N. Vijayakumar, *Determinants Of Implied Volatility Function On The Nifty Index Options Market: Evidence From India*, Asian Academy Of Management Journal Of Accounting And Finance, **4**, pp. 45-69 (2008)

[Ta 10] Chia Chiang Tan, *Demystifying Exotic Products: Interest Rates, Equities and Foreign Exchange*, John-Wiley and Sons (2010)

[To 94] R. G. Tompkins, *Options explained*^{2}, Macmillan Press (1994)

[To 01] R. Tompkins, *Implied Volatility Surfaces: Uncovering Regularities for Options on Financial Futures*, The European Journal of Finance, **7** No. 3 pp. 198-230 (2001)

[To 08] F. Tourrucôo, *Considerations on approximate calibration of the SABR smile*, Universidade Federal do Rio Grande do Sul Working Paper, (2008)

[We 05] G. West, *Calibration of the SABR Model in Illiquid Markets*, Applied Mathematical Finance, **12**, No. 4, pp. 371-385, (December 2005)

[We 06] T. Weithers, *Foreign Exchange: a Practical Guide to the FX Markets*, *Wiley Finance* (2006)

[Wi 98] P. Wilmott, *Derivatives: the theory and practise of financial engineering*, *John Wiley & Sons* (1998)

[Wy 06] U. Wystup, *FX Options and Structured Products*, *John Wiley & Sons* (2006)

[Wy 08] U. Wystop, *Vanna-Volga Pricing*, MathFinance AG Waldems, Germany (June 2008)

[Ya 01] Len Yates, *Buying and Selling Volatility*, http://www.optionvue.com/Articles/ArticlesDirectory.htm

[Zh 97] P. G. Zhang, *Exotic Options: A guide to second generation options*, *World Scientific* (1997)

[ZX 05] Jin Zhang and Yi Xiang, *Implied Volatility Smirk*, University of Hong Kong Working Paper (2005)