Option Valuation Under Stochastic Volatility

-An Empirical Investigation of the Weak GARCH Diffusion Model

by

Eivind Sars Veddeng

Master's thesis in Financial Economics

Trondheim, May 25, 2012

Norwegian University of Science and Technology

Faculty of Social Sciences and Technology Management

Department of Economics

Academic Supervisor: Professor Egil Matsen

NTNU – Trondheim
Norwegian University of
Science and Technology

Preface

This thesis is the completion of the two-year Master of Science program in Financial Economics at the Norwegian University of Science and Technology (NTNU). During the exciting but challenging process I have strengthened my interest and deepened my understanding of option pricing and financial modeling. The thesis has been written in LaTeX. Microsoft Excel with VBA, OxMetrics and MATLAB have been used extensively.

I would like to thank my supervisor Professor Egil Matsen for guidance and helpful suggestions over the past six months. I would also like to thank Dag Skreden for proof-reading.

Eivind Sars Veddeng
Trondheim, May 2012

Contents

List of Figures

List of Tables

1 Introduction

Suppose we use the standard deviation... of possible future returns on a stock... as a measure of its volatility. Is it reasonable to take that volatility as a constant over time? I think not. -Fischer Black (1976)

The option valuation problem is a very fascinating topic and has evolved significantly over the past decades. In the early 1900s Louis Bachelier suggested a fair game approach using a normal distribution for the underlying security price, almost the same as the 1973 Black and Scholes model. Black and Scholes let the underlying security follow a process called a random walk, where the change in the stock price over a pre-specified period is completely random. In continuous time the process is said to follow a Brownian motion.

As suggested in the quote in the beginning of this section, the assumption of constant volatility is rather debatable. Figure 1 which shows the daily returns on the Standard & Poor 500 index (SPX) calls for another approach. In the figure, periods of large fluctuation in the returns are followed by new periods of large fluctuations. Similarly, periods with small fluctuations also tend to follow each other. This tendency is called volatility *clustering* where the uncertainty of the returns i.e. the *volatility* tends to cluster. The volatility is usually measured by the standard deviation of the returns. This empirical fact can not be sufficiently captured by a model that assumes this measure of uncertainty as a constant.

In the existing literature there are many different models that try to correct for the clustering of returns. A common approach is to allow the volatility to be *stochastic*, meaning that it is allowed to fluctuate over time. Bakshi, Cao and Chen (1997) summarize some of the different approaches including stochastic volatility models, stochastic interest rate models, stochastic models with jumps in the stock price and constant elasticity of variance models.[1]

Bakshi et al. test which of the models described above reflects the market prices most precisely. The article gives empirical evidence that stochastic volatility is of first order importance in improving the Black-Scholes formula.

The motivation behind this thesis is a sincere interest in option pricing, implied volatility and financial modeling. Also, the thought of using a version of the Generalized Autoregressive Conditional Heteroskedasticity model (GARCH), a model designed to pick up volatility clustering, is intriguing. The GARCH model has been an interesting topic in the Master Program in Fi-

[1]For:
- Stochastic interest models see Merton(1973) and Amin and Jarrow(1992)
- Jump -diffusion models see Merton(1976), Bates(1991) and Madan and Chang (1996)
- Constant elasticity of variance see Cox and Ross(1976).

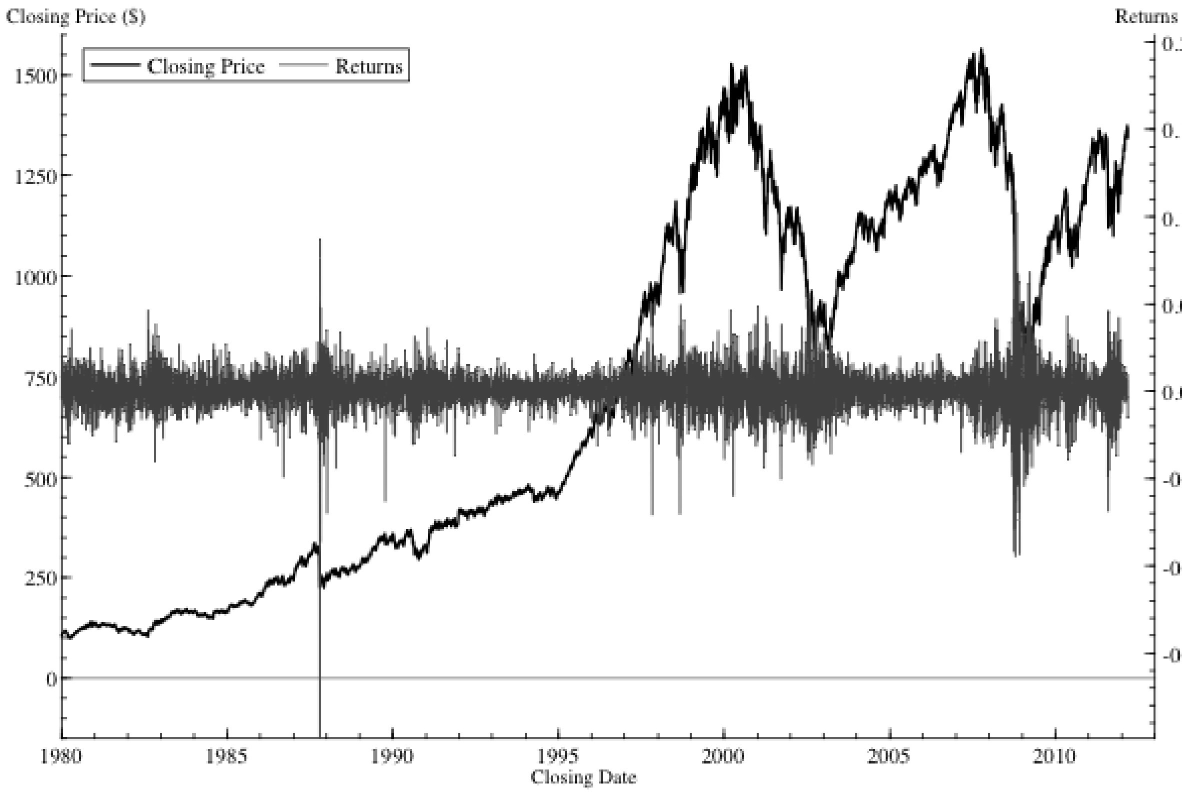

Figure 1: Standard & Poor 500 daily closing price and logarithmic returns.
The dataset consists of 8120 observations from January 1st 1980 to March 8th 2012.

nancial Economics. Other graduate topics covered in the program was financial modeling using stochastic differential equations. In Alexander (2008b) we were introduced to a model using a continuous version of the GARCH, with an option to include skewness and kurtosis, as the volatility process. This process is referred to as the *weak* GARCH. In the same book the author praised the properties of the model but issued a warning because of the difficult implementation. The thought of carrying out an analysis of a GARCH model using stochastic differential equations has been a key motivational factor. Also, to the best of our knowledge, no empirical investigation of the option pricing ability of the model has been carried out.

In Lehar, Scheicher and Schittenkopf (2002) an empirical comparison between the discrete time GARCH model and another continuous stochastic volatility model, the Heston model, is performed.[2]The authors conclude that the GARCH model clearly dominates the Black-Scholes model and the Heston model when it comes to out of sample valuation. The Heston model on the other hand, was superior when measuring Value at Risk. The results in the article, combined with the personal motivation described above, provided plenty motivation to analyze the performance of the continuous GARCH model.

The thesis is to perform an empirical investigation of the weak GARCH diffusion model. Cen-

[2]The Heston model is one of the most common stochastic volatility models used for option pricing. It is named after the author Steven Heston.

tral questions are:

(i) Is the weak GARCH diffusion model able to capture the properties of implied volatility?

(ii) Will the model price options significantly better than the benchmark Black-Scholes (BS) model?

In this thesis options will be valued using a Monte Carlo algorithm and directly compared to two Black-Scholes benchmark models. We use two different calibrations of the BS model, one using the historical standard deviation as a proxy for the volatility. The other use the reported implied volatility for options traded on the SPX index. For the GARCH diffusion we will use the limit of a weak discrete-time GARCH model, called the weak GARCH diffusion. This is the limit of a GARCH model with no strong assumption on the distribution of the residuals.

Section one of this thesis is an introduction to the option pricing framework and the underlying assumptions. This is followed by a description of the dataset and the methodology used in the analysis. The ending consists of an analysis of the pricing and volatility properties and concluding remarks.

2 Option Theory, Volatility and Stochastic Processes

2.1 Plain Vanilla Options

The following section is from Mcdonald (2006). Plain vanilla options are defined as the simple options such as a *call* or a *put*. A call is a right but not an obligation to *buy* the underlying security (e.g. the stock) at a predetermined price, *the strike price*. We denote the price of the underlying security at time t as S_t, and the strike price as K. If the stock price is higher than the strike price on the date of expiry, denoted T, a rational investor will choose to exercise the option. If the stock price is below the strike the investor will let the option expire. We define the payoff of a call at the expiration date as:

$$C[S(T), K, T] = max(S_T - K, 0) \tag{2.1}$$

A put option is a right to *sell* the stock to the strike price. Here the investor earns a payoff if the strike is higher than the spot price at expiry.

$$P[S(T), K, T] = max(K - S_T, 0) \tag{2.2}$$

Common option terminology often refers to the option moneyness. An option is *in- the- money* (ITM) if $S_t > K$ for a call, and $K > S_t$ for a put. Analogously the options are *out-of-the-money* (OTM) if the converse holds and *at-the-money* (ATM) if $S_t = K$. In tables and figures we define the degree of moneyness as $\frac{S}{K}$ for calls and $\frac{K}{S}$ for puts.

Options traded in the markets are unfortunately not quoted at one price, a derivative is typically quoted at a *bid*-price, an *ask*-price and in the last traded price. The ask price is the the lowest price a seller is willing to accept for the option, whereas the bid price is the highest price that the bidder (buyer) is willing to pay. The ask price always is higher than the bid price and the difference is called the bid-ask spread.

Using the last traded price often violates the no-arbitrage conditions for options. An arbitrage opportunity is when an investor can realize a positive cash flow with a net zero cost of investment and no risk involved. These opportunities may arise when using the last traded price since deep in the money and out of the money options typically are illiquid, thus making the quoted price stale. This is because of the dynamic pricing property of the option. The value of the underlying security and the time to expiry change over the life of the option making yesterdays quote irrelevant. Consider an ATM call option with six months to expiry at time t and $t + 1$. If the underlying security and the other market parameters does not change the option value will still deteriorate. This is because the time to expiry decrease and thus the time value (the conditional probability of the option ending ITM) of the option is reduced. If the last trade price was reported with the transaction date this is not a problem. However, in several of the option databases like EcoWin, Datastream, Marketwatch etc. this is not the case. Because of this, the price where an investor can acquire an option immediately, namely the ask price, will be used throughout the thesis.

Several types of options also exists, where European and American styled options are the most common. The difference is that European options may only be exercised at the expiry date, while an American option can be exercised anytime between the date of purchase and the date of expiry. Only European options will be considered in this thesis.

2.2 The Brownian Motion

This section explains one of the fundamental assumptions of option pricing and provide some inital motivation for an expansion of the existing framework. In option pricing stock prices are assumed to be *stochastic* and described by a process called a *geometric Brownian motion* which is modeled in continuous time. To understand the framework of which the model is built upon some initial theory is in order. The following is based on Hull (2003) and Mcdonald (2006).

The original option pricing paper by Black and Scholes (1973) and the majority of papers begin by assuming that the price of the underlying asset follows the stochastic differential equation (SDE):

$$\frac{dS(t)}{S(t)} = \alpha dt + \sigma dZ(t) \tag{2.3}$$

In equation (2.3) $S(t)$ is the stock price at time t and $dS(t)$ is the instantaneous infinitesimal change in the price of the stock. α_t is the continuously compounded expected price return of the stock at time t, σ is the standard deviation of the stock at time t, and Z(t) is a standard normally distributed random variable. $Z(t)$ is also referred to as a Wiener Process.

Another way of describing a continuous random process is the *arithmetic* Brownian motion where the SDE takes the form:

$$dS(t) = \alpha dt + \sigma dZ(t) \tag{2.4}$$

In the arithmetic Brownian motion the change in the stock price is not proportional to the initial price, thus allowing the process to become negative. In financial modeling of returns the geometrical Brownian motion is most common, but for other purposes like volatility modeling the arithmetic version is often used.

The stochastic process $Z(t)$ in the Brownian motions described above has the following characteristics (Williams (2006)):

(i) $Z(0) = 0$
The initial value of the process Z is 0.

(ii) $Z(t+s) - Z(t) = \Delta_s Z \sim \mathcal{N}(0, s)$ or
$E[Z_t|\mathcal{F}_s] = Z_s$ for all $s \leq t \in [0, T]$
In discrete time, a change in the process equal to s is normally distributed with mean zero and a standard deviation of s. The same property is valid for continuous time where the change in the process s limits to zero (that is, we let the length of each step become infinitely small). This is also a *martingale* property where the expectation of the stochastic process at time t, equals the value of Z at time s, where $s \leq t$.

(iii) $Z(t+s_1) - Z(t)$ is independent of $Z(t+s_2) - Z(t)$, where $s_1, s_2 > 0$
All draws are independent of each other, i.e. the value at time t is independent of the value at time $t+s$.

(iv) $E[|Z_T|] < \infty$ for all $t \in [0, T]$
The expectation of the process is a finite measure.

(v) Z_t is $\mathcal{F}_t$-measurable for each $t \in [0, t]$
Where $\mathcal{F}_t$ is the filtration of information available to the investor up to time t.

$\mathcal{F}_t$ can also be regarded as the probability space containing all possible events that ever might be observed, or to which we can assign probabilities. The filtration keeps track of which information is known at each time step where $t \in [0.T]$. Thus the information is increasing in time. $Z(t)$ is also called an innovation, whereas the α term in (2.3) is usually called the drift. If Z_t follows the properties above it is referred to as a *martingale*. When simulating random numbers throughout the thesis we let the processes follow the properties above.

2.3 Lognormality

One important implication of the geometric Brownian motion is the properties of $\ln S$ when S follows the process in equation (2.3). To find these properties we utilize Itô's Lemma. Itô's Lemma shows that a function G(S,t) (e.g. a contingent claim) follows the process when S follows (2.3):

$$dG = \Big(\frac{\partial G}{\partial S}\alpha S + \frac{\partial G}{\partial t} + \frac{1}{2}\frac{\partial^2 G}{\partial S^2}\sigma^2 S^2\Big)dt + \frac{\partial G}{\partial S}\sigma S dZ \tag{2.5}$$

This result is based on a Taylor series expansion integrated with respect to time and shows how the claim value changes over time when the underlying security changes.

We can now use the above result to evaluate the properties of $\ln S$. Let $G = \ln S$, since:

$$\frac{\partial G}{\partial S} = \frac{1}{S} \qquad \frac{\partial^2 G}{\partial S^2} = -\frac{1}{S^2} \qquad \frac{\partial G}{\partial t} = 0$$

then:

$$dG = \Big(\alpha - \frac{1}{2}\sigma^2\Big)dt + \sigma dZ \tag{2.6}$$

Because α and σ are constants, the equation indicate that $\ln S$ follows an arithmetic Brownian motion with a constant drift rate $\alpha - \frac{1}{2}\sigma^2$ and constant variance σ^2. By definition the change between $\ln S(t)$ and $\ln S(0)$ is normally distributed with mean $\big(\alpha - \frac{1}{2}\sigma^2\big)t$ and variance $(\sigma^2)t$: that is:

$$\ln[S(t)] - \ln[S(0)] \sim \mathcal{N}\Big(\alpha - \frac{1}{2}\sigma^2)t, \sigma\sqrt{t}\Big)$$

or:

$$\ln[S(t)] \sim \mathcal{N}\Big(\ln[S(0)] + (\alpha - \frac{1}{2}\sigma^2)t, \sigma\sqrt{t}\Big) \tag{2.7}$$

The integral representation of the change in $lnS(t)$ and $lnS(0)$ can be written:

$$\int_0^t \frac{dS(t)}{S(t)} = \ln S(t) - \ln S(0) = \int_0^t (\alpha - \frac{1}{2}\sigma^2)dt + \int_0^t \sigma dZ(t) \tag{2.8}$$

this implies:

$$\ln S(t) = \ln S(0) + (\alpha - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}z$$

or:

$$S(t) = S(0)e^{(\alpha - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}z} \tag{2.9}$$

Where $\sqrt{dt}z$ is the discretization of $dZ(t)$ and $z \sim \mathcal{N}(0,1)$. This is the link between the Brownian motion and lognormality. If a variable is distributed in such a way that instantaneous percentage changes follow a geometric Brownian motion, then over discrete periods of time, the variable is lognormally distributed.

The result is useful since lognormality implies that the price of a stock following a lognormal distribution cannot be negative. Further, the distribution has a long right tail, thus including more extreme values. Equation (2.9) allows us to simulate the stock price process and it will be used extensively throughout the text.

2.4 The Risk Neutral Process

There are several ways of simulating stock prices in such a way that they obey the martingale property. One approach is to find the expected value using the real probability measure and discount with the appropriate interest rate. Another approach is to discount with the risk-free interest rate and change the probabilities such that the expected return of the asset is equivalent to the risk-free rate. Either way we require a specific probability measure to ensure that the process is a martingale.

A key property of the model framework is that there is no involvement of any variable that is affected by risk preferences. If we include α, which is the expected return on the stock, then this would not be the case. This is because the increase in utility of a risk averse investor on a gain of \$1 is less than the utility decrease after a \$1 loss. To offset this, the investor demands a premium $\alpha - r$ in equilibrium, where r is defined as the risk-free interest rate. Therefore a closely related process to equation (2.3) is often used, namely the *risk neutral* process or the *equivalent martingale measure*, where we change the probability measure such that the martingale condition holds.

To arrive at the risk adjusted Brownian motion we expand equation (2.3) to include the continuously compounded dividend yield δ, and denote the true return of the stock as μ. That is, we let the process of (2.3) simulate the pure price evolution meaning that dividends has to be subtracted from the expected return. That is: $\alpha = \mu - \delta$, such that:

$$\frac{dS(t)}{S(t)} = (\mu - \delta)dt + \sigma dZ(t) \tag{2.10}$$

To derive the risk neutral process we use Girsanov's theorem based on Wiersema (2008) which states that a Brownian motion $dS(t)$ can be transformed into a new SDE using a Girsanov transformation. Let the original SDE be specified by (2.10) and rewrite the drift term based on a new coefficient μ_{new}. That is:

$$\begin{aligned}\frac{dS(t)}{S(t)} =& \mu_{new}dt - \mu_{new}dt + \big[\mu - \delta\big]dt + \sigma dZ(t) \\ =& \mu_{new}dt + \big[\mu - \delta - \mu_{new}\big]dt + \sigma dZ(t) \\ =& \mu_{new}dt + \sigma\Big[\big(\frac{\mu - \delta - \mu_{new}}{\sigma}\big)dt + dZ(t)\Big] \\ =& \mu_{new} + \sigma\big[\varphi dt + dZ(t)\big] \end{aligned} \tag{2.11}$$

where:

$$\varphi = \frac{\mu - \delta - \mu_{new}}{\sigma}$$

we then apply the Girsanov Transformation:

$$\tilde{Z}(t) \equiv \int_{s=0}^{t} \varphi ds + Z(t)$$

or:

$$d\tilde{Z}(t) = \varphi dt + dZ(t) \tag{2.12}$$

Where $d\tilde{Z}(t)$ is a Brownian motion under a new probability distribution, P_{new}. Defining μ_{new} as the risk-free drift rate, $r - \delta$, we see that the representation of φ is nothing but the Sharpe ratio, $\varphi = \dfrac{\mu - r}{\sigma}$. Using this and (2.12) in equation (2.10) we obtain:

$$\begin{aligned}\frac{dS(t)}{S(t)} =&(\mu-\delta)dt + \sigma dZ(t) = (\mu-\delta)dt + \sigma\big(d\tilde{Z}(t) - \frac{\mu - r}{\sigma}dt\big) \\ =&(r-\delta)dt + \sigma d\tilde{Z}(t)\end{aligned} \tag{2.13}$$

Equation (2.13) is the *risk-adjusted* process which follows the new probability measure P_{new}. The new probability distribution can be thought of as the continuos time equivalent of the discrete time example where we changed the probability measure such that the asset generated the risk free rate in expectation. To evaluate the new probability measure Wiersema (2008) apply the Radon-Nikodym derivative. For our purpose further investigation of the new probability measure is not necessary (and beyond the scope of this thesis) since the process is risk-neutral. We can perform valuation as if the investors were risk-neutral since we have transformed Z_t to $\tilde{Z}_t$ in order to make the investor behave risk-neutrally with respect to the revised process. This allows us to use equation (2.13) when performing the Monte Carlo valuation of derivatives with the stock as an underlying security. This is also a key property of the Black-Scholes-Merton differential equation since it allows for an option pricing framework without any variables affected by the risk preferences of investors (Hull (2003)).[3]

2.5 The Black-Scholes-Merton Differential Equation

We now turn to the pricing of plain vanilla options based on the work of the article by Black and Scholes (1973). Hull (2003) summarize the assumptions needed on the market where the assets are traded:

(i) The stock price follows a geometric Brownian motion as in equation (2.13).

(ii) Short selling of assets are permitted, meaning that it is possible of creating zero cost portfolios financed by lending money through borrowing and selling stock.

(iii) Frictionless market with no tax distortion. Models including tax distortion effects will not be discussed in this paper.

(iv) Perfectly divisible assets, meaning that one can hold any position preferable in the market.

(v) Efficient markets with no riskless arbitrage. This implies that all securities are priced fairly, and that any mispricing will be instantly exploited by arbitrageurs thus removing all arbitrage opportunities.

(vi) Continuous asset trading, meaning that assets are bought and sold non stop at all times.

[3] Since we use the risk adjusted process in the forthcoming analysis, we will for ease of notation write the Girsanov-transformed Brownian motions as $dZ(t)$ instead of $d\tilde{Z}(t)$.

(vii) The risk free rate, r, is constant and equal for all securities.

The option to be valued is a European call option with the following boundary condition on the maturity date, T, where K is the strike price.

$$G[S(T),T] = \begin{cases} S_T - K & \text{if } S_T \geq K, \\ 0 & \text{if } S_T \leq K, \end{cases} \tag{2.14}$$

Assume that the stock price follows a geometric Brownian motion, that is:

$$\frac{dS(t)}{S(t)} = (\mu - \delta)dt + \sigma dZ(t) \tag{2.15}$$

Since the Wiener process underlying both $dG(S,t)$ and dS are the same, we can choose a portfolio consisting of a short position in the derivative, and a delta hedged position in the stock to remove all source of randomness in the portfolio.[4] Defining the change in this portfolio as:

$$d\Pi = -dG(S,t) + \frac{\partial G}{\partial S}(dS + \delta S dt) \tag{2.16}$$

Where, $\delta S dt$ stems from the underlying stock paying dividends. By applying Itô's Lemma on equation (2.16) we obtain:

$$\begin{aligned} d\Pi = &-\Big(\frac{\partial G}{\partial S}(\mu - \delta)S + \frac{\partial G}{\partial t} + \frac{1}{2}\frac{\partial^2 G}{\partial S^2}\sigma^2 S^2\Big)dt - \frac{\partial G}{\partial S}\sigma S dZ \\ &+ \frac{\partial G}{\partial S}\big((\mu - \delta)S dt + \sigma S dZ + \delta S dt\big) \\ = &\Big(-\frac{\partial G}{\partial t} - \frac{1}{2}\frac{\partial^2 G}{\partial S^2}\sigma^2 S^2 + \frac{\partial G}{\partial S}\delta S\Big)dt \end{aligned} \tag{2.17}$$

This portfolio does not involve any uncertainty, and thus must yield a payoff equal to the risk free interest rate during the time interval. That is:

$$d\Pi = \Big(-\frac{\partial G}{\partial t} - \frac{1}{2}\frac{\partial^2 G}{\partial S^2}\sigma^2 S^2 + \frac{\partial G}{\partial S}\delta S\Big)dt = \Pi r dt \tag{2.18}$$

This implies:

$$\Big(\frac{\partial G}{\partial t} + \frac{1}{2}\frac{\partial^2 G}{\partial S^2}\sigma^2 S^2 - \frac{\partial G}{\partial S}\delta S\Big)dt = r\Big(G(S,t) - \frac{\partial G}{\partial S}S\Big)dt \tag{2.19}$$

or:

$$\frac{\partial G}{\partial t} + (r - \delta)S\frac{\partial G}{\partial S} + \frac{1}{2}\frac{\partial^2 G}{\partial S^2}\sigma^2 S^2 = rG(S,t) \tag{2.20}$$

[4] Delta hedging is defined as holding a position to offset changes in the portfolio price when the underlying security price changes. That is: $\alpha = \frac{\partial G}{\partial S}$.

Equation (2.20) is the famous BSM differential equation. Every contingent claim based on (2.15) must satisfy this equation. There are infinitely many solutions to this equation, but only one satisfying the boundary conditions proposed in equation (2.14). This is the Black-Scholes formula presented in Black and Scholes (1973) where the authors derive the following result for the price of a call option:

$$G[S(t), t] = C_{BS} = Se^{-\delta(T-t)}N(d_1) - Ke^{-r(T-t)}N(d_2) \tag{2.21}$$

where:

$$d_1 = \frac{ln\frac{S}{K} + (r - \delta + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}} \tag{2.22}$$

$$d_2 = \frac{ln\frac{S}{K} + (r - \delta - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}} \tag{2.23}$$

It is easy to verify that the BS call pricing formula follows the BSM equation. A European put option with a boundary condition opposite of (2.14) also follow this equation. The theoretical value of a European put option can be written as:

$$G[S(t), t] = P_{BS} = Ke^{-r(T-t)}N(-d_2) - Se^{-\delta(T-t)}N(-d_1) \tag{2.24}$$

With d_1 and d_2 defined in (2.22) and (2.23). The fact that a portfolio consisting of a short put and a long call yields the underlying instrument as a payoff at all possible states, the net value of this position has to equal a leveraged position in the underlying. That is:

$$\begin{aligned} C_{BS} - P_{BS} =& Se^{-\delta(T-t)}N(d_1) - Ke^{-r(T-t)}N(d_2) - \big(Ke^{-r(T-t)}N(-d_2) - Se^{-\delta(T-t)}N(-d_1)\big) \\ =& Se^{-\delta(T-t)} - Ke^{-r(T-t)} \end{aligned} \tag{2.25}$$

This is known as the *put-call parity*. Both (2.21) and (2.24) are important results that will be used to a large extent throughout the thesis. This framework has since the discovery been the benchmark of both option theory and practical valutation, but are the results above coherent with empirical option prices? In the next section we question a key assumption in the formula for pricing European options, namely the assumption of constant volatility.

2.6 The Implied Volatility Surface

To visualize the problem of constant volatility options on the Standard & Poor 500 index with strikes from \$950 to \$1700 and expiry dates ranging from March 16th to December 21st 2012 have been used. The closing price of the index was \$1358. The data was downloaded from Marketwatch (2012) at February 16th.

Consider the price of a European put option, the Black-Scholes formula (2.24) states the price of the option in terms of other quantities which are already known. All of the input variables are directly observable in the market, such as the interest rate, the dividend yield etc.. The only variable not observable is the standard deviation of the underlying price, i.e. the volatility σ.

Equation (2.24) is derived under the assumption of continuous trading where the time interval between each observation is infinitely small. This allows σ to be the standard deviation of the innovations in the geometric Brownian motion (2.13). According to this process, σ must remain constant over time. As stated earlier, this variable is not observable in the markets so it has to be estimated empirically.

Following Rouah and Vainberg (2007) the most frequently used graphical comparison between option model prices and market prices is done by using the *implied* volatility surface. Given the observed market price of the option, the implied volatility in our framework would be the volatility measure backed out of the option formula. For the Black-Scholes formula to hold the volatility backed out of the option prices must remain constant regardless of moneyness and time to expiration.

To investigate these properties we need a method of finding the implied volatilities of the options. This may only be done numerically since the Black-Scholes formula does not allow a closed form solution. We have used a numerical root finding algorithm of the objective function to find the σ_{iv} such that:

$$P_{obs}(K,T) = P_{BS}(\sigma_{iv}, K, T) \tag{2.26}$$

this is expressed as the root of the objective:

$$f(\sigma) = P_{BS}(\sigma_{iv}, K, T) - P_{obs}(K,T) = 0 \tag{2.27}$$

Where σ_{iv} denotes the implied volatility, $P_{obs}(K,T)$ denotes the observed market put price, and $P_{BS}(K,T)$ is the Black-Scholes theoretical price of the option with the same underlying parameter specifications.

The objective function of a root-finding algorithm takes the value zero when the Black-Scholes theoretical price is matched with the volatility priced in the market. Brooks (2008) argue that the Newton-Rhapson method or the Bisection root-finding method is well suited to seek the implied volatility.[5] A problem with the Newton-Rhapson is its sensitivity to the initial guess location where a remote guess could lead to a drastic deviation from the root. The Bisection method

[5]The Newton- Rhapson seeks the root by iterating $\sigma_{i+1} = \sigma_i - \dfrac{f(\sigma_i)}{f'(\sigma_i)}$, where $f'(\sigma_i) = \dfrac{\partial C_{BS}}{\partial \sigma_i}$ = Option Vega. The Bisection Method seeks the root by evaluating the sign of the root function with two boundary conditions, a and b, then $\sigma_i = 0.5(a+b) = m$.

If $P_{obs} - P_{BS}(m) \begin{cases} > 0 & \text{then } a = m \\ < 0 & \text{then } b = m \end{cases}$

avoids this problem but is less efficient in convergence. To ensure the quick convergence of the Newton-Rhapson combined with the steadiness of the Bisection a Newton-Rhapson Bisection (NRB) method will be applied to find the implied volatilities of the options discussed. The VBA code for the NRB-method can be found in Appendix B.1 and B.2. Below is a brief explanation of the iteration steps in the root-finding algorithm.

The Newton-Rhapson Bisection Volatility Iteration Algorithm (NRB)

(i) The NRB starts with a simple Newton-Rhapson estimation by guessing the midpoint m , of the initial volatility. This is the initial step of the iteration.

(ii) If the current estimate of volatility is outside the current volatility boundary conditions, then a Bisection step is performed.

(iii) Depending on the outcome a new boundary is set following the process of either Newton-Rhapson or the Bisection.

(iv) Proceed to next i.

(v) This process repeats until the objective function is sufficiently close to zero or i approaches the iteration limit, which is set to 1000.

The resulting implied volatilities for the Standard & Poor 500 index options can be found in Appendix A.1.

A problem with the options is that they are traded at strikes increasing with \$5-\$25 depending on the time to expiry. To create a smooth volatility surface we have used a simple interpolation method described in Rouah and Vainberg (2007) called the Natural Cubic Spline Method. The natural cubic spline is a series of cubic polynomials with starting point $\{(x_j, a_i)\}_{i=1}^{n}$defined as

$$S_j(x) = a_j + b_j(x - x_j) + c_j(x - x_j)^2 + d_j(x - x_j)^3 \tag{2.28}$$

With the boundary conditions $S''(x_1) = S''(x_n) = 0$. Finding the coefficients b_j, c_j and d_j is done by working backwards from the linear system of equations from (2.28). The VBA code for the interpolation is given in Appendix B.3.

Using the techniques described above and the option prices in table 13 we are now able to create the implied volatility surface as a function of moneyness and time to expiry. The graphical representation for the SPX options can be found in figure 2.

If the assumption of a constant volatility is correct, the NRB- algorithm should find the same implied volatility for all strikes and expiry dates, i.e. the surface in figure 2 should be flat. This is clearly not the case. When interpreting the implied volatility we observe some common structures, also discussed by Misra, Kannan and Misra (2006), Rubinstein (1998) and Alexander (2008b):

(i) **The Floating Smile Property:** At any fixed expiration, options with moneyness near 1 almost always trade at lower implied volatilities than deep in- and out-of- the- money options. This creates a smile or a smirk, known as the floating smile property. The property is referred to as floating because the volatility surface is very sensitive to the current price of the underlying security(Alexander (2008b)). Investigating the float on the SPX options we look at table 12 where an April call option with strike \$1200 has an implied volatility of 28.2%. If the underlying spot price were to change with 0.5%, ceteris paribus, the implied volatility decreases to 22.66% or with 24.41%. Recognizing this and identifying the local volatility minimum at the April option with strike around \$1450 we conclude that the SPX options has the floating smile property.

(ii) **The Strike Structure:** At any fixed expiration the implied volatility vary with the strike level. In almost every case, the implied volatility increase with decreasing strike on equity options. That is, out-of-the-money puts trade at higher implied volatilities than out-of-the-money calls. This property is often called the negative skew since it is comparable with a negatively skewed log price density. We find the same property in our options looking at figure 3 where the skewness is observable when looking at the relative difference of volatility on out-of-the-money call options and out-of-the-money puts.

(iii) **The Term Structure:** At any fixed strike price, the implied volatility vary with time to expiration. There is a clear tendency for implied volatility to decrease, regardless of moneyness, when time to expiration increases. This implies that the long term return have a distribution that looks more normal than the short term returns.The decreasing curvature is easily observable in figure 2.

As the volatility surface clearly shows, the market does not believe that the Black-Scholes option price is correct.

Alexander (2008b) argues that the clear violation of the constant volatility assumption opens up for risk preferences in the option valuation problem without the creation of arbitrage opportunities. When Black and Scholes derived the option pricing formula in 1973 one of the key properties where that the equation did not involve any variable that is affected by the risk preference of the investors. If volatility is not constant the pricing problem becomes more complex and the risk preferences of the representative investor can explain some of the features of the volatility surface.

The strike structure of implied volatility cannot be explained by a time varying standard deviation per se, but a stochastic volatility factor can affect pricing through hedging demand. From figure 3 we can see that ITM calls and OTM puts implied volatilities are much higher than ITM puts and OTM calls. This is because a deep OTM put option is insurance agains large drops in the index. When an investor already holds the share, the increase in portfolio cost of including

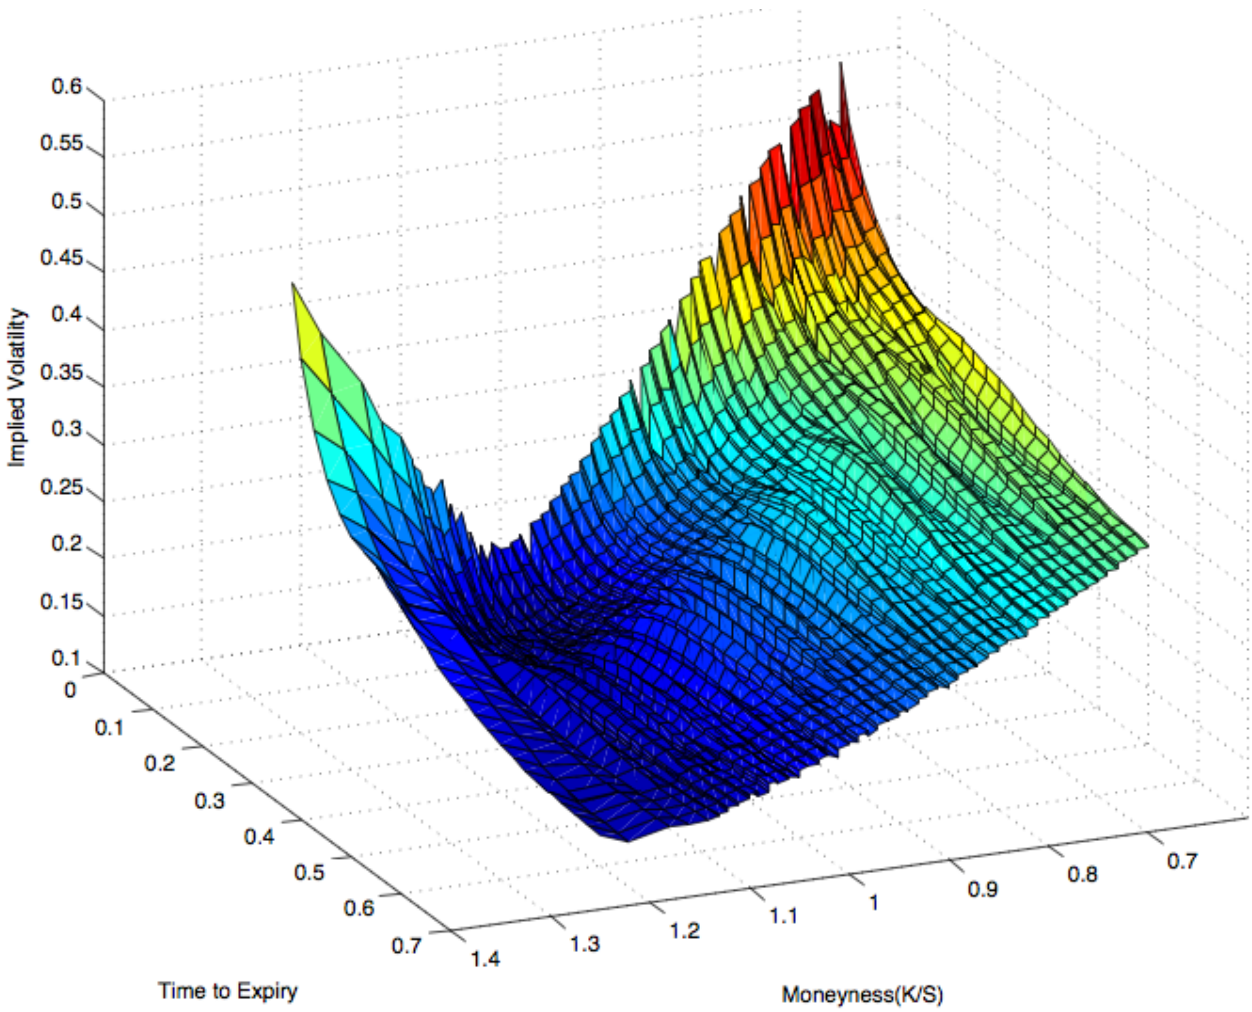

Figure 2: The implied volatility surface for S&P 500 put options.
The figure shows the level of implied volatility for different degrees of moneyness and different times to expiration (measured in years). The price of an index share S_t, is equal to \$1358. The option data has been smoothed by using a natural cubic spline interpolation method.

a deep OTM put is very low, making the investment a cheap hedge against big losses.[6] Deep ITM calls trade at higher implied volatilities could also be explained by the supply- demand situation for the option. An abundance of call buyers could drive the implied volatility up since the deep ITM option behaves similar to the stock and offers a higher return due to leverage (Alexander (2008b)).[7] This is closely related to the market crash of 1987, before this date the historical probability of extreme left tailed events was low and not considered in hedging strategies. The effect of the inclusion of this is discussed in an article by Rubinstein (1998) and is called *crashophobia*. This means that traders are concerned about the possibility of another crash similar to October 1987 and they price options accordingly.

Stringent capital requirements and risk averse investors can also cause option volatility to vary.

[6]The April put option with a \$1000 strike only costs \$2.40.

[7]A deep ITM has a *Delta* (the partial derivative of the option price with respect to the underlying security) around 1 implying a 1 for 1 change in the option price when the underlying price changes.

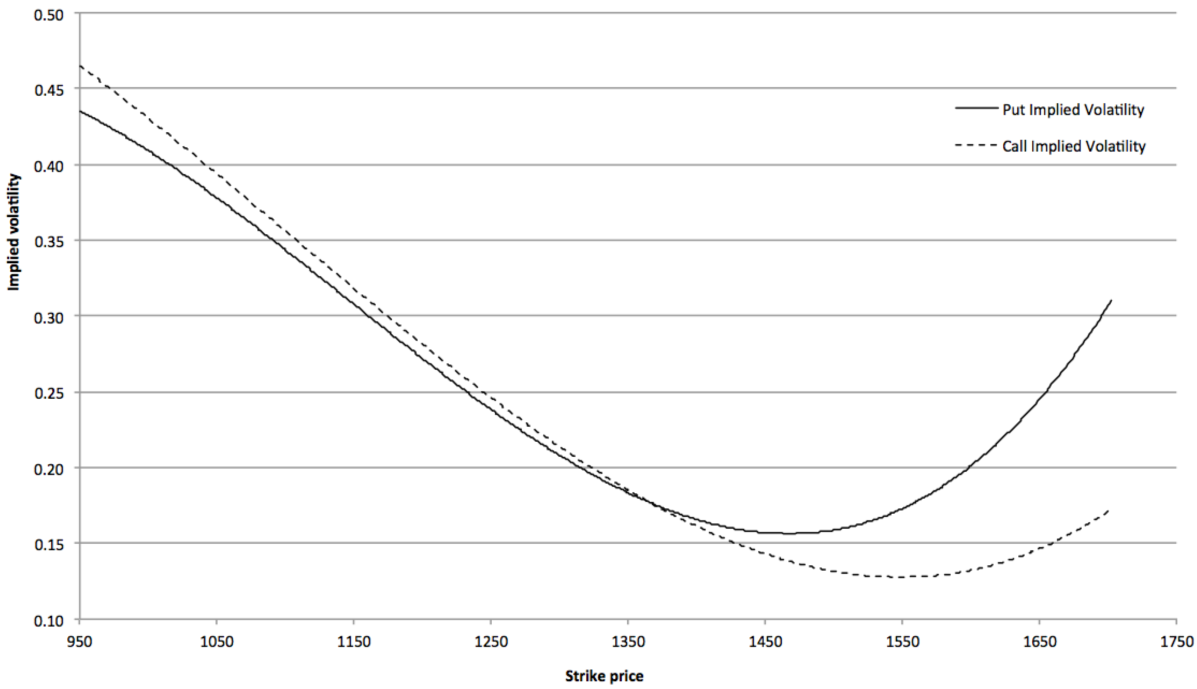

Figure 3: The implied volatility skew for April 2012 S&P 500 options.
Deep out-of- the-money puts trade at a higher volatility than deep out-of-the-money calls. The price of an index share S_t is equal to $1358. The same cubic interpolation technique described earlier has been applied to create a continuous smile.

Investors are prepared to pay a premium for the insurance value of the OTM put options, especially during periods of high uncertainty. This is reflected at higher option prices in the markets. On the other hand, market prices of OTM call options are sometimes lower than the Black-Scholes price because the market believes that the probability for the option to expire ITM is lower than the one predicted by the option pricing formula.

There are also other factors that contribute to the volatility curvature from a market maker perspective. Two important factors is the cost of hedging an option portfolio, and the credit premium implied by default risk. Black and Scholes (1973) assumed perfect and frictionless capital markets with no transaction or bankruptcy costs. With these two frictions apparent, an option trader will add a premium on the price to properly take the hedging costs and default risks into account. Both the hedging costs and the credit risk can be considerable, especially for long term options. Further, the hedging and credit factors are equal across all strikes. When factoring in the costs the effect will be relatively much greater on OTM calls and put compared to the more valuable ITM options. This could result in the smirky shape of figure 2.

The term structure in figure 2 is decreasing, this feature could stem from the fact that volatility tends to cluster. Volatility clustering has the effect of inducing a downward sloping term structure during relatively volatile periods, and an upward sloping structure in relatively calm periods. If volatility clustering is asymmetric, i.e. clusters of high volatility tends to be much stronger than its weak counterparts, the downward sloping term structures tends to be steeper than the upward sloping ones (Alexander (2008b)).

The Black-Scholes model also assume that the price follows a geometric Brownian motion (equation (2.10)) and that the risk neutral log price is normally distributed. But the premiums discussed above correspond to a price density with a heavier lower tail and a lighter upper tail compared to the normal density curve. This corresponds to a more negative view on stock returns since the probability of a value depreciation is higher than an appreciation. One interesting phenomenon first shown by Rubinstein (1994) is that the volatility smile was not observed in its current form in the markets before the "Black Monday" of October 19th 1987. On this day the S&P 500 index took the largest single drop of its history, plunging over 22.5%. Before this date the implied volatility was not shaped as a smile, but merely a constant when plotted over different strikes.

Rubinstein (1998) discuss additional reasons for the observed skew:

(i) **Leverage effect:** As stock prices fall, debt-equity ratio rises, which leads to a rise in volatility.

(ii) **Correlation effect:** Stocks become more highly correlated in down markets, which leads to volatility of the market index to rise since the benefits of diversification are reduced due to increasing correlation.

(iii) **Investor wealth effect:** As the market falls, investors feel poorer and become more risk averse so that any news lead to greater market reactions and more trading, which causes volatility to rise.

(iv) **Risk effect:** As volatility rises, risk premium increases, leading to declines in the market.

In summary the skew can be described by a technical supply-demand view. There is a strong demand for out- of -the- money puts created by portfolio hedgers, at- the- money puts does not share the same attractive hedging property. Further deep- in- the- money calls have a strong demand, thus a high implied volatility, because of the similar behavior to the underlying security. The same holds for in-the-money money puts.

Throughout this section we have seen evidence of an option pricing framework that do not sufficiently take into account risk-aversion and time varying volatility. In the following section we will try to expand the Black-Scholes framework to develop a more coherent option pricing model.

2.7 Generalized Autoregressive Conditional Heteroskedasticity (GARCH)

In this section we will introduce the concept of GARCH and its advantages when measuring volatility. The following draws on Brooks (2008) and Alexander (2008a). As mentioned in the introduction, GARCH models are designed to pick up the effects of clustering in the dataset. To understand volatility modeling consider a discrete-time econometric system given by:

$$\begin{aligned} y_t &= \phi_0 + \phi_1 x_t + u_t \qquad u_t \sim \mathcal{N}(0, \sigma_t^2) \\ \sigma_t^2 &= \omega_0 \end{aligned} \tag{2.29}$$

Where ϕ_0, ϕ_1 and ω_0 are constants. u_t is the residual and is the difference between y_t and the estimated relationship $\hat{y}_t = \hat{\phi}_0 + \hat{\phi}_1 x_t$. This system explains the linear relationship between y_t and x_t. ϕ_0 is the intercept term and ϕ_1 is the slope coefficient. The relationship may be estimated by the Ordinary Least Squares (OLS) method by minimizing the sum of squared residuals, $\sum u_t^2$. u_t is the noise of the estimation which is not captured by OLS and is defined as the residual at time t. We assume u_t to be normally distributed with mean zero and variance equal to σ_t^2. This system has constant variance and everything behaves just as in the Black and Scholes article. As we have shown, this system does not sufficiently model the observed market prices. To better model reality we expand the volatility process in system (2.29) above:

$$\begin{aligned} y_t &= \phi_0 + \phi_1 x_t + u_t \qquad u_t | \mathcal{F}_{t-1} \sim \mathcal{N}(0, \sigma_t^2) \\ \sigma_t^2 &= \omega_0 + \alpha_1 u_{t-1}^2 \end{aligned} \tag{2.30}$$

In (2.30) α_1 is the slope coefficient of the squared residual in time $t-1$ and $\mathcal{F}_t$ is the filtration of information at time t. This model is known as an Autoregressive Conditional Heteroskedasticity model or ARCH(1). Under the ARCH(1) model stochastic volatility is modeled by allowing the conditional variance of the error term, σ_t^2, to depend on one lagged value of the squared error term. ARCH(1) is the simplest model in the ARCH class and is usually expanded with more lagged values of the squared residual. Brooks (2008) argue that the use of ARCH models are rather limited since the number of lagged values might be very large, making the model computationally ineffective. Other issues are the determination of the optimal ARCH length, and the risk of violating the non-negativity constraints when including more lagged values.

A natural extension of the ARCH is the *Generalized* Autoregressive Conditional Heteroskedasticity (GARCH) and is based on Bollerslev (1986). GARCH is a generalization of the ARCH(1) model described in (2.30). The generalization consists of letting the conditional variance of the error term depend on lagged values of the conditional variance term itself. That is:

$$\begin{aligned} y_t &= \phi_0 + \phi_1 x_t + u_t \qquad u_t | \mathcal{F}_{t-1} \sim \mathcal{N}(0, \sigma_t^2) \\ \sigma_t^2 &= \omega_0 + \alpha_1 u_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \end{aligned} \tag{2.31}$$

New notation in (2.31) is β_1 which is the slope coefficient for σ_{t-1}^2. The symmetric normal GARCH(1,1) parameters above have a natural interpretation in terms of the reaction to market shocks and the mean reversion of the volatility. The GARCH(1,1) is a variance process, but the interpretation is the same for the square root of the variance, i.e. the volatility:

(i) α_1 measures the reaction of the conditional volatility to market shocks. When α_1 is relatively large (e.g. above 0.1) then volatility is very sensitive to market events.

(ii) β_1 measures the persistence in conditional volatility regardless of any market events. When β_1 is relatively large (e.g. above 0.9) then the volatility shocks diminish slowly after a high-volatility market event.

(iii) The sum $\alpha_1+\beta_1$ determines the rate of convergence of the conditional volatility to the long term average level. When $\alpha_1 + \beta_1$ is relatively large, then the term structure of volatility forecasts from the GARCH-model is relatively flat.

(iv) $\dfrac{\omega_0}{1-(\alpha_1+\beta_1)}$ determines the long term average volatility i.e. the *unconditional* volatility in the GARCH model.

To fully understand a GARCH model we must understand the difference between the unconditional variance and the conditional variance of a time series of returns. The *unconditional* variance is just the variance of the unconditional returns distribution, which is assumed to be constant over the entire data period. On the other hand, the *conditional* variance will change at every point in time, because it is conditional on the history of returns up to that point. Thus, we account for the dynamic properties of returns by regarding their distribution at any point in time as being conditional on all information up to that point. In section 2.2 we introduced the concept of continuous filtration. In (2.30) and (2.31) we write the conditional variance in the same manner.[8]

Taking the limit of the GARCH model lets the time interval between each lag approach 0. The limit allows us to represent the powerful GARCH-process as a SDE and utilize the explanatory power of the framework. The next section explains this concept.

2.8 The Weak GARCH Diffusion Limit

The first study on continuous GARCH modeling is the paper of Nelson (1990), where one of the most important approaches for GARCH option pricing was introduced. Nelson derived the continuous limit of GARCH using a theorem of weak convergence on the following system of equations:

$$\begin{aligned} ln\left(\frac{S_t}{S_{t-1}}\right) &= (r-\delta) + \sigma Z_t, \qquad Z_t|\mathcal{F}_{t-1} \sim \mathcal{N}(0,\sigma_t^2) \\ \sigma_t^2 &= \omega + \alpha Z_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned} \tag{2.32}$$

Where the expression for the stock movement is the discrete time version of the Brownian motion in (2.13). Z_{t-1}^2 has replaced U_{t-1}^2 and is the squared residual of the system. When

[8]For ease of notation we use the same notation for both discrete and continuous information filtering.

$\lim_{\Delta t \to 0}$ the following conditions are assumed to hold:

$$m = \lim_{\Delta t \to 0} \big(\frac{{}_\Delta\omega}{\Delta}\big); \qquad \alpha = \lim_{\Delta t \to 0} \big(\frac{{}_\Delta\alpha}{\sqrt{\Delta}}\big) \qquad \theta = \lim_{\Delta t \to 0} \Big(\frac{1 - \big({}_\Delta\alpha + {}_\Delta\beta\big)}{\Delta}\Big); \quad 0 < \omega, \alpha, \theta < \infty \tag{2.33}$$

Then by the convergence theorem for stochastic difference equations to stochastic differential equations by Nelson (1990) the symmetric *strong* GARCH diffusion can be written as:

$$\begin{aligned} \frac{dS(t)}{S(t)} =& (r - \delta)dt + \sqrt{V(t)}dW(t), \\ dV(t) =& \varphi(m - V(t))dt + \sqrt{2\alpha}V(t)dZ(t), \end{aligned} \tag{2.34}$$

Where $\langle dW(t), dZ(t) \rangle = 0$ implying zero correlation between the innovations, and

$$\varphi = 1 - \alpha - \beta, \qquad m = \frac{\omega}{\varphi}$$

In (2.34) $V(t)$ is the variance process governed by an arithmetic Brownian motion. New notation is φ and m, which are constants. The parameter interpretation is discussed below.

By comparing equation (2.34) to the simple Brownian motion in (2.13) the difference in the stock process is that the impact of randomness σ is no longer constant. It is now governed by a new process $V(t)$ which is the continuous GARCH(1,1).The continuous GARCH includes a new innovation with correlation zero to the original. The empirical fact that volatility comes in clusters means that volatility term structures converge to a long term average level. Thus, almost all stochastic volatility models assume that the drift in the diffusion process for volatility is specified in such a way that volatility can never drift too far away from its long term average. This process is called an Ornstein-Uhlenbeck process and it permits mean reversion. In (2.34) the drift parameters have the following interpretation:

(i) m is the long term average to which the variance converges.

(ii) φ is the continuous rate of mean reversion. φ governs the time taken for the drift term to revert volatility to its long term average, thus the greater the φ the greater speed of mean reversion. The parameter must be greater than 0 or the process would explode.

According to Alexander and Lazar (2005) several properties of the Nelson (1990) strong GARCH diffusion limit makes its applicability to equity options questionable:

(i) The diffusion is not time aggregating. This implies that resampling of the model at another frequency does not yield another GARCH process. Without time aggregation it makes less sense to reduce the time frequency to derive the diffusion limit.

(ii) The specific form in system (2.34) does only arise if one makes a specific assumption of the convergence behavior of the GARCH parameters. Corradi (2000) showed that if the limit of α converge with the convergence rate Δ instead of $\sqrt{\Delta}$ the result is not a stochastic volatility model, but a deterministic variance model of the form:

$$\begin{aligned}\frac{dS(t)}{S(t)} =& (r-\delta)dt + \sqrt{V}dW(t) \\ dV =& (\omega - \theta V)dt\end{aligned} \tag{2.35}$$

(iii) Any discretization of the continuous limit of the strong GARCH is not a GARCH model (Corradi (2000)).

(iv) The most apparent empirical problem with the model in (2.34) is due to the fact that the Brownian motions are independent, i.e. the price will have a zero correlation with the volatility, making the volatility surface symmetric.

In the case of the SPX options, figure 2 shows that a model that allows for a non-zero correlation should yield a better market price for the options. Further a model with a well defined continuous limit ensures that the limit of the GARCH is not deterministic as in equation (2.35). Alexander and Lazar (2005) proposed a solution to the problems of the strong symmetric GARCH model by deriving the diffusion limit of a *weak* GARCH(1,1). To understand the difference between the strong and the weak GARCH we return to the discrete-time discussion. In discrete time we used the following definition for the variance term in the GARCH framework:

$$\begin{aligned}E\big[u_{t+1}|\mathcal{F}_t\big] &= 0 \\ E\big[u_{t+1}^2|\mathcal{F}_t\big] &= \sigma_t^2\end{aligned}$$

This is the classical strong definition for the sequence $\{u_t, t \in \mathbb{Z}\}$ where the values of the residuals are assumed to be normally distributed. Drost and Nijman (1993) defines the symmetric weak GARCH as

$$\begin{aligned}P\big[u_t|u_{t-1}, u_{t-2}, \ldots\big] =& 0, \\ P\big[u_t^2|u_{t-1}, u_{t-2}, \ldots\big] =& \sigma_t^2,\end{aligned} \tag{2.36}$$

Where $P[X_t|u_{t-1}, u_{t-2}, \ldots\big]$ denotes the best linear predictor of X_t given the information available by filtration. Best linear predictors are linear functions of X that best predicts the values of Y on the form $f(Y) = \alpha + \beta X$ and does not necessary assume normality.

By using the weak definition, Nelson's convergence theorem and the limits described in (2.33) Alexander and Lazar (2005) derived the following weak GARCH diffusion:

$$\begin{aligned}\frac{dS(t)}{S(t)} &= (r-\delta)dt + \sqrt{V(t)}dZ_1(t) \\ dV(t) &= \varphi(m - V(t))dt + \sqrt{(\eta-1)}\alpha V(t)dZ_3(t)\end{aligned} \tag{2.37}$$

where: $\varphi = 1 - \alpha - \beta, \qquad m = \frac{\omega}{\varphi}, \qquad$ and $\varrho = \frac{\tau}{\sqrt{(\eta - 1)}}$

New notation includes a new innovation $dZ_3(t)$, a kurtosis parameter η, a skewness parameter τ and a correlation coefficient ϱ which is discussed below.

By relaxing the assumptions on the GARCH measure we are now able to include skewness and kurtosis in the distribution by allowing a correlation between the price and the volatility process. Skewness is a measure of asymmetry around the mean, and the kurtosis represents the height of the mean, i.e. the number of observations located around the mean value. The correlated innovations are easily computed by applying the Cholesky decomposition on the correlation matrix(Wiersema (2008)). Since we have only two correlated innovations the Cholesky-matrix is the same for all stages in the process. The decomposition yields the following:

$$\begin{bmatrix} 1 & 0 \\ \varrho & \sqrt{(1-\varrho^2)} \end{bmatrix} \begin{bmatrix} 1 & \varrho \\ 0 & \sqrt{(1-\varrho^2)} \end{bmatrix} = \begin{bmatrix} 1 & \varrho \\ \varrho & 1 \end{bmatrix}$$

then the Cholesky-matrix is just:

$$\begin{bmatrix} 1 & 0 \\ \varrho & \sqrt{(1-\varrho^2)} \end{bmatrix}$$

which means that we can obtain correlated simulations by using the sum of two independent Brownian motions as a new Brownian motion with ϱ as the conditional correlation coefficient. That is:

$$Z_3(t) \equiv \varrho Z_1(t) + \sqrt{(1-\varrho^2)} Z_2(t) \qquad (2.38)$$

where $\langle dZ_1(t), dZ_2(t) \rangle = 0$. Here equation (2.38) is the diffusion part of the variance process in (2.37). In the model, the skewness and kurtosis decides the model *volatility* of volatility, i.e. the degree of uncertainty in the volatility measure.

The symmetric strong GARCH is nested in the weak GARCH and appears as a special case when log returns are assumed to be normally distributed. The stochastic volatility introduces positive excess kurtosis into the log returns distributions, and if the log returns are also skewed the price-volatility,ϱ, will be non-zero. The limit of equation (2.37) has none of the problems of the strong GARCH diffusion outlined above and should therefore be a better tool of modeling volatility and financial returns. The main disadvantage of the weak GARCH diffusion is the lack of an easy calibration method, since there has been no assumption on the distribution of the variables. Thus, no closed form solution is yet available. However, we estimate the weak GARCH using the strong GARCH parameters. The use of strong GARCH parameters in the weak model will cause a small bias since we assumed no excess kurtosis and no skewness in the distribution (Drost and Nijman (1993)). Nevertheless, the weak GARCH opens up for a non-normal distribution and according to Lewis (2000) and Alexander (2008b) use of the strong GARCH parameters for weak estimation yields negligible errors.

The parameter α still determines the reaction to market shocks (i.e. the volatility of volatility) and β determines volatility persistence. $\alpha+\beta$ determines the rate of convergence and $\frac{\omega}{1-\alpha-\beta}$ determines the long run GARCH volatility.

The Weak GARCH diffusion limit should be very capable to capture the shape of the volatility surface. We now continue by looking on the properties of the historical returns of the S&P500 index, and the observed market prices of the SPX options.

3 Data and Methodology

The following section is divided into two parts. The first part explains the properties of the historical S&P data and the option market prices that will be used to compare the Black-Scholes model and the GARCH diffusion model. The last part will go into details about the calibration and implementation of historical data to the different models.

3.1 The Standard & Poor 500 Index and Market Parameters

The S&P 500 (SPX) index has been used for the comparison of the two models. The SPX is a capitalization weighted index of the prices of the 500 largest companies with public stock actively traded in the United States. The index covers approximately 75% of US equities (Standard&Poor (2012)).

According to Rubinstein (1994) the SPX index best replicates the conditions necessary for Black-Scholes option pricing. The index has actively traded options for a wide range of strikes and the options are quoted with a relatively small bid-ask interval which implies high liquidity. Since the index is an adjusted total return index, where the dividends paid are used to repurchase shares, we have used the dataset containing only the non-adjusted closing price. The dividends are included explicitly in the model. The sample period consists of several interesting periods such as the global equity crash of 1987, the outbreak of two Gulf wars (1990−91 and 2003), the Asian currency crisis (1997), the LTCM (1998), the dot-com bubble during the late 90's , the September 11th terrorist attacks (2001) and the recent credit and banking crisis. The dataset was downloaded from Reuters EcoWin and the range is from January 2nd 1980 to March 8th 2012. The returns are calculated based on the following formula:

$$r_t = ln\Big(\frac{SPX_t}{SPX_{t-1}}\Big) \tag{3.1}$$

In (3.1) we obtain the logarithmic return by taking the natural logarithm of the closing price of the Standard & Poor index at time t, divided by the closing price at time $t-1$. The properties of the time series are summarized in table 1.

When excluding "The Black Monday" the SPX returns still exhibit a negative skew and excess kurtosis, but not to the same extent as in the case when it is included. The properties of the complete dataset, including the Black Monday, will be used in the option analysis.[9] The kurtosis in a normal distribution is 3, implying that the correct kurtosis for the distribution is the reported value plus three (Brooks (2008)). Following Hull (2003) we use the US 12-month Treasury Bill interest as a proxy for the risk free interest rates. For the dividend yield we have used the

[9]The descriptive statistics on a sample from January 1st 1980 to March 8th 1986 shows positive skewness and a excess kurtosis of only 1.38.

Table 1: Descriptive statistics for the S&P 500 log returns.
The returns α_t in column 1 describes the complete dataset and the returns α_t^* in column 2 excludes the extreme left tail event of October 19th 1987. The results are based on 8120 observations from January 3rd 1980 to March 8th 2012.

	α_t	α_t^*
Mean (a)	0.0315%	0.034%
Annualized mean	11.5%	12.526%
Std.Devn (s)	1.157%	1.13%
Annualized Std. Devn.	22.104%	21.588%
Skewness (τ)	-1.1696	-0.2535
Kurtosis (η)	29.217	11.5122
Minimum	-22.83%	-9.469%
Maximum	10.95%	10.95%

last reported annualized dividend payouts on the SPX index divided by the closing price of the index at the relevant date. Both rates are transformed to continuous time using the appropriate formula from Alexander (2008b). For the conversion of discrete observations we use:

$$\begin{aligned} r_t^f &= \ln\left(1 + R_t^f\right) \\ \delta_t &= \ln\left(1 + D_t\right) \end{aligned} \tag{3.2}$$

The discrete time interest rate is denoted R_t^f and the continuous is denoted r_t^f. The dividends follow the same notation with D_t for discrete data and δ_t for the continuous. The data was downloaded from EcoWin Reuters. The summary of the used market parameters can be found in table 2.

As a measure of volatility we have used the implied volatility of the options traded on the SPX index. This is called the VIX index and the spot value can be interpreted as the market expectations for the annualized 30 day volatility (Chicago Board of Exchange (CBOE) (2012)).

Table 2: Estimated and observed market parameters.
S_t is the observed market price of the underlying at the time when the option prices were recorded. D_t and R_t^f denotes the discrete dividend yield and the risk free interest rate. δ_t and r_t^f is the continuous transformation. $\sigma_{\alpha,t}$ is the historical volatility and σ_{VIX} is the VIX implied volatility. The different dates in the sample are February 16th, March 8th and March 27th 2012.

Date	*Market Parameters*						
	S_t	D_t	δ_t	R_t^f	r_t^f	$\sigma_{\alpha,t}$	σ_{VIX}
February 16th	\$1358	2.15%	2.1272%	0.1584%	0.1583%	22.1040%	19.22%
March 8th	\$1346	2.10%	2.0783%	0.1558%	0.1557%	22.1040%	17.98%
March 27th	\$1413	2.06%	2.0390%	0.1669%	0.1668%	22.1040%	17.29%

In figure 4 the sample distribution is plotted against the Gaussian. Clearly, a model that assumes a normal distribution of the asset returns will fail to provide unbiased results. This is because the empirical distribution is *leptokurtic*, meaning that the returns tend to exhibit fat tails and excess peakedness at the mean (Brooks (2008)). Therefore, to ensure parameter generation consistent with empirical findings, it is important to allow the distribution to have nonzero skewness, τ and positive kurtosis, η.

3.2 The SPX Option Chains

To evaluate the Stochastic volatility models performance we will compare it to historical option quotes from Marketwatch (2012). The sample consists of 2608 ask quotes collected in the period February 16th to March 27th. We have chosen to use option ask quotes because the ask is the price where an investor immediately can acquire the option, and it does not suffer from stale quotes like the last traded quoted price which may be several days old. The open interest i.e. the number of contracts ranges 0 to 93400 in the sample. Zero open interest does not imply that the option can not be bought or sold, it implies no secondary market for the option. Whether open interest and liquidity affects the option price is not discussed. All option models seeks to find the market price for the option, not the bid or ask price. It would be optimal to compare the model price with a liquid market price, but due to lack of data we have chosen the ask price.

The full sample originally consisted of 4002 option prices, but we have chosen to exclude contracts with zero open interest that was extremely deep out-of-the money and extremely deep-in-the money. This was due to a very low probability of a transaction occurring (thus a low probability for a trade around the ask price) and large bid-ask intervals. Regardless of the filtration, some contracts with zero open interest are still in the sample. These options have in

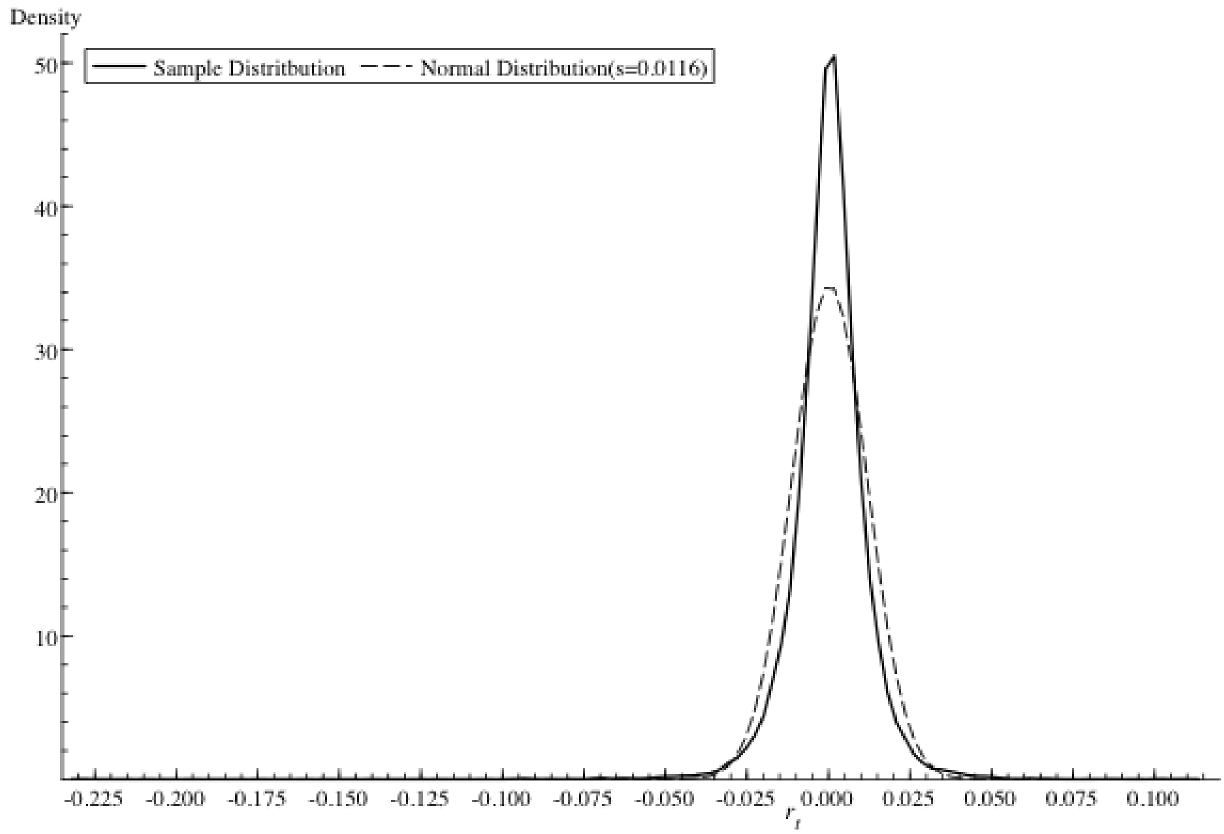

Figure 4: Actual return distribution for the S&P 500 index.
The density plot is leptokurtic with a negative skew and is the solid curve in the figure. The Normal distribution is the dotted curve and has zero skew and no excess kurtosis. The moments of the actual return distribution are described in table 1 above.

common that there are contracts with similar strikes and time to expiry traded with significant open interest. For example, the April put with $K = \$1335$ from February 16th have no active market. The surrounding puts, with strike $K = \$1330$ and $K = \$1340$ have an open interest of 4019 and 3789 respectively. With contracts actively trading nearby, one could also expect trading in the option currently not traded. The bid-ask spreads are also smaller compared to the deep OTM and ITM cases. We have therefore chosen not to exclude options of this type from the dataset and allow the ask price to serve as a proxy for the market price.

Following Misra et al. (2006) contracts with moneyness, $\frac{K}{S}$ for puts and $\frac{S}{K}$ for calls, greater than 1.05 are defined as in-the-money. Options with moneyness of 0.95 or less are defined as out-of-the money. Further deep-in-the money and deep-out-of-the money are options with moneyness over 1.15 and under 0.85. The expiry date of the option ranges from 10 to 485 days, where 365 days is defined as one year. There are limited amounts of deep-in the money and deep-out-of the money contracts listed at the exchange, therefore we only have a few prices available for comparison. The properties of the SPX options are summarized in table 3.

The evaluation criterias we have chosen is the *Root Mean Squared Error (RMSE)* and *Mean*

Table 3: Summary of the different S&P 500 options.
Time to expiry is measured in days and varies from 10 to 485. Total number of calls and puts is 1307 and 1301 respectively. The option prices was downloaded from Marketwatch (2012)at February 16th, March 8th and March 27th 2012.

Moneyness	*Time to expiry*			
	$T \leq 50$	$51 \leq T \leq 199$	$T \geq 200$	*Total*
i)	*Puts*			
≤ 0.85	166	178	56	400
$0.85 - 0.95$	110	130	40	280
$0.95 - 1.05$	109	144	41	294
$1.05 - 1.15$	95	110	35	240
≥ 1.15	20	39	28	87
Total	500	601	200	1301
ii)	*Calls*			
≤ 0.85	10	24	20	54
$0.85 - 0.95$	102	122	42	266
$0.95 - 1.05$	109	144	42	295
$1.05 - 1.15$	90	109	32	231
≥ 1.15	193	204	64	461
Total	504	603	200	1307
iii)	*Total*			
≤ 0.85	176	202	76	454
$0.85 - 0.95$	212	252	82	546
$0.95 - 1.05$	218	288	83	589
$1.05 - 1.15$	185	219	67	471
≥ 1.15	213	243	92	548
Total	1004	1204	400	2608

Absolute Percentage Error (MAPE) from Brooks (2008). The root mean squared error criteria measures the distance between the observed market price and the estimated model-price. The goal of the estimation should be to minimize this error. The root mean squared error is obtained by the following formula:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(P_{i,m} - P_i\right)^2} \tag{3.3}$$

Where n is the number of observations, P_i is the model price at time i and $P_{m,i}$ is the market price at time i. MAPE yields the relative difference between the market price and the predicted price. MAPE has the attractive property compared to the RMSE that it can be interpreted as a percentage error, implying that a large absolute difference is weighted similar to a smaller absolute difference. The mean absolute percentage error is given by:

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\Big|\frac{P_{i,m} - P_i}{P_{i,m}}\Big| \tag{3.4}$$

We have now established a framework to price options with stochastic volatility and developed an algorithm to find the implied volatility and the volatility surface of the options. The remaining part of this section consists of an explanation of the GARCH diffusion algorithm and implementation of market data.

3.3 Implementation of the Weak GARCH Diffusion Model

Following Alexander (2008b) we have chosen an intuitive method for the model calibration. Since the weak GARCH does not assume any distribution the estimated parameters from the strong GARCH have been used in the model. The parameters for the strong GARCH have been estimated using Maximum Likelihood which is a maximization problem assuming a normal distribution such that the parameter values chosen are most likely to produce the observed data. This is done by forming a likelihood function $L(\Theta)$, which will be a multiplicative function of the actual data. Using the properties of the normal distribution where each draw, $y_1, y_2, \ldots, y_T$, is independent from the previous we can write:

$$f(y_1, y_2, \ldots, y_T) = f(y_1)f(y_2)\ldots f(y_T) = \prod_{t=1}^{T} f(y_t) \tag{3.5}$$

we can write the likelihood function as:

$$\begin{aligned} L(\Theta) =& f(y_1, y_2, \ldots, y_T) = \prod_{t=1}^{T}\left(\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{1}{2}(\frac{y_t-\mu}{\sigma_t})^2}\right) \\ =& \frac{1}{(\sigma^2 2\pi)^{\frac{T}{2}}}e^{-\frac{1}{2}\sum_{t=1}^{T}(\frac{y_t-\mu}{\sigma_t})^2} \end{aligned} \tag{3.6}$$

Brooks (2008) suggests a monotonic transformation of the likelihood function using the natural logarithm to make the problem additive. The maximization then reduces to:

$$\max_{\omega,\alpha,\beta} \quad Ln\big(L(\Theta)\big) = -\frac{1}{2}\sum_{t=1}^{T}\Big(ln(\sigma_t^2) + \big(\frac{y_t-\mu}{\sigma_t}\big)^2\Big) \tag{3.7}$$
$$\text{where } \sigma_t = \sqrt{\omega + \alpha u_{t-1}^2 + \beta\sigma_{t-1}^2}$$

Further, we have used the daily price return observations from the Standard & Poor 500 index and solved the optimization problem using OxMetrics. The optimization is done numerically using a process similar to the NRB method described in the theory section. The results are summarized in table 4. In the theory section we defined a relatively high value of α as above 0.1, our dataset has an α of 0.0779. This implies that the volatility is relatively sensitive to market events implying that a market shock has a significant effect on the volatility. The sample β is 0.9125, this is defined as a relatively large degree of persistence meaning that the effect of a volatility shock slowly diminishes from the model. In other words, when market shocks occur, the effect is visible in the model for a multitude of periods. The sum of α and β is above 0.99 yielding a relatively flat term structure of the volatility and the long term daily average volatility is 1.19%.

We are now able to simulate a stock price process that follows the Weak GARCH diffusion process. The payoffs of plain vanilla options were discussed in the theory section, are nothing but the maximum of 0 and , $S_T - K$ for calls and $K - S_T$ for puts. Since a closed form solution for the weak GARCH model does not exist a Monte Carlo simulation will be used. A Monte Carlo simulation values the options by taking the discounted average of n option payoffs. The average of several thousand possible payoffs should yield a fair estimate for the price of the claim. Since we use the risk-neutral process for the stock the correct discount factor should be the risk free interest rate (Mcdonald (2006)). Below is an explanation of the algorithm used to price the options, the complete Matlab code can be found in appendix B.4.

The algorithm is a double loop creating a leptokurtic price path of specified length for the stock with a model-determined price variance correlation (ϱ). The process is then repeated several thousand times for the Monte Carlo Valuation. The system to be simulated is equation (2.37) from the theory section and is given by:

$$\begin{aligned} \frac{dS(t)}{S(t)} &= (r-\delta)dt + \sqrt{V(t)}dZ_1(t) \\ dV(t) &= \varphi(m - V(t))dt + \sqrt{(\eta-1)}\alpha V(t)dZ_3(t) \end{aligned} \tag{3.8}$$

where: $\varphi = 1 - \alpha - \beta, \qquad m = \frac{\omega}{\varphi}, \qquad \text{and } \varrho = \frac{\tau}{\sqrt{(\eta-1)}}$

The stock price process has an explicit solution given from equation (2.9). The variance process is modeled in the form described above yielding the following system of equations to be

Table 4: Estimated daily GARCH and model parameters.
The estimation sample is from January 1st 1980 to March 8th 21012. The results have small standard errors yielding statistically significant parameter estimates both at 5% and 1% level of significance.α is the reaction of conditional variance to market shocks, β is the persistence of volatility shocks. m is the long term variance, φ is the speed of adjustment towards the long run mean. ϱ is the correlation between the variance and the price of the stock. The total number of observations is 8120.

Parameters	*Estimated Values*			
	Coefficient	Std.Error	t-value	t-prob
ω	1.35677e-06	1.895e-07	4.11	0.000
α	0.0778921	0.005725	3.93	0.000
β	0.912546	0.005895	51.6	0.000
φ	0.0096	-	-	-
m	1.4189e-04	-	-	-
ϱ	-0.232911	-	-	-

evaluated numerically:

$$\begin{aligned} S(t+1) =& S(t)e^{\left(r_f - \delta - \frac{1}{2}V(t) + \sqrt{V(t)}Z_1(t)\right)} \\ dV(t+1) =& \varphi(m - V(t))dt + \sqrt{(\eta-1)}\alpha V(t)\left(\varrho dZ_1(t) + \sqrt{(1-\varrho^2)}dZ_2(t)\right) \end{aligned} \tag{3.9}$$

The model is calibrated to market data using the parameter values from table 2. The interest rate r_t^f and the dividend yield δ_t is divided by 365 to yield daily values. Since all parameters now are daily data the discretization of the wiener process is done using the following random draw from the standard normal distribution.

$$dZ_i \sim \sqrt{dt}\mathcal{N}\left(0,1\right) \tag{3.10}$$

In this case, because of limited computational power, we have allowed a single time step to equal one day. This is a rough assumption but the effects are dampened due to the fact that we repeat the simulations several thousand times to find the option payoffs. This causes one innovation to the stock price and variance process to count for a very small fraction of the option price. The steps of the algorithm are explained in more detail below:

The Weak GARCH Diffusion Option Pricing Model Algorithm

(i) for $t = 1, j = 1$. The algorithm generates two innovations, Z_1 and Z_2 with correlation -23.29%. The innovation generates a movement in the variance in (3.9). The variance is then used in the expression for the stock process. This generates a new stock price with stochastic volatility. *next* t.

(ii) The process repeats until $t = T$ and we have simulated a price path with a length of T days. The payoff of an option with T days to expiry and a strike price of K is recorded. *next* j.

(iii) The process repeats until $j = 100000$ and we have simulated *j* different price paths used to value the option.

(iv) The algorithm now takes the discounted average value of the *payoffs* for each option at each stock path at the maturity date. This process is repeated for strike prices ranging from \$950 to \$1700. The values recorded are the weak GARCH diffusion model option prices.

To illustrate the pricing algorithm we have carried out one simulation which is showed graphically in figure 5. The stock and volatility process is simulated over a period of 200 days. The starting price of the stock is \$100 and the starting volatility is 1.21% which is the volatility equivalent of an annualized VIX of 23.11%. The sample annualized volatility of volatility is 11.25% and the sample correlation is -27.26%. The figure clearly shows that periods with large fluctuations in volatility yields large price changes in the stock price. Further the sample volatility also yields a clear mean reverting tendency towards the long term equilibrium represented by the turquoise line. The diffusion model is also able to capture jump-like changes in the stock price without including jump parameters. This is an effect of the very flexible weak GARCH diffusion and could enable us to capture the empirical properties of the SPX index.

This completes the data and methodology chapter. The next section carries out the analysis of the distribution and the evaluation of the model option prices.

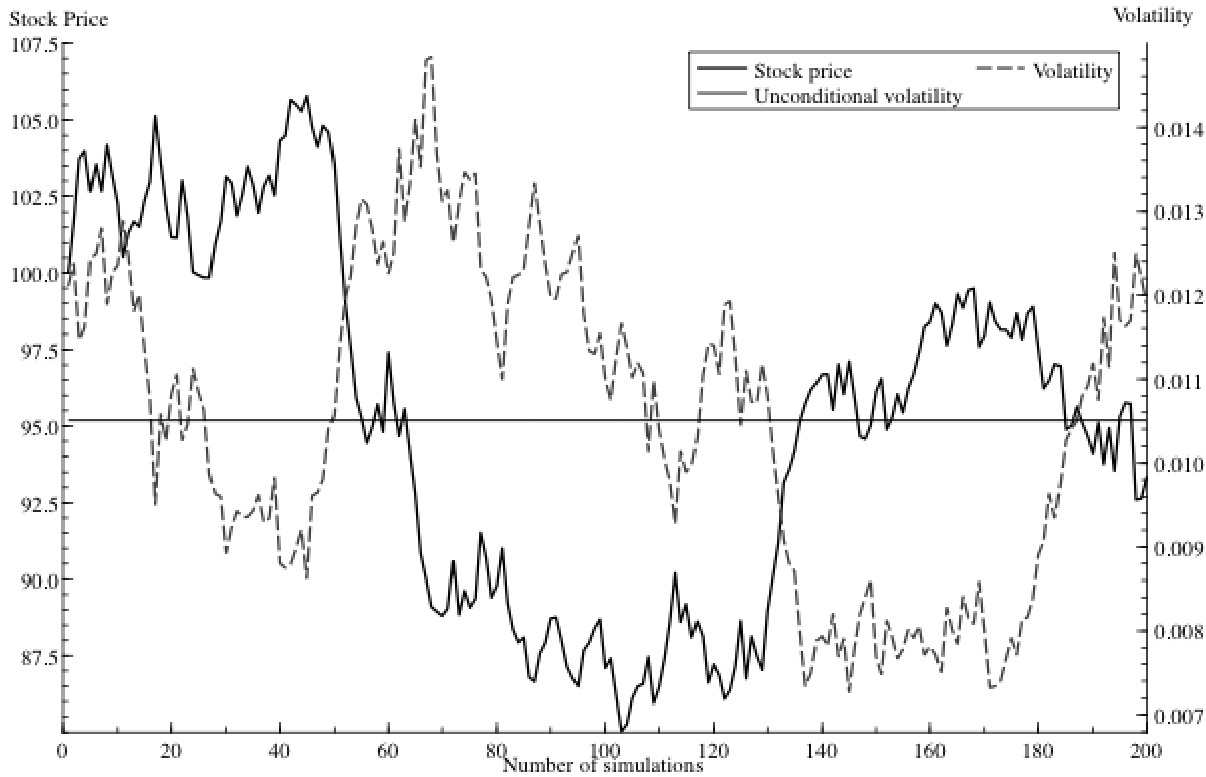

Figure 5: Stochastic volatility and price simulation.
The stock price is measured on the left axis with $S(0) = 100$ and the volatility, $\sqrt{V(t)}$, is measured on the right. Initial daily volatility is 1.21%. The number of simulations is 200.

4 Analysis

In this section we will analyze the distribution generated by the system of differential equations. Further, a discussion of the models ability to generate an empirically consistent volatility surface will follow before we turn to the comparison of the option prices.

4.1 Simulated Return Distribution

After calibrating the model to the market parameters, a test of the generated parameters was performed. We modeled the stock process of 8120 days, which is the length of the SPX return sample and recorded the returns of each price vector. This process was repeated 1000 times to evaluate the distribution descriptive statistics. The results are summarized in table 5. The mean of a lognormal variable is the drift term which is deterministic and thus not included in the table.

In table 5 we see that, on average, the model generates a leptokurtic return distribution. But the table also describes one of the main weakness of a weak GARCH diffusion model. Since we try to replicate the distribution of the SPX index we allow the model to generate jump like price paths, without including a jump parameter. This makes the parameters very flexible as reflected in the standard deviation on the distribution parameters. The returns generated are slightly skewed to the left, but not as much as the empirical distribution. Further the model generates a distribution with a too high degree of peakedness with a very large standard deviation. According to Rubinstein (1994) the Standard & Poor 500 index is preferred for analyzing option price behavior when it comes to structural parameters such as liquidity and market completeness. But according to Kaeck and Alexander (2011) the SPX index yields higher moments (i.e skewness and kurtosis) not easily generated by a stochastic volatility model because the shape of the distribution with large excess kurtosis and negative skewness is not easily captured.

Another problem with the model when calibrated with the method described above is the correlation between the price and variance, ϱ. The model calibration implies a correlation of $\varrho = \frac{\tau}{\sqrt{\eta - 1}} = -0.2202$. This correlation is not consistent with empirical studies and this could create a pricing error if the model correlation and the empirical correlation is significantly different. Most empirical studies agree that the correlation between variance and price is negative (see Andersen, Benzoni and Lund (2002);Bakshi et al. (1997); Rubinstein (1998)). All articles point out that the correlation coefficient is of great importance when dealing with asymmetric smiles. A correlation of zero yields symmetric smiles, while a negative correlation yields the smirk observed in the data.

To find the daily price variance correlation we need intraday data to capture the variations within each trading day. Since we do not have access to intraday data we are not able to generate the

Table 5: Descriptive statistics for simulated returns with model price-variance correlation. The results are based on a fixed number of days equal to 8120. The loop has been repeated 1000 times. The parameters are the mean values of the simulation moments with a total of 8.12 million simulations.

	Parameter Values	Standard Deviation
Daily Minimum	-15.55%	0.1379
Daily Maximum	15.32%	0.1135
Annualized standard deviation(s)	18.15%	0.0726
Skewness (τ)	-0.2549	1.5663
Kurtosis (η)	48.3137	67.4832
Price-Variance Corr(ϱ)	-0.2205	-

daily variance for the index. Instead we have used another common method which is to measure the correlation between the squared CBOE implied volatility (VIX) and the SPX index. The squared VIX represents the market 30 day annualized implied *variance* for the options on the Standard & Poor index. We have used daily time series from January 2nd 1990 to March 6th 2102 to calculate the sample correlation. The empirical sample correlation between the change in variance and the log returns was found to be -71.99%.

To correct for the correlation error, we have modified the model to allow the price variance correlation to be exogenous instead of the kurtosis.[10] Using the empirical correlation and skewness we find the kurtosis to be:

$$\eta = \left(\frac{\varrho}{\tau}\right)^2 + 1 = 3.6389 \tag{4.1}$$

Equation (4.1) is a simple manipulation of the model correlation in equation (2.37). The reparametrization of the model implies the descriptive statistics reported in table 6.

An important lesson from the re-parametrization is the sensitivity of the model with respect to the kurtosis-parameter. When increasing η the impact of the innovations in the model also increase. This affects the stability of the parameters resulting in larger standard deviations and less stability.[11] Due to the relatively large moments of the SPX returns stochastic volatility models often fail to take these effects into account. As pointed out in Kaeck and Alexander (2011) it is easier to calibrate the model to value options in other markets where the skewness and kurtosis is relatively lower. This is why all parameters increase significantly in stability when we use the empirical correlation instead of the implied correlation. Another effect is that the distribution turned more normal when looking at the third and fourth moments. In addition

[10] A special thanks to Prof. Carol Alexander for valuable input on how to strengthen the model.

[11] Increasing the kurtosis parameter to 90 resulted in output kurtosis above 200. Moment standard deviations also increased rapidly.

Table 6: Descriptive statistics for simulated returns with empirical price-variance correlation. The results are based on a fixed number of days equal to 8120. The loop has been repeated 1000 times. The parameters are the mean values of the simulation moments with a total of 8.120 million simulations.

	Parameter Values	Standard Deviation
Daily Minimum	-6.97%	0.0199.
Daily Maximum	7.19%	0.0210
Annualized standard deviation(s)	17.33%	0.0150
Skewness (τ)	-0.0017	0.1024
Kurtosis (η)	5.6731	2.0532
Variance Return Corr(ϱ)	-0.7202	-

to the kurtosis, the skewness is also significantly lower compared to the method above. The least favorable consequence of the weak GARCH is the lack of ability to sufficiently calibrate the model to generate returns with correct higher moments and an empirically correct correlation between returns and volatility. Using the correct correlation the model should generate returns with kurtosis equal to 3.6389. The model overshoots this measure and generates kurtosis closer to the observed values, on the other hand the model is not able to capture the skewness of the observed data. Nevertheless, the existing literature emphasize that it is the price volatility correlation that creates the asymmetric smile effect and therefore we focus the analysis on option pricing with the empirical correlation.[12] Using the real correlation we will now turn to the models ability to produce implied volatilities consistent with the observed surface in figure 2 in the theory section.

4.2 The Generated Volatility Surface

In this section we present the volatility surface generated by the model and discuss the differences between the observed surface from section 2.6.

As observed in figure 2 in the theory section the volatility surface for the SPX options takes an asymmetric smirky shape where the local minimum is observed where option moneyness is close to 1. Further, the observed implied volatility increases for both decreasing and increasing moneyness, with the largest observed values for deep out-of-the money options.

In figure 6 the data generated volatility surface is shown. We have generated a 33X76 matrix consisting of a total of 2508 option prices for 76 different degrees of moneyness ranging from 0.07 to 1.25. Time to expiry ranges from 0.082 to 0.7 years (30 to 256 days). The implied

[12]We implicitly say the same, since the kurtosis and skewness generates the price volatility correlation.

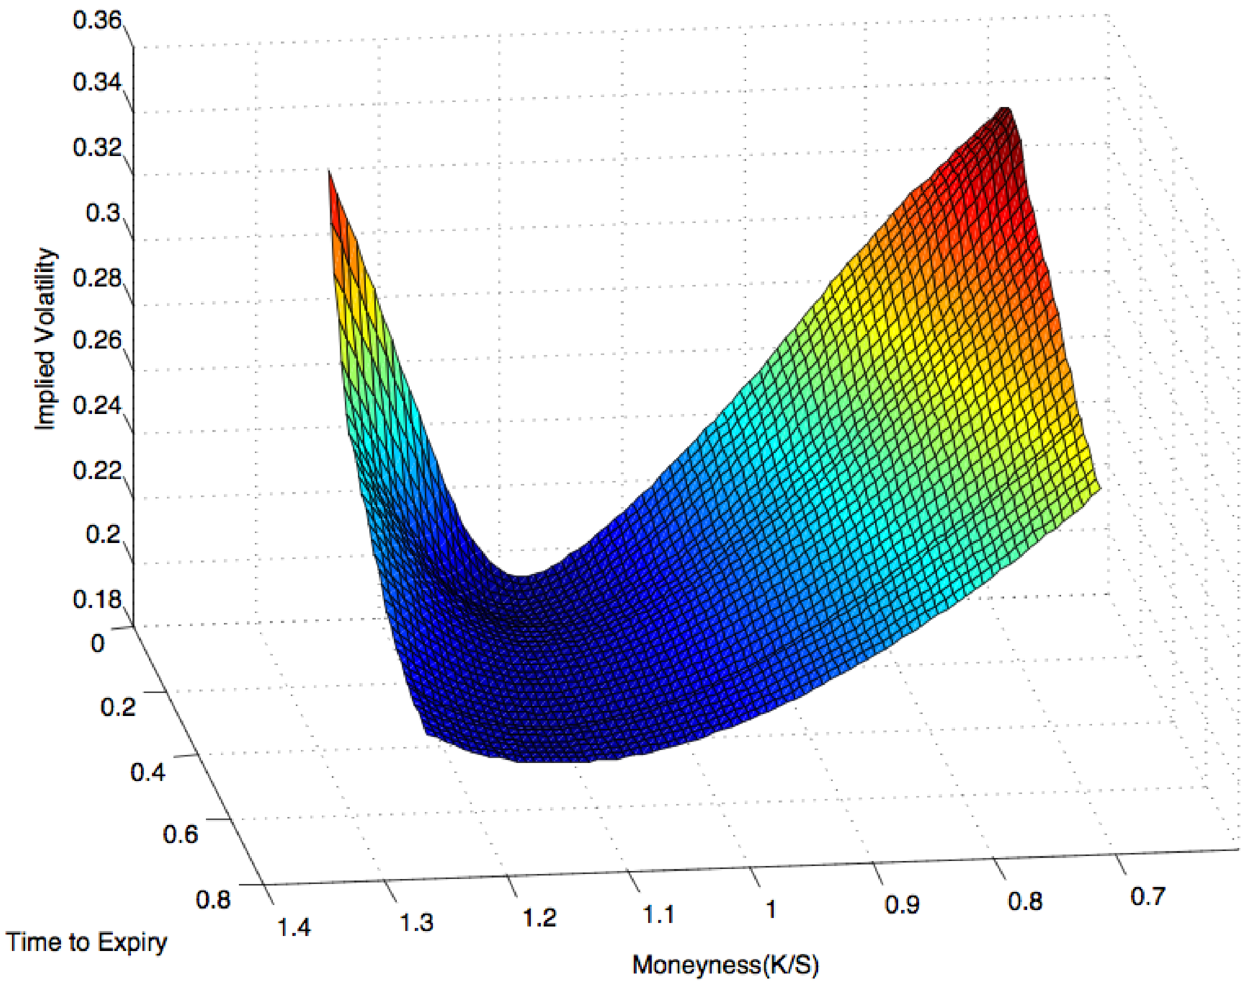

Figure 6: Simulated implied volatility surface for S&P 500 put options.
The figure shows the level of implied volatility for different degrees of moneyness and different times to expiration (measured in years). The option data has been smoothed by using a natural cubic spline interpolation method described in the theory section.

volatilities was extracted using the Newton- Rhapson Bisection method discussed in the theory section. With the correct empirical price volatility correlation we are able to generate a surface with a very similar same shape as the empirical counterpart. The empirical local minimum volatility is at 14.75% around moneyness 1, whereas the sample maximum is at 57.45% deep out of the money. The generated surface suffers from the fact that the underlying return distribution is less skewed and has significantly lower kurtosis. This implies that the probabilities for extreme left tail events are lower than what we observe in the markets, thus creating a more symmetric flatter surface. The local minimum for the generated surface is at 19.91% and the maximum is at 34.54%. Although there is a difference between the extreme values of the surface we have generated a surface with a very similar shape as the empirical, which again could lead to a better volatility framework for option pricing. The generated surface is, regardless of the extreme value deviations, a much better proxy for the empirical surface than the plane suggested by the Black and Scholes framework.

We have now seen the data generating abilities of the model and studied the implied volatility surface generated in closer detail. The main part of the analysis follows in the next section where we investigate the prices generated by the stochastic volatility model presented and compare them to the Black and Scholes price.

4.3 Black and Scholes vs. The Weak GARCH Diffusion Model

In this section we present the pricing errors from the three models discussed. First we present the pricing error for the put prices using RMSE and MAPE from two calibrations of the Black Scholes formula, one using the historical standard deviation, and the other using the closing VIX as the standard deviation.[13] The Black and Scholes prices are then compared directly to the estimated weak GARCH price. Further we repeat the analysis for call prices before we look at the total pricing error induced by the models. The results are given in table 7, 8 and 9.

For the put pricing the weak GARCH model outperforms the other two models when it comes to total put pricing performance. The weak GARCH has a total put pricing MAPE of 33.62% and a RMSE of 5.26. This is clearly lower than the BS prices. Further, the weak GARCH is clearly superior when pricing deep-out-of-the and out-of-the money puts where the only model coming even close is the BS with historical standard deviation. The reason for this is that the weak GARCH model is able to price the options with higher implied volatility than the two other models. The model with the historical standard deviation is closer than the BS VIX because the historical standard deviation is higher than the one implied by VIX. This has the largest impact for options with a long time to expiry and is visible in the upper right corner where both the RMSE and MAPE for BS HIST are much lower compared to the BS VIX. This effect is also the cause for the BS VIX and BS HIST outperformance of the weak GARCH at-the money pricing, where the volatility implied from the weak GARCH is not low enough to price the puts correctly and thus overpricing the options compared to the market price.

It is also worth to notice brackets where one model outperforms when using RMSE and is outperformed when using the MAPE. The reason for this is that the MAPE looks at the relative difference in pricing, so when the price is relatively low, a deviation could be much more significant in relative terms than in absolute. This may cause ambiguous results for the comparison and which model that is superior is a matter of preferences towards relative or absolute pricing errors.

For all the three columns measuring the total pricing errors with respect to moneyness and a constant time to expiry the weak GARCH price is more accurate than the two BS models. This shows that the volatility smile generated from the weak GARCH is more accurate than the plane implied by Black and Scholes. The weak GARCH prices are clearly dominant in relative

[13]Except for the volatility, all the other parameters are equal for the two models.

Table 7: The BS_{VIX}, BS_{HIST} and the Weak GARCH Diffusion Model put pricing errors. The table shows the Root Mean Squared Error and the Mean Absolute Percentage Error for the put prices generated by the three models for different degrees of moneyness and time to expiration. The highlighted figures in the table represents the total values where the Weak GARCH diffusion model outperforms the Black and Scholes framework and vice versa. The time to expiry is measured in days and varies from 10 to 485 days. Total number of puts is 1301.The prices have been generated using the Black and Scholes formula, the Weak GARCH diffusion Monte Carlo Algorithm and was calibrated with the market data described in the data section.

i) Puts		*Time to expiry*							
Moneyness		$T \leq 50$		$51 \leq T \leq 199$		$T \geq 200$		*Total*	
		RMSE	MAPE	RMSE	MAPE	RMSE	MAPE	RMSE	MAPE
≤ 0.85	BS_{HIST}	1.43	107.93%	5.71	99.73%	22.95	44.05%	6.27	95.53%
	BS_{VIX}	1.46	108.95%	6.69	109.58%	32.30	62.28%	8.00	102.78%
	$GARCH$	1.33	103.98%	3.12	54.95%	16.20	28.95%	**4.16**	**72.22%**
$0.85 - 0.95$	BS_{HIST}	2.87	64.71%	4.49	29.41%	11.34	12.05%	4.89	40.63%
	BS_{VIX}	3.95	82.23%	9.02	55.88%	27.00	30.59%	9.75	62.39%
	$GARCH$	2.49	57.11%	3.92	31.64%	10.62	12.12%	**4.37**	**38.68%**
$0.95 - 1.05$	BS_{HIST}	6.25	27.92%	11.73	24.61%	9.74	8.70%	9.40	23.75%
	BS_{VIX}	3.94	18.52%	5.02	10.37%	13.60	11.46%	**5.75**	**13.59%**
	$GARCH$	2.38	12.38%	9.47	21.53%	6.14	8.14%	6.35	16.29%
$1.05 - 1.15$	BS_{HIST}	2.47	2.23%	9.56	7.51%	18.48	27.53%	8.06	8.28%
	BS_{VIX}	1.49	1.26%	3.78	3.04%	4.44	9.41%	**2.98**	**3.25%**
	$GARCH$	2.99	2.20%	6.85	5.08%	8.08	8.54%	5.52	4.45%
≥ 1.15	BS_{HIST}	0.78	0.30%	2.74	1.05%	15.15	106.10%	5.89	31.55%
	BS_{VIX}	0.62	0.23%	1.00	0.38%	5.65	53.37%	**2.26**	15.82%
	$GARCH$	2.32	3.35%	4.05	2.06%	9.41	24.79%	5.20	**9.01%**
Total	BS_{HIST}	2.96	53.85%	7.38	39.80%	15.89	37.80%	6.99	44.89%
	BS_{VIX}	**2.50**	55.75%	**5.70**	43.69%	18.08	34.72%	6.37	46.95%
	$GARCH$	2.61	**47.78%**	5.73	**27.37%**	**10.49**	**17.00%**	**5.26**	**33.62%**

measures and has a lower absolute error in one of the three cases. Again, this represents that the model outperforms in relative measure, but the result is opposite in absolute terms. That being said, the weak GARCH is never the worst pricing alternative and it is always very close to the best model in the absolute cases.

Further, the CBOE does not price options lower than $0.05 which for deep-out- of- the- money options, where the prices are very low, causes high relative errors compared to the absolute which is the case in the upper left bracket of table (7).

For call pricing errors the weak GARCH has a more dominant position. The pricing errors are larger but the weak GARCH prices are more accurate in more cases compared to the put pricing.

Table 8: The BS_{VIX}, BS_{HIST} and the Weak GARCH Diffusion Model call pricing errors. The table shows the Root Mean Squared Error and the Mean Absolute Percentage Error for the call prices generated by the three models for different degrees of moneyness and time to expiration. The highlighted figures in the table represents the total values where the Weak GARCH diffusion model outperforms the Black and Scholes framework and vice versa. The time to expiry is measured in days and varies from 10 to 485 days. Total number of calls is 1307. The prices have been generated using the Black and Scholes formula, the Weak GARCH diffusion Monte Carlo Algorithm and was calibrated with the market data described in the data section.

ii) Calls		*Time to expiry*							
Moneyness		$T \leq 50$		$51 \leq T \leq 199$		$T \geq 200$		*Total*	
		RMSE	MAPE	RMSE	MAPE	RMSE	MAPE	RMSE	MAPE
≤ 0.85	BS_{HIST}	0.06	87.59%	2.60	403.11%	11.23	237.98%	5.24	290.04%
	BS_{VIX}	0.05	60.95%	0.87	136.50%	4.13	99.46%	1.89	110.32%
	$GARCH$	0.06	68.60%	0.55	70.99%	2.79	119.92%	**1.26**	**88.06%**
$0.85 - 0.95$	BS_{HIST}	2.30	428.95%	9.66	431.05%	18.72	90.55%	8.22	379.00%
	BS_{VIX}	0.96	143.95%	4.15	187.11%	5.58	30.39%	**3.15**	147.07%
	$GARCH$	0.55	76.02%	4.50	178.91%	6.79	32.55%	3.34	**117.65%**
$0.95 - 1.05$	BS_{HIST}	6.45	79.82%	12.03	53.89%	13.97	15.40%	10.18	58.55%
	BS_{VIX}	4.16	41.84%	6.22	24.07%	21.66	18.23%	**7.49**	**29.98%**
	$GARCH$	3.24	30.92%	8.73	38.23%	19.98	18.04%	8.15	32.81%
$1.05 - 1.15$	BS_{HIST}	3.48	2.93%	8.28	11.53%	12.22	6.67%	**7.03**	7.44%
	BS_{VIX}	4.76	4.07%	11.80	9.92%	27.39	15.21%	11.48	8.47%
	$GARCH$	3.20	2.71%	6.99	6.60%	19.13	10.55%	7.39	**5.70%**
≥ 1.15	BS_{HIST}	2.02	0.79%	6.06	2.22%	23.17	7.56%	6.62	2.32%
	BS_{VIX}	2.07	0.81%	7.21	2.71%	33.14	11.14%	8.49	3.03%
	$GARCH$	1.90	0.72%	5.81	2.03%	24.11	7.93%	**6.59**	**2.26%**
Total	BS_{HIST}	3.24	117.01%	8.48	137.41%	17.01	56.37%	7.77	117.14%
	BS_{VIX}	2.69	43.98%	6.71	59.08%	20.12	**28.57%**	7.21	48.59%
	$GARCH$	**2.06**	**26.19%**	**6.10**	**56.12%**	**15.89**	29.85%	**6.04**	**40.56%**

The fact that a Monte Carlo simulation is more accurate for empirical put pricing than for calls is an empirical result also observed in other option studies (see for instance Misra et al. (2006)). For the call errors the weak GARCH model clearly outperforms the Black Scholes models in all cases of a constant time to expiry for different degrees of moneyness. As discussed earlier this stems from the implied model volatility. In the call error case the weak GARCH is dominant in the deep-in-the-money case, where the Black Scholes price was more accurate in the put price analysis. Thus, even though the pricing errors are larger, the weak GARCH option price is more accurate compared to the BS model.

The total RMSE for the weak GARCH is 6.04, almost 20% lower than BS VIX, the MAPE is

Table 9: The BS_{VIX}, BS_{HIST} and the Weak GARCH Diffusion Model total pricing errors. The table shows the Root Mean Squared Error and the Mean Absolute Percentage Error for all prices generated by the three models for different degrees of moneyness and time to expiration. The highlighted figures in the table represents the total values where the Weak GARCH diffusion model outperforms the Black and Scholes framework and vice versa. The time to expiry is measured in days and varies from 10 to 485 days. Total number of prices is 2608. The prices have been generated using the Black and Scholes formula, the Weak GARCH diffusion Monte Carlo Algorithm and was calibrated with the market data described in the data section.

iii) Total		*Time to expiry*							
Moneyness		$T \leq 50$		$51 \leq T \leq 199$		$T \geq 200$		*Total*	
		RMSE	MAPE	RMSE	MAPE	RMSE	MAPE	RMSE	MAPE
≤ 0.85	BS_{HIST}	1.33	106.46%	5.17	152.14%	19.04	108.69%	6.11	127.05%
	BS_{VIX}	1.36	105.48%	5.69	114.23%	22.92	74.68%	7.02	104.00%
	$GARCH$	1.24	101.42%	2.68	57.72%	11.73	59.28%	**3.70**	**74.78%**
$0.85 - 0.95$	BS_{HIST}	2.59	250.91%	7.20	240.00%	15.16	52.65%	6.62	215.97%
	BS_{VIX}	2.43	113.79%	6.47	124.69%	15.92	30.49%	6.33	106.27%
	$GARCH$	1.50	66.77%	4.23	108.86%	8.64	22.68%	**3.84**	**79.60%**
$0.95 - 1.05$	BS_{HIST}	6.35	53.87%	11.88	39.30%	11.86	12.05%	9.79	41.18%
	BS_{VIX}	4.06	30.18%	5.62	17.24%	17.63	14.84%	**6.62**	**21.80%**
	$GARCH$	2.81	21.65%	9.10	29.91%	13.07	13.09%	7.26	24.56%
$1.05 - 1.15$	BS_{HIST}	2.94	2.55%	8.98	9.35%	15.44	17.39%	7.58	7.89%
	BS_{VIX}	3.01	2.56%	7.46	6.19%	15.60	12.23%	6.93	5.67%
	$GARCH$	3.09	2.43%	6.91	5.78%	13.45	9.52%	**6.39**	**5.03%**
≥ 1.15	BS_{HIST}	1.86	0.72%	5.32	1.96%	20.23	43.69%	**6.47**	8.47%
	BS_{VIX}	1.88	0.74%	5.83	2.19%	23.06	26.62%	7.18	5.72%
	$GARCH$	2.93	1.06%	5.71	2.04%	18.72	14.11%	6.81	**3.68%**
Total	BS_{HIST}	3.10	85.56%	7.93	88.68%	16.45	47.08%	7.38	81.10%
	BS_{VIX}	2.60	49.84%	6.21	51.40%	19.11	31.64%	6.80	47.77%
	$GARCH$	**2.34**	**36.94%**	**5.92**	**41.77%**	**13.19**	**23.42%**	**5.66**	**37.10%**

40.56%, 18.03 percentage points lower than BS VIX. For the call price errors the BS HIST has very large pricing errors for out of the money options, causing the total pricing error to be very large. Again, some of the high relative errors stems from the price floor set by CBOE on the options.

In table 9 the total pricing errors from the models are summarized. The table shows the same picture as the put and call price errors. The weak GARCH is dominant in capturing the moneyness effect, prices deep-out-of-the money, out-of-the-money, in-the-money and deep-in-the-money options significantly better than the BS model. The total pricing error of the weak GARCH is, measured in RMSE and MAPE respectively, 5.66 and 37.10%. This implies a total outperformance of BS VIX of 20.14% and 10.67 percentage points. Compared to the BS HIST the RMSE and MAPE are 30.38% and 44 percentage points lower.

Using a difference in mean hypothesis test where we test the hypothesis that the average RMSE and MAPE is lower for the weak GARCH diffusion model than for the BS HIST and BS VIX. We conclude that with the 2608 observations in our sample we reject the null hypothesis of no difference both at 95% and 99% significance. In light of statistical inference it is therefore valid to conclude that the weak GARCH pricing is significantly better than the BS pricing.

When comparing the estimated option prices it is possible to observe a pattern in the pricing. All three models, with different magnitude, underprices deep in the money calls and underprice when moneyness decrease to below 0.85 with below 100 days to expiry. When time to expiry increases the models still underprice the deep- in- the money calls but price the rest of the options rather fairly. For put options the models all underprice deep- out- of- the money options with below 50 days to expiry. When time to maturity increases all pricing converges to the market price, again with a significant difference in magnitude. A possible reason for this is the lack of an exact distribution for the underlying security. The probability of extreme events are too low and the probabilities of a left tail event is not significantly larger than for a right tail event. Further we have assumed a negative drift in our analysis, since the current dividend yield which we have used as a proxy for δ is larger than the 12-month treasury, which is our proxy for the risk free interest rate. This implies that, everything else equal, one should expect the underlying stock to depreciate over time. If this is not the belief of the option market one could experience underpricing of the options since the expected value of the stock becomes too low. Another issue is the fact that we have used the ask price as a proxy for the market price due to lack of a dataset consisting of arbitrage free and consistent option prices. In the existing literature the ask price, an average of the bid and the ask and an arbitrage filtered dataset of the last traded price have been used. All models are designed to capture the market price of the option, not the ask price. Since the market price is somewhere between the bid and the ask this could be a reason for the underpricing of several options.

The liquidity of the options also has an impact on the pricing. The most illiquid options are excluded from the dataset, but we still have some options where there is no market. This usually increases the bid-ask spread and could also be a cause for pricing bias when we use the ask price. That being said, options very similar to the ones with zero open interest are trading actively which ensures that the ask price is not too inaccurate.

The overall picture is that all models struggle with short time options, a reason for this could be that the implied volatility surface for observed options is most extreme when time to expiry is close. It is in this area where both the local minimum and maximum are achieved. Further, the implied volatility converges to normality when time to expiry increases. In GARCH terminology the volatility also converges to the long term equilibrium. The equilibrium volatility is lower than the short term volatility in our case and it should be no surprise that the model assuming the lowest standard deviation, the BS VIX, has the lowest observed pricing error for options with the longest time to maturity.

The results above answers the questions raised. We have shown by including stochastic volatility in the option valuation that we are able to generate the empirical properties of volatility with respect to moneyness and term structure. We are also able to reduce the pricing errors of the traditional framework with 20.14% to 30.38%. Further we have reduced the extremum pricing errors from 428.95%, which is the case for out-of-the-money options with less than 50 days to expiry priced by the BS HIST, to 119.92% for deep-out-of the money options with over 200 days to expiry. It feels safe to suggest the weak GARCH diffusion model as a better option pricing tool compared to the Black Scholes formula when it comes to pricing accuracy.

Several other factors may also bias the pricing that is outside the scope of this thesis. We have not excluded options where the price floor is reached. As discussed above this leads to irregular pricing for options reaching this barrier. Excluding these observations, or including a correction for the floor in the pricing algorithm could help fix this issue. Further, since there is no weak GARCH estimation technique available we have chose to use the strong GARCH parameters as a proxy. The effect of this proxy is rather uncertain and is beyond the scope of the analysis, but a theoretically correct estimation technique could be of significant value to this type of analysis.

Another technical assumption that could give rise to inaccurate pricing is the discretization method for the differential equations that we base our valuation on. We allowed one draw from the standard normal distribution to represent one time step, i.e. one day, and then averaged this several thousand times. This is a rather crude method but necessary to limit the computation time. An interesting analysis could be to repeat the procedure with each time step as several hundred draws from the standard normal distribution. This could also correct the model for the parameter stability issues discussed above. Even though we decrease the dependency to each draw by repeating the process, and limit the process by using the empirical price variance correlation which implied lower kurtosis for the distribution, the discretization process is still debatable and could give rise to inaccuracy and instability.

Regardless of the potential sources of bias above the option pricing method implemented has a higher pricing accuracy than the Black and Scholes framework. Another approach to test the performance qualities of the model could be to analyze the pricing accuracy of the weak GARCH compared to the popular Heston stochastic volatility model

5 Conclusion

In this thesis we have tried to answer several questions within the topic of option valuation under stochastic volatility. We have questioned the assumption of taking the standard deviation of returns as a constant and an investigation of the pricing abilities of the existing framework has been performed. Further we have implemented another option pricing tool using a weak GARCH diffusion for the stochastic volatility. The main question that we tried to answer was which option pricing tool minimized the deviation from the observed market price. We used the Standard & Poor 500 index as the underlying security due to the favorable properties of the index. Further we used the ask price of the options on this index as a proxy for the market price.

To answer the questions raised we implemented a weak GARCH diffusion option pricing model using Monte Carlo valuation. The analysis consisted of an implementation of strong GARCH parameters and market data to a self written option pricing algorithm created in MATLAB. We obtained option prices for a total of 2608 options using the weak GARCH option pricing algorithm and two different calibrations of the Black and Scholes formula. The difference between the two was the measure of volatility, in the BS HIST we used the historical standard deviation obtained from the observed data, and in BS VIX we used the spot value of the VIX index as a proxy for volatility.

In the calibration we encountered problems with the price-variance correlation implied by the model. This correlation was not consistent with empirical findings, thus allowing for larger pricing errors. We proposed a solution where we used the historical correlation between the squared VIX and the returns from the SPX index as input instead of the sample kurtosis.

In the analysis we concluded that the Weak GARCH diffusion model is clearly dominant over the two BS alternatives. We reject the null hypothesis of no RMSE and MAPE difference between the models both at a 95% and 99% level of confidence. Over the sample we found that the RMSE of the weak GARCH was 20.14% lower compared to the BS VIX. Compared to the BS HIST the difference was 30.38%. Further the MAPE was 10.67 and 44 percentage points lower compared to BS VIX and BS HIST respectively.

To our knowledge, there has not been any empirical analysis of the weak GARCH diffusion model earlier, but the results are in line with other studies comparing Black Scholes prices to prices found with stochastic volatility models.

The results of the analysis was then questioned with the validity of the assumptions made in the Weak GARCH algorithm. Especially the discretization process is debatable, in addition to the use of strong GARCH estimates in a weak distribution. Further, an interesting discussion is to compare the results we obtain with a similar analysis using other proxies for the risk free interest rate and dividend yield. Another interesting approach would be to apply the algorithm on high frequency observed market prices to see if some of the structural biases disappear.

References

Alexander, C. (2008a): *Market Risk Analysis II : Practical Financial Econometrics*. John Wiley & Sons, Inc.

Alexander, C. (2008b): *Market Risk Analysis III : Pricing, Hedging and Trading Financial Instruments*. John Wiley & Sons, Inc.

Alexander, C. and Lazar, E. (2005): "On The Continuous Limit of GARCH." *ICMA Center Discussion Papers in Finance 13*, 1–19.

Andersen, T. G., Benzoni, L. and Lund, J. (2002): "An Empirical Investigation of Continuous-Time Equity Return Models." *The Journal of Finance 57*, 1239–1284.

Bakshi, G., Cao, C. and Chen, Z. (1997): "Empirical Performance of Alternative Option Pricing Models." *The Journal of Finance 52*, 2003–2049.

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

Bollerslev, T. (1986): "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics 31*, 307–327.

Brooks, C. (2008): *Introductory Econometrics for Finance*. Cambridge University Press.

Chicago Board of Exchange (CBOE) (2012): "THE CBOE VOLATILITY INDEX- VIX." Accessed at March 29th. URL `http://www.cboe.com/micro/vix/vixwhite.pdf`.

Corradi, V. (2000): "Reconsidering the Continuous Time Limit of the GARCH(1,1) Process." *Journal of Econometrics 96*, 145–153.

Drost, F. and Nijman, T. (1993): "Temporal Aggregation of GARCH Processes." *Econometrica 61*(4), 909–927.

Fischer Black (1976): "The pricing of commodity contracts." *Journal of Financial Economics 3*, 167–169.

Hull, J. C. (2003): *Options, Futures, and Other Derivatives*. Pearson Education, Inc.

Kaeck, A. and Alexander, C. (2011): "Stochastic Volatility Jump-Duffusions for European Equity Index Dynamics." *European Financial Management 17*, 1–26.

Lehar, A., Scheicher, M. and Schittenkopf, C. (2002): "GARCH vs. stochastic volatility: Option pricing and risk management." *Journal of Banking & Finance 26*, 323–235.

Lewis, A. L. (2000): *Option Valuation Under Stochastic Volatility With Mathematica Code*. Finance Press

Marketwatch (2012): “Standard & Poor 500 Option Chains.” Accessed at 16.02. URL `http://www.marketwatch.com/investing/index/spx/options`.

Mcdonald, R. L. (2006): *Derivatives Market*. Pearson Education, Inc.

Misra, D., Kannan, R. and Misra, S. D. (2006): “Implied Volatilty Surfaces: A Study of NSE NIFTY Options.” *International Research Journal of Finance and Economics 6*, 1–17.

Nelson, B. (1990): “ARCH Models as Diffusion Approximations.” *Journal of Econometrics 45*, 7–38.

Rouah, F. D. and Vainberg, G. (2007): *Option Pricing Models & Volatility Using Excel-VBA*. John Wiley & Sons, Inc.

Rubinstein, M. (1994): “Implied Binomial Trees.” *Journal of Finance 49*, 771–818.

Rubinstein, M. (1998): “Edgeworth Binomial Trees.” *Journal of Derivatives 5*(3), 20–27.

Standard&Poor (2012): “S&P500 fact sheet.” Accessed at March 21st. URL `http://www.standardandpoors.com/indices/main/en/eu`.

Wiersema, U. F. (2008): *Brownian Motion Calculus*. John Wiley & Sons Ltd.

Williams, R. J. (2006): *Introduction to the Mathematics of Finance*. American Mathematical Society.

A Option Data

A.1 Option Prices and Implied volatility

The tables consists of selected prices from the options traded on the SPX index. Implied volatility is found by applying the NRB-method.

Table 10: S&P 500 call option ask prices.
The option prices was downloaded from Marketwatch (2012)at February 16th.

Expiry	Option strike												
	950	1000	1050	1100	1150	1200	1250	1300	1350	1400	1500	1600	1700
16-Mar	406.20	356.80	307.60	258.70	210.00	162.00	115.50	72.10	35.30	11.00	1.85	0.10	0.05
20-Apr	406.70	357.60	308.90	260.40	212.50	165.70	120.70	79.10	43.40	17.60	4.50	0.25	0.20
18-May	406.60	358.10	309.90	262.20	215.50	170.30	127.20	87.40	52.90	26.40	9.50	0.40	0.20
15-Jun	407.70	359.90	312.60	266.10	220.90	177.40	136.30	98.50	65.30	38.60	19.20	0.95	0.35
21-Sep	409.20	363.20	318.20	274.40	231.80	191.10	152.80	117.30	85.30	57.90	36.60	4.80	0.85
21-Dec	411.50	367.20	324.00	282.00	241.50	202.70	165.90	131.70	100.80	73.90	50.70	11.20	2.65

Table 11: S&P 500 put option ask prices.
The option prices was downloaded from Marketwatch (2012)at February 16th.

Expiry	Option strike												
	950	1000	1050	1100	1150	1200	1250	1300	1350	1400	1500	1600	1700
16-Mar	0.90	1.45	2.30	3.20	4.60	6.50	10.10	16.80	30.30	56.80	147.80	246.80	346.70
20-Apr	2.20	2.40	3.50	5.00	7.90	10.10	16.20	24.80	39.20	64.20	148.60	248.00	347.70
18-May	4.30	5.80	7.60	9.90	13.30	18.00	24.90	35.40	51.10	75.10	152.60	251.10	350.70
15-Jun	7.90	10.00	12.70	16.30	21.20	27.80	36.70	48.80	65.70	89.20	159.50	253.30	352.50
21-Sep	16.00	19.70	24.50	30.50	37.90	47.00	58.60	73.10	91.50	114.90	177.80	262.30	358.60
21-Dec	24.70	30.10	36.80	44.30	53.60	64.60	77.90	93.30	113.30	136.70	196.70	274.90	366.30

Table 12: S&P 500 call option implied volatilities.
The implied volatilies are reported in percentages, and are computed with the NRB-method. For the calculations we used the following: $S = 1358$,$r = 0.17\%$,$\delta = 1.95\%$.

Expiry	Option strike												
	950	1000	1050	1100	1150	1200	1250	1300	1350	1400	1500	1600	1700
16-Mar	47.23	48.96	46.77	43.40	38.84	34.15	29.59	25.21	21.17	17.83	15.90	20.10	23.83
20-Apr	46.69	42.92	39.39	35.55	31.79	28.20	24.72	21.48	18.44	15.78	13.76	14.61	18.62
18-May	42.41	39.41	36.28	33.12	30.15	27.33	24.56	21.85	19.29	17.00	14.83	13.66	16.08
15-Jun	41.05	38.36	35.63	32.96	30.45	28.04	25.70	23.44	21.26	19.30	17.45	15.11	15.48
21-Sep	36.13	34.17	32.29	30.46	28.57	26.76	25.00	23.26	21.52	19.87	18.50	15.49	14.30
21-Dec	34.20	32.62	31.08	29.54	28.01	26.48	24.93	23.41	21.97	20.63	18.17	16.22	14.77

Table 13: S&P 500 put option implied volatilities.
The implied volatilies are reported in percentages, and are computed with the NRB-method. For the calculations we used the following: $S = 1358$,$r = 0.17\%$,$\delta = 1.95\%$.

Expiry	Option strike												
	950	1000	1050	1100	1150	1200	1250	1300	1350	1400	1500	1600	1700
16-Mar	57.45	53.74	49.98	45.03	40.24	35.01	30.31	25.82	21.88	19.29	25.80	35.27	44.88
20-Apr	44.12	39.16	36.25	33.17	30.86	26.66	24.25	21.20	18.25	15.96	15.81	22.00	27.54
18-May	41.41	38.66	35.71	32.68	29.88	27.07	24.36	21.81	19.34	17.25	15.54	20.47	25.45
15-Jun	41.11	38.34	35.64	33.05	30.60	28.21	25.87	23.56	21.41	19.54	17.01	18.53	22.45
21-Sep	36.15	34.06	32.13	30.26	28.42	26.58	24.83	23.11	21.48	20.04	17.46	16.12	16.95
21-Dec	34.34	32.66	31.11	29.43	27.88	26.33	24.83	23.26	22.01	20.75	18.49	16.88	16.21

B VBA and Matlab Code

B.1 The Black-Scholes Formula

The following code is self written based on the formulas in Mcdonald (2006).

```
Function BSCall(S As Double, K As Double, T As Double, _
r As Double, d As Double, v As Double) As Double
Dim d1 As Double, d2 As Double
d1 = (Log(S / K) + (r - d + 0.5 * v ^ 2) * T) / (v * Sqr(T))
d2 = d1 - v * Sqr(T)

BSCall = Exp(-d * T) * S * Application.NormSDist(d1)  _
- Exp(-r * T)* K *Application.NormSDist(d2)

End Function

Function BSPut(S As Double, K As Double, T As Double, _
 r As Double,d As Double, v As Double) As Double
BlackScholes = K * Exp(-r * T) * Application.NormSDist(-d2)_
 - Exp(-d * T) * S * Application.NormSDist(-d1)
End Function



Function VegaCall(S As Double, K As Double, T As Double,_
 r As Double, d As Double, v As Double) As Double

VegaCall = S * Exp(-d * T) * Fz(d1) * Sqr(T)

End Function

Function VegaPut(S As Double, K As Double, T As Double, _
 r As Double, d As Double, v As Double) As Double
VegaPut = VegaCall(S, K, T, r, d, v)
End Function

Function Fz(x)
Fz = Exp(-0.5 * -x ^ 2) / Sqr(2 * Application.Pi())
```

```
End Function
```

B.2 The Newton-Rhapson Bisection Method

The following code is self written based on Rouah and Vainberg (2007). The code has been simplified and expanded to take into account the dividend yield. The VBA code for a put option is not included since it is identical except the replacement of BSCall with BSPut.

```
Function ImpliedVolCall(Stock As Double, Exercise As Double,_
 Time As Double, Interest As Double, _
Dividends As Double, realC As Double)


Dim ACC As Double  ' Deviation from zero allowed
Dim a As Double 'lowest boundary for volatility
Dim b As Double 'highest boundary for volatility
Dim MidP As Double

ACC = 1e-10
a = 1e-06
b = 5

    MidP = 0.5 * (a + b)
    dxold = (b - a)
    dx = dxold
    midCdif = realC - BSCall(Stock, Exercise, Time, Interest,_
     Dividends, MidP)
    midCvega = VegaCall(Stock, Exercise, Time, _
    Dividends, Interest, MidP)
    For i = 1 To 1000
    'The following are the boundary conditions:
        If ((((MidP - b) * midCvega - midCdif) * ((MidP - a)_
         * midCvega - midCdif) > 0) Or (Abs(2 * midCdif) >_
          Abs(dxold * midCvega))) Then
            dxold = dx
            dx = 0.5 * (b - a)
            MidP = a + dx
        Else
```

```
            dxold = dx
            dx = midCdif / midCvega
            temp = MidP
            MidP = MidP - dx
        End If
        midCdif = realC - BSCall(Stock, Exercise, Time,_
         Interest, Dividends, MidP)
        If (Abs(midCdif) < EPS) Then
             Exit For
        End If
        midCvega = VegaCall(Stock, Exercise, Time, _
         Dividends, Interest, MidP)
        If (midCdif < 0) Then
            b = MidP
        Else
            a = MidP
        End If
    Next i
ImpliedVolCall = MidP
End Function
```

B.3 Natural Cubic Spline Interpolation

The following code is from Rouah and Vainberg (2007)

```
Option Base 1
Function NSpline(S, x, a)
n = Application.Count(x)

Dim h() As Double, alpha() As Double
ReDim h(n - 1) As Double, alpha(n - 1) As Double

For I = 1 To n - 1
    h(I) = x(I + 1) - x(I)
Next I

For I = 2 To n - 1
    alpha(I) = 3 / h(I) * (a(I + 1) - _
a(I)) - 3 / h(I - 1) * (a(I) - a(I - 1))
```

```
Next I

Dim l() As Double, u() As Double, z() As Double,_
 c() As Double, b() As Double, d() As Double
ReDim l(n) As Double, u(n) As Double, z(n) _
  As Double, c(n) As Double, _
  b(n) As Double, d(n) As Double

l(1) = 1: u(1) = 0: z(1) = 0
l(n) = 1: z(n) = 0: c(n) = 0

For I = 2 To n - 1
    l(I) = 2 * (x(I + 1) - x(I - 1)) - h(I - 1) * u(I - 1)
    u(I) = h(I) / l(I)
    z(I) = (alpha(I) - h(I - 1) * z(I - 1)) / l(I)
Next I

For I = n - 1 To 1 Step -1
    c(I) = z(I) - u(I) * c(I + 1)
    b(I) = (a(I + 1) - a(I)) / h(I) - h(I) * (c(I + 1) + _
     2 * c(I)) / 3
    d(I) = (c(I + 1) - c(I)) / 3 / h(I)
Next I

For I = 1 To n - 1
    If (x(I) <= S) And (S <= x(I + 1)) Then
         NSpline = a(I) + b(I) * (S - x(I)) + c(I) * _
   (S - x(I)) ^ 2 + d(I) * (S - x(I)) ^ 3
    End If
Next I

End Function
```

B.4 The Weak GARCH Diffusion Option Pricing Model

The following MATLAB code is self written and is a Monte-Carlo simulation for option pricing. The parameters implemented may vary from those used in the actual calculation.

```
clear all;
```

```
n=100000; %simulations
T=30; %days to expiry

%Distribution
tau=-1.1696; %skewness
eta=3.63890827; %kurtosis
%Input

S = zeros(T,n);
V = zeros(T,n);
X = zeros(T,n);
Y = zeros(T,n);
Z = zeros(T,n);
Vol = zeros(T,n);
drift = zeros(T,n);
Price_T = zeros(1,n);
Price_T = zeros(T,n);
% daily GARCH parameters
Gw=1.35677e-06;
Ga=  0.0778921; %Reaction to market shock
Gb=   0.912546;%volatility of volatility

%Market Parameters
Mr=0.0014; %Risk free rate
Md=0.0202; %Dividends
S(1,:)=1358; %initial price for SV model

%Model
r=Mr/(365);
d=Md/(365);
w=Gw;
a=Ga;
b=Gb;
phi=1-a-b;
m=(w/phi); %long term average variance
rho=tau/(sqrt(eta-1));

V(1,:)= 0.1920^2/(365); %initial variance
```

```
Vol(1,:) = (0.1920/(sqrt(365)));

for i=1:n
    for t=1:T-1
        %Generates two standard normal random variables
     X(t,i)=randn;
     Y(t,i)=randn;
     Z(t,i)=(rho*X(t,i)+sqrt(1-rho^2)*Y(t,i));
     %Mean reverting volatility function
V(t+1,i)=V(t,i)+(phi*(m-V(t,i))+sqrt(eta-1)*a*V(t,i)*Z(t,i));
        if V(t+1,i)<0;
           V(t+1,i)=V(t,i);
        end
     Vol(t+1,i)=sqrt(V(t+1,i));
     drift(t+1,i)=(r-d-0.5*V(t+1,i));
     %SV stock process
     S(t+1,i)=S(t,i)*exp(drift(t,i)+Vol(t,i)*X(t,i));
        if S(t+1,i)<0.01;
            S(t+1,i)=1;
        end
   end
Price_T(1,i)=S(T,i);
end

strikes = 751;
Payoffs_C = zeros(strikes, n);
Payoffs_P = zeros(strikes, n);
K=zeros(strikes,1);
for i = 1:strikes
    for j = 1:n
        K(i)=949+i;
        Payoffs_C(i,j) = max(Price_T(1,j) - K(i),0);
        Payoffs_P(i,j) = max(K(i)-Price_T(1,j),0);
    end
end
 Call_GARCH = zeros(1,strikes);
Put_GARCH = zeros(1,strikes);
for i = 1:strikes
    Call_GARCH(1,i) = exp(-T*r)*mean(Payoffs_C(i,:));
```

```
Put_GARCH(1,i) = exp(-T*r)*mean(Payoffs_P(i,:));
```