MASTER

Trading the difference between realised and implied volatility

de Kreuk, Anne C.C.

Award date:
2005

TECHNISCHE UNIVERSITEIT EINDHOVEN

Department of Mathematics and Computing Science

MASTER'S THESIS

Trading the Difference between Realised and Implied Volatility

by

A.C.C. de Kreuk

Supervisor	:	dr. C.A.J. Hurkens
Advisors	:	dr. B. Oldenkamp (ABN Amro Asset Management)
	:	drs. J. Tolenaar (ABN Amro Asset Management)
	:	prof. dr. ing. G. Woeginger
	:	dr. F.H.J. Redig

Eindhoven, January 2005

Summary

In literature on econometrics and investment strategies is often stated that the investors have a biased perception of the variance (risk) of stock values. On average the perceived standard deviation (which is called *implied volatility*, since it is implied by the prices of options on the stocks) is larger than the true standard deviation (which is called *realised volatility*). Therefore it seems profitable to find an investing strategy with which the difference between implied and realised volatility can be traded. The objective of the research presented in this thesis is to find an optimal strategy for building portfolios that minimise risk and maximise return to exploit the phenomenon that implied volatility overestimates realised volatility. To define a strategy like this, the instruments that can be used to trade the difference between implied and realised volatility are analysed and a robust optimisation model that can give the desired optimal strategy is formulated.

The straddle is the most basic instrument that can be used to exploit the phenomenon that implied volatility overestimates realised volatility. The straddle has the disadvantage that it is highly sensitive to the level of the stock price, since the profit pattern is very much dependent on the level of the stock price at maturity, compared to the strike price. To prevent this sensitivity, the straddle can be delta-hedged. The delta-hedged straddle has an expected payoff that is linear in the difference between implied and realised volatility, but the payoff is dependent on the timing of the volatility. An over-the-counter product, that has the same expected payoff as the delta-hedged straddle, is the volatility swap. Volatility swaps are, however, not traded in practice. The offering party has to replicate an over-the-counter instrument with products that are traded on exchanges, such as options and futures. Since the only known way to replicate the volatility swap is to keep a position in a delta-hedged straddle, which gives the exposure to the timing of the volatility, the replication of volatility swaps does not seem attractive. Another over-the-counter product is the variance swap. A variance swap has a payoff that is linear in the difference between implied and realised *variance*. The variance swap has the advantage that it can be replicated more reliably and can therefore be valued more reliably.

The decision supporting model that is developed gives a distribution of an investment over the different markets that exist in the world (i.e. the Asian market, the European market, and the U.S.A. market). Furthermore, the model finds an optimal distribution over the volatility instruments (within a market) such that the risk is minimised and the performance is as defined by the customer. There is a lot of uncertainty about the processes that describe the underlying indices and therefore the optimisation model is made robust to the input parameters that define the processes of the underlying indices by using the technique of *robust optimisation*.

Preface

The research presented in this thesis has been performed at ABN Amro, department Structured Asset Management, as a part of the Master's degree program of Applied and Industrial Mathematics at the department of Mathematics and Computer Science, Eindhoven University of Technology. It was a very interesting project for me, since I studied an application of mathematics that was totally new to me. It was a great opportunity for me to do an internship at ABN Amro Structured Asset Management and I had the possibility to learn a lot during this internship.

The aim of the research, which I have performed during the last nine months, is to find an optimal strategy for building portfolios that minimise risk and maximise return to exploit the phenomenon that implied volatility overestimates realised volatility. I am aware of the fact that not all readers are familiar with the theory of financial derivatives. To be able to understand this thesis, some knowledge of financial derivatives is required and therefore I have tried to summarise some important theory in Appendix A. If the reader is not familiar with the theory of derivatives it is advisable to read this appendix first, before the rest of the report is read.

I am very grateful to my supervisors Dr. C.A.J. Hurkens (TU/e, Department of Mathematics and Computer Science), Dr. B. Oldenkamp (ABN Amro, Structured Asset Management), and Drs. J. Tolenaar (ABN Amro, Structured Asset Management). They have given many insightful comments, helpful discussions, and the opportunity to do an interesting graduation project. Furthermore, I would like to thank all my 'colleagues' at ABN Amro Structured Asset Management for answering my questions and providing a pleasant work environment.

Anne de Kreuk

Contents

CHAPTER 1

Introduction

There are many ways to invest money. It can, for example be invested in real estate or stocks. Many variations are possible. It is, for example, very common to use products that are derived from stocks (*derivatives*). Examples of derivatives are *futures* (contracts that state that a certain underlying product has to be sold or bought for a certain price at a certain time in the future) and *options* (contracts that state that a certain underlying product *can* be sold or bought for a certain price at a certain time in the future). Many complex strategies can be constructed by using derivatives. A particular investment strategy that ABN Amro Structured Asset Management wants to know more about is investing in combinations of derivatives such that a position is obtained that has an exposure to the movements of the stock price (*volatility*). The objective of the research presented in this thesis is to find an optimal strategy, using some characteristics of volatility.

The research described in this thesis has taken place at ABN AMRO Asset Management in the department Structured Asset Management in Amsterdam. Therefore, in this chapter first a general introduction is given to the ABN AMRO Bank and the department Structured Asset Management in particular. After that, a description of the problem of trading volatility is given, including an overview of some relevant literature. This chapter concludes with an overview of the rest of the report.

1.1 ABN AMRO Bank N.V.

The ABN AMRO Bank is divided in three global Strategic Business Units (SBU's):

- *Consumer & Commercial Clients (C&CC)* - for individuals and small to medium-sized enterprises requiring day-to-day banking.
- *Wholesale Clients (WCS)* - for major international corporations and institutions.
- *Private Clients & Asset Management (PCAM)* - for high net-worth individuals and institutional investors.

1.1.1 ABN AMRO Asset Management

The business unit *Asset Management* provides mutual funds and handles investment mandates. Investment management products are distributed through financial intermediaries and offered to institutional clients directly. Asset Management operates in more than 30 countries across Europe, North America, South America, Asia and Australia. Portfolios are often managed locally, and local knowledge is used as input for international co-ordination of the investment policy.

ABN AMRO Asset Management is able to effectively combine global resources to deliver specialised investment products. Global and regional equity and fixed income investment products benefit from the valuable source of local expertise.

Structured Asset Management

The *Structured Asset Management* department is a separate portfolio management department within ABN AMRO Asset Management. Structured Asset Management (SAM) specialises in the pragmatic use of quantitative techniques and models in investing. The main activities of the Structured Asset Management department include the management of guaranteed products, optimised derivative portfolios and quantitative equity products.

At present SAM has two units, one in Amsterdam and one in Hong Kong. Currently, there are three product groups within SAM. There is one product group that deals with equity derivatives solutions, one product group that deals with quantitative equity and the last product group is structured interest rates solutions.

1.2 Problem Description

In this section the problem is described. The first subsection gives an introduction to the framework and background of the project. After that, an overview of previous research is given and in Section 1.2.3 the goal of the project is described.

1.2.1 Background

Derivatives such as options are traded all over the world. The prices of derivatives are determined by different components. The option price is determined considering the strike price, the remaining time to maturity, the current stock price, the return, and the *volatility.* The volatility is the standard deviation of the return of the underlying asset. All the components of which the option price consists are perceptible except volatility. The volatility, however, can be determined by extracting it from a given price, since all the other components are known. This is called the *implied volatility.* The *realised volatility* is the actual volatility and can only be determined over history and not over future periods. The implied volatility is often considered an estimator for realised volatility, even though implied volatility and realised volatility usually differ.

Some traders consider it interesting to trade volatility. There are several reasons why it can be profitable to invest in volatility. Firstly, the return of volatility is usually negatively correlated with the price return of the underlying asset. This can give the investor either a diversifier to equity exposure or a tool to use his expectations considering future volatility. Furthermore, volatility appears to be mean reverting. This could give the investors some indication about an expected decrease or increase of the volatility in the future. It is also possible that the investor has a view on the difference between implied and realised volatility.

The difference between implied and realised volatility is tradable by different strategies (usually combinations of options). Examples of strategies that are interesting for investors that would like to trade this difference are straddles and variance swaps. Different strategies arise, for example by hedging or not hedging and holding a position until maturity or not.

The research described in this thesis is included in a project in the department Structured Asset Management that is called the 'straddle project'. The main goal of the 'straddle project' is to find a robust investment strategy that profits from the phenomenon that implied volatility usually overestimates realised volatility, as is shown in Figure 1.1. Figure 1.1 shows the difference between the implied and realised volatility on the S&P 500 from 1994 until 2004. The figure shows that the implied volatility is on average 5% larger than the realised volatility.

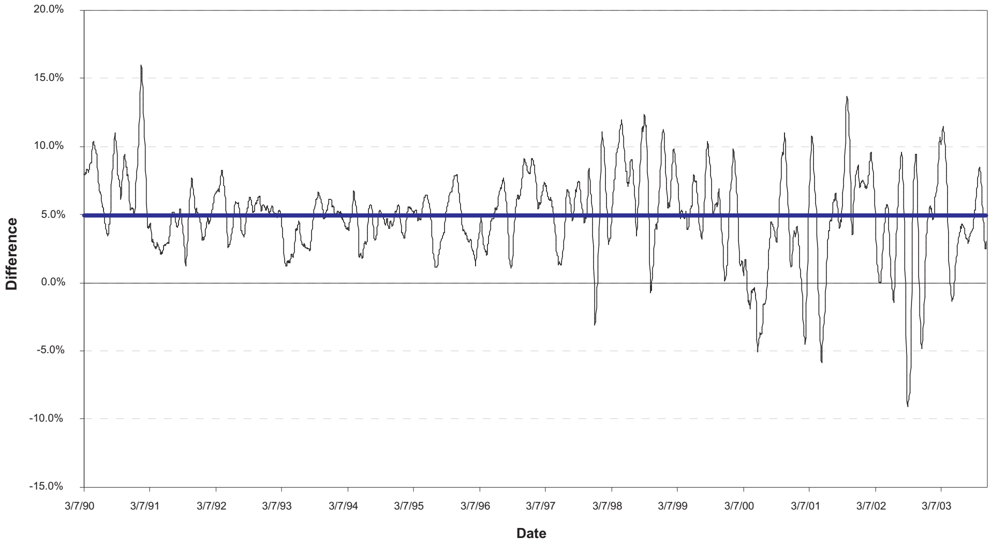

Figure 1.1: Implied volatility − Realised volatility on the S&P 500 1994–2004

1.2.2 Previous Research

There is some literature focussed on the difference between implied and realised volatility using empirical research, for example in Canina and Figlewski [6]. Since it is showed by empirical studies that implied volatility often overestimates realised volatility, is seems interesting to profit from this phenomenon. For example Coval and Shumway [8] and Driessen and Maenhout [11] discuss some trading strategies by which the difference between implied volatility and realised volatility is sold and test these with data sets from history.

Coval and Shumway [8] considered returns on zero-beta index straddles (combinations of long positions in call and put options that have offsetting covariance with the underlying index), at-the-money (ATM) zero-beta 'crash-neutral' straddles (a straddle position with an offsetting short position in a deeply out-of-the-money (OTM) put) and selling ATM 'crash-neutral' straddles (not zero-beta, but an equivalent number of calls and puts are sold). Coval and Shumway's conclusions are that the long positions in these straddles give rise to negative returns. Hence, by selling these straddles, a positive return can be generated. Therefore they conclude that besides market risk, it is important to price the risk associated with option contracts and that stochastic volatility may be an important factor for pricing assets. Driessen and Maenhout [11] considered returns on OTM put options and ATM straddles. Their conclusion is that it is always optimal to have a short position on the OTM put options and the ATM straddles, except for the case where investors were extremely loss-averse and where distorted probabilities (extreme stock market outcomes that are overweighed) were used. A difference between these two papers is that they use different methods to calculate the option returns. The main difference is that Coval and Shumway hold the options until maturity and Driessen and Maenhout do not.

Different option pricing models are developed that have different assumptions and give an indication of how to price options given these assumptions. There is some literature focussed on the option pricing models. For every model discussed in this report, there are papers that describe the specific model, like the Black and Scholes model in Black and Scholes [3], the Heston model in Heston [14], and implied trees in Jackwerth [19] and Rubinstein [22]. Other literature, like Bates [1], gives an overview of the different option pricing models and compares these.

Some literature can be found that is focussed on the pricing and hedging of volatility instruments, like straddles, volatility swaps, and variance swaps. Carr and Madan [7] discuss three instruments that can be used to trade realised volatility (i.e. a long position in a straddle, delta-hedging an option position, and buying or selling over-the-counter products of which the payoff is an explicit function of volatility). Windcliff, Forsyth, and Vetzal [25] also discuss some volatility instruments (instruments that can be used to trade volatility). In both papers, the return of the instruments is considered as well as the possibilities and performance of replication strategies. On the theory that deals with volatility and variance swaps specifically, more can be found in [12], Howison [16] and Demeterfi et al. [10].

1.2.3 Goal

Since most of the results found in literature considering the gap between implied and realised volatility draw conclusions from empirical research, it is interesting to get more insight in the theoretical background. This can be obtained by analysing some different pricing models and the behaviour of some different volatility instruments. After this insight is gained, a decision-supporting model can be made that finds an optimal portfolio using the information about how to use difference between the realised and implied volatility. The main goal of this research can therefore be formulated as follows:

> *Find an optimal strategy for building portfolios that minimise risk and maximise return to exploit the phenomenon that implied volatility overestimates realised volatility.*

This strategy should be robust to the assumptions that have been made considering the asset return generating process. Furthermore, the model should give the user an indication of the optimal diversification of capital over the different markets in the world and the optimal investment strategy considering the expected difference between implied and future realised volatility.

1.3 Overview of the Report

To reach the goal described in the previous section, first some insight in the theoretical background has to be gained. In the Chapters 2 and 3 the findings considering the theoretical background of option pricing models are described. In the Chapters 4 and 5 the results of the research considering the performance of the volatility instruments are described. Chapter 4 describes the findings from theory and Chapter 5 results from simulations.

After that, this information has to be used to make a decision supporting model that can be used to build an optimal portfolio. This model and the results following from implementing the decision supporting model are presented in Chapter 6. Finally, the conclusions and recommendations following from the research described in this thesis are presented in Chapter 7.

In Appendix A, some information is given for readers who are not familiar with the theory of derivatives. In the other appendices, some derivations and calculations are presented.

CHAPTER 2

Black and Scholes Model

In 1973 Black, Scholes, and Merton developed the Black and Scholes model, which was a major breakthrough in the pricing of options (more information in [3]). Although the Black and Scholes model is often quite successful in explaining option prices and still used by traders, it has some biases. The assumption of constant volatility has been questioned particularly.

Black, Scholes, and Merton (more information in Hull [17] and references in [17]) use a differential equation to determine the options pricing formulas. The parameters that are input to the differential equation are the spot price of the underlying, the strike price of the option, the interest rate, the time until expiration, the dividend rate, and the volatility of the underlying. The Black and Scholes model is widely used due to its simplicity and consistency with the capital asset pricing model (CAPM model, more information in [5]). The weakness of the model, however, is that some assumptions are (too) restrictive and therefore not very realistic.

The assumptions that Black, Scholes, and Merton use to derive their option pricing formulas are as follows:

1. The stock price (S) follows the following process: $dS = \mu S dt + \sigma S dW$
   σ is the constant volatility of the stock price, μ is the constant expected rate of return and W is a Wiener process.

2. There are no transaction costs or taxes. All securities are perfectly divisible.

3. There are no riskless arbitrage opportunities (more information can be found in Appendix A.6).

4. Security trading is continuous.

5. European exercise terms are used for the options (more information can be found in Appendix A.6).

6. The risk-free interest rate, r, is constant and the same for all maturities.

Especially the assumptions about the constant volatility and the constant risk-free interest rate are often questioned.

The pricing formulas are as follows (C is the price of a call option, P is the price of a put option):

$$C = S_0 e^{-qT} \mathrm{N}(d_1) - K e^{-rT} \mathrm{N}(d_2), \tag{2.1}$$

$$P = K e^{-rT} \mathrm{N}(-d_2) - S_0 e^{-qT} \mathrm{N}(-d_1), \tag{2.2}$$

with $d_1 = \frac{\ln(S_0/K)+(r-q+\sigma^2/2)T}{\sigma\sqrt{T}}$,

$d_2 = \frac{\ln(S_0/K)+(r-q-\sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$,

$\mathrm{N}(x)$ is the cumulative probability distribution function for a standardised normal distribution,

S_0 is the price of the underlying at time 0,

K is the strike price,

q is the dividend rate.

Even though the Black and Scholes model is still used widely by traders, other pricing models are developed that deal with some of the restrictions that are used in the Black and Scholes model. An overview of these models is given in the next chapter.

CHAPTER 3

Other Option Pricing Models

The Black and Scholes model assumes constant volatility. This is a very restrictive assumption, which is not very close to reality, as can be seen in Figure 3.1, which shows the realised volatility of the S&P 500 from May 17th 1994 until March 23rd 2004. Several models have been developed to forecast volatility and price options using this forecast. In this chapter, some option pricing models are discussed that have the feature that they recognise that volatility is not constant. There are several models that use deterministic volatility. Some deterministic volatility models are discussed in Section 3.1. Other models use stochastic volatility. Some stochastic volatility models are described in Section 3.2. There are also numerical procedures to determine option prices. These are discussed in Section 3.3.

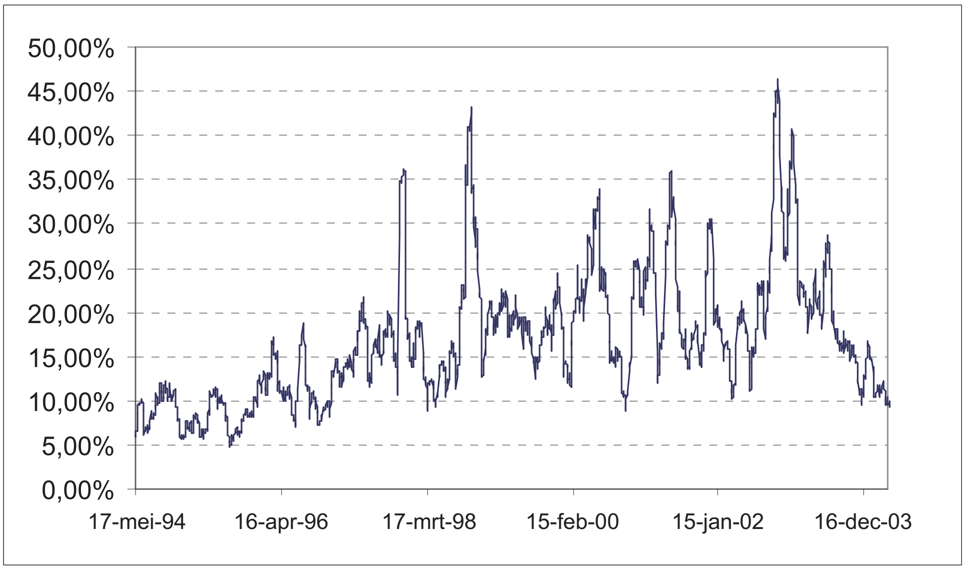

Figure 3.1: Realised volatility of the S&P 500 from May 1994 to March 2004

3.1 Deterministic Volatility Approach

3.1.1 The Constant Elasticity of Variance Model

The constant elasticity of variance model is developed by Cox and Ross in 1976 (See Cox and Ross [9]). It was the first model that relaxed the assumption of constant volatility of the Black and Scholes model. In this model the volatility changes depending on the level of the stock price. The constant elasticity of variance diffusion process of the underlying is described by:

$$\frac{dS}{S} = \mu dt + \sigma S^{\rho-1} dW. \tag{3.1}$$

W is a Wiener process. The variance of the percentage stock price change (return) is given by $\sigma^2 S^{2(\rho-1)}$. ρ is the correlation between the volatility and the stock price. When $\rho = 1$, the process reduces to the process used in the Black and Scholes model (the variance becomes independent of the asset price). The value of a European call option is:

$$C = S(\sum_{n=0}^{\infty} g(\frac{s^*}{n+1})\sigma(\frac{K^*}{n+1} + \frac{1}{2-2\rho})) - Ke^{-rT}(\sum_{n=0}^{\infty} g(\frac{s^*}{n+1} + \frac{1}{2-2\rho})\sigma(\frac{K^*}{n+1})), \tag{3.2}$$

with
$$\begin{aligned} S^* &= S^{2-2\rho}(\tfrac{2re^{rT(2-2\rho)}}{\sigma^2(2-2\rho)e^{rT(2-2\rho)}-1}), \\ K^* &= K^{2-2\rho}(\tfrac{2r}{\sigma^2(2-2\rho)e^{rT(2-2\rho)}-1}), \\ g(X,m) &= \tfrac{e^{-x}x^{m-1}}{\Gamma(m)} \text{ (the Gamma density function).} \end{aligned}$$

3.1.2 Displaced Diffusion Option Pricing

Rubinstein (in [23]) has derived the displaced diffusion option pricing formula by modeling a firm as the holder of two types of assets, namely the risky and the riskless assets. The former are assumed to follow a geometric Brownian motion and the latter grow at the risk-free interest rate. Let α be the proportion of the total value (V) of the firm that is invested in the risky asset and let $1-\alpha$ be the proportion invested in the risk-free assets. The value of the firm after t years is:

$$\lfloor \alpha e^y + (1 - \alpha r^t) \rfloor V, \tag{3.3}$$

where y is a normally distributed random variable with volatility $\sigma\sqrt{t}$. The price of a call option of this firm is now:

$$C = aS\mathrm{N}(x) - (X - bS)r^{-t}\mathrm{N}(x - \sigma\sqrt{t}), \tag{3.4}$$

with
$$\begin{aligned} x &= \tfrac{\ln(\frac{\alpha S}{k-bS}r^{-t})}{\sigma\sqrt{t}} + \tfrac{1}{2}\sigma\sqrt{t}, \\ a &= \alpha(1+\beta)\textstyle\prod_{k\le t}(1-\delta_k), \\ b &= (1-\alpha-\alpha\beta)r^t - \textstyle\sum_{k\le t}\delta_k r^{t-k}, \end{aligned}$$

S is the current market value of the stock,
β is the debt-equity ratio (current market value),
k is the time until the ex-dividend date,
σ is the volatility of the underlying,
$\mathrm{N}(x)$ is the cumulative normal distribution function.

The Black and Scholes model is a special case of this Rubinstein model. The Rubinstein model provides a more realistic approach to dividend payout than the Black and Scholes model. A disadvantage of the model, however, is that the stock price can become negative in extreme cases.

3.2 Stochastic Volatility Approach

3.2.1 Discrete Time Models

Examples of discrete time models that forecast volatility are the exponentially weighted moving average (EWMA) model, the autoregressive conditional heteroscedasticity (ARCH) model and an extension to ARCH, the generalised autoregressive conditional heteroscedasticity (GARCH) model (more information in Hull [17]). The volatility forecasts obtained from these models are often used as input in pricing models such as the Black and Scholes model. For example Heston and Nandi in [15] propose a closed-form option pricing model that uses the GARCH model. The difference of these models with the models described in the previous sections (the Black and Scholes model and the models that use a deterministic non-constant volatility) is that the models described here make use of volatility-data from history and give a specified weight to these data. In this section the GARCH model is described.

The GARCH model is proposed by Bollerslev in 1986 (described in [4]). The model assumes the following formula for updating volatility estimates:

$$v(t) = \omega + \sum_{i=1}^{p} \beta i v(t - i\Delta) + \sum_{i=1}^{q} (\alpha i z(t - i\Delta) - \gamma i \sqrt{v(t - i\Delta)})^2, \tag{3.5}$$

where r is the continuously compounded interest rate for the time interval Δ and $z(t)$ is the standard normal disturbance. $v(t)$ is the conditional variance of the log return between $t - \Delta$ and t and is known from the information at time $t - \Delta$. α and β are the weight parameters.

3.2.2 Continuous Time Models

The early literature on stochastic volatility assumed zero risk premia. For example the model of Hull and White (see [18]), who assumed the following processes for the stock price and its volatility:

$$dS = \mu S dt + \sigma S dW_1, \tag{3.6}$$

$$d\sigma^2 = \phi\sigma^2 dt + \delta\sigma dW_2. \tag{3.7}$$

Here W_1 and W_2 independent Wiener processes and μ depends on S, σ, and t. ϕ and δ may depend on σ and t. Since the variance is not associated to any systematic risk, the pricing formula of Black and Scholes can be used for the stock price process, using the average variance. Compared to this model, the Black and Scholes model overvalues the At-the-Money (ATM) options, while it undervalues the deep In-the-Money (ITM) and Out-of-the-Money (OTM) options.

Stein and Stein (see [24]) provide the first closed-form solution for the value of a call option when the variance changes stochastically and the volatility risk premium is non-zero. Stein and Stein assume the following processes for the stock price and its volatility (the volatility follows a mean reverting process):

$$dS = \mu S dt + \sigma S dW_1, \tag{3.8}$$

$$d\sigma = -\delta(\sigma - \vartheta)dt + k dW_2. \tag{3.9}$$

Here δ, ϑ and k are fixed constants (δ is the speed of the mean reversion and ϑ the mean of the mean reversion) and W_1 and W_2 are independent Wiener processes. Stein and Stein use characteristic functions to derive the value of a call option.

One of the most fundamental models in option pricing is derived by Heston (see [14]). Heston derives a solution for a call option when both systematic volatility risk and arbitrary correlation between volatility and returns are allowed. The stock price and its volatility are described by the following processes:

$$\frac{dS}{S} = \mu dt + \sigma dW_1, \tag{3.10}$$

$$dv = -\delta(v - \vartheta)dt + \sigma\sqrt{v} dW_2. \tag{3.11}$$

Here W_1 and W_2 are *dependent* Wiener processes with correlation coefficient ρ, v the variance of the stock price, ϑ the mean of the mean reversion and δ the speed of the mean reversion. σ is the volatility of the volatility.

The value of a call option is then given by:

$$C_t = S_t P_1 - K e^{-r(T-t)} P_2, \tag{3.12}$$

with: $P_j = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{re}[\frac{e^{i\phi \ln(K)} f_j(\ln(S), \sigma^2, T, \phi)}{i\phi}] d\phi$, $j = 1, 2$.

Here f_j is the characteristic function of the logarithm of the stock price.

Heston (in [14]) concludes that the Black and Scholes model, compared to the Heston-model, overvalues OTM options and undervalues ITM options when the correlation between volatility and asset returns is negative and that the Black and Scholes model undervalues OTM options and overvalues ITM options when the correlation between volatility and asset returns is positive.

3.2.3 Models Implementing Jumps

The underlying price or index can also be modeled by a process with jumps. The jumps are often added to a stochastic volatility model and the jump process can, for example, be a Poisson process. The size of the jumps can either be random or not. Often the size of the jumps is assumed to be lognormally distributed. When the jump process is a Poisson process with arrival intensity λ and this jump process is added to the Heston model described in the previous section, the underlying follows the following process:

$$dS = (\mu - \lambda m) S dt + \sigma S dW_1 + (J-1) S dq, \tag{3.13}$$

$$dv = -\delta(v - \vartheta) dt + \sigma \sqrt{v} dW_2. \tag{3.14}$$

Here $m = \mathbb{E}[J-1] = e^{\mu_J + \frac{1}{2}\gamma_J^2} - 1$ and $\mathbb{P}(dq = 1) = \lambda dt$ and $\mathbb{P}(dq = 0) = 1 - \lambda dt$.

3.3 Numerical Procedures

When exact formulas are not available, numerical procedures can be used to price derivatives. Examples of numerical procedures are binomial trees and an extension to binomial trees: implied trees. Another example of a numerical procedure and an extension to implied trees is generalised binomial trees. Numerical procedures are particularly useful when the trader has to make early exercise decisions prior to maturity (American options).

3.3.1 Implied Trees and Generalised Binomial Trees

The basic idea of implied trees (see Rubinstein [22]) is to construct a (binomial) tree that can fit currently traded derivatives prices, either exactly or fit in some ways. The tree can then be used to price other derivatives on the same underlying asset and with the same or earlier maturity. Generalised binomial trees (see Jackwerth [19]) generalise an assumption of implied trees, namely that all paths in the tree that lead to the same end node are equally likely to be taken. By using these generalised binomial trees it is possible to price derivatives, considering more than one other derivatives on the same underlying asset with the same or earlier maturity. An advantage of these models is that American options (i.e. options that can be exercised before the expiration date) can be priced by them as well.

3.4 Conclusions

After te most well-known option pricing model, the Black and Scholes mode, is discussed in the previous chapter, in this chapter some other option pricing models are discussed. The models described in this chapter have in common that they recognise that volatility is not constant.

Models that use deterministic volatility were discussed in Section 3.1 and models that use stochastic volatility in Section 3.2. Furthermore, we have seen some numerical procedures to determine option prices in Section 3.3.

The models discussed in this chapter can be used to analyse the instruments that exploit the phenomenon that implied volatility is on average larger than realised volatility (*volatility instruments*). In the Chapters 4 and 5 an analysis of the performance of these instruments is presented.

CHAPTER 4

Volatility Instruments - Theory

4.1 Introduction

Literature often states that the market is hard to predict. The market prediction falls into two categories: 1) the prediction of the short-term movement of prices, and 2) the prediction of volatility of the underlying. It seems like a much easier task to predict volatility than to predict prices, since volatility is considered mean reverting. Therefore it seems more attractive to trade volatility than other elements of the market.

There are several ways to trade volatility. Firstly, it is possible to take a position in the implied volatility by buying or selling futures on a volatility index, like the VIX. A volatility index gives measures of market expectations of the volatility conveyed by certain option prices. The VIX, for example, gives these measures by using the options on the S&P 500 index. In this situation, the trader will always be concerned with buying volatility near the bottom of its range and sell it when it gets back to the middle or the high of the range.

Another way to trade volatility is to use the difference between implied and realised volatility. If one expects that the realised volatility will be higher than the implied volatility, it is profitable to buy the difference between implied and realised volatility (realised minus implied volatility) and if the realised volatility is expected to be lower, it is more profitable to sell this difference. This strategy would work best if the difference between implied and realised volatility can be isolated perfectly. In practice however, it is nearly impossible to isolate this, but it is still possible to establish positions in which the direction of the stock price movements is irrelevant to the return of the position. Instruments that can be used to trade the difference between implied and realised volatility distinguish themselves by the way the return depends on other factors than volatility and volatility itself. These instruments are called *volatility instruments* in the remainder of this report. Since implied volatility is usually higher than realised volatility (as is shown in Figure 1.1) according to the literature (Canina and Figlewski in [6]), we focus on short positions in the trading of the difference between implied and realised volatility.

In this chapter, some instruments that can be used to trade the difference between implied and realised volatility are described theoretically. Some of these instruments are traded in practice and some are usually not (like volatility swaps). It is interesting to analyse how these instruments perform in different situations and therefore some simulations have been performed, which are discussed in the next chapter.

4.2 Straddles

A straddle (See Hull [17] and Appendix A.4.2) involves buying a European call option and a European put option with the same strike price and expiration date. The profit pattern is shown in Figure 4.1. A long position in an at-the-money (ATM) straddle is appropriate when an investor is expecting a large move in a stock price, but does not know in which direction this move will be (the investor expects high volatility). An investor who writes a straddle has the opposite profit pattern, so by writing straddles with the strike price (close to) the current stock price (ATM straddles), one can profit from a low volatility. Therefore, writing a straddle can be seen as the most basic instrument that profits from the phenomenon that implied volatility usually overestimates realised volatility.

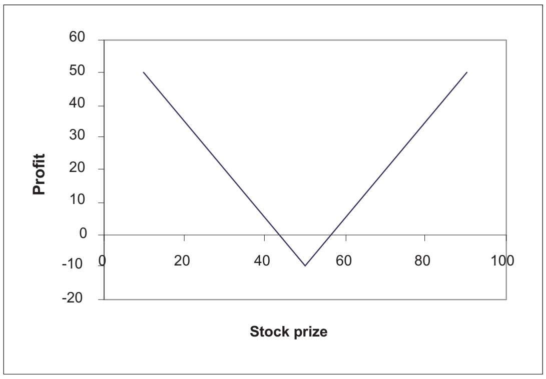

Figure 4.1: The payoff of a long postition in a straddle

A straddle is sensitive to the level of the stock price, since the profit pattern is very much dependent on the level of the stock price at maturity, compared to the strike price. To avoid this, the straddle can be hedged. This is discussed in Section 4.2.1.

4.2.1 Delta-hedged Straddle

It is common to neutralise the influence of the level of the stock price on the return pattern of the straddle. This can be done by delta-hedging[1] the straddle, for example by keeping the delta of the position zero by buying or selling futures on the underlying (in Appendix A.5, more information about (delta-) hedging is given). Unfortunately, the delta cannot be kept zero continuously, because continuous trading is not possible in practice and also not profitable because of transaction costs.

Hedging Continuously

The expected return of a short position in a continuously delta-hedged straddle (no transaction costs) is approximated by vega $*\,(\sigma_{impl} - \sigma_{real})$, when the return of the underlying is lognormally distributed and the hedging is performed with stocks. This can be explained intuitively, because vega is the change of the value of the portfolio due to the change of the value of volatility (Appendix A.3.4).

The expected return of a continuously delta-hedged straddle is only dependent on the future realised volatility. The expected costs of the hedge (to keep the straddle delta-neutral) are equal to the premium that is received for the straddle (one call and one put) at time 0. This premium is the Black and Scholes price of the call and put option. The actual costs of the hedge, however, depend on the future realised volatility and are therefore determined by the Black and Scholes price of the straddle, using the future realised volatility instead of the implied volatility. The return of a straddle is now determined by the difference between the expected costs of the hedge (the premium received for the straddle) and the actual costs of the hedge. To see this, take two portfolios with the same parameters and underlying process, except for the volatility. The first portfolio, P_1, is the portfolio that determines the expected costs of hedging a straddle and uses the implied volatility. The second portfolio, P_2, determines the realised costs of hedging a straddle and uses the realised volatility. The difference between the values of these portfolios

[1] Usually the delta of the model that is assumed to be the underlying model for the option prices is used. This is often the Black and Scholes delta.

determines the return of a continuously delta-hedged straddle. The value of P_1 is given by the Black and Scholes value of a call option and a put option:

$$P_1 = S_0\mathrm{N}(d_{1,impl}) - Ke^{-rT}\mathrm{N}(d_{2,impl}) + Ke^{-rT}\mathrm{N}(-d_{2,impl}) - S_0\mathrm{N}(-d_{1,impl}) =$$
$$2S_0\mathrm{N}(d_{1,impl}) - 2Ke^{-rT}\mathrm{N}(d_{2,impl}) + Ke^{-rT} - S_0,$$

in which:

$$d_{i,impl} = \frac{\ln(S_0/K) + (r - q \pm \sigma^2_{impl}/2)T}{\sigma_{impl}\sqrt{T}}, i = 1, 2.$$

Similarly, the value of P_2 is given by:

$$P_2 = 2S_0\mathrm{N}(d_{1,real}) - 2Ke^{-rT}\mathrm{N}(d_{2,real}) + Ke^{-rT} - S_0,$$

in which:

$$d_{i,real} = \frac{\ln(S_0/K) + (r - q \pm \sigma^2_{real}/2)T}{\sigma_{real}\sqrt{T}}, i = 1, 2.$$

So the difference between P_1 and P_2 is given by:

$$\begin{aligned} P_1 - P_2 = \quad & 2S_0\mathrm{N}(d_{1,impl}) - 2Ke^{-rT}\mathrm{N}(d_{2,impl}) + Ke^{-rT} - S_0 \\ & - \big(2S_0\mathrm{N}(d_{1,real}) - 2Ke^{-rT}\mathrm{N}(d_{2,real}) + Ke^{-rT} - S_0\big) = \\ & 2S_0(\mathrm{N}(d_{1,impl}) - \mathrm{N}(d_{1,real})) - 2Ke^{-rT}(\mathrm{N}(d_{2,impl}) - \mathrm{N}(d_{2,real})) \end{aligned}$$

and is therefore only depending on the difference between realised and implied volatility. The difference between the portfolios (the return of a delta-hedged straddle) can be written as the following Taylor series expansion:

$$\begin{aligned} P_1 - P_2 = \delta P = \frac{\delta P}{\delta\sigma}\delta\sigma + \frac{1}{2}\frac{\delta^2 P}{\delta\sigma^2}\delta\sigma^2 + \ldots = \mathrm{vega}(\sigma_{impl} - \sigma_{real}) + \frac{1}{2}\frac{\delta\mathrm{vega}}{\delta\sigma}(\sigma_{impl} - \sigma_{real})^2 + \ldots \quad (4.1) \\ \approx \mathrm{vega}(\sigma_{impl} - \sigma_{real}) \end{aligned}$$

because $\frac{\delta\mathrm{vega}}{\delta\sigma}$ and higher order derivatives of vega to the volatility become very small.

Since an expected return for a discretely hedged straddle is derived for options on indices, the hedging is done with futures and an extra term for the interest on the start-value is added. The expected return becomes: $\mathrm{vega}(\sigma_{impl} - \sigma_{real}) + (e^{(r-q)T} - 1)$.

Hedging Discretely

The expected return of a short position in a discretely delta-hedged straddle is derived. The expected return is given by:

$$\mathbb{E}[\mathit{TotalReturn}] = \mathbb{E}[W_T] - \mathbb{E}[\mathit{OptionSettle}] + \mathbb{E}[(S_0 + C_0 + P_0)e^{rT}], \quad (4.2)$$

where $\mathbb{E}[W_T]$ is the expected value of the hedge-portfolio at maturity (time T), $\mathbb{E}[\mathit{OptionSettle}]$ is the expected settlements price of the short position of the put and call option (the straddle) at time T, and $(S_0 + C_0 + P_0)e^{rT}$ is the value of the portfolio at time 0 (S_0 is the initial value of the underlying, and C_0 and P_0 are the prices of the call and put respectively that the holder of the short position receives).

The total expected return of a straddle is:

$$\begin{aligned} \mathbb{E}[\mathit{TotalReturn}] = (S_0 + C_0 + P_0)e^{rT} \quad (4.3) \\ + \sum_{j=0}^{n-1}[e^{r(T-\frac{j}{n}T)}(-2S_0e^{\frac{j}{n}T\mu}\mathrm{N}(\frac{a^{(j)} + b^{(j)}(\sigma_{impl} - \kappa)j\sqrt{\frac{1}{n}T}}{\sqrt{1 + b^{(j)2}j}}) + S_0e^{\frac{j}{n}T\mu})] \\ - \sum_{j=0}^{n-1}[e^{r(T-\frac{j+1}{n}T)}(-2S_0e^{\frac{j+1}{n}T\mu}\mathrm{N}(\frac{a^{(j)} + b^{(j)}(\sigma_{impl} - \kappa)j\sqrt{\frac{1}{n}T}}{\sqrt{1 + b^{(j)2}j}}) + S_0e^{\frac{j+1}{n}T\mu})] \\ -S_0\mathrm{N}(d_1^*)e^{\mu T} - K\mathrm{N}(d_2^*) - K\mathrm{N}(-d_2^*) - S_0\mathrm{N}(-d_1^*)e^{\mu T}, \end{aligned}$$

where

N(x) is the cumulative probability function of the standardised normal distribution,

$d_1^* = \frac{\ln(\frac{S_0}{K})+(\mu+\frac{1}{2}\sigma_{real}^2)T}{\sigma_{real}\sqrt{T}}$

$d_2^* = d_1^* - \sigma_{real}\sqrt{T}$

The derivation of the above formula can be found in Appendix B. The derived expectation gives results that are close to vega$(\sigma_{impl} - \sigma_{real}) + (e^{(r-q)T} - 1)$. The short position in a discretely delta-hedged straddle is determined by simulations as well (see Petit et al. [21] and Chapter 5). The results of these simulations showed a distribution of the return with average close to vega$(\sigma_{impl} - \sigma_{real}) + (e^{(r-q)T} - 1)$. In both the theoretical derivation and the simulations, the expected return is a linear function of the difference between implied and realised volatility, unless for a small realised volatility (smaller than 5%).

The results of the simulations performed by Petit et al. [21] also show a non-symmetric distribution around the average (vega$(\sigma_{impl} - \sigma_{real}) + (e^{(r-q)T} - 1)$). It would therefore be interesting to have some information about the second moment of the distribution of the return, which is simulated. The results are presented in Chapter 5.

Figure 4.2 shows the expected return of a discretely hedged straddle and a continuously hedged straddle for an implied volatility of 25% and hedged every hour. If the delta-hedge is performed every day instead of every hour (which is more realistic), no significant changes in the expected return can be seen. The distribution around the expected return is expected to be wider around the average when the delta-hedge is performed less often. This can be derived from the second and higher moments of the distribution of the return of a discretely delta-hedged straddle.

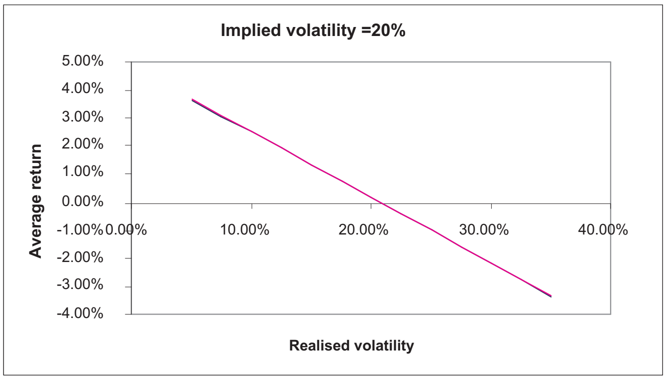

Figure 4.2: The payoff of the discretely delta-hedged straddle (short position)

4.2.2 Variations on Straddles

Not only 'pure' straddles can be used as volatility trading instruments, but also variations on straddles. One example is the β-neutral straddle which is, for example, formed by a normal straddle with an opposite position in an out-of-the money (OTM) call and an OTM put. β is defined as the measure of systematic risk of a portfolio; When the portfolio exactly mirrors the market, $\beta = 1$, when the excess return on the portfolio tends to be exactly as much as the excess return on the market.

Another example of a variation on the straddle is a position consisting of an ATM straddle with a long position in a deeply OTM put to prevent crashes (a crash neutral straddle).

4.2.3 Influence of Different Index Processes on the Return

As we have shown in this section, under the Black and Scholes assumptions the return of a delta-hedged straddle is vega$(\sigma_{impl} - \sigma_{real}) + (e^{(r-q)T} - 1)$ when the rebalancing is done continuously and approximately the same when this is done discretely. As soon as the volatility is not constant anymore (like in the Heston model), the return of a delta-hedged straddle is dependent on the path of the volatility, as will be shown in Section 4.5. Adding jumps to the model for the index process can have a hold on the return. The hedging is often based on a model that does not consider a possibility for jumps, like the Black and Scholes model. This gives an imperfect replication whenever jumps occur. Furthermore, the rebalancing is often performed discretely. Due to this, the jumps that occur between two rebalancing moments can influence the quality of the hedge portfolio (and therefore the return) a lot.

4.3 Strangles

A strangle (see Hull [17] and Appendix A.4.4) involves buying a put with a low strike price and a call with a high strike price. The put and the call option have the same expiration date. The possible strategies concerning volatility using a strangle are similar to those using a straddle, since the payoff patterns of strangles (shown in Figure 4.3) are similar to those of straddles.

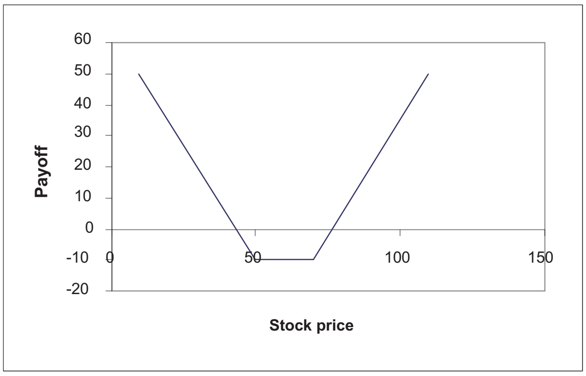

Figure 4.3: The payoff of a long position in a strangle

4.4 Butterflies

A call-butterfly (see Hull [17] and Appendix A.4.1) is created by buying call options with strike prices K_1 and K_3 and selling two call options with strike price K_2, halfway between K_1 and K_3. When the volatility is low, it is profitable to have a butterfly with K_2 close to the current stock price. The profit pattern is shown in Figure 4.4. When the stock movements are large, the profit of the butterfly will be (close to) zero, since the settlement is zero and the investment costs of buying the two calls are a little higher than the premium received for selling the two calls. The same profit pattern can be created by using puts in the same way as the calls are used in the call-butterfly that is described above. When an investor expects the volatility to be high, he can create the opposite butterfly position by selling call options with strike prices K_1 and K_3 and buying two call options with strike price K_2.

4.5 Volatility Swap

An over-the-counter (OTC) product that has the same payoff pattern as the delta-hedged straddle is the volatility swap. A volatility swap is a forward OTC contract on annualised volatility. Its payoff at expiration is equal to $N * (\sigma_{impl} - \sigma_{real})$. σ_{real} is the realised volatility (given by the average of the volatility over the life of the contract), σ_{impl} is the implied volatility and N is the notional amount of the swap in dollars per annualised volatility point.

The volatility swap seems to have the advantage (compared to the delta-hedged straddle) that the payoff only depends on the difference between implied and realised volatility (since it is an OTC product),

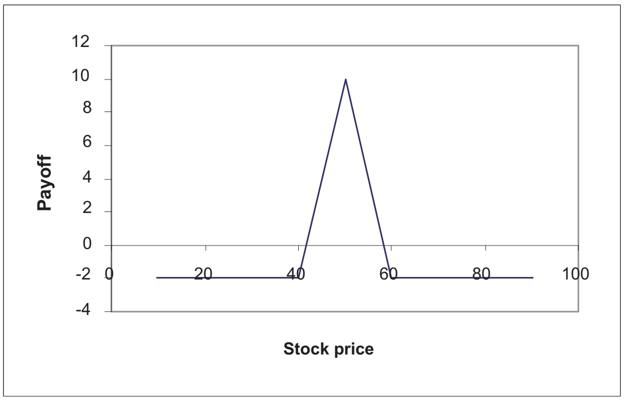

Figure 4.4: The payoff of a long position in a butterfly spread

whereas the return of the delta-hedged straddle is influenced by the level of the underlying. Furthermore, the timing of the volatility influences the return of a delta-hedged straddle, while it does not influence the return of the volatility swap, which is shown in the following example:

An example of the influence of the timing of volatility on the return of a delta-hedged straddle, compared to a volatility swap is given in this example:

Example:
Assume the interest rate is zero, the implied volatility is 15% and there are two indices X and Y as follows:

T	X	Y	$\ln(\frac{X_{i-1}}{X_i})$	$\ln(\frac{Y_{i-1}}{Y_i})$	Return hedge long straddle on X	Return hedge long straddle on Y
0	100	100			101.51	101.51
1	100	110	0.00	0.10	101.51	101.58
2	120	120	0.18	0.09	101.64	111.58
3	110	100	-0.09	-0.18	91.64	91.58
4	100	100	-0.10	0.00	81.64	91.58

The total volatility of X and Y is the standard deviation of the values $\ln(\frac{S_{i-1}}{S_i})$. This is 12.90% for both X and Y, so the return of a short volatility swap is $Notional * (\sigma_{impl} - \sigma_{real}) = Notional * (2.10\%)$ for both X and Y. The return of a short straddle is different for X and Y, since the returns of the hedge portfolios are different (namely 18.36% on X and 8.42% on Y) and the other components that determine the return are the same for X and Y. In this example, the return of the delta-hedged straddle is influenced by the timing of the volatility, whereas the return of the volatility swap is not.

Volatility swaps are, however, not traded in practice. Volatility swaps are products that are traded over-the-counter (OTC) instead of on an exchange. Therefore, the offering party has to replicate these instruments with products that are traded on exchanges, such as options and futures. Since the only known way to replicate this instrument is to keep a position in a dynamically delta-hedged straddle, which gives the exposure to the underlying index and timing of the volatility as is mentioned in this section.

4.6 Variance Swap

A variance swap is a forward contract on annualised variance. Its payoff at expiration is equal to $N * (\sigma^2_{ref} - \sigma^2_{real})$. σ^2_{real} is the realised variance, σ^2_{ref} is the reference variance of the contract and N is the notional amount of the swap in dollars per annualised variance point. The fair level of σ_{ref}^2 is determined by valuing the portfolio that replicates the variance. A variance swap can be replicated more reliably (i.e. the replicating portfolio shows a smaller standard deviation on the returns) than a volatility swap and can therefore be valued more reliably.

A variance swap can be replicated by a log contract and a dynamic position in futures (described in Section 4.6.1). The replication that makes use of the log contract gives a good approximation of a variance swap, but holding far out-of-the-money puts and calls is required. Especially when there are transaction costs, it can be better to avoid holding these far out-of-the-money options until they actually are important for the replication strategy. A strategy that makes use of this is delta-gamma hedging that is described in Section 4.6.2.

4.6.1 Log Contract Hedging

In Demeterfi et al. [10] the replication of a variance swap with a static position in a log contract is proposed. To value the variance swap, in [10] a continuous process of the underlying is assumed (no jump are allowed), which is given by the following process:

$$\frac{dS_t}{S_t} = \mu(t, \ldots)dt + \sigma(t, \ldots)dZ_t, \tag{4.4}$$

with μ the drift and σ the volatility, which are both arbitrary functions of time and other parameters. The value for a forward contract on realised variance (a variance swap) is now given by:

$$F = \mathbb{E}[e^{-rT}(V - K)]. \tag{4.5}$$

The fair delivery value of a variance swap is the strike that gives an expected return of zero. The strike $K_{var} = \mathbb{E}[V]$ is the fair delivery value of future realised variance. Because $V = \frac{1}{T}\int_0^T \sigma^2(t, \ldots)dt$, we have that $K_{var} = \mathbb{E}[V] = \frac{1}{T}\mathbb{E}[\int_0^T \sigma^2(t, \ldots)dt]$. By applying Ito's lemma (see Hull [17]) to $\ln S_t$ one can find: $d(\ln S_t) = (\mu(t, \ldots) - \frac{1}{2}\sigma^2(t, \ldots))dt + \sigma dZ_t$ and by subtracting this from 4.4 the following can be found: $\frac{dS_t}{S_t} - d(\ln S_t) = \frac{1}{2}\sigma^2 dt$. So for the variance we have:

$$V = \frac{1}{T}\int_0^T \sigma^2 dt = \frac{2}{T}[\int_0^T \frac{dS_t}{S_t} - d(\ln S_t)]. \tag{4.6}$$

This indicates that the replication strategy for variance is to have a continuously rebalanced stock position of $\frac{1}{S_t}$ shares (this position is worth 1) and a static position in a contract which at expiration pays the logarithm of the total return (a log contract). We find for the fair variance: $K_{var} = \frac{2}{T}\mathbb{E}[\int_0^T \frac{dS_t}{S_t} - d(\ln S_t)]$.

A log contract is a contract of which the payoff at time T is $\ln(\frac{S_T}{S_0})$. A portfolio of European options with all possible (infinite) strike prices with a forward contract on the stock is a log contract (this is shown below). In practice, no log contracts are produced by an exchange and not all strike prices are available but a certain (discrete) range of strike prices around the reference price (S_*) can be taken to approximate the log contract.

The log payoff can be rewritten as follows: $\ln(\frac{S_T}{S_0}) = \ln(\frac{S_T}{S_*}) + \ln(\frac{S_*}{S_0})$. The second term is constant, so only the first term has to be replicated:

$$-\ln(\frac{S_T}{S_*}) = -\frac{S_T - S_*}{S_*} + \int_0^{S_*} \frac{1}{K^2}\max(K - S_T, 0)dK + \int_{S_*}^{\infty} \frac{1}{K^2}\max(S_T - K, 0)dK. \tag{4.7}$$

The portfolio consists of:

- a short position in $\frac{1}{S_*}$ forward contracts at S_* (which gives the term $-\frac{S_T - S_*}{S_*}$),
- a long position in $\frac{1}{K^2}$ put options at K , for all strikes from 0 to S_* (which gives the term $\int_0^{S_*} \frac{1}{K^2}\max(K - S_T, 0)dK$),

- a long position in $\frac{1}{K^2}$ call options at K, for all strikes from S_* to ∞ (which gives the term $\int_{S_*}^{\infty} \frac{1}{K^2} \max(S_T - K, 0) dK$).

In practice however, this replication is imperfect, since not all strikes are available. If we have a finite number strikes of call options as follows:

$$K_0 = S_* < K_{1c} < K_{2c} < \dots,$$

and a finite number of strikes of put options as follows:

$$K_0 = S_* > K_{1p} > K_{2p} < \dots,$$

then the number of call options of strike K_{nc} is:

$$w_c(K_{nc}) = \frac{f(K_{n+1c}) - f(K_{nc})}{K_{n+1c} - K_{nc}} - \sum_{i=1}^{n-1} w_c(K_{ic}), \tag{4.8}$$

and the number of put options of strike K_{np} is:

$$w_p(K_{np}) = \frac{f(K_{n+1p}) - f(K_{np})}{K_{np} - K_{n+1p}} - \sum_{i=1}^{n-1} w_p(K_{ip}) \tag{4.9}$$

as is derived by Demeterfi et al. in [10].

The replication strategy of holding a static position in a log contract and a dynamic position in the underlying is a good replication for all possible processes of the underlying, as long as the the processes evolve without jumps. When jumps occur, the log contract does not give an expected value that is equal to the variance anymore ($\frac{1}{T}\int_0^T \sigma^2 dt \neq \frac{2}{T}\int_0^T \frac{dS_t}{S_t} - \ln(\frac{S_T}{S_0})$), so the replication is not perfect anymore.

4.6.2 Delta-Gamma Hedging

As mentioned before, it makes sense to avoid holding too many far out-of-the-money options, like in the replicating strategy proposed in the previous section. The delta-gamma hedging strategy can be thought of as constructing only the part of the log contract that is near the current asset price. A delta-gamma hedging strategy holds the following position: x_1 futures on the underlying index and x_2 in short term options with x_1, x_2 as follows:

$$x_1 = V_S - x_2 I_S, \tag{4.10}$$

$$x_2 = \frac{V_{SS}}{I_{SS}}, \tag{4.11}$$

where V is the price of the variance swap, V_S is the first derivative of the variance swap to S, V_{SS} is the second derivative of the variance swap to S. I_S and I_{SS} are the delta and gamma respectively of the secondary instruments that are used for the hedging (options).

Assume, for example, that the strike price of options on the index are spaced by ΔK. The Delta-Gamma hedging strategy is performed by constructing straddle or strangle positions at each volatility observation, using the strike prices close to the current asset price. If the index is between the two strikes K_i and K_{i+1} ($K_{i+1} - K_i = \Delta K$): $K_i < S < K_{i+1}$, then the following position is constructed:

- A straddle position with strike K_i if $S - K_i < 0.2 * \Delta K$;
- A straddle position with strike K_{i+1} if $K_{i+1} - S < 0.2 * \Delta K$;
- An out-of-the-money strangle position using put options with strike K_i and a call option with strike K_{i+1} otherwise.

If an out-of-the-money strangle position is created, the position is only changed when the index moves beyond the strike prices of either the call or the put options. In Windcliff et al. [25], this method is tested and compared to the log contract. Windcliff et al. found that the log contract hedging strategy is similar to the delta-gamma hedging strategy, but the number of positions that need to be taken is much higher for the log contract-method than for the delta-gamma method. The disadvantage of delta-gamma hedging is that the replication is less perfect than the replication that is obtained using the log contract replication, especially when large jumps occur in the index process.

4.6.3 Convexity Bias

In the Sections 4.6.1 and 4.6.2, replication strategies for a variance swap are shown. These replication strategies make no assumptions about the level of future volatility, other than that the stock price evolves without jumps. The replication of volatility swaps, however, is different. It is affected by changes in volatility and the value of a volatility swap therefore depends on the volatility of future realised volatility. The volatility swap can be replicated by holding a variance contract, using the following approximation (for a long position in the contracts):

$$N * (\sigma_{real} - K_{vol}) \approx \frac{N}{2K_{vol}}(\sigma_{real}^2 - K_{vol}^2). \tag{4.12}$$

K_{vol} is the strike of the volatility swap. So $\frac{1}{2K_{vol}}$ variance contracts with strike K_{vol}^2 and notional N can approximate a volatility contract with notional N. The convexity bias is the mismatch between the variance and volatility swap payoffs. If $K_{vol} = \sqrt{K_{var}}$, the payoff of the variance swap compared to the volatility swap is as shown in Figure 4.5. In Figure 4.5, an implied volatility of 30% is used and a long position in a variance swap is shown. If $K_{vol} = \sqrt{K_{var}}$, the convexity bias is given by for a certain realised volatility (σ_{real}):

$$N * (\sigma_{real} - \sqrt{K_{var}}) - \frac{N}{2\sqrt{K_{var}}}(\sigma_{real}^2 - K_{var}) = \frac{N}{2\sqrt{K_{var}}}(\sigma_{real} - \sqrt{K_{var}})^2. \tag{4.13}$$

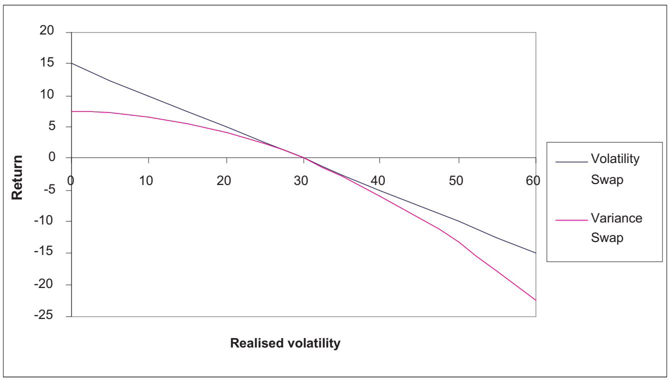

Figure 4.5: Payoff of a volatility swap and a variance swap ($K_{vol} = \sqrt{K_{var}}$)

Usually the variance swap is priced with a reference volatility higher than the implied volatility ($K_{vol} < \sqrt{K_{var}}$) due to the construction of the replication of a variance swap, which is influenced by the skew. In this case the payoff of the long position in a volatility swap will be higher than the payoff of a long position in a variance swap close to $\sigma_{impl} = \sigma_{real}$. An example is shown in Figure 4.6.

The difference between K_{var} and K_{vol} influences the convexity bias as well (as can be seen in Figure 4.6). The difference between K_{vol} and $\sqrt{K_{var}}$ is approximated by:

$$K_{vol} - \sqrt{K_{var}} = \frac{\text{var}(\text{var})}{8 * (K_{var})^{3/2}}$$

which is derived in Appendix C. So when $K_{vol} < \sqrt{K_{var}}$ the convexity bias is approximately given by:

$$N * (\frac{1}{2\sqrt{K_{var}}})(\sigma_{real} - \sqrt{K_{var}})^2 - \frac{\text{var}(\text{var})}{8 * (K_{var})^{3/2}}. \tag{4.14}$$

var(var) is the variance of the future realised variance (the future realised volatility squared). So the size of the convexity bias is influenced by the variance of the future realised variance. The variance of the future realised volatility gives the following relation between $\sqrt{K_{var}}$ and K_{vol} under the risk-neutral measure:

$$\text{var}(\sigma_{real}) = \mathbb{E}[\sigma_{real}^2] - (\mathbb{E}[\sigma_{real}])^2 = K_{var} - K_{vol}^2. \tag{4.15}$$

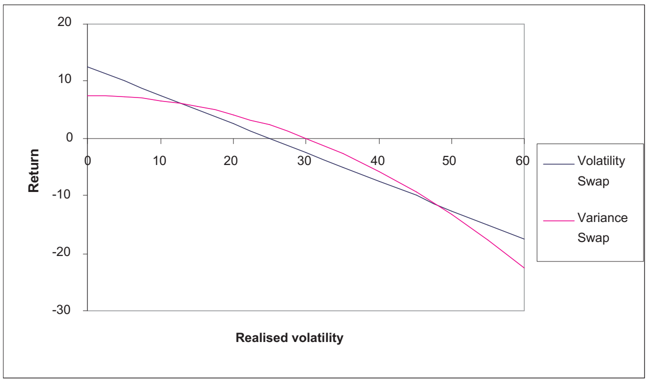

Figure 4.6: Payoff of a volatility swap and a variance swap, $K_{vol} < \sqrt{K_{var}}$

Since K_{var} is determined by the expected value of the future realised variance ($\mathbb{E}[\sigma^2_{real}]$) and K_{vol} by the expected value of the future realised volatility ($\mathbb{E}[\sigma_{real}]$, given by the implied volatility: σ_{impl}).

So the difference between the variance swap rate (K_{var}) and the volatility swap rate squared (K^2_{vol}) is a measure for the volatility in the return volatility.

The skew (the effect of the skew is described in Section 4.6.4) of the implied volatility determines the reference level of the variance swap, K_{var}, so the skew canbe seen as an estimator of the market of the volatility of volatility (due to equation 4.15).

The magnitude of the convexity bias depends a lot on the model that has been chosen, since this magnitude depends on the volatility of realised volatility. When there are no jumps in the process that is assumed for the underlying, the magnitude of the convexity bias is however not so model-dependent. In the next section an example is given for the Heston model.

Convexity Bias in the Heston Model

In the Heston model, the instantaneous variance v follows the process:

$$dv(t) = -\delta(v(t) - \vartheta)dt + \sigma\sqrt{v(t)}dW. \tag{4.16}$$

The total variance is $\mathrm{var}_T = \int_0^\infty v_t dt$. The Laplace transform of the total variance is given by (as can be found in [12]):

$$\mathbb{E}[e^{-\psi \mathrm{var}_T}] = Ae^{-\psi v B}, \tag{4.17}$$

with: $A = \{\frac{2\phi e^{(\phi+\delta)\frac{T}{2}}}{(\phi+\delta)(e^{\phi T}-1)+2\phi}\}^{2\delta\vartheta/\sigma^2}$,

$B = \frac{2(e^{\phi T}-1)}{(\phi+\delta)(e^{\phi T}-1)+2\phi}$,

and $\phi = \sqrt{\delta^2 + 2\psi\sigma^2}$.

This gives for the expected total variance in the Heston model:

$$\mathbb{E}[\mathrm{var}_T] = -\frac{\partial}{\partial\psi}\mathbb{E}[e^{-\psi \mathrm{var}_T}] \mid_{\psi=0} = \frac{1 - e^{-\delta T}}{\delta}(v - \vartheta) + \vartheta T. \tag{4.18}$$

When $v = \vartheta$,

$$\mathrm{var}[\mathrm{var}_T] = \mathbb{E}[\mathrm{var}^2_T] - (\mathbb{E}[\mathrm{var}_T])^2 = \frac{\partial^2}{\partial\psi^2}\mathbb{E}[e^{-\psi \mathrm{var}_T}] \mid_{\psi=0} -(\frac{\partial}{\partial\psi}\mathbb{E}[e^{-\psi \mathrm{var}_T}] \mid_{\psi=0})^2 = \tag{4.19}$$

$$\vartheta T\frac{\sigma^2}{\delta^2} + O(T^0).$$

So, as $T \to \infty$, the standard deviation of the annualised variance has the leading order behaviour $\sqrt{\frac{\vartheta}{T}}\frac{\sigma}{\delta}$, so this is not very model dependent, but quite dependent on the choice of parameters.

4.6.4 Effects of Skew

Because the skew influences the prices of the set of put and call options that replicate a log contract, the skew has influence on the fair value of the variance swap. Demeterfi et al. [10] derived the effects of a skew linear in strike and the effects of a skew linear in delta. In this thesis the effect of a more general skew is derived.

The skew linear in the strike is given by:

$$\Sigma(K) = \Sigma_0 - b_1 \frac{K - S_F}{S_F}, \tag{4.20}$$

where $S_F = S_0 e^{rT}$ the forward index value, Σ_0 is the at-the-money forward implied volatility ($= \Sigma(K = S_F)$) and b_1 is the slope of the skew. In [10] is derived that the fair variance with a skew linear in the strike is:

$$K_{var} = \Sigma_0^2(1 + 3Tb_1^2 + \ldots). \tag{4.21}$$

The skew increases the value of the fair variance (it becomes larger than the at-the-money forward level of volatility), and the size of the increase is proportional to the time to maturity times the square of the skew slope.

The skew linear in the delta is given by:

$$\Sigma(\Delta) = \Sigma_0 - b_2(\Delta_p + \frac{1}{2}), \tag{4.22}$$

where Δ_p is the Black and Scholes delta of a put options which is given by $\Delta_p = -\mathrm{N}(-d_1)$ with $d_1 = \frac{\ln(\frac{S_F}{K}) + \frac{1}{2}\Sigma_0^2 T}{\Sigma_0\sqrt{T}}$ and b_2 is the slope of the skew. In [10] is derived that the fair variance with a skew linear in the strike is:

$$K_{var} = \Sigma_0^2(1 + \frac{1}{\sqrt{\pi}} b_2\sqrt{T} + \frac{1}{12}\frac{b_2^2}{\Sigma_0^2} + \ldots). \tag{4.23}$$

Here, the first order correction (on the at-the-money forward volatility) is of magnitude $b_2\sqrt{T}$. This is due to the fact that a variation in delta is not equivalent to a variation linear in the strike.

In Demeterfi et al. [10], the effects of the skew linear in the strike and linear in the delta are derived. In this thesis, a the influence of a skew which is given by a more general function of the strike is derived in a similar way as Demeterfi et al. did in [10].

The skew is now given by:

$$\Sigma(K) = \Sigma_0 + b_3 f(K), \tag{4.24}$$

where f is an arbitrary function of K (e.g. quadratic) and b_3 an arbitrary (small) value. In Appendix D the fair value of the reference variance is derived and the following expression is found:

$$K_{var} = \Sigma_0^2 + b_3(\frac{2}{T}e^{rT})\{\int_0^\infty \frac{1}{K^2} f(K) \frac{S\sqrt{T}}{\sqrt{2\pi}} e^{-d_1^2/2} dK\} \tag{4.25}$$

$$+\frac{1}{2}b_3^2(\frac{2}{T}e^{rT})\{\int_0^\infty \frac{1}{K^2} f^2(K)(\frac{\ln(\frac{S_F}{K})}{\Sigma_0^2\sqrt{T}} - \frac{1}{2}\sqrt{T})\frac{S\sqrt{T}}{\sqrt{2\pi}} d_1 e^{-d_1^2/2} dK\} + \ldots.$$

For $f(K) = \frac{K-S_F}{S_F}$, the same derivation as in [10] is obtained. For $f(K) = -\frac{K-S_F}{S_F} + (\frac{K-S_F}{S_F})^2$ we have the type of skew that is shown in Figure 4.7.

$f(K) = -\frac{K-S_F}{S_F} + (\frac{K-S_F}{S_F})^2$ gives the following fair variance for a variance swap:

$$\begin{aligned}
K_{var} = \; & \Sigma_0^2 - b_3(\tfrac{2}{T}e^{rT})\left\{\int_0^\infty \tfrac{1}{K^2} f(K) \tfrac{S\sqrt{T}}{\sqrt{2\pi}} e^{-d_1^2/2} dK\right\} \\
& +\tfrac{1}{2}b_3^2(\tfrac{2}{T}e^{rT})\left\{\int_0^\infty \tfrac{1}{K^2} f^2(K)(\tfrac{\ln(\frac{S_F}{K})}{\Sigma_0^2\sqrt{T}} - \tfrac{1}{2}\sqrt{T})\tfrac{S\sqrt{T}}{\sqrt{2\pi}} d_1 e^{-d_1^2/2} dK\right\} + \ldots \\
= \; & \Sigma_0^2 - b_3(\tfrac{2}{T}e^{rT})\left\{\int_0^\infty \tfrac{1}{K^2}(\tfrac{K-S_F}{S_F} - (\tfrac{K-S_F}{S_F})^2)\tfrac{S\sqrt{T}}{\sqrt{2\pi}} e^{-d_1^2/2} dK\right\} \\
& +\tfrac{1}{2}b_3^2(\tfrac{2}{T}e^{rT})\left\{\int_0^\infty \tfrac{1}{K^2}((\tfrac{K-S_F}{S_F})^2 - 2(\tfrac{K-S_F}{S_F})^3\right. \\
& \left. + (\tfrac{K-S_F}{S_F})^4)(\tfrac{\ln(\frac{S_F}{K})}{\Sigma_0^2\sqrt{T}} - \tfrac{1}{2}\sqrt{T})\tfrac{S\sqrt{T}}{\sqrt{2\pi}} d_1 e^{-d_1^2/2} dK\right\} + \ldots.
\end{aligned} \tag{4.26}$$

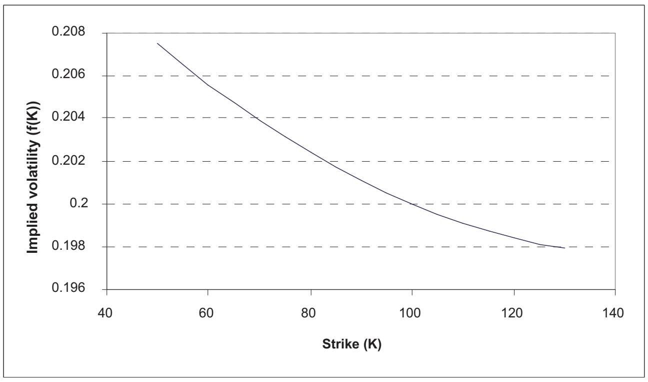

Figure 4.7: Volatility skew for $f(K) = -\frac{K-S_F}{S_F} + (\frac{K-S_F}{S_F})^2$

By changing the integration variable to $z = \frac{\ln(\frac{S_F}{K})+\frac{1}{2}v_0}{\sqrt{v_0}} (= d_1)$, where $v_0 = \Sigma_0^2 T$, the equation for the fair variance becomes:

$$\begin{aligned}
K_{var} = \;& \Sigma_0^2 - b_3 \left\{ \int_{-\infty}^{\infty} 2\Sigma_0^2 \frac{1}{\sqrt{2\pi}} e^{-z^2/2} (e^{z\sqrt{v_0}-\frac{v_0}{2}} - 1) dz \right. \\
& \left. + \int_{-\infty}^{\infty} 2\Sigma_0^2 \frac{1}{\sqrt{2\pi}} e^{-z^2/2} (e^{-z\sqrt{v_0}+\frac{v_0}{2}} + e^{z\sqrt{v_0}-\frac{v_0}{2}} - 2) dz \right\} \\
& + \frac{1}{2} b_3^2 \left\{ - \int_{-\infty}^{\infty} \frac{2}{\sqrt{2\pi}} (z - \frac{\sqrt{v_0}}{2}) (e^{z\sqrt{v_0}-\frac{v_0}{2}}) z e^{z^2/2} \right. \\
& \left((e^{-z\sqrt{v_0}+\frac{v_0}{2}} - 1)^2 - 2(e^{-z\sqrt{v_0}+\frac{v_0}{2}} - 1)^3 + (e^{-z\sqrt{v_0}+\frac{v_0}{2}} - 1)^4 \right) dz \\
& + \int_{-\infty}^{\infty} \frac{\sqrt{v_0}}{\sqrt{2\pi}} (e^{z\sqrt{v_0}-\frac{v_0}{2}}) z e^{z^2/2} \\
& \left. \left((e^{-z\sqrt{v_0}+\frac{v_0}{2}} - 1)^2 - 2(e^{-z\sqrt{v_0}+\frac{v_0}{2}} - 1)^3 + (e^{-z\sqrt{v_0}+\frac{v_0}{2}} - 1)^4 \right) dz \right\}.
\end{aligned} \tag{4.27}$$

This gives: $K_{var} = \Sigma_0^2(1 - 2b_3\Sigma_0 T - \frac{11}{2} b_3^2 T + \ldots)$.

So the skew decreases the value of the fair variance compared to the at-the-money forward level of volatility and the size of the decrease is proportional to the time to maturity, the skew slope and the at-the-money implied volatility.

4.6.5 Effects of Jumps

In many stochastic volatility models in which jumps are added to the process, the jumps are assumed to follow a Poisson process with arrival intensity λ and the size of the jump lognormally distributed with mean α and standard deviation δ. Gatheral and Lynch [12] show that the error caused by this jump-process is very small and given by:

Jump error= $\frac{1}{3}\lambda T\alpha(\alpha^2 + 3\delta^2)$+*higher order terms.*

For $\lambda = 0.02$, $\alpha = -0.45$ and $\delta = 0.45$ this error is 0.0024 on a one-year variance swap (this corresponds with 0.4 vol points if the volatility level is 30%), so, even with jumps, the log contracts seem to be a good average hedge for a variance swap.

4.6.6 Effects of Index Processes

As is shown in this section, under the Black and Scholes model with a constant volatility the return of a (long) volatility swap is $N * (\sigma_{real} - K_{vol})$ and the return of a comparable variance swap is $\frac{N}{2\sqrt{K_{var}}} *$

$(\sigma^2_{real} - K_{var})$. Not only under Black and Scholes assumptions, but for all continuously evolving processes of the underlying, the replication of the variance swap can be done by a log contract. When jumps occur, this replication is not perfect anymore, since the log contract does not give an expected value equal to the variance $(\frac{1}{T}\int_0^T \sigma^2 dt \neq \frac{2}{T}\int_0^T \frac{dS_t}{S_t} - \ln(\frac{S_T}{S_0}))$. The influence of the volatility skew is discussed in Section 4.6.4 and the influence of jumps in Section 4.6.5. The difference between K_{vol} and $\sqrt{K_{var}}$, the convexity bias, is independent of the chosen model, assuming there are no jumps.

4.7 Conclusions

The most basic volatility instrument is the straddle (described in Section 4.2), on which variations are possible such as the strangle and butterflies. The straddle has the disadvantage that it is very sensitive for movements in the price of the underlying. This can be neutralised by delta-hedging the straddle. The delta-hedged straddle has an expected return that is linear in the difference between implied and realised volatility. The delta-hedged straddle, however, has the disadvantage that the hedge is dynamic and therefore very sensitive for the timing of the volatility. An over-the-counter instrument that has the same return as the expected return of the delta-hedged straddle is the volatility swap. The volatility swap is not traded in practice, since it has to be replicated by the dynamically delta-hedged straddle, which is sensitive for the timing of the volatility.

Another over-the-counter instrument is the variance swap. The variance swap can be replicated by a static set of options (i.e. a log contract) and is therefore not sensitive to the timing of the volatility, but only to the realised volatility at maturity. The variance swap has a return that is linear in the *variance*.

Furthermore we have seen that the skew implies the size of the convexity bias and gives a market estimator of the volatility of volatility.

CHAPTER 5

Volatility Instruments - Simulations

In the previous chapter, a theoretical analysis of the volatility instruments is done and a lot of characteristics of the volatility instruments were derived. More performance characteristics of the volatility instruments can be analysed by performing simulations. It is interesting to obtain information about the performance of the volatility instruments when different processes are assumed for the underlying. Therefore, two different index processes are simulated by Monte Carlo simulation. One is the Black and Scholes process, which is described in Chapter 2. The other is the Heston process (described in Section 3.2.2) with a component added to simulate jumps in the index process. The arrival process of the jumps is a Poisson process with a certain arrival intensity. The jumpsize is lognormally distributed.

The main difference between the processes is that the Black and Scholes process assumes deterministic volatility and the Heston process has stochastic volatility. The simulations are performed to obtain more information about the influence of the different processes on the return of the volatility instruments. Furthermore, it is important to know more about the performance of the different volatility instruments compared to each other in different situations. In the simulations that have been performed, no transaction costs are assumed and a mid-price is taken (no bid-ask spread is taken into account).

The parameters that are used to simulate these indices are as follows:

Number of rebalancings per day:	1,
Simulation time:	22 days (252 days per year),
The number of simulations that have been done per run:	1000,
The initial index level:	100,
The strike price:	100,
The expected return on a stock:	10%,
The interest rate:	0,
The linear skewness:	1%,
The implied volatility:	20%,
The expected realised volatility:	15%
More simulated levels of realised volatility:	5%, 10%, 15%, 20%, 25%, 30%, 35%, 40%, 45% and 50%,

The notional for the volatility swap:	1,
The notional for the variance swap:	1/(2*Reference volatility) (as is described in Section 4.6.3),
The dividend yield:	0,
The difference between the strikes of the options that are used to replicate a log contract:	5%,
The strikes:	in the range from 40% to 200%.

The parameters that are specific for the Heston model with jumps are chosen such that the process is realistic:

The arrival intensity of the jumps:	1/20 (one jump in 20 years),
The size of the jumps	20%,
The correlation coefficient of the Wiener processes determining the index process and the volatility process:	-0.01,
The speed of the mean reversion:	10,
The volatility of volatility	10%.

The return of some volatility instruments is analysed using these simulated indices. The most basic volatility instrument is the volatility swap, which was described theoretically in Section 4.5. The payoff of this instrument is based directly on the volatility of the underlying index. The volatility swap is an over-the-counter product. Another instrument that has the same expected return as the volatility swap (linear in the realised volatility) and is therefore used to replicate the volatility swap, is the delta-hedged straddle. The delta-hedged straddle is a combination of options and futures and therefore not an over-the-counter product such as the volatility swap. A theoretical description of the delta-hedged straddle can be found in Section 4.2.1.

Another over-the-counter volatility instrument is the variance swap. The variance swap has a payoff that is not linear in the volatility, but is easier to replicate and therefore easier to understand theoretically. More theoretical information on the variance swap is given in Section 4.6. The variance swap is an over-the-counter product, like the volatility swap. There is, however, a possibility to find a good replication of the variance swap using a static position in options and a dynamic position in futures. This is described in Section 4.6.1. The replication can be used to find the fair strike price (reference variance) of the variance swap.

5.1 General Results in the Black and Scholes World

5.1.1 Return of the Instruments in the Black and Scholes World

The returns of the instruments are analysed using a simulation of a Black and Scholes process with an implied volatility of 20% and a realised volatility varying between 5% and 50%. The expected returns of the volatility swap and of the delta-hedged straddle are equal and look exactly like we expect, considering the theory (Sections 4.2.1 and 4.5). The expected returns of the volatility swap and of the delta-hedged straddle are given in Figure 5.1.

The replication of the volatility swap (the dynamically delta-hedged straddle) is not perfect, since the dynamic hedging cannot be performed continuously. Therefore, given the level of realised volatility the delta-hedged straddle has a standard deviation that is larger than zero, whereas the standard deviation of the volatility swap is zero by definition. In Figure 5.2 the distribution of returns of the delta-hedged straddle is given for different levels of realised volatility.

Figure 5.2 shows that the distribution of returns is wider when the realised volatility is further from the implied volatility. This is due to the quality of the hedge-portfolio, which will be worse when the realised volatility differs more from the implied volatility, since the implied volatility is used to determine the hedge-portfolio. Furthermore, the quality of the hedge will be worse when the realised volatility is large, due to the large range of possible index-levels. This gives a wider distribution of returns for a larger realised volatility.

When the level of volatility changes (but the difference between the implied volatility and realised volatility stays the same), the standard deviations of the returns of the replication of the volatility swap

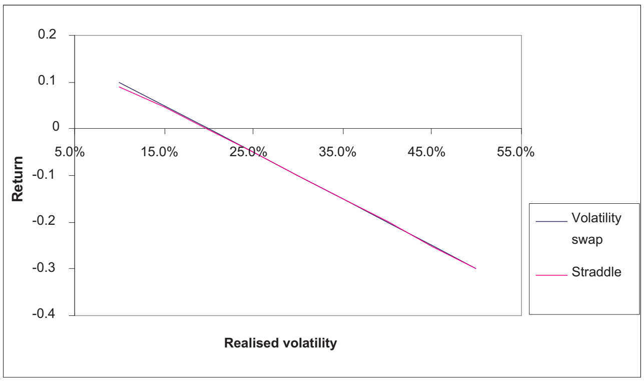

Figure 5.1: Average return of a delta-hedged straddle and a volatility swap

(the delta-hedged straddle) and the replication of the variance swap increase slowly. In Table 5.1 the standard deviations of the delta-hedged straddle and the replication of the variance swap are given for different levels of volatility.

Volatility Level	Standard deviation delta-hedged straddle	Standard deviation replicated variance swap
Implied volatility=20% Realised volatility=15%	2.50 %	0.71 %
Implied volatility=25% Realised volatility=20%	2.74 % (increase: 0.24%)	0.76 % (increase: 0.05%)
Implied volatility=30% Realised volatility=25%	3.33 % (increase: 0.59%)	0.92 % (increase: 0.16%)
Implied volatility=35% Realised volatility=30%	3.71 % (increase: 0.38%)	1.00 % (increase: 0.08%)
Implied volatility=40% Realised volatility=35%	4.18 % (increase: 0.47%)	1.19 % (increase: 0.19%)
Implied volatility=45% Realised volatility=40%	4.57 % (increase: 0.39%)	1.47 % (increase: 0.28%)

Table 5.1: Increase of standard deviation for different levels of volatility

Table 5.1 shows that the standard deviation of the delta-hedged straddle is larger than the standard deviation of the replication of the variance swap. Furthermore, Table 5.1 shows that the standard deviation of the delta-hedged straddle increases more than the standard deviation of the replication of the variance swap when the level of the volatility increases. This is due to the fact that the level of the index return influences the return of the delta-hedged straddle more than the return of the replication of the variance swap. The return of a short position in a variance swap is determined by $\frac{N}{2\sigma_{ref}}(\sigma_{ref}^2 - \sigma_{real}^2)$.

Demeterfi et al. in [10] do a proposition to give a good reference variance (σ_{ref}^2). In [10], the fair value of future realised variance (the reference variance) is estimated using a static position in a perfect log contract (a contract that pays the logarithm of the total return at expiration) and a dynamic position in futures (more information about this replication can be found in Section 4.6.1 and in [10]). A perfect log contract does not exist and the replication of the log contract consists of a set of options and a forward contract. The reference variance is therefore determined using a replication that is as close to a perfect log contract as possible. Using the replication of the log contract instead of a perfect log contract does not have much influence on the average return and the confidence interval around the average return.

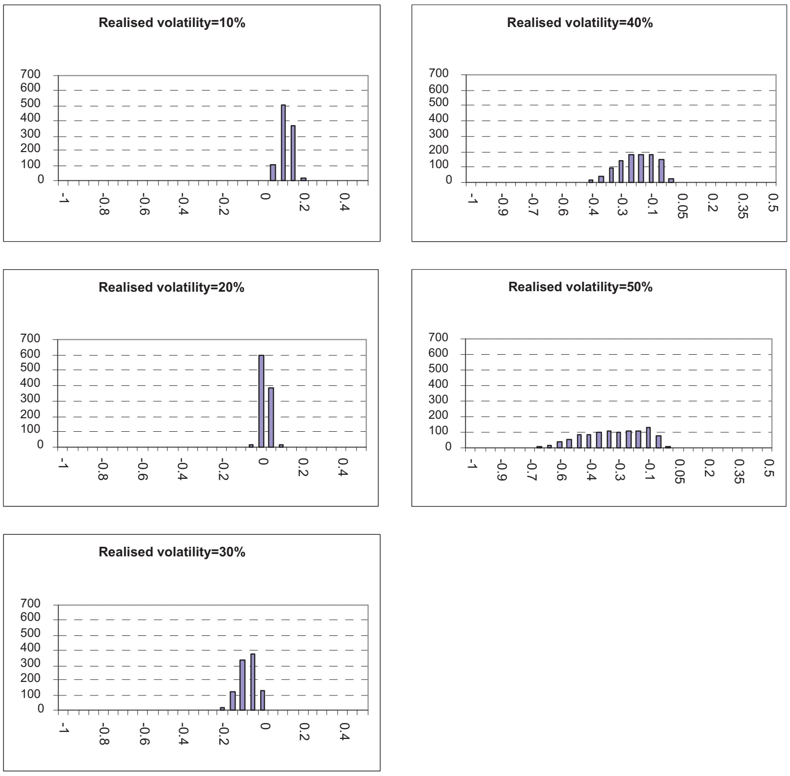

Figure 5.2: The distribution of the return of a delta-hedged straddle

The expected return of the variance swap and the replication of the variance swap are equal to the return that is expected from the theory (Section 4.6). Figure 5.3 shows the returns of the variance swap and the replication of the variance swap. In Figure 5.3 the return of the volatility swap is also presented to show the difference between volatility and variance swaps. Furthermore, the confidence intervals of the delta-hedged straddle and the replication of the variance swap are given in Figure 5.3. The confidence interval of the delta-hedged straddle is much larger than the confidence interval of the replication of the variance swap. This indicates a more reliable replication for the variance swap than for the volatility swap.

The replication of the variance swap has an expected return equal to the return of the variance swap, but has a standard deviation that is not equal to zero. The variance swap has a standard deviation of zero, by definition. In Figure 5.4, the probability distribution of the variance swap is given. Like the distribution of the delta-hedged straddle, the distribution of the replication of the replication of the variance swap is wider when the realised volatility is larger, due to the quality of the hedge when the realised volatility is large, which is worse when the range of possible index-levels is large.

5.1.2 Influence of the Skew

The simulations showed that due to the effect of the skew that implies the convexity adjustment, the variance swap (or the replication of the variance swap) performs better than the volatility swap (or the replication of the volatility swap) when the realised volatility is between 10% and 35%. When the realised volatility is smaller than 10% or larger than 35%, the performance of the variance swap is worse than the performance of the volatility swap. The skew, which is implied from the market prices of options, can be seen as an estimator of the volatility of volatility, as is described in Section 4.6.3. It is interesting to find out if a certain realistic skew gives a good estimator for the volatility of volatility. When the skew gives an implied volatility of volatility (volatility of volatility that is estimated by the market) that is

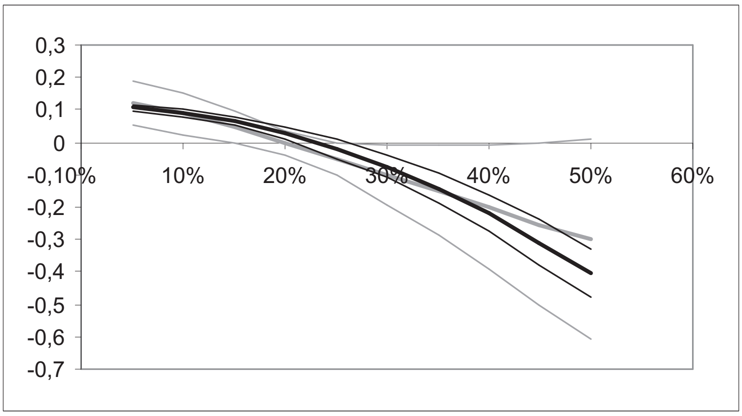

Figure 5.3: The average return and confidence intervals of the replication of the variance swap (dark lines) and volatility swap (light lines)

larger than the realised volatility of volatility, the variance swap will on average be more profitable than the volatility swap, because a risk premium is paid for the larger risk due to volatility. To find out if the implied volatility of volatility is larger than the realised volatility of volatility, the difference between the old calculation of the VIX[1] and the new calculation of the VIX can be compared to the realised volatility. The calculation of the old VIX finds the implied volatility, whereas the new VIX calculates the reference volatility (as is described by Demeterfi et al. [10]). The difference between the reference volatility and the implied volatility gives an estimator (as is described in Section 4.6.3) for the implied volatility of volatility. From the realised volatility, the realised volatility of volatility can be derived by calculating the standard deviation of the realised volatility. This is left for future research.

5.1.3 Influence of the Frequency of Rebalancing

In Section 4.2.1, the expected return of the discretely delta-hedged straddle is determined. In this section was stated that the frequency of rebalancing did not influence the average return, but it was expected to influence higher moments of the distribution of return. No expressions were derived for these higher moments, but using the results of the simulations more can be said about the influence of the frequency of rebalancing.

In the simulation of the Black and Scholes index, the influence of the frequency of rebalancing is as expected from the theory (Section 4.2.1). The frequency of rebalancing does not have any significant influence on the average return of any instrument. For the instruments for which rebalancing is used (the delta-hedged straddle and the replication of the variance swap), however, the standard deviation increases as the frequency of rebalancing decreases, especially for the higher values of realised volatility. This is shown in Figure 5.5, which depicts the standard deviation of the delta-hedged straddle and the replicated variance swap for three different frequencies (the darker the line, the lower the frequency).

5.2 Performance in the Heston World Compared to the Black and Scholes world

As is mentioned before, two different index processes are simulated. The return characteristics of the different instruments in the Black and Scholes world are discussed in the previous section. The other process that is simulated is the Heston process (described in Section 3.2.2). The performance of the different volatility instruments in both the Heston world and the Black and Scholes world is analysed and

[1]The VIX is a volatility index for the Chicago Board Options Exchange.

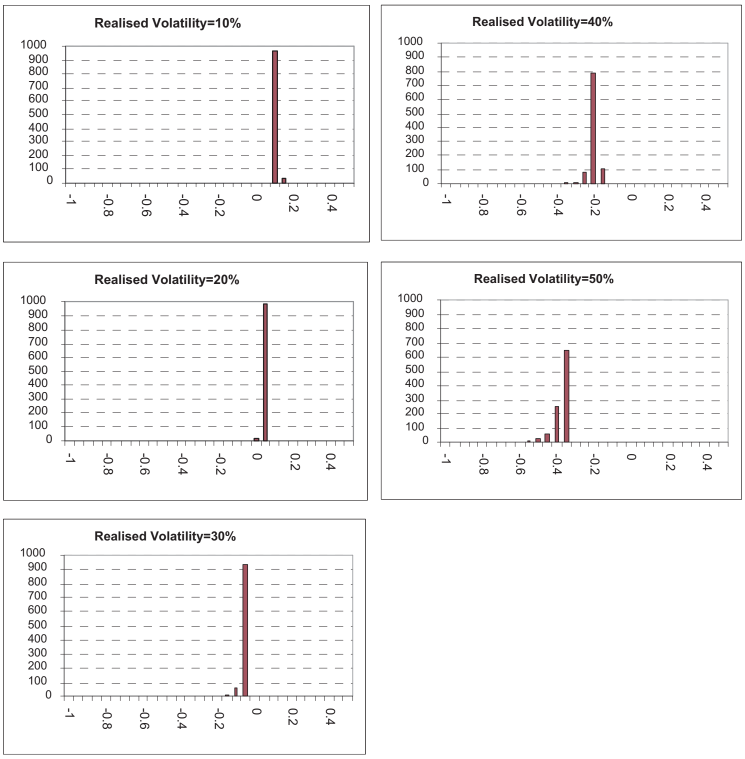

Figure 5.4: The Distribution of the Return of a Replication of a Variance Swap

compared by analysing some performance measures and analysing the influence of the characteristics of the simulated index processes.

5.2.1 Performance Measures

Both given a Black and Scholes process and given a Heston process, the performance of the volatility instruments is analysed and compared using several performance measures. The performance measures that are used are the average return, the median of the return, the standard deviation of the return, the downside risk, and the Sharpe ratio.

The average return of the 1000 simulations that are performed given the Black and Scholes index process is shown in Table 5.2. In Table 5.3 the average returns of the volatility instruments given the Heston process are depicted. The average returns of all volatility instruments are smaller for the Heston process than for the Black and Scholes assumptions. The average returns of the variance swap and its replication however are more different than the average return of the volatility swap and the delta-hedged straddle, due to the non-linear payoff pattern of the variance swap and its replication. Furthermore, the medians do not show many differences between the Heston model and the Black and Scholes model. This suggests more skewness in the returns when there is stochastic volatility and jumps (the Heston assumptions are used). Especially for the variance swap and its replication, the standard deviation of the returns is much higher in Table 5.3 than in Table 5.2. This is due to the stochasticity in the Heston index process. The standard deviations of the variance swap and its replication increase more when there is more stochasticity than the standard deviations of the volatility swap and its replication. This is due to the non-linear payoff pattern of the variance swap.

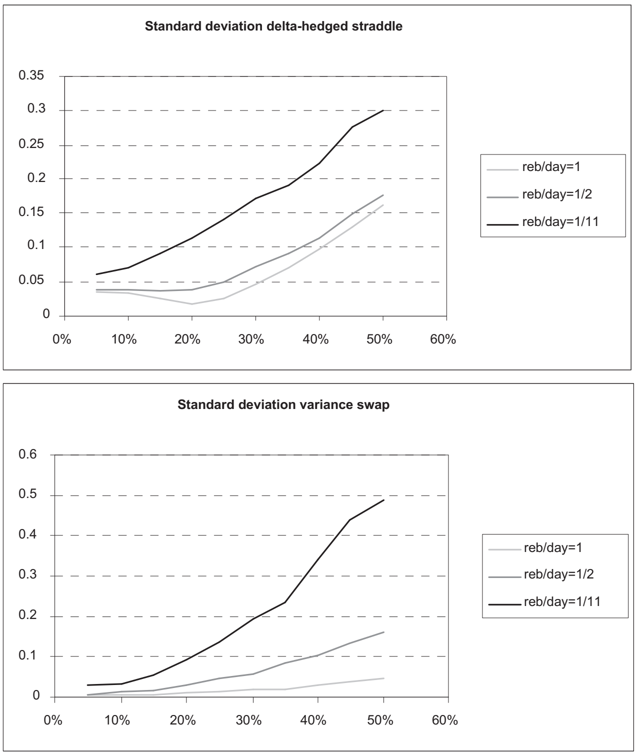

Figure 5.5: The influence of the frequency of rebalancing on the standard deviation

The downside risk of the volatility instruments is the fraction of the total number of simulations for which the returns are negative. The downside risk is shown in the third column of Table 5.2 and Table 5.3. The downside risk is lower for the variance swap and its replication than for the volatility swap and the delta-hedged straddle, both when the Heston simulation and when the Black and Scholes simulation is considered. Furthermore, the downside risk in the Black and Scholes simulation is not very different from the downside risk in the Heston simulation.

The last performance measure we consider is the Sharpe ratio. The Sharpe ratio (reward-to-volatility ratio) is defined by the excess return divided by the standard deviation of the return[2]. The Sharpe ratios are lower in the Heston process for all instruments, due to the higher standard deviation of the return and the lower average return in the Heston process. Furthermore, the Sharpe ratio gives better results for the variance swap and its replication in the Black and Scholes process compared to the volatility swap and the delta-hedged straddle. This is due to the higher average returns and the smaller standard deviations for the variance swap and its replication, which is dependent on the size of the skew compared to the realised number of jumps (which is zero in the Black and Scholes process). Table 5.3 shows exactly the opposite result for the Heston process. This is due to the smaller average returns and the higher standard deviations for the variance swap and its replication in the Heston model, caused by more realised jumps compared to the same skew.

5.2.2 Influence Skew

In the simulations a linear skew is used. Both the simulations that make use of the Black and Scholes assumptions and the simulations that make use of the Heston assumptions show larger average returns for the variance swap and its replication when the skew is chosen larger. In the Black and Scholes simulations,

[2]The Sharpe ratio assumes that returns are fully described by their mean and standard deviation (i.e. normal). This is not the case in the returns of the volatility instruments, but still the Sharpe ratio gives information. On top of that the Sharpe ratio is used here since it is used very often in practice.

	Average return	Median	Standard deviation	Downside risk	Sharpe ratio
Volatility swap	5.14%	5.21%	2.36%	1.7%	2.182
Straddle	4.97%	4.66%	3.16%	4.3%	1.572
Variance swap	6.83%	6.99%	1.52%	0%	4.479
Replication variance swap	6.52%	6.60%	1.48%	0.1%	4.397

Table 5.2: Results of 1000 runs in the Black and Scholes world (the implied volatility is 20% and the realised volatility is 15% on average)

	Average return	Median	Standard deviation	Downside risk	Sharpe ratio
Volatility swap	4.91%	5.22%	4.51%	2.1%	1.087
Straddle	4.69%	4.75%	6.44%	4.6%	0.729
Variance swap	6.36%	7.00%	7.75%	0.4%	0.821
Replication variance swap	6.12%	6.72%	6.97%	0.4%	0.878

Table 5.3: Results of 1000 runs in the Heston world (the implied volatility is 20% and the realised volatility process is stochastic, but mean reverting to 15%)

the average returns of all instruments are equal when there is no skew and become larger for the variance swap and its replication when the skew is larger. Furthermore, the standard deviation of the return of the instruments becomes smaller for the variance swap and its replication when the skew is larger. This gives a larger Sharpe ratio for these instruments for a nonzero skew.

The skew gives an indicator for the volatility of volatility, estimated by the market (as is described in Section 4.6.3). When the risk of volatility of volatility estimated by the market is larger than the actual risk, the performance of the variance swap and its replication is better than the performance of the volatility swap and delta-hedged straddle.

5.2.3 Influence of the Chosen Heston-Parameters

The influence of volatility of volatility is distinguishable in the simulation results, whereas the influence of the correlation coefficient and the speed of the average reversion is not clear from the results. A larger volatility of volatility gives an average return of the variance swap (and its replication) that is closer to the average return of a volatility swap (and a delta-hedged straddle). When the volatility of volatility is small, the average return of a variance swap (and its replication) is much larger than the average return of the volatility swap (and the delta-hedged straddle). This can be explained by the return pattern. Figure 5.6 shows that a variance swap will give a much larger average return compared to the volatility swap when the volatility is between a% and b% (smaller volatility of volatility) than when the volatility is between c% and d% (larger volatility of volatility).

5.2.4 Influence of the Characteristics of the Jump Process

The jumps that occur have influence on the return. Both when more jumps occur and when the mean jump size is larger, the average return of all instruments, but especially the average return of the variance swap and its replication (due to the non-linear payoff pattern) decreases. Furthermore, the standard

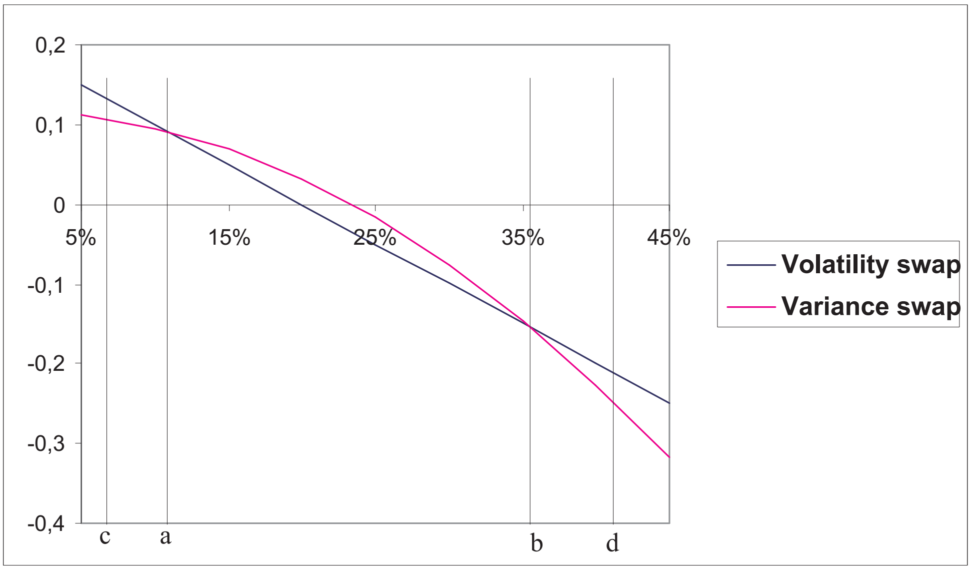

Figure 5.6: The influence of the volatility of volatility on expected return of volatility swap and variance swap

deviation of the return of all instruments (especially of the variance swap and its replication) increases. This gives better Sharpe ratios for the variance swaps when no jumps or only very small jumps occur and worse results when jumps occur. In the Tables 5.4 and 5.5, the average returns and standard deviations of the returns for four different arrival intensities are given respectively. For an arrival intensity of 15 jumps per 40 years (1 jump in $2\frac{2}{3}$ year), the expected returns of the volatility swap and the variance swap are equal. From this can be concluded that if the implied volatility has a skew like the one assumed here, the estimation of the market considering the volatility of volatility is such that on average one jump in $2\frac{2}{3}$ year occurs.

Arrival Intensity	0	1/20	4/20	15/40
Volatility swap	5.20%	4.93%	3.89%	3.49%
Straddle	5.03%	4.66%	3.38%	2.70%
Variance Swap	6.87%	6.37%	4.31%	3.49%
Replication of Variance swap	6.58%	6.10%	4.30%	3.53%

Table 5.4: Average return of the volatility instruments for different arrival intensities

A cap can be added to a volatility or variance contract to prevent great losses. A cap on a volatility or variance swap is very cheap in practice. The capped volatility swap and capped variance swap are simulated as well. The results of these instruments are compared to the normal volatility and variance swap. The results are depicted in Table 5.6. Table 5.6 shows no difference between the capped and the normal instruments for an arrival intensity of the jumps of 1/20 (on average one jump per 20 years).

Arrival intensity	0	1/20	4/20	15/40
Volatility swap	2.34%	4.58%	9.19%	10.48%
Straddle	3.22%	6.74%	12.66%	14.56%
Variance swap	1.51%	8.01%	17.88%	20.78%
Replication of variance swap	1.51%	7.29%	15.90%	18.67%

Table 5.5: Standard deviation of the return of the volatility instruments for different arrival intensities

Therefore the theoretical price of a cap is almost zero. If the arrival intensity of the jumps is larger, the cap will influence the average return more. In that situation, the cap will be worth more.

Arrival intensity	0	1/20	4/20	15/40
Volatility swap	5.20%	4.91%	3.99%	3.49%
Capped Volatility swap	5.20%	4.91%	3.99%	3.49%
Difference	0%	0%	0%	0%
Variance swap	6.87%	6.36%	4.82%	3.49%
Capped Variance swap	6.87%	6.42%	5.02%	3.99%
Difference	0%	0.055%	0.204%	0.502%

Table 5.6: Average return of the capped volatility instruments compared to the normal volatility instruments

5.3 Conclusions

Summarising, the most important conclusions that can be drawn from the simulations are:

- From the simulations can be concluded that the confidence interval of the replication variance swap is much smaller than the confidence interval of the replication of the volatility swap (i.e. the delta hedged straddle).
- In Chapter 4 is concluded that the skew gives an estimator for the volatility of volatility. From the simulations is concluded that when the skew gives a larger estimator than the realised volatility of volatility, the variance swap gives higher returns than the volatility swap and when the skew gives a smaller estimator, the variance swap gives smaller returns than the volatility swap.
- Jumps in the process of the underlying have more influence on variance swaps than on volatility swaps, due to the (non-linear) payoff-pattern of variance swaps.
- The price of a cap on the volatility or variance swap is very low, unless many large jumps occur.

Since the variance swap usually gives larger expected returns, but also gives a larger standard deviation than the volatility swap (due to the non-linear payoff pattern of variance swaps), it is best to find an optimal combination to the two instruments for a desired risk-return-ratio.

CHAPTER 6

Decision Supporting Model

As was mentioned in the previous chapter, a certain combination of volatility instruments can give the risk-return-ratio that is desired by the customer. It can also be useful to divide an investment over different markets that exist in the world (i.e. the Asian market, the European market and the U.S.A. market). This can spread the risk, since all markets have different characteristics.

In this chapter a model is presented that can support in investment decisions considering volatility instruments. The decision supporting model that is developed is supposed to give a distribution of an investment over the different markets that exist in the world. Furthermore, an optimal distribution of volatility instruments (within a market) has to be found such that the risk is minimised and the performance is as defined by the customer. It is important that the optimisation model is robust to the input parameters that define the process of the underlying indices. The underlying index is assumed to follow a certain process, given the input parameters. The uncertainty in this process can be approximated by a number of scenarios for given input parameters.

In the next section, the general approach to portfolio optimisation is described and the concept of robust optimisation is introduced. In Section 6.2, first the assumptions that have been made to model the problem are given. After that the model itself is described in Section 6.3. In Section 6.3.4, the results of the optimisation model are presented. Furthermore, in Section 6.4 some extensions to the model are given.

6.1 Portfolio Optimisation

The first mathematical model for portfolio selection was developed by Markowitz in 1952 (Markowitz [20]). In the Markowitz portfolio selection model, the return of a portfolio is measured by the expected value of the random portfolio return, and the associated risk is quantified by the variance of the portfolio return. Markowitz showed that, given either an upper bound on the risk that the investor is willing to take or a lower bound on the return the investor is willing to accept, the optimal portfolio can be obtained by solving a convex quadratic programming problem. In practice, the Markowitz model and comparable models are often not useful, since these are very sensitive to the values of the input parameters.

Several techniques have been proposed to reduce the sensitivity of the portfolios to input uncertainty. Goldfarb and Iyengar give an overview in [13]. Examples of these techniques are the introduction of constraining portfolio weights, the introduction of other estimating techniques like Bayesian estimation of means and covariances. These techniques reduce the sensitivity of the portfolio composition to the parameter estimates, but they are not able to provide any guarantees on the risk-return performance of the portfolio. Another suggestion is to resample the mean returns and the covariance matrix of the assets from a confidence region around a nominal set of parameters, and then aggregate the portfolios obtained by solving a Markowitz problem for each sample. Another proposition to reduce the uncertainty is scenario based stochastic programming. Both of these approaches do not provide any hard guarantees on the performance of the portfolio and become very inefficient as the number of assets grows.

Another method that can reduce the sensitivity of the portfolio to input uncertainty is robust optimisation. This method generates portfolios that are robust to parameter uncertainty and estimation errors. More information on different variations on robust optimisation can be found in Ben-Tal and Nemirovski [2]. In this framework, the market parameters are modeled as unknown, but bounded, and optimisation problems are solved assuming worst-case behaviour of these parameters. Since worst-case situations can be important in the problem that is modeled here (great losses can be obtained when a portfolio of short volatility instruments experiences high realised volatility levels), a robust optimisation model seems the best approach in this situation.

The problem of finding an optimal distribution of an investment over different exchanges and an optimal distribution of the investment over the volatility instruments on these exchanges is somewhat more difficult than the portfolio model proposed by Markowitz in [20]. In the problem described in this thesis, two portfolio optimisations have to take place instead of one (as proposed by Markowitz), namely the optimisation over the exchanges and the optimisation over the different volatility instruments within the exchanges.

6.2 Assumptions

6.2.1 Risk and Performance

The objective of the problem described above is to minimise the risk. Markowitz proposed the variance of the return as a risk measure in [20]. This has become a standard and will also be used here.

The minimal performance is measured by the expected rate of return. The Sharpe ratio, described in Section 5.2.1, is derived from Markowitz's idea. As mentioned in Section 5.2.1, the Sharpe ratio assumes normally distributed returns and so does the risk-reward curve described by Markowitz. To adjust the objective of the problem for distributions of returns that are not completely determined by the mean and standard deviation, third (and higher moments) of the distribution can be added.

The risk-reward-curve is expected to be non-decreasing and convex, since higher returns are expected to give higher variance.

6.2.2 Correlation

In the problem described in this chapter, the optimal distribution of an investment over the volatility instruments within an exchange has to be found. The returns of the volatility instruments are highly correlated within one exchange, since they all depend on the same index process.

There is no correlation assumed between the different exchanges. It can be introduced in the model by assuming a correlation and integrating this into the models for the index processes. This is discussed in Section 6.4.

6.2.3 Robustness

The model has to be as robust as possible to the assumptions that have been made considering the process of the underlying index. To make the model more robust, the parameters that determine the index process have to be estimated in such a way that the model is suitable for many situations. There are several parameters that determine the process of the underlying. When for example the Heston process (Section 3.2.2) with jumps is assumed to model the underlying index (jumps that have a deterministic size and arrive according to a Poisson process), 5 parameters have to be estimated. These parameters are the correlation coefficient between the volatility process and the index process (ρ), the speed of the mean reversion (δ), and the volatility of volatility (σ) for the Heston model and the arrival intensity of the jumps, and the size of the jumps for the jump process. It can be the case that some parameter values depend on other values. The jumpsize is for example expected to be larger when fewer jumps occur and therefore the jumpsize can be dependent on the arrival intensity of the jumps. There are some other parameters that characterise an exchange and influence the return of the volatility instruments. Examples of these parameters are, for example, the skew, the at-the-money implied volatility and the expected return of a stock.

To make the model more robust, a 'worst-case' analysis is performed by choosing more values per parameter instead of fixing the parameters to one value. The set of parameter values that gives the worst results is used in the remainder of the problem, to make sure the solution is acceptable for all parameter combinations.

When, for example, 3 levels are chosen for the Heston parameters and 2 levels for the arrival intensity of the jumps (and the jumpsize dependent on the arrival intensity), $3^3 * 2 = 54$ possible parameter combinations are obtained for every exchange. The possible parameter values are determined by using historical data for every exchange.

It could be useful to check if all parameter combinations are realistic. It could for example be that the values of the different parameters are dependent (such as described for the jumpsize and arrival intensity of jumps).

6.2.4 Scenarios

Given an exchange and a combination of parameter values, the index process is still uncertain. The uncertainty is modeled by a set of scenarios. These scenarios are used to estimate the expected return and the variance of the return of the volatility instruments. K scenarios are defined per exchange and per combination of parameter values to estimate the expected return and the variance of return. In total $N * P * K$ scenarios have to be generated (N exchanges, P combinations of parameter values, and K scenarios). The scenarios can for example be generated by the Monte Carlo method.

6.3 Optimisation Model

6.3.1 Input

- Set of M volatility instruments, denoted by index $i \in \{1, \dots, M\}$;
- Set of N exchanges, denoted by index $j \in \{1, \dots, N\}$;
- Set of combinations of parameter values $p \in P^j$. (P^j is the set of possible parameter combinations for exchange j). This set is different for every exchange j;
- $g_{i\cdot}^{jp}$ is the return of volatility instrument i, uniquely determined by exchange j and combination of parameter values p;
- Set of K scenarios per exchange j and combination of parameter values p, denoted by $k \in K^{jp}$. (K^{jp} is the set of scenarios for exchange j and parameter combination p);
- β is the minimal outperformance of the total investment, for example compared to cash. This is determined in advance by the user of the optimisation model and implied by the risk-averseness of the user (the more risk-averse, the lower β is chosen);

6.3.2 Decision Variables

The decisions that have to be made in this problem are how to divide an investment over different volatility instruments and over different exchanges. This is denoted by a_{ij}, which is the fraction of the investment that is invested in volatility instrument i on exchange j.

6.3.3 Model

The optimisation model for this problem now becomes:

P=

Minimise	$\sum_{j=1}^{N} \max_{p \in P^j} \{\text{var}[\sum_{i=1}^{M} g_{i\cdot}^{jp} a_{ij}]\}$	
Subject to	$\sum_{j=1}^{N} \min_{p \in P^j} \{\mathbb{E}[\sum_{i=1}^{M} g_{i\cdot}^{jp} a_{ij}]\} \geq \beta$	(Performance)
	$\sum_{i=1}^{M} \sum_{j=1}^{N} a_{ij} = 1$	(All available money is invested somehow)
	$a_{ij} \in [0,1]\ \forall_i \forall_j$	

In this formulation, the expected return and the variance of the return are given by:

$$\mathbb{E}[\sum_{i=1}^{M} g_{i\cdot}^{jp} a_{ij}] = \sum_{i=1}^{M} (\frac{1}{K} \sum_{k=1}^{K} g_{ik}^{jp}) a_{ij};$$

$$\text{var}[\sum_{i=1}^{M} g_{i\cdot}^{jp} a_{ij}] = \sum_{i=1}^{M} \sum_{i'=1}^{M} \text{cov}[g_{i\cdot}^{jp} a_{ij}, g_{i'\cdot}^{jp} a_{i'j}]$$

$$= \sum_{i=1}^{M} \sum_{i'=1}^{M} a_{ij} a_{i'j} (\frac{1}{K} \sum_{k=1}^{K} g_{ik}^{jp} g_{i'k}^{jp}) - \sum_{i=1}^{M} \sum_{i'=1}^{M} a_{ij} (\frac{1}{K} \sum_{k=1}^{K} g_{ik}^{jp}) a_{i'j} (\frac{1}{K} \sum_{k=1}^{K} g_{i'k}^{jp}).$$

Due to the fact that the variance is not linear in the decision variables a_{ij}, the model becomes a non-linear programming model.

The objective is the sum of the worst-case variances (maximum variance over all parameter combinations) over the exchanges j. The first constraint makes sure that the expected return (for all parameter combinations) is better than the required expected return (β).

The optimisation problem is rewritten and the corresponding robust optimisation formulation is obtained:

P'=

Minimise	$\sum_{j=1}^{N} v_j^2$	
Subject to	$\text{var}[\sum_{i=1}^{M} g_{i\cdot}^{jp} a_{ij}] \leq v_j^2$	$\forall_j \forall_p$
	$\mathbb{E}[\sum_{i=1}^{M} g_{i\cdot}^{jp} a_{ij}] \geq \alpha_j$	$\forall_j \forall_p$
	$\sum_{j=1}^{N} \alpha_j \geq \beta$	$\forall_j \forall_p$
	$\sum_{i=1}^{M} \sum_{j=1}^{N} a_{ij} = 1$	
	$a_{ij} \in [0,1]$	$\forall_i \forall_j$
	$v_j^2 \in \mathbb{R}^+$	$\forall_j$
	$\alpha_j \in \mathbb{R}$	$\forall_j$

In this model the vectors $\underline{a}$, $\underline{v}$, and $\underline{\alpha}$ are the decision variables. The input parameters are g_{ik}^{jp} and β.

6.3.4 Results

The model has been implemented in AIMMS. Three exchanges are considered ($N = 3$) and two volatility instruments ($M = 2$), namely the delta-hedged straddle (described in Section 4.2.1) and the variance swap (described in Section 4.6). A one-month period is considered. The expected return and confidence interval of the returns of the one-month delta-hedged straddle is shown in Figure 6.1. The expected return and confidence interval of a one-month variance swap is shown in Figure 6.2.

1000 scenarios are used to estimate the expected return and the variance of the return for every instrument i, exchange j, and parameter combination p. 16 parameter combinations are being considered, consisting of the following parameter values:

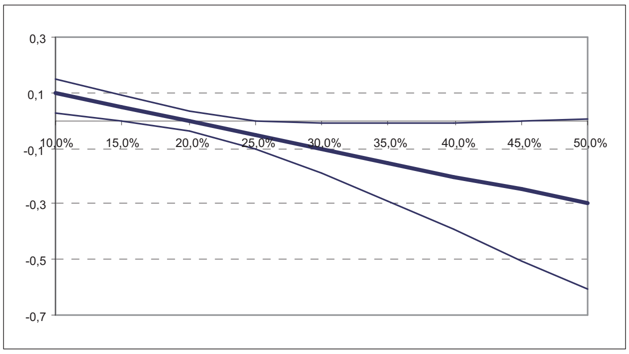

Figure 6.1: The expected return and confidence interval of the return of the one-month delta-hedged straddle

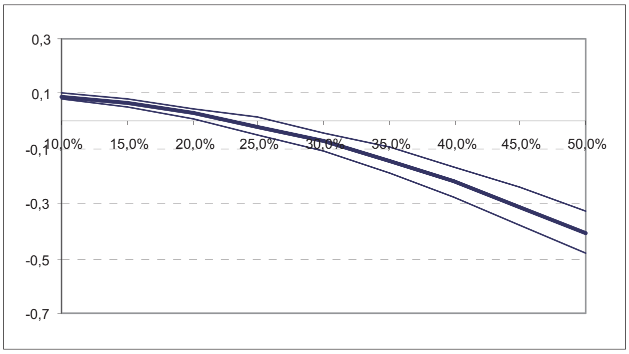

Figure 6.2: The expected return and confidence interval of the return of the one-month variance swap

ρ:	-0.01 or 0
δ:	0 or 5
σ:	0 or 10%
Arrival intensity of jumps:	0 (0 jumps per year, with mean jumpsize 0) or 0.05 (1 jump per 20 years, with mean jumpsize 0.2)

Furthermore, a realistic skew is implemented (fitted on data of the S&P 500 of the fourth quarter of 2004). The at-the-money implied volatility is assumed to be 20% and the expected realised volatility (without considering the jumps) is 15%. The optimisation model is solved for the parameter combinations that are given above and some other values, to obtain more information on the influence of these parameter values on the robustness of the solution.

Standard Case

When all three exchanges have the parameter combinations given above (the *standard case*), a required expected return of 3% ($\beta = 3\%$) gives an investment strategy of investing everything in the delta-hedged straddle (equally divided over the exchanges). When the required expected return is larger (maximal 4.2%, otherwise the problem becomes infeasible), almost everything is invested in the variance swap and the standard deviation of the return becomes larger. The variance swap gives a larger expected return and a larger standard deviation compared to the volatility swap, due to the non-linear payoff pattern of the variance swap. When the required expected return is 4.2%, the maximum standard deviation is 7.8%, whereas the standard deviation is 4.8% when the required expected return is 3%. In Figure 6.3, the standard deviation is compared to the required expected return. Figure 6.3 shows that the standard deviation of the return of the portfolio increases when the required expected return increases. This is due to the fact that variance swaps give a higher expected return, but a higher standard deviation as well. This is again caused by the non-linear payoff pattern of the variance swap. Therefore, more has to be invested in variance swaps when the required expected return is larger and this gives a higher standard deviation. In de first point of Figure 6.3, everything is invested in the delta-hedged straddles. The fractions of the investment in delta-hedged straddles and variance swaps are denoted in the circle diagrams for the other two points. In the circle diagrams, the light color denotes the fraction of the investment invested in delta-hedged straddles and the dark color denotes the fraction invested in variance swaps. The investments are equally divided over the three exchanges, since all exchanges have the same characteristics. The circle diagrams therefore denote the investment-distribution of straddles and variance swaps over the three exchanges, but also per exchange.

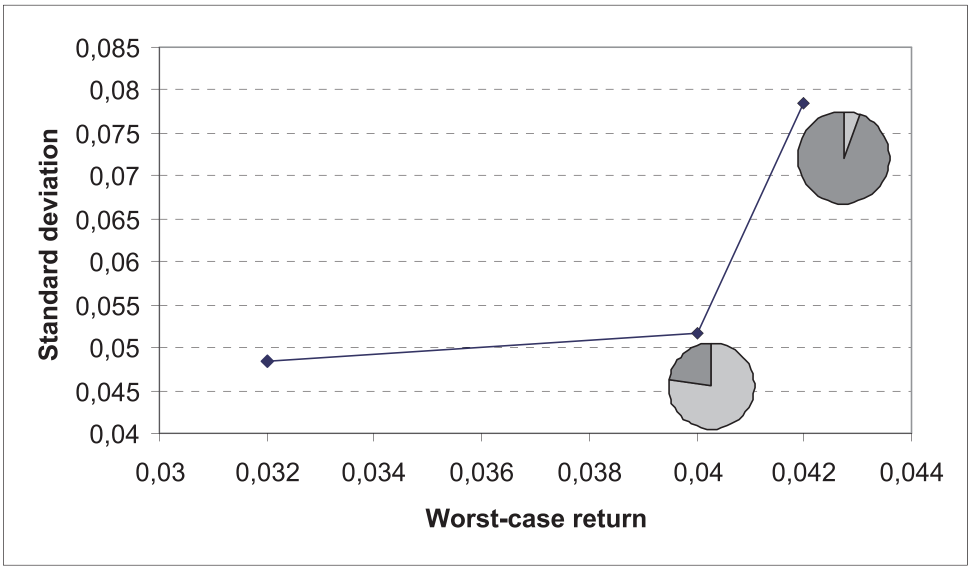

Figure 6.3: The standard deviation (vertical axis) compared to the required expected return (β, horizontal axis) per month and the distribution over the volatility instruments (circle diagrams), total over all exchanges

Summarising, this gives the following results:

1. The larger the required expected return is, the larger is the fraction invested in variance swaps, since variance swaps give a larger expected return (but also a larger standard deviation due to the non-linear payoff pattern of the variance swaps).

2. When a larger expected return is required (β is larger), the standard deviation of the returns increases, since more is invested in variance swaps (variance swaps give a larger standard deviation of the returns).

When the realised volatility is larger than expected (for example larger than the implied volatility instead of 5% smaller as is expected), the return of the investment strategy is worse. In Table 6.1 the results of the strategies proposed in the data points of Figure 6.3 are described in the first column. In the second column the expected return of the standard case is given. When the realised volatility is 1% larger than the implied volatility, the results are worse, which is given in the third column. In the last column, the expected returns that are obtained when the realised volatility appears to be 10% larger than the implied volatility are depicted.

% in variance swaps	Expected return standard case	Expected return realised volatility 1% larger (no jumps)	Expected return realised volatility 10% larger (no jumps)
0%	3.2%	−1%	−10%
22%	4%	−0.79%	−10.19%
94%	4.2%	−0.10%	−10.81%

Table 6.1: Results for a realised volatility that is larger than the implied volatility

Table 6.1 shows that when the realised volatility is only 1% larger than the implied volatility, the results become better when the strategy is to invest more in variance swaps, whereas for a realised volatility that is 10% larger than the implied volatility, the results become worse when more is invested in variance swaps. This can be explained by the fact that the variance swap gives better results (compared to the delta-hedged straddle) when the realised volatility is close to the implied volatility and worse results when the realised volatility is further from the implied volatility. This can be seen in the Figures 6.1 and 6.2.

Average Case

Not only the 'worst-case' variance (maximal variance over the parameter combinations) is optimised, but also the *average case*, which is constructed by assuming that all possible parameter combinations are equally likely to happen. This average case is compared to the standard case, by looking at the *expected* variance (average over all parameter combinations, instead of maximum over all parameter combinations) given by the optimal values of the decision variables of the standard case (optimal distribution of the investment). In Figure 6.4, the standard deviations, given a certain required expected return, are given for both the expected results of the standard case as the average case.

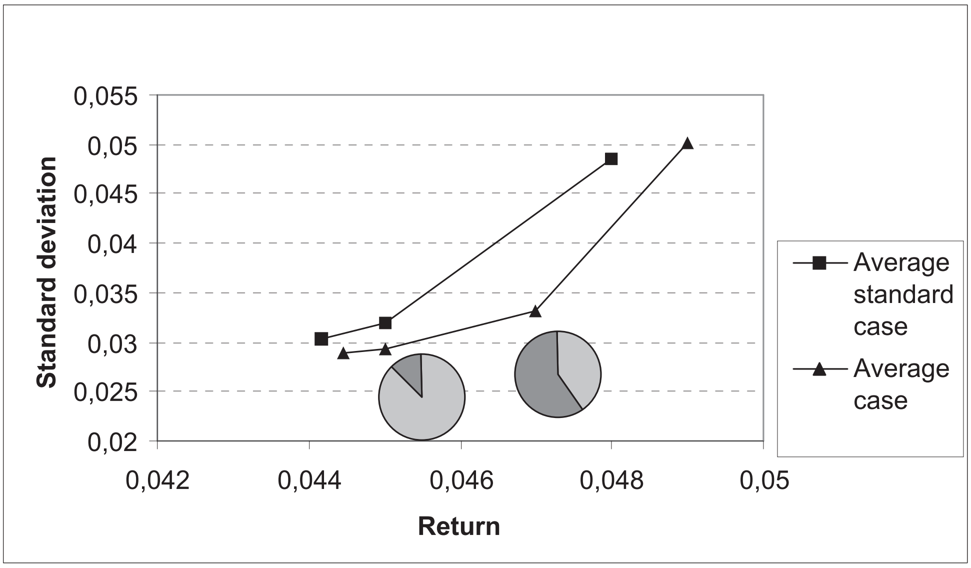

Figure 6.4: The standard deviation (vertical axis) compared to the required expected return (β, horizontal axis) of the average of the standard case and the average case and the distribution over the volatility instruments (circle diagrams)

From Figure 6.4 can be concluded that when the standard case is optimised, slightly larger expected standard deviations are obtained. This is caused by the fact that the standard case gives optimal results for the worst parameter combination, but these results might not be optimal for the average of the parameter combinations. Furthermore, a larger required expected return (4.9% in the average case, compared to 4.8% in the standard case) is still feasible for the average case (4.9% in the average case, compared to 4.8% in the standard case). In Figure 6.4, the distribution over the volatility instruments is also shown. In the first data point, everything is invested in delta-hedged straddles. In the last point, everything is invested in variance swaps. The fractions invested in delta-hedged straddles and variance swaps in the two points in between are shown in the circle diagrams in Figure 6.4. In the circle diagrams, the light color denotes the fraction of the investment invested in delta-hedged straddles and the dark color denotes the fraction invested in variance swaps. Since all exchanges have the same characteristics, the investment is divided equally over the exchanges. The fractions depicted in the circle diagrams therefore give information on the derivation of the volatility instruments over the sum of all exchanges and per exchange.

Summarising, this gives the following results:

1. A similar investment strategy is chosen when the average case is optimised as when the standard case is optimised.
2. The strategy chosen in the average case optimisation gives slightly better expected returns.

More and Smaller Jumps

When on one of the exchanges more and smaller jumps occur, the investment strategy changes. On exchange 3, the arrival intensity of the jumps is, for example, either $\frac{1}{5}$ (one jump per 5 years, with mean size of the jumps of 0.05) or $\frac{1}{3}$ (one jump per 3 years, with mean size of the jumps of 0.03), while on the other exchanges, the arrival intensity of the jumps is as in the standard case. More smaller jumps make it more attractive to invest in variance swaps, since there is less risk involved. This can be seen in the proposed investment strategy for this situation, which proposes to invest 87% of the investment in variance swaps on exchange 3 and the rest of the investment in delta-hedged straddles in exchange 1 and 2. This gives a required expected return of 4.4% (even for a required expected return (β) of 3%) with a standard deviation of 1.90%. When even 91% is invested in variance swaps on exchange 3, the required expected return is 4.9% and the standard deviation is 1.91%.

Summarising, more and smaller jumps on one exchange give the following results:

1. The largest part of the investment is invested on the exchange that has more and smaller jumps, since more and smaller jumps give less risk than fewer and larger jumps;
2. The required expected return is larger and the standard deviation is smaller, compared to the standard case. This is due to the fact that the overall risk decreases due to the exchange that has more and smaller jumps;
3. A larger part of the investment is done in variance swaps, compared to the standard case, since there is less risk for the variance swaps on an exchange that has smaller jumps.

Larger Volatility of Volatility

The investments in variance swaps are expected to give larger standard deviations when the volatility of volatility is larger. When the volatility of volatility in exchange 3 is either 0 or 30% (instead of either 0 or 10%, which is still valid for the other exchanges), an investment strategy which gives a required expected return of 4.4% gives a standard deviation of 9.5%. To obtain a required expected return that is larger than β, the investment strategy proposes to invest 95% of the investment in *variance swaps* on the exchanges 1 and 2. This gives the larger standard deviation (compared to the standard case). In the analysis described above, the skew and reference variance are not adjusted to the larger volatility of volatility, so in this situation, the market estimation of the volatility of volatility (derived from the skew) is smaller than the realised volatility of volatility and therefore the results are worse than the standard case. When the skew (and the reference variance implied by the skew) is adjusted for the larger volatility of volatility, however, the results do not differ much from the standard case described above, since the return of the variance swaps is now adjusted to the risk of the larger volatility of volatility.

Summarising, a larger volatility of volatility on one exchange gives the following results when the reference level of the variance swap is not adjusted to the level of the volatility of volatility:

1. The largest part of the investment is invested on the exchanges that have the smallest volatility of volatility, since the risk on these exchanges is smaller than the risk on the exchange with a larger volatility of volatility;

2. Compared to the standard case, a larger standard deviation is obtained. This is due to the fact that the overall risk invreases since the risk on one exchange increases (due to a larger volatility of volatility).

When the reference level of the variance swap is adjusted to the larger volatility of volatility, however, the larger volatility of volatility does not influence the results.

6.4 Extensions

Two extensions have been added to the model described in the previous section. First, a correlation between the exchanges is added. This is described in Section 6.4.1. After that, a multi-period model is introduced to analyse the influence and size of a cash component. This is described in Section 6.4.2.

6.4.1 Correlation

As is mentioned before, the correlation between the different exchanges was not included in the model so far. The correlation can play an important role and therefore a different model for the processes that describe the underlying that includes this correlation is included in the standard case (as discussed in the previous section). In this model, the processes of the underlying are modeled by Black and Scholes processes. 3 exchanges are considered (for example U.S.A., Europe and Asia) and exchange 1 and 2 are assumed to be highly correlated. Furthermore, a jump process is assumed for the exchanges, in which the timing of the jumps of exchange 1 and 2 is again highly correlated.

Positive Correlation

When there is a positive correlation between either the index processes of exchange 1 and 2 or the jump processes of exchange 1 and 2 (or both), the following results are obtained. When the desired minimal expected return (β) is small, the proposed investment strategy is to invest equal amounts in all exchanges. When the required expected return is larger, 40% of the investment is done in the index that is not correlated with the other two indices and 60% is done in the other two indices. This strategy can be explained by the fact that when the desired return is larger, a larger part of the investment is done in variance swaps. Since variance swaps are more sensitive to extreme events, the correlation becomes more important (all extremes occur at the same time).

Negative Correlation

A negative correlation between two exchanges makes it more attractive to invest in the exchanges that are negatively correlated, since the risk is opposite when the investment is done in these exchanges.

6.4.2 Cash Component

When more periods are considered, the influence of a cash component on the return and variance of the return becomes interesting. The cash component is added to prevent great losses in case of large realised volatility (for example when large jumps occur).

The influence of the cash component can be analysed by analysing the following multi-period model in which c is the return of the cash investment:

$$
\begin{array}{lll}
\mathbf{P}^{cash} = & & \\
\text{Minimise} & \sum_{t=1}^{T}(\sum_{j=1}^{N} v_{jt}^2)\lambda^2 & \\
\text{Subject to} & \text{var}[\sum_{i=1}^{M} g_{i\cdot}^{jpt} a_{ij}] \le v_{jt}^2 & \forall_j \forall_p \forall_t \\
& \mathbb{E}[\sum_{i=1}^{M} g_{i\cdot}^{jpt} a_{ij}] \ge \alpha_{jt} & \forall_j \forall_p \forall_t \\
& \sum_{t=1}^{T}(\sum_{j=1}^{N} \alpha_{jt}\lambda + c(1-\lambda)) \ge B & \forall_j \forall_p \\
& \sum_{i=1}^{M}\sum_{j=1}^{N} a_{ij} = 1 & \\
& a_{ij} \in [0,1] & \forall_i \forall_j \\
& v_{jt}^2 \in \mathbb{R}^+ & \forall_j \forall_t \\
& \alpha_{jt} \in \mathbb{R} & \forall_j \forall_t \\
& \lambda \in [0,1] &
\end{array}
$$

5 years (60 months) are simulated. The fraction invested in volatility instruments (λ) increases linearly with the total required expected return per month (β), until this fraction is equal to 1 (everything invested in volatility instruments). This is shown in Figure 6.5.

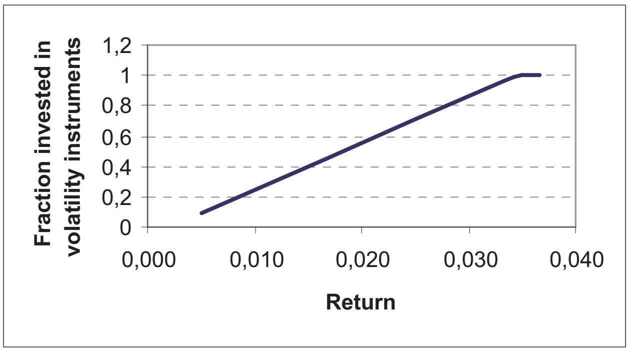

Figure 6.5: The fraction invested in volatility instruments (λ, vertical axis) compared to the required expected return per month (β, horizontal axis)

When everything is invested in volatility instruments and nothing in cash, the model proposes to invest more in variance swaps when a larger expected return per month (β) is required. Therefore, the standard deviation of the portfolio also increases linearly with the required expected return, until the fraction invested in volatility instruments is 1, which is shown in Figure 6.6.

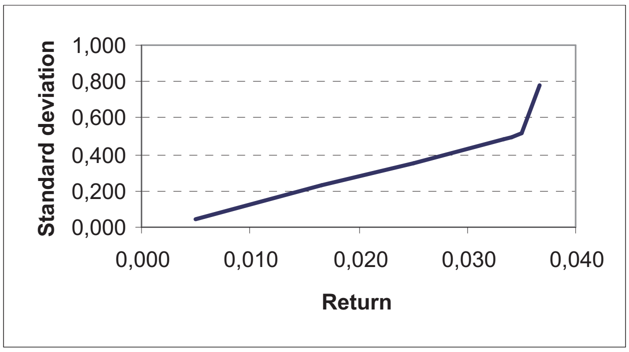

Figure 6.6: The standard deviation of the return (vertical axis) compared to the required expected return per month (β, horizontal axis)

When a capped variance swap is used instead of a variance swap without a cap, similar results are obtained considering the fraction invested in volatility instruments and the standard deviation as when a 'normal' variance swap would have been used. When a capped variance swap is used, however, more is invested in variance swaps (since the capped variance swaps give less risk than the variance swaps without a cap). Furthermore, the maximum expected return that can be required (and still gives a feasible solution) is larger for a given level of the standard deviation. This is due to a larger investment in variance swaps.

CHAPTER 7

Conclusions and Recommendations

7.1 Conclusions

The objective of the research presented in this thesis was to find an optimal strategy for building portfolios that minimise risk and maximise return to exploit the phenomenon that implied volatility overestimates realised volatility. To define a strategy like this, first the instruments that can be used to trade the difference between implied and realised volatility were analysed. After that a robust optimisation model was formulated that was able to give the desired optimal strategy. The conclusions of this research can therefore be divided into two parts: The conclusions considering the volatility instruments (given in Section 7.1.1) and the conclusions considering the portfolio strategy (given in Section 7.1.2).

7.1.1 Volatility Instruments

There are several instruments that can be used to trade volatility. The most basic instrument is the straddle, which has the disadvantage that the losses are big when the price of the underlying moves away from the strike. Another volatility instrument is an over-the-counter instrument, namely the volatility swap. The volatility swap has a payoff linear in the difference between implied and realised volatility and can be replicated by a delta-hedged straddle. The volatility swap has the disadvantage that its replication is a dynamic position in options. An instrument that does not have this disadvantage is the variance swap. The variance swap is an over-the-counter instrument that can be replicated by a static position in options and a dynamic position in futures, which gives a smaller standard deviation of the return of a variance swap compared to the volatility swap.

The volatility skew implies the size of the convexity bias (which is defined by the difference between the payoff of a volatility swap and a variance swap for a certain realised volatility). he skew therefore gives a market estimator of volatility of volatility. When the skew gives a larger market estimator of the volatility of volatility than the realised volatility of volatility is, the variance swap gives higher returns than the volatility swap.

Furthermore, due to the non-linear payoff pattern of the variance swap, jumps in the process that describes the underlying have more influence on the return of variance swaps than on the return of volatility swaps.

7.1.2 Portfolio

An optimisation model is developed that has the objective to minimise the variance of the portfolio and requires a minimal performance. Depending on the value of the required expected performance of the portfolio, the model proposes an investment in volatility swaps and/or variance swaps. Furthermore, a division of the investment over the different exchanges is proposed.

A larger fraction of the investment is done in variance swaps as soon as more performance is required. Furthermore, investing in an exchange with many small jumps seems more attractive than investing in an exchange with not so many large jumps. A larger part of the investment is done on the exchanges with a smaller volatility of volatility.

By adding correlation to the processes that describe the underlying, the investment strategy changes slightly. More is invested in the exchange that is not correlated with the other exchanges.

The influence and size of a cash component of the portfolio are analysed by considering a multi-period model. Depending on the required level of expected return, an investment strategy is proposed that includes a fraction of the investment in cash, a fraction in volatility instruments and a division over the exchanges and over the volatility instruments.

7.2 Recommendations

In this thesis an optimisation model is presented by which an investment strategy in volatility instruments can be determined, of which the required expected performance is still acceptable. To develop this model, an extensive analysis of the volatility instruments is performed. Furthermore, two extensions to the optimisation models are discussed. The recommendations for future research consider both the analysis of the volatility instruments and the optimisation model.

Considering the analysis of the volatility instruments:

- It is interesting to investigate the performance of the volatility instruments with parameters that are estimated from the historical data of different exchanges. Especially a realistic skew and volatility of volatility are interesting to consider (as is proposed in Section 5.1.2);
- As is mentioned in Section 5.2.1, the Sharpe ratio assumes normally distributed returns of the portfolio. Since the returns of the volatility instruments do not seem normally distributed (as can be seen in Figure 5.2 and Figure 5.4), it is better to find another performance measure than the Sharpe ratio to analyse the return. Furthermore the risk and performance are included in the portfolio optimisation model from the same point of view as the Sharpe ratio (as is described in Section 6.2.1). The objective of the optimisation model could be adjusted to another performance measure as well.

Two propositions have been done for expanding the optimisation model (in Section 6.4), namely adding correlation of the processes describing the indices and adding a cash component. These propositions can be expanded and investigated further:

- Considering the correlation of the processes describing the underlying, only Black and Scholes processes with jumps are used to analyse the influence of the correlation. It is interesting to look at more realistic processes that include correlation in future research.
- Considering the cash component, a multi-period model is introduced to analyse the influence of the cash component. The multi-period model can be improved. For example, results from the previous period (considering i.e. realised volatility) can be used to estimate parameter settings in each period and correlation between the exchanges can be added as well.

APPENDIX A

Background Information

In this appendix, an overview of the most important definitions and terms for understanding this report are given. This chapter is meant to give the readers who are not involved in option theory more understanding in the subject and the possibility the read and understand the rest of the report.

A.1 Derivatives

A derivative is a financial instrument of which the value depends on the value of other, more basic underlying variables, for example stocks or indices (like AEX, DAX, FTSE). Examples of derivatives are forward contracts, futures contracts and option contracts. In every type of derivative you can have a long position or a short position. Usually the person with the long position buys something and the person with the short position sells something. 'Writing' a derivative denotes taking a short position in the contract.

A market that is going down is called a bear-market and a market that is going up is called a bull-market.

A.1.1 Forward Contracts

A forward contract is an agreement to buy or sell an asset on a certain future time for a certain price. One of the parties to a forward contract assumes a long position and agrees to buy the underlying asset on a certain specified future date for a certain specified price. The other party assumes a short position and agrees to sell an asset on the same date for the same price.

The payoff from a long position in a forward contract on one unit of an asset is $S_T - K$. K is the delivery price; S_T is the spot price of the asset at maturity of the contract. The payoff from a short position in a forward contract on one unit of an asset is $K - S_T$.

A.1.2 Futures Contracts

Futures contracts are an agreement between to parties to buy or sell an asset at a certain time in the future for a certain price. The difference with forward contracts is that futures contracts are traded on an exchange.

A.1.3 Options

Put and Call Options

Options are traded both on exchanges and in the over-the-counter market. There are two basic types of options. A call option gives the holder the right to buy the underlying asset by a certain date for a certain price. A put option gives the holder the right to sell the underlying asset by a certain date for a certain price.

American and European Options

American options can be exercised at any time up to the expiration date and European options can be exercised only on the expiration date.

Specifications of Option Contract

In an option contract several things are specified. The exercise price or strike price is the price in the contract. The expiration date or maturity is the date on which the contract expires.

Payoff

There are two sites to every option contract. One side is the investor who has taken the long position (has bought the option). The other side is the investor who has taken a short position (has sold (i.e. written) the option). The payoff from a long position in a European call option is $\max(S_T - K, 0)$. The payoff to the holder of a short position in the European call option is: $-\max(S_T - K, 0) = \min(K - S_T, 0)$. The payoff from a long position in a European put option is: $\max(K - S_T, 0)$. The payoff to the holder of a short position in the European put option is $-\max(K - S_T, 0) = \min(S_T - K, 0)$. These payoffs are shown in Figure A.1.

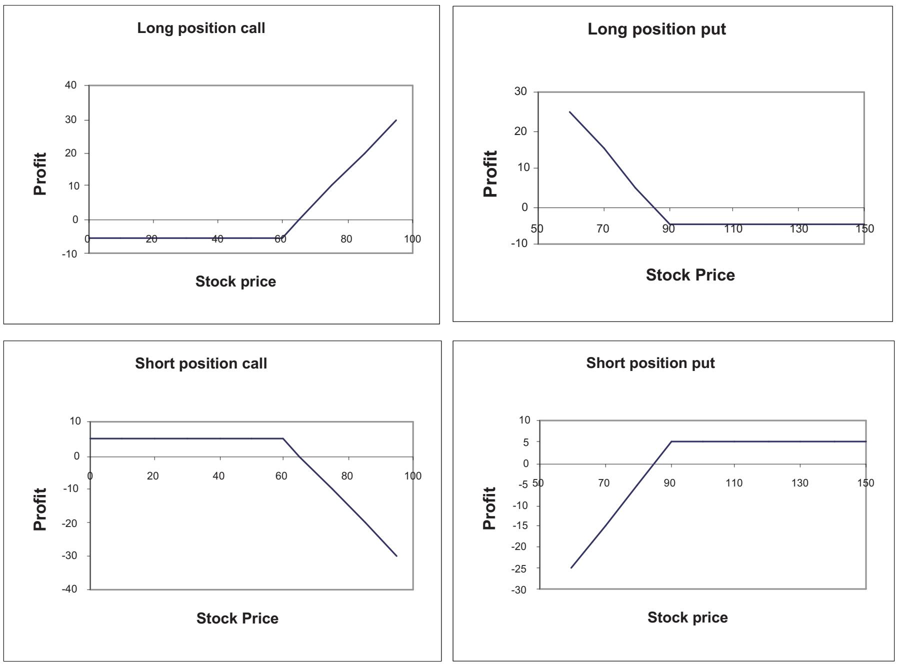

Figure A.1: The payoff pattern of long and short positions in call options and put options

Option Price

The price of an option is determined by the strike price, the price of the underlying asset, the time to maturity and the volatility (uncertainty in the future stock price, measured by the standard deviation).

Volatility
Unlike the other components of the option price, the volatility is not known in advance. Since all the other components of the price are known, the volatility can be calculated, because the option price is known (the market agrees on a certain price). This is called the implied volatility and can be seen as the estimator of the market for the future realised volatility.

Usually there is a difference between the implied volatility and the future realised volatility and generally the implied volatility overestimates the realised volatility.

Volatility Skew
The volatility used to price equity options (on individual stocks and on stock indices) has the general form shown in Figure A.2. This is usually referred to as the volatility skew. The volatility decreases as the strike price increases. This can be explained by the fact that there is more willingness to pay for risk (volatility) when the strike price is low than when the price is higher.

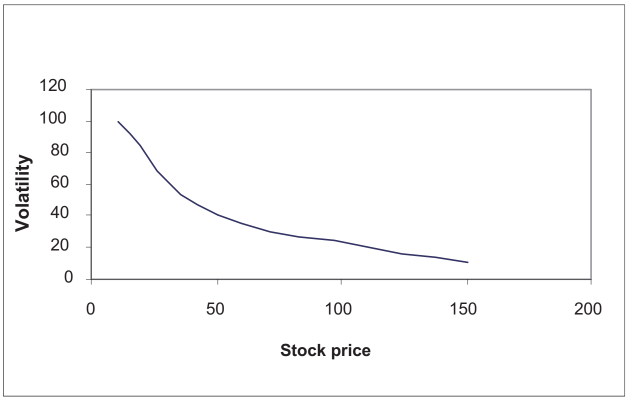

Figure A.2: The volatility skew

A.2 Option Pricing Models

An option pricing model that is very well known is the Black and Scholes model. It assumes that percentage changes in the stock price in a short period of time are normally distributed. From this, it follows that the stock price at time T has a lognormal distribution.

There are many more pricing models, that all have different assumptions and therefore different advantages and disadvantages. Most option pricing models take the Black and Scholes model as a basis and relax assumptions of the Black and Scholes model. In many models the assumption of constant volatility is relaxed, in Figure A.3, a scheme of the relation of these models is given.

Another assumption that is often relaxed, is the assumption of no jumps in the process of the underlying. This can be added in many of the models given in Figure A.3. Numerical procedures are also often used to price options. Examples of numerical methods are implied trees and generalised binomial trees.

A.3 The Greek Letters

The Greek letters help an investor to get more insight in the risk of a certain portfolio. Each Greek letter gives a measure for a different dimension to the risk in an option position. When an investor

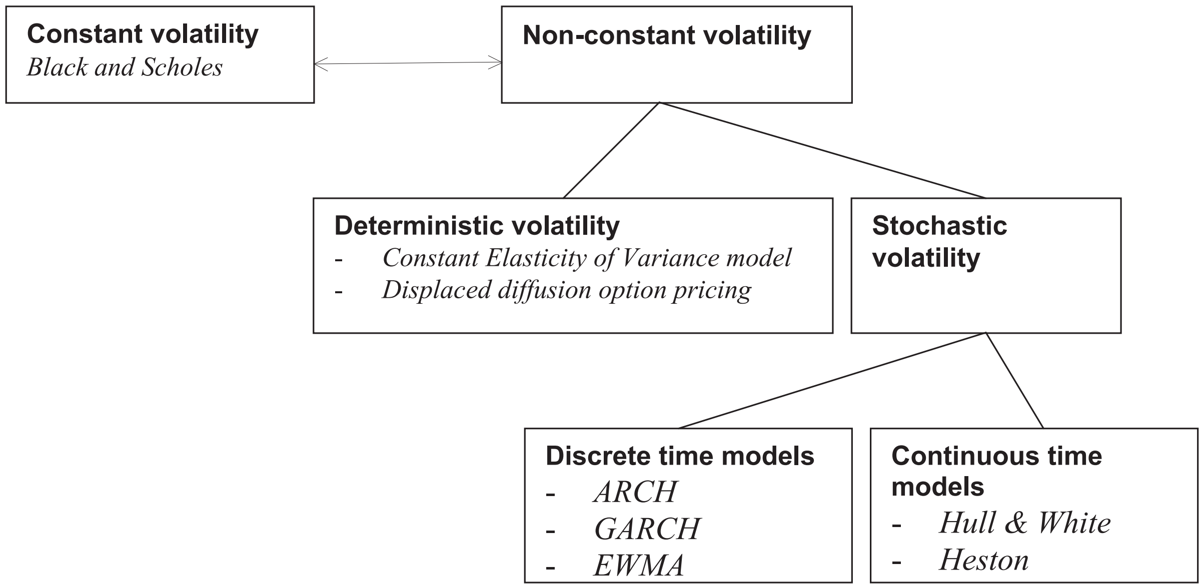

Figure A.3: Scheme of Option Pricing Models

hedges his portfolio, he considers 'the Greeks' and tries to manage the Greek letters such that all risks are acceptable.

A.3.1 Delta

Delta is the rate of change of the option price with respect to the underlying asset. Usually, traders want to keep a delta-neutral position (delta=0). Since the delta changes continuously, the investor's position remains delta-hedged (delta-neutral) only for a short period of time. The hedge has to be adjusted periodically. This is called re-balancing.

A.3.2 Theta

Theta is the rate of change of the value of the portfolio with respect to the passage of time with all else remaining the same. Theta can be called the time decay of the portfolio.

A.3.3 Gamma

The gamma of a portfolio of options on an underlying asset is the rate of change of the portfolio's delta with respect to the price of the underlying asset. A portfolio can be made gamma-neutral; Delta-neutrality provides protection against relatively small stock price moves in between re-balancing and gamma-neutrality provides protection against larger movements in this stock price between hedge re-balancing.

A.3.4 Vega

Vega is not really a Greek letter, but it is a common term in option theory. The vega of a portfolio of derivatives is the rate of change of the value of the portfolio with respect to the volatility of the underlying asset. Vega-neutrality protects for a variable volatility.

A.3.5 Rho

The rho of a portfolio of options is the rate of change of the value of the portfolio with respect to the interest rate.

A.4 Trading Strategies Involving Options

A.4.1 Spreads

Buying a call option on an asset with a certain strike price and selling a call option on the same asset with a higher strike price creates a bull-spread. Both options have the same expiration date. In the left figure in Figure A.4, the dashed lines show the profit from the two option positions separately and the solid line shows the profit from the bull-spread. The value of the option bought is always less than the value of the option sold, because the call price decreases as the strike price increases. A bull spread created from calls therefore requires an initial investment. Buying a put with a low strike price and selling a put with a high strike can also create bull-spreads. Bull-spreads created from puts involve a positive cash flow up front.

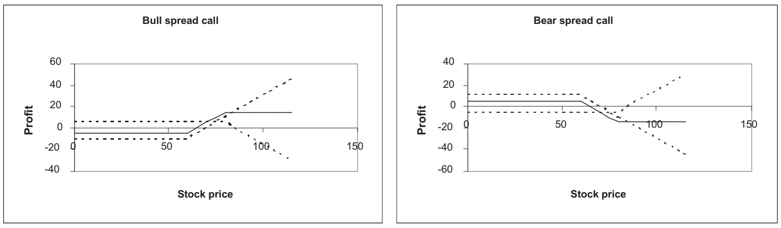

Figure A.4: The payoff of a bull spread and a bear spread using call options

Buying a call with one strike price and selling a call with a lower strike price creates a bear-spread. Both options have the same expiration date. In the right figure of Figure A.4, the dashed lines show the profit from the two option positions separately and the solid line shows the profit from the bear-spread. A bear-spread created from calls involves an initial cash inflow for the investor. Buying a put with a high strike price and selling a put with a low strike price can also create bull-spreads. Bull-spreads created from puts involve a negative cash flow up front.

An investor who enters into a bull-spread is hoping that the price of the underlying asset will increase (bull market) and an investor who enters a bear-spread is hoping that the price of the underlying asset will decrease (bear market).

A butterfly spread involves positions in options with three different strike prices. It can be created by buying a call option with a relatively low strike price K_1; buying a call option with a relatively high strike price K_3; and selling two call options with strike price $K_2 = (K_1 + K_3)/2$ (generally K_2 is close to the stock price). The payoff is shown in Figure 4.4 (in Section 4.4). A butterfly spread can also be created using put options. The investor buys a put with a low strike price, buys a put with a high strike price and sells two puts with an intermediate strike price.

A.4.2 Straddles

A straddle involves buying a call option and a put option with the same strike price and expiration date. The profit pattern is shown in Figure 4.1 (in Section 4.2). A straddle is appropriate when an investor is expecting a large move in a stock price, but does not know in which direction this move will be. An investor who writes a straddle has the opposite profit pattern, so by writing straddles with the strike price (close to) the current stock price, one can use the fact that the realised future volatility is usually smaller than implied volatility.

A.4.3 Strips and Straps

A strip consists of a long position in one call and two puts with the same strike price and expiration date. A strap consists of a long position in two calls and one put with the same strike price and expiration date.

A.4.4 Strangle

A strangle involves buying a put with a low strike price and a call with a high strike price. The put and the call have the same expiration date.

A.4.5 Volatility and Variance Swaps

A volatility swap is a forward contract on annualised volatility. Its payoff at expiration is equal to N($\sigma_{real} - K_{vol}$). σ_{real} is the realised volatility, K_{vol} is the annualised volatility delivery price and N is the notional amount of the swap in dollars per annualised volatility point.

Volatility and variance swaps are traded by portfolios of options that replicate them. The variance swap can be replicated more reliably (by portfolios of options of varying strikes) than the volatility swap. A variance swap is a forward contract on annualised variance, the square of the realised volatility. Its payoff at expiration is equal to N($\sigma^2_{real} - K_{var}$), σ^2_{real} is the realised variance, K_{var} is the reference variance and N is the notional amount of the swap in dollars per annualised volatility point squared.

A.5 Hedging

The aim of hedging is to reduce a particular risk by trading for example futures on the underlying. An example of the risk that has to be reduced is the level of the index. The hedge should replicate the portfolio as close as possible in connection to the risk that has to be reduced. Often this is not done perfectly.

> **Example:**
> The portfolio consists of a short position in a call option on index A that has to be hedged against changes in index A. This call option has to be delta-hedged, with delta defined as the rate of change of the option price with respect to the index rate:
> $\Delta = \frac{\delta C}{\delta S}$, with C the price of the call option and S the level of the index. If, for example, the delta of the call option is 0.6, the value of the call option changes by 0.6*x if the index changes by x. The hedge portfolio of the short position in the call option consists of a long position in futures on index A.

Often index futures are used to hedge options, because the transaction costs associated with the trades in the index futures are generally lower than the costs associated with the underlying stocks.

It is usually not possible to hedge against all risk. A delta-neutral portfolio (neutral against changes in the level of the underlying) is often not vega-neutral (neutral against volatility changes).

A.6 The Principle of No Arbitrage

The assumption of no arbitrage is often done in the theory of derivatives. The principle of no arbitrage states that two equivalent goods in the same competitive market must have the same price. This assumption only holds in a totally efficient market, i.e. a market where the following conditions satisfy:

- Liquidity: all players are small relative to the size of the market and there are no transaction costs.
- Price Continuity: the price is set only by the underlying supply-demand situation and temporary imbalances do not affect the price.
- Fairness: there is no discrimination between different players in the market.
- Equal access to information: all market players have equal access to all available information.

Since these conditions are never totally satisfied, no market is really efficient.

A.7 Performance Measures

To evaluate the performance of a certain portfolio, certain measures are developed. In this section some frequently used performance measures are discussed.

A.7.1 Beta

The performance measurement theory is rooted in the Capital Asset Pricing Model (CAPM, [5]). The market risk measure in the CAPM model is given by:

$$\mathbb{E}[r_p] = r_f + \beta * (\mathbb{E}[r_{mkt}] - r_f)$$

where	r_p	is the return on portfolio p,
	r_f	is the risk free interest rate,
	r_{mkt}	is the return on the market,
	β	is the risk measure.

A.7.2 Downside Risk

The performance measure that is called downside risk is given by the probability of negative return.

A.7.3 Sharpe Ratio

The Sharpe ratio is given by the difference between returns that a portfolio provides and the risk free rate of return divided by the standard deviation of the portfolio:
$SR_p = \frac{\mathbb{E}[r_p]-r_f}{\sigma_p}$

where	r_p	is the return on portfolio p,
	r_f	is the risk free interest rate,
	σ_p	is the standard deviation of the return on portfolio p.

APPENDIX B

Derivation of the Expected Return of a Short Straddle that is Hedged Discretely

B.1 Assumptions and Notation

The assumptions that have been made to derive the expectation of the return in a 'Black and Scholes world'. The following assumptions have been made:

- The implied volatility and the realised volatility are constant (like in the Black and Scholes model).
- The implied and the realised volatility differ a constant factor: $\sigma_{impl} - \sigma_{real} = \kappa$.
- The dividend rate is zero: $q = 0$.
- Hedging is done discretely and the rebalancing is done n times with intervals of the same length.
- The position ends at time T.
- There are no transaction costs.
- The index is assumed to have a lognormal distribution. The increments in the index are independent.

The following notation is used:

- W_t is the value of the hedge-portfolio at time t.
- S_t is the value of the index at time t.
- C_0 and P_0 are the values of the call and put at time 0.
- K is the strike price.
- r is the risk-free interest rate.
- Δ_t is the delta of the position at time t.
- N(x) is the cumulative probability function of the standardised normal distribution function.

B.2 Total Return

The total return of a discretely hedged straddle is given by the value of the portfolio at time T plus the value of the original position at time 0 (a call and a put are sold at time 0) minus the settlement price of the call and put at time T:

$$\mathit{TotalReturn} = W_T - \mathit{OptionSettle} + (S_0 + C_0 + P_0)e^{rT}. \tag{B.1}$$

The expected return is therefore given by:

$$\mathbb{E}[\mathit{TotalReturn}] = \mathbb{E}[W_T] - \mathbb{E}[\mathit{OptionSettle}] + \mathbb{E}[(S_0 + C_0 + P_0)e^{rT}]. \tag{B.2}$$

In this equation:

- $\mathbb{E}[W_T]$, this is derived in the next section,
- $\mathbb{E}[\mathit{OptionSettle}] = \mathbb{E}[-\max\{0, S_T - K\} - \max\{0, K - S_T\}$, this is derived in Section B.2.2
- $(S_0 + C_0 + P_0)e^{rT}$ is deterministic.

B.2.1 Value at Time T

The value of the portfolio at time t is given by the value of the portfolio at time $t-1$ plus the change in portfolio due to the delta-hedge at time t.

$$W_t = W_{t-1} + \Delta_{t-1}(S_{t-1}e^{r(T(1-\frac{t-1}{n}))} - S_t e^{r(T(1-\frac{t}{n}))}) \Rightarrow$$

$$W_t = W_0 + \sum_{j=0}^{t-1} \Delta_{j-1}(S_{j-1}e^{r(T(1-\frac{j-1}{n}))} - S_j e^{r(T(1-\frac{j}{n}))}). \tag{B.3}$$

In the Black and Scholes model, the delta of a call option is given by $\Delta_{i,call} = \mathrm{N}(d_1^{(i)})$ and the delta of a put option is given by $\Delta_{i,put} = \mathrm{N}(d_1^{(i)}) - 1$, so the delta of a short position in a call and a put option is given by $\Delta_i = -2\mathrm{N}(d_1^{(i)}) + 1$

The expectation of W_T is therefore:

$$\mathbb{E}[W_T] = \mathbb{E}[W_0] + \sum_{j=0}^{t-1} \mathbb{E}[\Delta_{j-1}(S_{j-1}e^{r(T(1-\frac{j-1}{n}))} - S_j e^{r(T(1-\frac{j}{n}))})]$$

$$= 0 + \sum_{j=0}^{t-1} e^{r(T(1-\frac{j-1}{n}))}\mathbb{E}[(-2\mathrm{N}(d_1^{(j-1)}) + 1)S_{j-1}] - e^{r(T(1-\frac{j}{n}))}\mathbb{E}[(-2\mathrm{N}(d_1^{(j-1)}) + 1)S_j]$$

$$= 0 + \sum_{j=0}^{t-1} e^{r(T(1-\frac{j-1}{n}))}\mathbb{E}[-2\mathrm{N}(d_1^{(j-1)})S_{j-1} + S_{j-1}] - e^{r(T(1-\frac{j}{n}))}\mathbb{E}[-2\mathrm{N}(d_1^{(j-1)})S_j + S_j]. \tag{B.4}$$

To find $\mathbb{E}[W_T]$, first the following expectations have to be found:

1. $\mathbb{E}[S_j]$ for all $j = 1, \ldots, n$,
2. $\mathbb{E}[\mathrm{N}(d_1^{(j)})S_j]$ for all $j = 1, \ldots, n-1$,
3. $\mathbb{E}[\mathrm{N}(d_1^{(j)})S_{j+1}]$ for all $j = 1, \ldots, n-1$.

1. $\rightarrow \mathbb{E}[S_j] = S_0 e^{\mu \frac{j}{n} T}$ because the index is lognormally distributed.
2. $\rightarrow \mathbb{E}[\mathrm{N}(d_1^{(j)})S_j] = S_0 e^{\mu \frac{j}{n} T} E_{\eta_j}[\Phi(a^{(j)} + b^{(j)}\eta_j)]$,
   where:

- $\eta_j \sim \mathrm{N}((\sigma_{impl} - \kappa) j \sqrt{\frac{1}{n}T}, j)$,
- $E_{\eta_j}[\Phi(a^{(j)} + b^{(j)}\eta_j)] = \mathrm{N}(\frac{a^{(j)}+b^{(j)}(\sigma_{impl}-\kappa)j\sqrt{\frac{1}{n}T}}{\sqrt{1+b^{(j)2}j}})$,
- $a^{(j)} = \frac{\ln(\frac{S_0}{K})+fracjn\mu T - \frac{j}{2}(\sigma_{impl}-\kappa)^2(\frac{1}{n}T)+(r+\frac{1}{2}\sigma^2_{impl})(\frac{n-j}{n}T)}{\sigma_{impl}\sqrt{\frac{n-j}{n}T}}$,
- $b^{(j)} = (1 - \frac{\kappa}{\sigma_{impl}})\sqrt{\frac{1}{n-j}}$,
- $d_1^{(j)} = a^{(j)} + b^{(j)}(\varepsilon_1 + \ldots + \varepsilon_j)$, where $\varepsilon_i \sim \mathrm{N}(0,1)$.

This is derived below.

3. $\rightarrow \mathbb{E}[\mathrm{N}(d_1^{(j)})S_{j+1}] = S_0 e^{\mu\frac{j+1}{n}T} E_{\eta_j}[\Phi(a^{(j)} + b^{(j)}\eta_j)]$.
This is derived below.

Derivation of Expectations

In this section is shown that

$$\mathbb{E}[\mathrm{N}(d_1^{(j)})S_j] = S_0 e^{\mu\frac{j}{n}T} E_{\eta_j}[\Phi(a^{(j)} + b^{(j)}\eta_j)]$$

and

$$\mathbb{E}[\mathrm{N}(d_1^{(j)})S_{j+1}] = S_0 e^{\mu\frac{j+1}{n}T} E_{\eta_j}[\Phi(a^{(j)} + b^{(j)}\eta_j)].$$

S_j is given by:

$$S_j = S_{j-1} e^{(\mu - \frac{1}{2}(\sigma_{impl}-\kappa)^2)(\frac{1}{n}T)+(\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}\varepsilon_j} = \ldots = S_0 e^{j(\mu - \frac{1}{2}(\sigma_{impl}-\kappa)^2)(\frac{1}{n}T)+(\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}(\varepsilon_1+\ldots+\varepsilon_j)}$$

$d_1^{(j)}$ is given by:

$$d_1^{(j)} = \frac{\ln(\frac{S_j}{K}) + (r + \frac{1}{2}\sigma^2_{impl})\frac{n-j}{n}T}{\sigma_{impl}\sqrt{\frac{n-j}{n}T}} =$$

$$\frac{\ln(\frac{S_0}{K}) + \frac{j}{n}\mu T - \frac{1}{2}(\sigma_{impl}-\kappa)^2\frac{j}{n}T + (\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}(\varepsilon_1+\ldots+\varepsilon_j)}{\sigma_{impl}\sqrt{\frac{n-j}{n}T}} + \frac{(r+\frac{1}{2}\sigma^2_{impl})\frac{n-j}{n}T}{\sigma_{impl}\sqrt{\frac{n-j}{n}T}}$$
$$= a^{(j)} + b^{(j)}(\varepsilon_1 + \ldots + \varepsilon_j),$$

with:

$$a^{(j)} = \frac{\ln(\frac{S_0}{K}) + \frac{j}{n}\mu T - \frac{1}{2}(\sigma_{impl}-\kappa)^2\frac{j}{n}T + (r+\frac{1}{2}\sigma^2_{impl})\frac{n-j}{n}T}{\sigma_{impl}\sqrt{\frac{n-j}{n}T}},$$

$$b^{(j)} = \frac{(\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}(\varepsilon_1+\ldots+\varepsilon_j)}{\sigma_{impl}\sqrt{\frac{n-j}{n}T}},$$

and

$$\varepsilon \sim \mathrm{N}(0,1).$$

Furthermore, we have for the general situation:

$$E_\varepsilon[\Phi(a+b\varepsilon)\exp(c+d\varepsilon)] = \int_0^\infty \Phi(a+b\varepsilon)\exp(c+d\varepsilon)f(\varepsilon)\mathrm{d}\varepsilon$$

and

$$\varepsilon \sim \mathrm{N}(\mu^*, \sigma^{*2}),$$

$$E_\varepsilon[\Phi(a+b\varepsilon)\exp(c+d\varepsilon)] = \int_0^\infty \Phi(a+b\varepsilon)\exp(c+d\varepsilon)\frac{1}{\sqrt{2\pi\sigma^{*2}}}\exp(-\frac{1}{2}(\frac{\varepsilon-\mu^*}{\sigma^*})^2)\mathtt{d}\varepsilon =$$
$$\exp(c)\int_0^\infty \Phi(a+b\varepsilon)\frac{1}{\sqrt{2\pi\sigma^{*2}}}\exp(\frac{-\frac{1}{2}\varepsilon^2+\varepsilon\mu^*+d\varepsilon\sigma^{*2}-\frac{1}{2}\mu^{*2}}{\sigma^{*2}})\mathtt{d}\varepsilon =$$
$$\exp(c+d(\mu^*+\frac{1}{2}d\sigma^{*2}))\int_0^\infty \Phi(a+b\varepsilon)\frac{1}{\sqrt{2\pi\sigma^{*2}}}\exp(-\frac{1}{2}(\frac{\varepsilon-(\mu^*+d\sigma^{*2})}{\sigma^*})^2)\mathtt{d}\varepsilon =$$
$$\exp(c+d(\mu^*+\frac{1}{2}d\sigma^{*2}))E_\eta[\Phi(a+b\eta)] = \exp(c+d(\mu^*+\frac{1}{2}d\sigma^{*2}))\mathrm{N}(\frac{a+b(\mu^*+d\sigma^{*2})}{\sqrt{1+b^2\sigma^{*2}}}),$$

with

$$\eta \sim \mathrm{N}(\mu^*+d\sigma^{*2},\sigma^{*2}).$$

Therefore we have:

$$\mathbb{E}[\mathrm{N}(d_1^{(j)})S_j] =$$
$$S_0e^{\frac{j}{n}T\mu}E_{\varepsilon_1+\ldots+\varepsilon_j}[\Phi(a^{(j)}+b^{(j)}(\varepsilon_1+\ldots+\varepsilon_j))\exp(-\frac{1}{2}(\sigma_{impl}-\kappa)^2(\frac{j}{n}T)+(\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}(\varepsilon_1+\ldots+\varepsilon_j))] =$$
$$S_0e^{\frac{j}{n}T\mu}E_{\eta_j}[\Phi(a^{(j)}+b^{(j)}\eta_j)] = S_0e^{\frac{j}{n}T\mu}\mathrm{N}(\frac{a^{(j)}+b^{(j)}(\sigma_{impl}-\kappa)j\sqrt{\frac{1}{n}T}}{\sqrt{1+b^{(j)2}j}})$$

and

$$\mathbb{E}[\mathrm{N}(d_1^{(j)})S_{j+1}] =$$
$$S_0e^{\frac{j+1}{n}T\mu}E_{\varepsilon_1+\ldots+\varepsilon_j}[\Phi(a^{(j)}+b^{(j)}(\varepsilon_1+\ldots+\varepsilon_j))\exp(-\frac{1}{2}(\sigma_{impl}-\kappa)^2(\frac{j+1}{n}T)+(\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}(\varepsilon_1+\ldots+\varepsilon_{j+1}))] =$$
$$S_0e^{\frac{j+1}{n}T\mu-\frac{1}{2}(\sigma_{impl}-\kappa)^2(\frac{j+1}{n}T)}*$$
$$E_{\varepsilon_1+\ldots+\varepsilon_j}[\Phi(a^{(j)}+b^{(j)}(\varepsilon_1+\ldots+\varepsilon_j))\exp((\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}(\varepsilon_1+\ldots+\varepsilon_j))\exp((\sigma_{impl}-\kappa)\sqrt{\frac{1}{n}T}\varepsilon_{j+1})] =$$
$$S_0e^{\frac{j+1}{n}T\mu-\frac{j+1}{2}(\sigma_{impl}-\kappa)^2(\frac{1}{n}T)+\frac{j}{2}(\sigma_{impl}-\kappa)^2(\frac{1}{n}T)}E_{\eta_j}[\Phi(a^{(j)}+b^{(j)}\eta_j)]\mathbb{E}[e^{(\sigma_{impl}-\kappa)(\sqrt{\frac{1}{n}T})\varepsilon_{j+1}}] =^1$$
$$S_0e^{\frac{j+1}{n}T\mu}E_{\eta_j}[\Phi(a^{(j)}+b^{(j)}\eta_j)] = S_0e^{\frac{j+1}{n}T\mu}\mathrm{N}(\frac{a^{(j)}+b^{(j)}(\sigma_{impl}-\kappa)j\sqrt{\frac{1}{n}T}}{\sqrt{1+b^{(j)2}j}})$$

B.2.2 Option Settle

The expected option settle is determined by:

$$\mathbb{E}[\mathit{OptionSettle}] = \tag{B.5}$$
$$\mathbb{E}[-\max\{0,S_T-K\}-\max\{0,K-S_T\}] = -\mathbb{E}[\max\{0,S_T-K\}]-\mathbb{E}[\max\{0,K-S_T\}] =$$
$$-\mathbb{E}[S_0\mathrm{N}(d_1^*)e^{rT}-K\mathrm{N}(d_2^*)]-\mathbb{E}[K\mathrm{N}(-d_2^*)-S_0\mathrm{N}(-d_1^*)e^{rT}] =$$
$$-S_0\mathrm{N}(d_1^*)e^{\mu T}-K\mathrm{N}(d_2^*)-K\mathrm{N}(-d_2^*)-S_0\mathrm{N}(-d_1^*)e^{\mu T}$$

where:

- $d_1^* = \frac{\ln(\frac{S_0}{K})+(\mu+\frac{1}{2}\sigma_{real}^2)T}{\sigma_{real}\sqrt{T}}$
- $d_2^* = d_1^* - \sigma_{real}\sqrt{T}$

[1] $\mathbb{E}[e^{(\sigma_{impl}-\kappa)(\sqrt{\frac{1}{n}T})\varepsilon_{j+1}}]$ is a moment generating function, so: $\mathbb{E}[e^{(\sigma_{impl}-\kappa)(\sqrt{\frac{1}{n}T})\varepsilon_{j+1}}] = M((\sigma_{impl}-\kappa)(\sqrt{\frac{1}{n}T})) = e^{\frac{1}{2}(\sigma_{impl}-\kappa)^2(\frac{1}{n}T)}$

B.3 Results

The total expected return of a straddle is:

$$\begin{aligned}\mathbb{E}[\text{TotalReturn}] &= (S_0 + C_0 + P_0)e^{rT} \\
&+ \sum_{j=0}^{n-1} [e^{r(T-\frac{j}{n}T)}(-2S_0 e^{\frac{j}{n}T\mu}\mathrm{N}(\frac{a^{(j)} + b^{(j)}(\sigma_{impl} - \kappa)j\sqrt{\frac{1}{n}T}}{\sqrt{1 + b^{(j)2}j}}) + S_0 e^{\frac{j}{n}T\mu})] \\
&- \sum_{j=0}^{n-1} [e^{r(T-\frac{j+1}{n}T)}(-2S_0 e^{\frac{j+1}{n}T\mu}\mathrm{N}(\frac{a^{(j)} + b^{(j)}(\sigma_{impl} - \kappa)j\sqrt{\frac{1}{n}T}}{\sqrt{1 + b^{(j)2}j}}) + S_0 e^{\frac{j+1}{n}T\mu})] \\
&- S_0\mathrm{N}(d_1^*)e^{\mu T} - K\mathrm{N}(d_2^*) - K\mathrm{N}(-d_2^*) - S_0\mathrm{N}(-d_1^*)e^{\mu T}\end{aligned} \tag{B.6}$$

APPENDIX C

Derivation of the Difference between the Reference Variance and the Reference Volatility

In this appendix the difference between $K_{var} = \mathbb{E}[Variance]$ and $K_{vol} = \mathbb{E}[\sqrt{Variance}]$ is given.

Notice that the Taylor expansion of $F(x) = x^{1/2}$ around x_0 is:

$$F(x) \approx F(x_0) + F'(x_0)(x - x_0) + \frac{1}{2}F''(x_0)(x - x_0)^2 + \ldots =$$

$$x_0^{1/2} + \frac{(x - x_0)}{2x_0^{1/2}} - \frac{(x - x_0)^2}{8x_0^{3/2}} + \ldots =$$

$$\frac{(x + x_0)}{2x_0^{1/2}} - \frac{(x - x_0)^2}{8x_0^{3/2}} + \ldots .$$

So for $x = v = Variance$ and $x_0 = \mathbb{E}[v] = K_{var}$, we have:

$$\sqrt{v} \approx \frac{(v + \mathbb{E}[v])}{2\sqrt{\mathbb{E}[v]}} - \frac{(v - \mathbb{E}[v])^2}{8(\mathbb{E}[v])^{3/2}} + \ldots .$$

By taking expectations on both sides:

$$K_{vol} = \mathbb{E}[\sqrt{v}] \approx \frac{(\mathbb{E}[v] + \mathbb{E}[v])}{2\sqrt{\mathbb{E}[v]}} - \frac{\mathbb{E}[(v - \mathbb{E}[v])^2]}{8(\mathbb{E}[v])^{3/2}} + \ldots = \tag{C.1}$$

$$\sqrt{\mathbb{E}[v]} - \frac{\mathrm{var}[v]}{8(\mathbb{E}[v])^{3/2}} + \ldots = \sqrt{K_{var}} - \frac{\mathrm{var}[v]}{8(K_{var})^{3/2}} + \ldots .$$

APPENDIX D

Derivation of the Fair Variance for the Skew Function of the Strike

In this appendix, the derivation of the fair variance for the skew a general function of the strike is given. The skew is given by:

$$\Sigma(K) = \Sigma_0 + b_3 f(K) \tag{D.1}$$

The fair variance of a variance swap is given by (Section 4.6.1):

$$K_{var} = \frac{2}{T}(rT - (\frac{S_0}{S_*}e^{rT} - 1) - \ln(\frac{S_*}{S_0}) + e^{rT}\int_0^{S_*}\frac{1}{K^2}P(K,\Sigma(b_3))dK + e^{rT}\int_{S_*}^{\infty}\frac{1}{K^2}C(K,\Sigma(b_3))dK) \tag{D.2}$$

The option prices for the call and the put options, $C(K, \Sigma(b_3))$ and $P(K, \Sigma(b_3))$, can be expanded as power series in b_3 around $b_3 = 0$ (the implied volatility is Σ_0 for all K):

$$C(K,\Sigma(b_3)) = C(K,\Sigma_0) + b_3\frac{\partial C}{\partial b_3}|_{b_3=0} + \frac{1}{2}b_3^2\frac{\partial^2 C}{\partial b_3^2}|_{b_3=0} + \ldots \tag{D.3}$$

$$P(K,\Sigma(b_3)) = P(K,\Sigma_0) + b_3\frac{\partial P}{\partial b_3}|_{b_3=0} + \frac{1}{2}b_3^2\frac{\partial^2 P}{\partial b_3^2}|_{b_3=0} + \ldots \tag{D.4}$$

When

$$\Sigma_0^2 = \frac{2}{T}(rT - (\frac{S_0}{S_*}e^{rT} - 1) - \ln(\frac{S_*}{S_0}) + e^{rT}\int_0^{S_*}\frac{1}{K^2}P(K,\Sigma(b_3))dK + e^{rT}\int_{S_*}^{\infty}\frac{1}{K^2}C(K,\Sigma(b_3))dK),$$

this gives the following formula for the fair variance:

$$\begin{aligned} K_{var} = \Sigma_0^2 + b_3(\frac{2}{T}e^{rT})\{\int_0^{S_*}\frac{1}{K^2}\frac{\partial P}{\partial b_3}|_{b_3=0}dK + \int_{S_*}^{\infty}\frac{1}{K^2}\frac{\partial C}{\partial b_3}|_{b_3=0}dK\} + \\ \frac{1}{2}b_3^2(\frac{2}{T}e^{rT})\{\int_0^{S_*}\frac{1}{K^2}\frac{\partial^2 C}{\partial b_3^2}|_{b_3=0}dK + \int_{S_*}^{\infty}\frac{1}{K^2}\frac{\partial^2 C}{\partial b_3^2}|_{b_3=0}dK\} + \ldots \end{aligned} \tag{D.5}$$

The derivatives that have to be found to solve equation D.5 are:

- $$\frac{\partial P}{\partial b_3}|_{b_3=0} = \frac{\partial P}{\partial \Sigma}|_{\Sigma_0}\frac{\partial \Sigma}{\partial b_3}|_{b_3=0},$$

- $$\frac{\partial C}{\partial b_3}\mid_{b_3=0} = \frac{\partial C}{\partial \Sigma}\mid_{\Sigma_0} \frac{\partial \Sigma}{\partial b_3}\mid_{b_3=0},$$

- $$\frac{\partial^2 P}{\partial b_3^2}\mid_{b_3=0} = \frac{\partial^2 P}{\partial \Sigma^2}\mid_{\Sigma_0} (\frac{\partial \Sigma}{\partial b_3})^2\mid_{b_3=0} + \frac{\partial P}{\partial \Sigma}\mid_{\Sigma_0} \frac{\partial^2 \Sigma}{\partial b_3^2}\mid_{b_3=0},$$

- $$\frac{\partial^2 C}{\partial b_3^2}\mid_{b_3=0} = \frac{\partial^2 C}{\partial \Sigma^2}\mid_{\Sigma_0} (\frac{\partial \Sigma}{\partial b_3})^2\mid_{b_3=0} + \frac{\partial C}{\partial \Sigma}\mid_{\Sigma_0} \frac{\partial^2 \Sigma}{\partial b_3^2}\mid_{b_3=0},$$

The derivatives to volatility are calculated using the Black and Scholes formula:

$$\frac{\partial P}{\partial \Sigma}\mid_{\Sigma_0} = \frac{\partial C}{\partial \Sigma}\mid_{\Sigma_0} = \frac{S\sqrt{T}}{\sqrt{2\pi}} e^{-d_1^2/2}$$

and

$$\frac{\partial^2 P}{\partial \Sigma^2}\mid_{\Sigma_0} = \frac{\partial^2 C}{\partial \Sigma^2}\mid_{\Sigma_0} = -\frac{S\sqrt{T}}{\sqrt{2\pi}} d_1 \frac{\partial d_1}{\partial \Sigma_0} e^{-d_1^2/2}$$

$$d_1 = \frac{\ln(\frac{S_F}{K}) + \frac{1}{2}\Sigma_0^2 T}{\Sigma_0\sqrt{T}}$$

For the skew that is considered here:

$$\frac{\partial \Sigma}{\partial b_3}\mid_{b_3=0} = f(K)$$

and

$$\frac{\partial^2 \Sigma}{\partial b_3^2}\mid_{b_3=0} = 0$$

Furthermore:

$$\frac{\partial d_1}{\partial \Sigma_0} = \frac{1}{2}\sqrt{T} - \frac{\ln(\frac{S_F}{K})}{\Sigma_0^2\sqrt{T}}$$

Filling in these derivatives gives the following formula for the fair variance:

$$K_{var} = \Sigma_0^2 + b_3(\frac{2}{T}e^{rT})\{\int_0^\infty \frac{1}{K^2} f(K) \frac{S\sqrt{T}}{\sqrt{2\pi}} e^{-d_1^2/2} dK\} \tag{D.6}$$
$$+\frac{1}{2}b_3^2(\frac{2}{T}e^{rT})\{\int_0^\infty \frac{1}{K^2} f^2(K)(\frac{\ln(\frac{S_F}{K})}{\Sigma_0^2\sqrt{T}} - \frac{1}{2}\sqrt{T})\frac{S\sqrt{T}}{\sqrt{2\pi}} d_1 e^{-d_1^2/2} dK\} + \ldots.$$

Bibliography

[1] Bates, D.S., Empirical Option Pricing: A Retrospection, *Journal of Econometrics* 116 (2003), 387–404.

[2] Ben-Tal, A. and Nemirovski, A., Robust convex optimization, *Mathematics of Operations Research* 23 (1998), Vol. 4, 769–805.

[3] Black, F. and Scholes, M., The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81 (1973), 637–654.

[4] Bollerslev, T., Generalised Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* 31 (1986), 307–327.

[5] Brealey, R.A., Myers, S.C. and Marcus, A.J., *Fundamentals of Corporate Finance* (1999), second edition, Irwin McGraw-Hill (First edition 1995).

[6] Canina, L. and Figlewski, S., The Informational Content of Implied Volatility, *Review of Financial Studies* 6-3 (1993), 659–81.

[7] Carr, P. and Madan, D., Towards a Theory of Volatility Trading, in "Volatility: New Estimation Techniques for Pricing Derivatives", *Risk Publications* (1998), R. Jarrow, 417–427.

[8] Coval, J.D. and Shumway, T., Expected Option Returns, *Journal of Finance* 56 (2000), 983–1009.

[9] Cox, J.C. and Ross, S.A., The Valuation of Options for Alternative Stochastic Processes, *Journal of Financial Economics* 3 (1976), 145–166.

[10] Demeterfi, K., Derman, E., Kamal, M. and Zou, J., More Than You Ever Wanted to Know about Volatility Swaps, *Quantitative Strategies Research Note*, Goldman Sachs (1999).

[11] Driessen, J. and Maenhout, P., A Portfolio Perspective on Option Pricing Anomalies, *AFA 2004 San Diego Meetings; EFA 2003 Annual Conference Paper* No. 916 (2003).

[12] Gatheral, J. and Lynch M., Lecture 6: Volatility and Variance Swaps, *Case Studies in Financial Modelling Course Notes* (2002).

[13] Goldfarb, D. and Iyengar, G., Robust Portfolio Selection Problems, *Mathematics of Operations Research* 28 (2003), Vol. 1, 1–38.

[14] Heston, S.L., A Closed-form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *The review of financial studies* 6 (1993), 327–343.

[15] Heston, S.L. and Nandi, S., A Closed-form GARCH Option Pricing Model, *Working paper* 97-9 (1997), Federal Reserve Bank of Atlanta.

[16] Howison, S.D., Rafailidis, A. and Rasmussen, H.O., A Note on the Pricing and Hedging of Volatility Derivatives, Oxford Financial Research Centre, *OFRC Working Paper Series*, Number 2001mf09 (2001).

[17] Hull, J.C., *Options, Futures, and Other Derivatives*, Fifth Ed. (2003), Prentice Hall.

[18] Hull, J.C. and White, A., The Pricing of Options on Assets with Stochastic Volatility, *Journal of Finance* 42 (1987), 281–300.

[19] Jackwerth, J.C., Generalised Binomial Trees, *Research Program in Finance Working Paper* (1997), University of California, Berkeley.

[20] Markowitz, H. M., Portfolio selection, *Journal of Finance* 7 (1952), 77–91.

[21] Petit, M.E., Pergamentsev, S. and Tolenaar, J., Volatility Product, *Internal Memorandum* (2004).

[22] Rubinstein, M., Implied Binomial Trees, *Journal of Finance* 49 (1994), 771–818.

[23] Rubinstein, M., Displaced Diffusion Option Pricing, *Journal of Finance* 38 (1983), 213–217.

[24] Stein, E.M. and Stein, J.C., Stock Price Distributions with Stochastic Volatility: an Analytic Approach, *review of financial studies* 4 (1991), 727–752.

[25] Windcliff, H., Forsyth, P.A. and Vetzal, K.R., Pricing Methods and Hedging Strategies for Volatility Derivatives, *University of Waterloo working paper* (2003).