



Pricing Derivatives:

Implementing Heston and Nandi's (2000) Model on the Swedish Stock Index

Abstract

This thesis is based on Heston and Nandi's (2000) paper. The aim is to check how their closed-form discrete-time GARCH option pricing model performs on Swedish data, and if there are any significant changes to its performance when estimating it via maximum likelihood using the Normal- and the Student-t distribution. The model is compared with the Black-Scholes- and the ad hoc Black-Scholes model of Dumas, Fleming, and Whaley (1998).

The results show that when the model is estimated under the Student-t distribution, its out-of-sample valuation performance increases in general when its parameters are updated. However, it is also shown that the model suffers from significant mispricing errors as it is greatly outperformed by both the BS- and the ad hoc BS model. This error is caused by poor estimates of two of its most important parameters, namely the volatility of volatility- and the skewness parameter. It is also shown that the model generally performs better when it is estimated using the Normal- than the Student-t distribution.

Master Degree Thesis in Economics

D Level ECTS: 15 points

Fall term 2007

Authors: Bujar Huskaj 820719

Sharlett Hanna 820301

Supervisor: Birger Nilsson

Acknowledgments

While remaining responsible for any errors in this thesis, the authors are thankful to their supervisor Birger Nilsson, who made this thesis possible thanks to his helpful comments, useful remarks, and support.

Table of Contents

1. Introduction	1
2. Theoretical Framework	4
2.1. Risk-Neutral Valuation.....	4
2.2. Implied Volatility.....	5
2.3. The Black-Scholes Model.....	5
2.4. The Implied Volatility Curve.....	7
2.5. The ad hoc Black-Scholes Model.....	8
2.6. Stochastic Volatility Models.....	10
2.7. The Heston-Nandi Closed-Form GARCH Option Valuation Model.....	11
2.8. Mispricing Error Measures	15
3. Results	16
3.1. Data Information.....	16
3.2. Estimations.....	16
3.3. Model Comparisons.....	18
3.3.1. In-sample Model Comparison	19
3.3.2. Out-of-Sample Model Comparison.....	22
4. Conclusions and Discussion	28
5. References	29

1. Introduction

The first scientific approach to value options dates back to 1900 when the French mathematician Louis Bachelier created the well-known Brownian motion to model the evolution of stock prices (which became the norm in the 1960s). Option pricing formulas were at that time derived by taking the discounted expectation. However, the major break-through within the field of option pricing was made by Black and Scholes (1973) and Merton (1973). Like previous pricing models, the Black-Scholes (1973) (henceforth BS) model assumes that the evolution of the stock price follows a geometric Brownian motion. Another assumption made in the BS model is that of a constant volatility, which according to Ross (1989) implies a constant flow of information as he means that volatility can be considered as a measure of information flow. The correctness of this constant volatility assumption has however been questioned by many since it is known that the BS model has some pricing biases (Rubinstein (1985)). Hull and White (1987) and Wiggins (1987), among many, suggest that this constant volatility assumption might be a reason for the failure of the BS model to value options exactly. A common way to express the deficiencies of the BS model is by plotting BS implied volatilities against strike price (K) or moneyness (S/K or K/S). This plot is also known as the volatility smile, where a constant volatility should result in “a neutral facial expression” (Dumas, Fleming, and Whaley (1998)). However, evidence against this has been presented numerous times. Two examples are Rubinstein (1994), who examines the S&P 500 index option market, and Heynen (1993), who examines the European Options Exchange. These deficiencies are naturally what induce researchers to pursue the development of more realistic models, incorporating empirical features such as stock price volatilities and interest rates as stochastic processes. Extensions of the BS model to incorporate the fact that the volatility of stock prices varies stochastically have for example been made by Scott (1987), Hull and White (1987), and Wiggins (1987). These so-called stochastic volatility (SV) models are divided into two groups, considering the volatility in continuous- and discrete-time. One alternative approach to characterize a stochastic volatility model in discrete-time is to use the generalized autoregressive conditional heteroskedasticity (GARCH) model (see Bollerslev (1986)).

This thesis is based on Heston and Nandi (2000), who present a closed-form discrete-time GARCH (henceforth HN GARCH(p,q)) model for pricing European options. The model applies for assets whose variances follow GARCH(p,q) processes. It should however be mentioned that only the GARCH(1,1) case is considered in this thesis, which according to a survey made by Bollerslev et al. (1992) is the simplest and most robust one of the family of

volatility models. The HN GARCH(p,q) model values options by using the volatilities computed directly from the history of asset prices, and it incorporates the correlation between spot prices and their volatility. Heston (1993) shows that this correlation is important for explaining the skewness seen in stock returns, making in this aspect the model superior to the BS model (which assumes that (continuously compounded) stock returns are normally distributed with known mean and variance). In addition to this, the HN GARCH(p,q) model combines also the cross-sectional information contained in options with the information in the time series of the underlying asset. In other words, it incorporates the volatility smile. Numerical methods such as simulations are valuable methods for pricing options and are for that reason widely used. However, despite their increasing popularity led by the growing capacity of computing power, they can be both time consuming and computationally intensive. This drawback is however an advantage of the HN GARCH(p,q) model, since its implementation is analytical.

The purpose of this thesis is to implement the HN GARCH(p,q) model for the Swedish stock index (OMXS30) using the Normal- and the Student-t distribution and compare its performance with the BS- and the ad hoc BS model.

The methodology is organized such that the BS- and the ad hoc BS model are estimated first, followed by two versions of the HN GARCH(1,1) model; one with constant- and one with updated parameters. The HN GARCH(1,1) model is also estimated using the Normal- and the Student-t distribution. All the models are estimated and compared in both in-sample and out-of-sample.

There are four main restrictions imposed in this thesis. The first one is that closing prices are used instead of intra-daily. This restriction should however not affect the results in a significant way since OMXS30 options are frequently traded, implying that the closing prices of the options and the stock index should still be reasonably synchronous. The second restriction considers the estimation of the HN GARCH(1,1) model parameters. While Heston and Nandi (2000) use non-linear least squares (NLLS) (which requires the usage of option prices) for this purpose, the method used here is maximum likelihood (ML). Heston and Nandi (2000) point out that the NLLS procedure is preferable since the information in option prices is more forward looking than the price of the underlying asset. This argument has however not been proven empirically, which is why the easier method, the ML procedure, is used here. The third restriction is that dividends are not taken into account. This should in general not have any significant effects on the results since many of the stocks in the OMXS30 index pay dividends

only once a year (mostly in Mars or April). The fourth and final restriction is that only call options are valued. However, put options can easily be valued using the put-call parity.

The outline of this thesis is as follows: Section 2 presents a theoretical framework discussing the BS-, the ad hoc BS-, and the HN GARCH(p,q) model. Section 3 presents the empirical results, and finally, section 4 gives a conclusion and a discussion.

2. Theoretical Framework

This section is the foundation of the theory behind all the aspects considered in the empirical part, and with that, constitutes a major corner stone of the thesis. It is comprised of various subsections considering discussions about risk-neutral valuation, implied volatility, different option pricing models, etc. The first subsection gives a brief discussion about the concept of risk-neutral valuation.

2.1. Risk-Neutral Valuation

Risk-neutrality implies that agents are indifferent to risk, meaning that they do not require any compensation for the exposure of it. A result introduced by Cox and Ross (1976) is that the expected return on all assets in a so-called risk-neutral world is equal to the risk-free interest rate. This follows that the price of an asset is equal to its expected payoff, discounted at the risk-free rate. Hence the value of a European call option for instance is given by

$$C = E \left[\exp \left\{ - \int_0^T r(t) dt \right\} \text{Max}(S(T) - K, 0) \right] = e^{-rT} E[\text{Max}(S(T) - K, 0)]. \quad (1)$$

Clearly, the second equality is an easier expression to grasp, but its drawback is that it only holds if the interest rate is assumed to be constant. The first equality however holds always. This result is known as risk-neutral valuation.¹

Risk-neutral valuation arises from the fact that none of the variables in the BS formulas (discussed in a later section), i.e. the current stock price, the stock price volatility, the time to maturity, and the risk-free interest rate, are dependent of individual risk preferences. This fact allows for the assumption that all investors are risk-neutral, stating that it is always possible to assume the world to be risk-neutral when pricing options. This assumption simplifies the analysis of options, or derivatives in general, since the resulting prices are valid in all worlds and not just in the assumed one. The reason to this is that there are only two changes that occur when moving from a risk-neutral world to a risk-averse one, and luckily, these two always offset each other exactly. These are the expected growth rate in the stock price and the discount

¹ For more about risk-neutral pricing, see for example Duffie (1996), Hull (1997), and Wilmott (1998).

rate that must be used for any of the derivative's payoffs (Hull (2006)). The significant simplification followed by using this assumption makes it one of the most important ones in explaining option prices. Hence the fact that all option pricing models are founded on this assumption is therefore not so surprisingly. Hull (2006) describes it as being "...without doubt the single most important tool for the analysis of derivatives."²

2.2. Implied Volatility

The most popular type of implied volatility σ is the so-called BS implied volatility. An option's BS implied volatility is that volatility obtained when equating the option's market value to its BS value, given the same strike price and time to maturity. It is extracted numerically due to the fact that the BS formula cannot be solved for σ in terms of the other parameters. Hence implied volatilities are embedded in option prices, which in turn reflect the future expectations of the market participants. This means that implied volatilities are important because they form a forward-looking estimate of the volatility of the underlying asset, and can therefore be used to monitor the market's opinion of it. Traders for example use the implied volatilities from actively traded options in order to estimate an appropriate volatility to use to price a less actively traded option on the same asset (Hull (2006)). Heston and Nandi (2000) warn however that "... using implied volatilities to value an option requires the use of other contemporaneous options that may not always be feasible if one does not have reliable option prices such as in cases of thinly traded or illiquid markets."³ Noteworthy is that the prices of deep-in-the-money- and deep-out-of-the-money options are relatively insensitive to volatility, meaning that the implied volatilities calculated from these options tend to be unreliable (Hull (2006)).

2.3. The Black-Scholes Model

Fischer Black and Myron Scholes developed in the early 1970s a method to price options on non-dividend paying stock, which has turned to be one of the most successful and widely used models in financial economics.⁴ This section presents only its formulas, since the derivation of

² See Hull (2006) pp. 293.

³ See Heston and Nandi (2000) pp. 586.

⁴ The first extension of this model was made shortly after by Merton (1973), who extended the model to consider the case of dividend-paying stocks.

the model and its underlying assumptions can be found in every literature in financial economics.

Assuming that the model's underlying assumptions hold, the BS formula for a European call option on a non-dividend-paying stock trading at time t is given by:⁵

$$C_t(S_t, T - t) = S_t N(d_1) - K e^{-r(T-t)} N(d_2), \quad (2)$$

where

$$d_1 = \frac{\log\left(\frac{S_t}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T - t)}{\sigma\sqrt{T - t}}, \quad (3)$$

$$d_2 = d_1 - \sigma\sqrt{T - t}. \quad (4)$$

C_t is the call price, $N(d_i)$ is the cumulative probability distribution value for a standard normal random variable with value d_i , S_t is the price of the underlying asset at time t , K is the strike price, r is the risk-free interest rate, and T is the option's time to maturity.

Two underlying assumptions of the BS model are that the underlying asset's volatility is constant and that its returns are normally distributed (equivalently, that the prices of the underlying asset are log-normally distributed). These two assumptions are major drawbacks of the BS model. Black (1976) for instance points out that “if the volatility of a stock changes over time, the option formulas that assume a constant volatility are wrong.”⁶ Empirical evidence has shown to be in accordance with Black's argument. A strong negative correlation between stock's current prices and their future volatilities has been seen. In other words, volatility tends to rise in response to “bad news” (i.e. when an unexpected price drop occurs) and to fall in response to “good news” (i.e. when an unexpected price rise occurs). This feature is called the leverage effect and was first noted by Black (1976), who also writes: “I have believed for a long time that stock returns are related to volatility changes. When stocks go up, volatility seem to go down; and when stocks go down, volatilities seem go to up.”⁷ Figure 1 illustrates this effect.

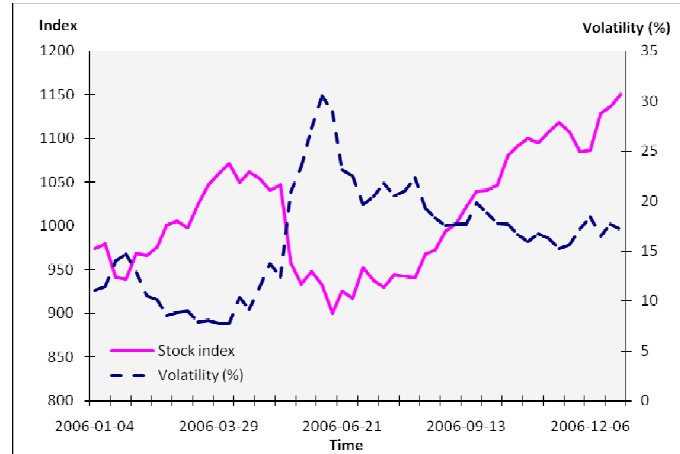
⁵ The BS formula for a European put option on a non-dividend-paying stock trading at time t is given by

$$P_t(S_t, T - t) = -S_t N(-d_1) + K e^{-r(T-t)} N(-d_2).$$

⁶ See Black (1976) pp. 177.

⁷ See Black (1976) pp. 177.

Figure 1 Leverage effect



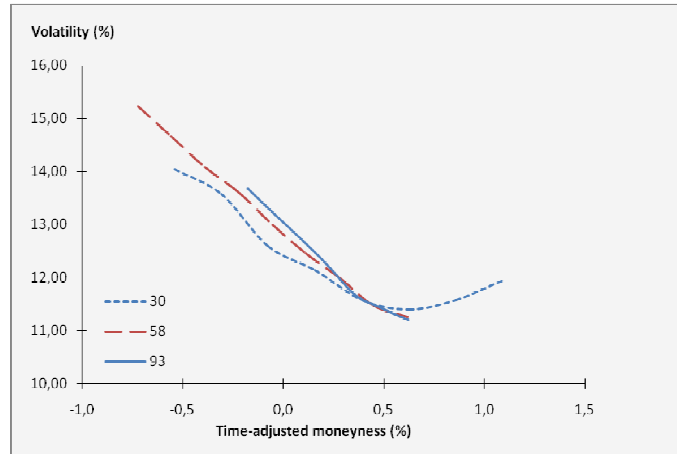
Note: OMXS30 stock index level and BS implied volatility each Wednesday during the period 2006-01-04 – 2006-12-27.

Hence Figure 1 clearly shows the leverage effect, with a correlation in the first difference of the series equalling -0.3977. Empirical evidence has also shown that financial data is generally not normally distributed, but is instead more skewed and has a greater kurtosis. The most known graphical evidence against the constant volatility assumption is the implied volatility curve, which is discussed in the next section.

2.4. The Implied Volatility Curve

A well-known and important feature is the implied volatility curve, also known as the volatility smile or the volatility smirk. This implied volatility curve is simply a plot of the BS implied volatilities (which are obtained by using a cross-section of option prices with a variety of strike prices and maturities) against strike price (K) or moneyness (S/K or K/S). The two factors that are believed to have the greatest effect on the shape of the implied volatility curve are moneyness and maturity (Rouah and Vainberg (2007)). Figure 2 illustrates the pattern in the OMXS30 BS implied volatilities.

Figure 2 Volatility smile



Note: BS implied volatilities on May 25, 2005. Implied volatilities are computed from OMXS30 index call option prices for the June, July, and August 2005 option expirations. The option prices are equal to the average of the bid- and ask prices. Time-adjusted moneyness is defined as $[(K/S) - 1]/\sqrt{T}$, where S is the index level, K is the option's exercise price, and T is the option's time to maturity (in days).

The wrong assumption that financial data is normally distributed is thought to explain the existence of smiles and smirks. More specifically, the data's relatively higher skewness and greater kurtosis is thought to explain the existence of smiles and smirks, respectively (Rouah and Vainberg (2007)). These are generally more evident for short-term options than for long-term options, which is equivalent with saying that long-term returns are more normally distributed than short-term returns (Rouah and Vainberg (2007)). The relatively greater kurtosis implies that extreme returns are more likely to occur, meaning that deep in-the-money- and deep out-of-the-money options are more expensive relative to their BS value. The relatively higher skewness in returns is also often shown to be negative, meaning that large negative returns are more likely to occur (which is a feature that is not allowed by the normal distribution). This greater likelihood of large negative returns leads in turn to higher implied volatilities for in-the-money- than for out-of-the-money calls.

2.5. The ad hoc Black-Scholes Model

The existence of the implied volatility curve and the two factors believed to have the greatest effect on its shape, namely moneyness and maturity, have resulted to the development of what

Dumas, Fleming, and Whaley (1998) (henceforth DFW) call the ad hoc BS model.⁸ This model involves using what DFW call a deterministic volatility function (DVF) to model the implied volatility, which is nothing more than a simple quadratic function of moneyness and maturity. Hence the underlying assumption of a constant volatility made in the BS model is dropped by making the volatility dependent on these two.

DFW considers four different structural forms for the DVF, namely

$$\text{Model 0: } \sigma = \max(0.01, a_0);$$

$$\text{Model 1: } \sigma = \max(0.01, a_0 + a_1K + a_2K^2);$$

$$\text{Model 2: } \sigma = \max(0.01, a_0 + a_1K + a_2K^2 + a_3T + a_5KT);$$

$$\text{Model 3: } \sigma = \max(0.01, a_0 + a_1K + a_2K^2 + a_3T + a_4T^2 + a_5KT),$$

where $a_0, a_1, a_2, a_3, a_4, a_5$ are the model parameters. A minimum value of the local volatility rate equalling 0.01 is imposed in each model in order to eliminate possible negative values of fitted volatility. Model 0 is obviously equivalent to the BS constant volatility, followed by gradually more complicated forms. A final model is a cross between Model 1, 2, and 3, depending on the number of different option maturities in the sample on that day. DFW (1998) show that Model 0 leads to the largest valuation errors, which is a result that is consistent with the fact that volatility is not constant across moneyness and maturity.

Christoffersen and Jacobs (2004) compare the pricing errors of the ad hoc BS- and the Heston (1993) model (which is discussed in the next section). They conclude that the ad hoc BS model is the superior one of the two. In addition to this advantage of the ad hoc BS model, it is made even more attractive when considering the ease it requires to implement it. The model parameters are estimated via ordinary least squares (OLS) on a series of BS implied volatilities.

⁸ This same model is referred to as the Practitioner BS model by Christoffersen and Jacobs (2004).

2.6. Stochastic Volatility Models

The behaviour documented in Figures 1 and 2 above can be explained by some option pricing models. The stochastic volatility model of Heston (1993) for example can explain them when a negative correlation exists between the asset price and its volatility (DFW (1998)).

Heston (1993) proposed the following two-factor asset pricing model;

$$\frac{dS(t)}{S(t)} = (\mu - \delta)dt + \sqrt{h(t)}dW(t), \quad (5)$$

$$dh(t) = \kappa(\theta - h(t))dt + \sigma(h)\sqrt{h(t)}dW(t)^h, \quad (6)$$

where W_t and W_t^h are two correlated Brownian motion processes with constant correlation ρ , i.e. $E[dW(t)dW(t)^h] = \rho dt$.⁹ This expression implies that the correlation between the returns $dS(t)/S(t)$ (given by equation (5)) and the changes in the conditional variance $h(t)$ (given by equation (6)) is incorporated in the option pricing model. Equation (6) models the conditional variance $h(t)$ of the percentage change of the stock price as a mean-reverting process. The first term on the right-hand side of the equation, $\theta - h(t)$, is a drift that pulls the conditional variance back to its long-run mean θ at rate κ . The latter parameter gives the speed of the mean reversion in the variance process. Hence the drift becomes negative when the variance at time t is higher than its long-run mean ($h(t) > \theta$), making the variance more likely to decrease over time towards its long-run mean, and vice versa when it is lower than its long-run mean ($h(t) < \theta$). Due to the $\sqrt{h(t)}$ term however, the variance cannot become negative since its own volatility $\sigma(h)$ approaches zero as $h(t)$ decreases (Heston (1993)). Equation (6) is intended to fit the non-constant variance by taking the past into account in a way that the size of the current volatility relates to the size of the past volatilities. This feature is called volatility clustering, meaning that high-volatility periods are followed by high-volatility periods and vice versa. The equation takes also into account a positive autocorrelation of squared log-returns, which is a well-known characteristic known as ARCH-effects.

One alternative approach to characterize a stochastic volatility model is to use the generalized autoregressive conditional heteroskedasticity (GARCH) model (see Bollerslev (1986)).¹⁰

⁹ See Hull (2006) for a discussion about the Brownian motion process.

¹⁰ See Bollerslev et. al. (1992) for a literature review of some important academic studies on ARCH and GARCH modelling in finance.

2.7. The Heston-Nandi Closed-Form GARCH Option Valuation Model

The HN GARCH(p, q) model has two basic assumptions. The first one concerns the process of the logarithmic spot price $\log(S(t))$ of the underlying asset. The second one is needed in order to transform this process to obtain the final risk-neutral version of the HN GARCH(p, q) model (explained in more detail below). Thus the first assumption is that the logarithmic spot price $\log(S(t))$ of the underlying asset (including accumulated interest or dividends, if they are considered) follows the following GARCH process over time steps of length Δ :

$$\log(S(t)) = \log(S(t - \Delta)) + r + \lambda h(t) + \sqrt{h(t)}z(t), \quad (7)$$

$$h(t) = \omega + \sum_{i=1}^p \beta_i h(t - i\Delta) + \sum_{i=1}^q \alpha_i \left(z(t - i\Delta) - \gamma_i \sqrt{h(t - i\Delta)} \right)^2, \quad (8)$$

where

$$\begin{aligned} p, q &\geq 0, \\ \omega &> 0, \\ \beta_i &\geq 0, i = 1, 2, \dots, p, \\ \alpha_i &\geq 0, i = 1, 2, \dots, q. \end{aligned}$$

Equation (7) and (8) is the mean model and the stochastic volatility model, respectively. r is the continuously compounded interest rate for the time interval Δ , λ is a parameter, $h(t)$ is the conditional variance of the log-return between time $t - \Delta$ and t , and $z(t)$ is a standard normal disturbance. Despite the fact that the specification of the volatility model above might seem a bit complex, it is still in the end nothing more than a version of the ordinary GARCH(p, q) model. The appearance of the conditional variance $h(t)$ in the mean model is interpreted as a return premium since it allows the average spot return to depend on the risk level.¹¹ Moreover, the term $\lambda h(t)$ implies that the expected spot return is assumed to exceed the risk-free rate by an amount proportional to the variance $h(t)$.

Since this thesis focuses only on the single lag version of the HN GARCH model, it follows naturally to discuss some of its process properties. The first property is that the first-order

¹¹ λ is assumed to be constant since option prices are very insensitive to this parameter. The functional form of this risk premium, i.e. $\lambda h(t)$, prevents arbitrage by ensuring that the spot asset earns the risk-free interest rate when it itself is risk-free, which is the case when the variance equals zero.

process is stationary if $\beta_1 + \alpha_1 \gamma_1^2 < 1$. Moreover, the conditional variance $h(t + \Delta)$ can be observed as a function of the spot price at time t , i.e.

$$h(t + \Delta) = \omega + \beta_1 h(t) + \alpha_1 \frac{\left(\log(S(t)) - \log(S(t - \Delta)) - r - \lambda h(t) - \gamma_1 h(t)\right)^2}{h(t)}. \quad (9)$$

α_1 and γ_1 determines the kurtosis and the skewness of the distribution of the log returns, respectively. The α_1 parameter being zero implies a deterministic time varying variance, whereas a zero value of the γ_1 (and λ) parameter implies a symmetric distribution. The variance process $h(t)$ and the spot return $\log(S(t))$ are in general correlated as

$$\text{Cov}_{t-\Delta}[h(t + \Delta), \log(S(t))] = -2\alpha_1 \gamma_1 h(t). \quad (10)$$

Given a positive α_1 parameter, a positive value for γ_1 results in a negative correlation between asset price and volatility.

It is not possible to use equations (7) and (8), or (9), above to value options due to the fact that the risk-neutral distribution of the spot price is unknown. This is where the second assumption comes in, which enables the transformation of equation (7) and (8) to their respective final risk-neutral versions. Heston and Nandi (2000) formalize this assumption as Proposition 1, which ensures that the risk-neutral process has the same form as the real process, i.e. as equations (7) and (8) but with λ replaced by $-1/2$ in the mean model and γ_1 replaced by $\gamma_1^* = \gamma_1 + \lambda + 1/2$ in the volatility model. Hence the final risk-neutral versions of equation (7) and (8) above are

$$\log(S(t)) = \log(S(t - \Delta)) + r - \frac{1}{2} h(t) + \sqrt{h(t)} z^*(t), \quad (11)$$

$$h(t) = \omega + \sum_{i=1}^p \beta_i h(t - i\Delta) + \sum_{i=2}^q \alpha_i \left(z(t - i\Delta) - \gamma_i \sqrt{h(t - i\Delta)}\right)^2 + \alpha_1 \left(z^*(t - \Delta) - \gamma_1^* \sqrt{h(t - \Delta)}\right)^2, \quad (12)$$

where

$$z^*(t) = z(t) + \left(\lambda + \frac{1}{2}\right) \sqrt{h(t)},$$

$$\gamma_1^* = \gamma_1 + \lambda + \frac{1}{2}.$$

$z^*(t)$ denotes now a standard normal risk-neutral disturbance.

The price of a European call option at time t with strike price K that expires at time T is given by

$$C = e^{-r(T-t)} E_t^* [\text{Max}(S(T) - K, 0)] = S(t)P_1 - Ke^{-rT}P_2, \quad (13)$$

where $E_t^*[\cdot]$ is the expectation at time t under the risk-neutral distribution and P_1 and P_2 are the risk-neutral probabilities.

P_1 is simply the delta of the call option whereas P_2 is the risk-neutral probability of the asset price being greater than the strike price at maturity, i.e. $P_2 = \text{Pr}[S(T) > K]$. The values of these two risk-neutral probabilities must be calculated in order to obtain the call price. The generating function of the underlying asset price, denoted $f(\emptyset) = E_t[S(T)^\emptyset]$, is also the moment generating function of the logarithm of the spot price $S(T)$. In other words, $E_t[S(T)^\emptyset] = E_t[\emptyset \log S(T)]$. Denoting the generating function for the risk-neutral process in (11) and (12) as $f^*(\emptyset)$, it takes the following log-linear form;

$$f(\emptyset) = S(t)^\emptyset \exp \left(A(t; T, \emptyset) + \sum_{i=1}^p B_i(t; T, \emptyset) h(t + 2\Delta - i\Delta) + \sum_{i=1}^{q-1} C_i(t; T, \emptyset) z(t + \Delta - i\Delta) - \gamma_i \sqrt{h(t + \Delta - i\Delta)} \right)^2, \quad (14)$$

where

$$A(t; T, \emptyset) = A(t + \Delta; T, \emptyset) + \emptyset r + B_1(t + \Delta; T, \emptyset) \omega - \frac{1}{2} \ln(1 - 2\alpha_1 B_1(t + \Delta; T, \emptyset)), \quad (15)$$

$$B_1(t; T, \emptyset) = \emptyset(\lambda + \gamma_1) - \frac{1}{2} \gamma_1^2 + \beta_1 B_1(t + \Delta; T, \emptyset) + \frac{1/2(\emptyset - \gamma_1)^2}{1 - 2\alpha_1 B_1(t + \Delta; T, \emptyset)}, \quad (16)$$

$$B_i(t; T, \emptyset) = \beta_i B_1(t + \Delta; T, \emptyset) + B_{i+1}(t + \Delta; T, \emptyset), \text{ for } 1 < i \leq p, \quad (17)$$

$$C_i(t; T, \emptyset) = \alpha_{i+1} B_1(t + \Delta; T, \emptyset) + C_{i+1}(t + \Delta; T, \emptyset), \text{ for } 1 < i \leq q - 1. \quad (18)$$

By using the following conditions as starting values:

$$A(T; T, \phi) = B_i(T; T, \phi) = C_i(T; T, \phi) = 0, \quad (19)$$

the coefficients for the generating function (14), given by equation (15), (16), (17), and (18), can be derived recursively from time T to t .

Due to that the generating function of the spot price is also the moment generating function of the logarithm of the spot price, it follows that $f(i\phi)$ is the characteristic function of the latter price. In order to use this function, ϕ in equations (15), (16), (17), and (18), must be replaced by $i\phi$ everywhere. Once the generating function is derived, the risk-neutral probabilities (P_1 and P_2) required for the call price can be calculated by inverting this characteristic function.¹² This requires numerical integration since the integrals representing these risk-neutral probabilities cannot be derived analytically. Heston and Nandi (2000) show that these probabilities have the following form;

$$P_1 = \frac{1}{2} + \frac{e^{-rT}}{\pi S_t} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\phi} f^*(i\phi + 1)}{i\phi} \right] d\phi, \quad (20)$$

$$P_2 = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\phi} f^*(i\phi)}{i\phi} \right] d\phi, \quad (21)$$

where $\operatorname{Re}[\cdot]$ denotes the real part of a complex number. The inclusion of these probabilities in equation (13) completes the option valuation formula. Hence the value of a European call option at time t is given by

$$C = e^{-r(T-t)} E_t^* [\operatorname{Max}(S(T) - K, 0)] = \frac{1}{2} S(t) + \frac{e^{-r(T-t)}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\phi} f_{i\phi+1}^*}{i\phi} \right] d\phi - K e^{-r(T-t)} \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\phi} f_{i\phi}^*}{i\phi} \right] d\phi \right), \quad (22)$$

where $E_t^*[\cdot]$ denotes the expectation under the risk-neutral distribution.¹³

¹² See Feller (1971), Kendall and Stuart (1977), or Heston (1993).

¹³ European put options can easily be valued by using the put-call parity.

2.8. Mispricing Error Measures

Three different pricing error measures are implemented in order to compare the different pricing models. The mean outside pricing error (MOE) measures the average pricing error outside the bid-ask spread. It is equal to the difference between the model value and the ask price if the model value exceeds the ask price, and to the difference between the model value and the bid price if the bid price exceeds the model value. In cases where the model price is within the bid-ask spread, the MOE is simply set equal to zero. Thus

$$MOE = \frac{1}{N} \begin{cases} \sum p_o - p_a & \text{if } p_o > p_a \\ \sum p_o - p_b & \text{if } p_o < p_b \\ 0 & \text{if } p_a \leq p_o \leq p_b \end{cases} \quad (23)$$

where N is the number of errors, p_o is the model value, p_a is the ask price and p_b is the bid price. MOE measures whether the model overprices as much as it underprices the option price, or if there are any systematic biases in the mispricing compared to the bid-ask spread. The mean absolute error (MAE) is the average absolute value of the MOE, i.e.

$$MAE = \frac{1}{N} \begin{cases} \sum |p_o - p_a| & \text{if } p_o > p_a \\ \sum |p_o - p_b| & \text{if } p_o < p_b \\ 0 & \text{if } p_a \leq p_o \leq p_b \end{cases} \quad (24)$$

meaning that it does not consider the direction of the pricing error, i.e. whether the price is over- or underestimated. The root mean squared error (RMSE) measures the squared errors between the model- and the market price.

$$RMSE = \sqrt{\frac{1}{N} \sum (p_o - p_m)^2}, \quad (25)$$

where p_m is the market price. The errors are squared before they are averaged, meaning that higher weights are given to larger errors. Hence the RMSE is always larger than the two other measures.

The next section considers the obtained empirical results.

3. Results

3.1. Data Information

The data used in this thesis is daily on the OMXS30 index call options traded on the Stockholm Stock Exchange and the risk-free rate is computed using the 90-day Treasury Bill SSVX. The stock- and the bond index are secondary and obtained from *DataStream* whereas the index options are obtained directly from the exchange. Dividends are not considered since many of the stocks in the OMXS30 index pay dividends only once a year (mostly in Mars or April), meaning that their exclusion should in general not have significant affects on the results.

The daily data set comprising the call options is sampled every Wednesday (or the next trading day if Wednesday is a holiday) for the years 2005 and 2006.¹⁴ The mid-point of the bid-ask quote is used as the option price.

Three criterions are used as filters when sampling the call options: The first one is to only include options with an absolute moneyness less than or equal to ten percent, i.e. $|K/S - 1| \leq 0.1$. The second one is to only include options whose time to maturity is between six- to 100 days.¹⁵ The third and final restriction is that a transaction must satisfy the no-arbitrage relationship (Merton (1973)), meaning that $C \geq S - Ke^{-rT}$ must hold.¹⁶

The data set comprising the call options consists of 2,046 observations. The average number of options per day is 20 with a minimum of three and a maximum of 31. The average bid-ask-spread is 1.70 SEK.

3.2. Estimations

Similar to Heston and Nandi (2000), this thesis focuses only on the single lag version of the HN GARCH model. Unlike their estimations however, the ones here are done by not only using the Normal distribution, but also the Student-t distribution. The GARCH(1,1) process is estimated via the maximum likelihood estimation (MLE) using the daily ($\Delta = 1$) index stock returns

¹⁴ The reason to why Wednesdays are used is that fewer holidays fall on a Wednesday than on any other trading day. In the sample used here, all but one day are Wednesdays.

¹⁵ See DFW (1998) for a further explanation of the exclusionary criteria about moneyness and maturity.

¹⁶ The no-arbitrage relationship (Merton (1973)) for put options implies that $P \geq Ke^{-rT} - S$.

during the period 2005-01-04 – 2006-12-29.¹⁷ Two specifications of the GARCH(1,1) model are estimated in order to analyze the significance of the skewness parameter γ_1 (which was said to capture leverage effects). First an unrestricted model is estimated, followed by a restricted one in which γ_1 is set equal to zero (which is equivalent to a symmetric GARCH). Both these processes are plotted and presented in the figures below. Figure 3A illustrates the process for the restricted model while Figure 3B illustrates the one for the unrestricted.

Figure 3A Restricted GARCH

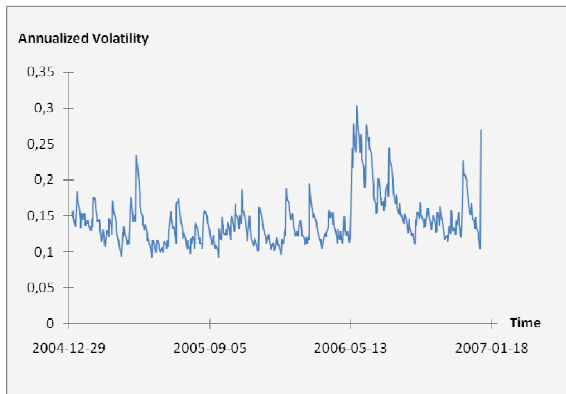



Figure 3B Unrestricted GARCH



Note: Figure 3A and 3B illustrate the daily annualized spot volatility from the restricted-/symmetric- and the unrestricted/asymmetric GARCH model, respectively, during the period 2005-01-04 – 2006-12-29 using daily OMXS30 index stock returns.

These plots indicate a difference between the two processes, emphasizing thereby the significance of the skewness parameter γ_1 . However, since “eyeballing” the data is not a substitute for formally testing for the significance of the parameter, a likelihood ratio (LR) test is made for each respective year and for the whole two-year period. The outcome is given in Table 1 below.

¹⁷ The starting value of the conditional variance $h(0)$ is set to equal the sample variance of the stock returns. The results are quite indifferent to the starting value of $h(0)$ due to the strong mean reversion in volatility.

Table 1

Maximum likelihood estimations using the Normal distribution and likelihood ratio tests

Year		α_1	β_1	γ_1	ω	λ	θ	$\beta_1 + \alpha_1 \gamma_1^2$	Log-Likelihood	LR test (p-value)
2005	GARCH	2.23E-06	0.4828	343.0002	1.05E-05	17.009	11.20%	0.7449	1157.9187	0.0002
	GARCH, $\gamma_1=0$	3.96E-06	0.6215		1.55E-05	0.2996	11.37%	0.6215	1150.7243	
2006	GARCH	1.24E-05	0.5757	124.4700	1.62E-05	8.5926	17.62%	0.7679	1049.9289	0.0002
	GARCH, $\gamma_1=0$	1.85E-05	0.8663		8.79E-07	0.9950	19.10%	0.8663	1041.2404	
2005-2006	GARCH	8.08E-06	0.7782	106.4889	2.83E-06	8.4110	14.53%	0.8698	2193.1971	0.0532
	GARCH, $\gamma_1=0$	1.15E-05	0.8490		1.79E-06	19.0708			2190.2635	

Note: Maximum likelihood estimations of the HN GARCH(1,1) model with $\Delta = 1$ (day) using the OMXS30 stock index returns during the period 2005-01-04 – 2006-12-29 for the unrestricted ($\gamma_1 \neq 0$) and the restricted ($\gamma_1 = 0$) model.

$$\log(S(t)) = \log(S(t - \Delta)) + r + \gamma h(t) + \sqrt{h(t)}z(t),$$

$$h(t) = \omega + \beta_1 h(t - \Delta) + \alpha_1 \left(z(t - \Delta) - \gamma_1 \sqrt{h(t - \Delta)} \right)^2.$$

The log-likelihood function is $\sum_{t=1}^T -0.5 \left(\log(h(t)) + z(t)^2 \right)$, where T is the number of days in the sample used. The likelihood ratio test statistic is computed as $\xi_{LR} = 2 \left[\log L(\hat{\theta}) - \log L(\tilde{\theta}) \right]$, where $\hat{\theta}$ and $\tilde{\theta}$ is the unrestricted- and the restricted ML estimator, respectively. The test statistic has a Chi-squared distribution with J degrees of freedom under the null hypothesis. $\theta = \sqrt{252(\omega + \alpha_1)/(1 - \beta_1 - \alpha_1 \gamma_1^2)}$ is the annualized (252 days) long run volatility implied by the parameter estimates. $\beta_1 + \alpha_1 \gamma_1^2$ measures the degree of mean reversion, where $\beta_1 + \alpha_1 \gamma_1^2 = 1$ implies that the conditional variance process is integrated.

The LR test easily rejects the null hypothesis of a symmetric GARCH for each respective year, implying that the negative correlation between returns and variance is an obvious feature of the stock index during each single year. However, the LR test does not reject the null hypothesis at a five percent significance level for the whole two-year period, i.e. 2005-2006. Apart from these results, the degree of mean reversion (given by $\beta_1 + \alpha_1 \gamma_1^2$) for the unrestricted GARCH process is 0.8698, the volatility of volatility (given by α_1) is 8.08e-06, and the annualized long-run mean of volatility (given by θ) is 14.53%.

3.3. Model Comparisons

This section presents the ML estimates of the HN GARCH(1,1) model using both the Normal- and the Student-t distribution. It gives also a discussion of the in-sample and the out-of-sample differences between the HN GARCH(1,1)-, the BS-, and the ad hoc BS model.

3.3.1. In-sample Model Comparison

The BS model is implemented using a single implied volatility that is estimated across all strikes and maturities on a given day. As was mentioned earlier, the implied volatility of the ad hoc BS model on a given day is modelled using the DVF.¹⁸ Both the BS- and the ad hoc BS model are re-estimated every week while the parameters of the HN GARCH(1,1) model are held constant over the entire estimation period, which are the first six months of each year. This implies that the model is constrained to use the variance from the history of asset prices. The ML estimates under each distribution and the in-sample valuation errors are given in Table 2 below, for both 2005 and 2006.

Table 2
In-sample estimations of the models in the first half of each year

		α_1	β_1	γ_1	ω	λ	θ	$\beta_1 + \alpha_1 + \gamma_1^2$	RMSE	Average price	Observations
2005										11.28	462
BS									5.12		
Ad hoc BS									3.01		
GARCH	Normal	9.91E-06	6.86E-15	63.5136	4.41E-05	11.3403	11.91%	0.0400	14.01		
	Student-t	6.09E-06	6.21E-17	100.7116	4.63E-05	19.8725	11.86%	0.0618	14.11		
2006										15.16	435
BS									7.14		
Ad hoc BS									2.69		
GARCH	Normal	2.00E-05	0.5827	96.2208	1.36E-05	-2.9090	19.09%	0.7676	44.28		
	Student-t	1.98E-05	0.5787	88.1739	1.62E-05	4.5113	18.42%	0.7325	46.03		

Note: BS is the Black-Scholes model which is implemented using a single implied volatility that is estimated across all strikes and maturities on a given day. The implied volatility of the ad hoc BS model on a given day is modelled using the deterministic volatility function. The parameters of the non-updated HN GARCH(1,1) model are estimated via maximum likelihood using the Normal- and the Student-t distribution. Both the BS- and the ad hoc BS model are re-estimated every week while the parameters of the HN GARCH(1,1) model are held constant over the entire estimation period, i.e. for the first six months of each year. RMSE is the root mean squared error (in SEK). Average price is the average option price in the used sample.

Although the long-run annualized volatility is quite similar under the two distributions for both years, it is higher in 2006 (around 19%) than in 2005 (around 12%).¹⁹ This variation in volatility between the two years is verified by the significant difference between the degree of mean reversion, being around five percent in 2005 and as high as around 75% in 2006. Hence this implies that the variation in volatility is much lower in 2005 than in 2006. The average option price is 11.28 SEK and 15.16 SEK in 2005 and 2006, respectively. The RMSE of the BS model is around 5 SEK in 2005 and around 7 SEK in 2006. With a RMSE of around 3 SEK in

¹⁸ See section 2.5. The ad hoc Black-Scholes Model.

¹⁹ These numbers should be treated with caution since options of up to only 100 days to maturity are used.

both years, the ad hoc BS- outperforms the BS model in every year. The reason to the ad hoc BS model's superiority lies in its flexibility to fit both the volatility smile in strike prices and the term structure of implied volatilities. Noteworthy is that the RMSE of the HN GARCH(1,1) model is around the alarming value of 14 SEK in 2005 and 45 SEK in 2006. A possible explanation to this model's larger in-sample RMSE could be given by examining the two figures below. Figure 4A and 4B illustrate the volatility process over the in-sample period of 2005 and 2006, respectively.

Figure 4A In-sample volatility 2005





Figure 4B In-sample volatility 2006



Note: Figure 4A and 4B illustrate the volatility process over the in-sample period of 2005 and 2006, respectively.

While the volatility process is relatively tranquil over the entire sample in 2005, a significant increase in volatility is seen in the last two months in the first half of 2006. This significant increase may affect the estimation of the volatility of volatility parameter α_1 , which is shown to have a great impact on the outcome of the HN GARCH(1,1) model. Heston and Nandi (2000) point out that "...option values are more sensitive to α_1 (that measures the volatility of volatility), and γ_1 (that controls the skewness of index returns) than they are to the other parameters. This stability is important for the GARCH model to fit the data reasonably well even with constant parameters."²⁰

According to Heston and Nandi (2000), the above comparison of the GARCH model with the ad hoc BS model is somewhat "unfair" due to that the latter model is updated every week. They also point out that "it is not clear whether the improved in-sample fit of the ad hoc model stems from a more flexible functional form or from the instability of the functional form of the GARCH process over a long enough time period."²¹ They estimate therefore also an "updated" HN GARCH(1,1) model, allowing its parameters to change every week. Worth noting here

²⁰ See Heston and Nandi (2000) pp. 605.

²¹ See Heston and Nandi (2000) pp. 604.

though is that despite this weekly update of the parameters, the variance $h(t + 1)$ is still drawn from the history of asset prices at time t . Hence the variance is estimated by using the previous 252 days of stock returns.²² This updating is only done in the second half of each year since the final objective is to compare the models' out-of-sample valuation errors in these same periods. The outcome of this "fair" in-sample comparison of the ad hoc BS model with the updated version of the HN GARCH(1,1) model is presented in Table 3.

Table 3
In-sample comparison of the ad hoc BS model and the updated GARCH model

		RMSE	Average Option price	Number of observations
2005			15.14	651
Ad hoc BS		1.92		
GARCH (updated)	Normal	9.25		
	Student-t	9.81		
2006			25.74	541
Ad hoc BS		2.20		
GARCH (updated)	Normal	17.65		
	Student-t	17.87		

Note: RMSE from the weekly estimation using option prices in the second half of each year (including the last Wednesday of the first half of each year) for the ad hoc BS- and the HN GARCH(1,1) model. The ad hoc BS is implemented using ordinary least squares whereas the HN GARCH(1,1) model is estimated via maximum likelihood. The conditional variance $h(t + 1)$ used in the HN GARCH(1,1) model is estimated using the daily history (last 252 days) of the OMXS30 stock index returns.

The average option price in the sample is around 15 SEK and 26 SEK in 2005 and 2006, respectively. Updating the parameters weekly results in a great improvement of the HN GARCH(1,1) model in 2006. However, the mispricing errors are still large in both years. Hence once again, the model is greatly outperformed by the ad hoc BS model, whose RMSE is around 2 SEK in each respective year. One way to check the quality of the updated HN GARCH(1,1) model's estimated coefficients is to analyze and compare their mean and standard deviations with each other. Table 4 presents these calculations, which are estimated via ML using the Normal- and the Student-t distribution.

²² According to Heston and Nandi (2000), the results are basically indifferent whether one uses longer time intervals or not when estimating the conditional variance. This is due to its strong mean reversion. Hence a period of 252 days is therefore satisfying to use here as well.

Table 4**Mean estimates from the updated HN GARCH(1,1) model using maximum likelihood under each distribution**

		α_1	β_1	γ_1	ω
Normal	Mean	1.67E-05	0.5832	106.9314	1.46E-05
	Standard Deviation	3.11E-06	0.0119	10.1059	4.33E-06
Student-t	Mean	1.61E-05	0.5866	105.2341	1.55E-05
	Standard Deviation	3.49E-06	0.0168	13.7136	5.56E-07

Note: The mean and standard deviation of the HN GARCH(1,1) model parameters from the weekly estimation via maximum likelihood using option prices in the second half of each year (including the last Wednesday of the first half of each year).

A striking result under both distributions is that the least stable parameter happens to be α_1 , which is then followed by γ_1 . Recalling Heston and Nandi's (2000) argument concerning these parameters, its large in-sample mispricing error is obviously due to the poor estimates of α_1 and γ_1 .

The next section discusses the out-of-sample differences between all the models.

3.3.2. Out-of-Sample Model Comparison

Since the interest now lies in the models' forecasting performances, the models are implemented using information at time t to value options at time $t + 1$. In other words, the BS model is implemented such that the estimated implied volatility from the current week is used to value options in the next week. The current week's estimated model parameters of both the ad hoc BS- and the updated HN GARCH(1,1) model are used to value options in the next week. A non-updated- and an updated version of the HN GARCH(1,1) model is once again estimated. The implementation of the non-updated version implies keeping its parameters fixed at their in-sample estimates for the particular year. Moreover, the updating of the conditional variance $h(t + 1)$ (which is still drawn from the dynamics of the daily stock returns) is done by first using the same starting variance $h(0)$ as in the in-sample estimation, and then from the entire daily history of stock prices for that year, obtain $h(t + 1)$ for any given time t in the out-of-sample period. This implies that all the out-of-sample computations for this non-updated version of the HN GARCH(1,1) model are based on option prices from the first six months of that year. The computation of out-of-sample option values at time t using the updated HN GARCH(1,1) model implies using the conditional variance $h(t + 1)$ obtained at time $t - 1$. "The important distinction between the out-of-sample implementations is that the non-updated GARCH model predicts options values up to 26 weeks ahead, whereas the BS, ad hoc BS, and

updated GARCH models only predict one week ahead.”²³ The out-of-sample valuation errors for the various models aggregated across the two out-of-sample periods are presented in Table 5 (Panel A) below.

Table 5
Out-of-sample valuation errors

Panel A: Aggregate valuation errors across all years						
		RMSE	MAE	MOE	Average option price	Number of observations
					20.05	1149
BS		9.45	1.15	0.34		
Ad hoc BS		6.66	0.69	0.11		
GARCH (non-updated)	Normal	14.49	1.86	1.59		
	Student-t	15.92	2.06	1.85		
GARCH (updated)	Normal	14.41	1.90	1.64		
	Student-t	14.71	1.83	1.53		
Panel B: Valuation errors by years						
		RMSE	MAE	MOE	Average option price	Number of observations
2005					15.16	625
BS		8.28	0.88	0.25		
Ad hoc BS		6.43	0.67	0.01		
GARCH (non-updated)	Normal	9.72	1.34	1.03		
	Student-t	9.81	1.36	1.11		
GARCH (updated)	Normal	9.48	1.23	0.79		
	Student-t	9.99	1.32	0.99		
2006					25.88	524
BS		10.48	1.42	0.42		
Ad hoc BS		6.89	0.71	0.22		
GARCH (non-updated)	Normal	18.05	2.38	2.15		
	Student-t	20.26	2.77	2.60		
GARCH (updated)	Normal	18.03	2.43	2.27		
	Student-t	18.25	2.48	2.38		

Note: Panel A presents the aggregated out-of-sample RMSE for all models. Panel B presents the out-of-sample RMSE for all models by each year. The BS model is implemented using the estimated implied volatility from the current week to value options in the next week. The current week’s estimated model parameters of both the ad hoc BS- and the updated HN GARCH(1,1) model are used to value options in the next week. The implementation of the non-updated HN GARCH(1,1) model implies keeping its parameters fixed at their in-sample estimates for the particular year and updating the variance from the daily OMXS30 stock index returns. MOE is the mean outside error (in SEK). MAE is the mean absolute error (in SEK).

²³ See Heston and Nandi (2000) pp. 606.

The average aggregated option price is 20.05 SEK. The aggregated RMSE of the BS- and the ad hoc BS model is around 9 SEK and 7 SEK, respectively. The mispricing errors of both versions of the HN GARCH(1,1) model are slightly improved when estimated under the Normal distribution. However, with both versions having an aggregated RMSE of around 15 SEK, they are still outperformed by both the BS- and the ad hoc BS model. Table 5 (Panel B) presents the out-of-sample RMSE for each out-of-sample period. The BS model improves slightly in 2005 and worsens in 2006, whereas the ad hoc BS model stays quite stable in both years. Both versions of the HN GARCH(1,1) model, who are still outperformed in both years, improve in 2005 and worsen in 2006. Hence the poor estimates of α_1 and γ_1 obviously affect the HN GARCH(1,1) model's forecasting performance too. In addition to the RMSE, Table 5 (Panel A and B) also present the MOE and the MAE. Nothing changes concerning the rankings of the models when looking at these error measures, except that it can be concluded that the two versions of the HN GARCH(1,1) model perform better when estimated under the Normal- than under the Student-t distribution. Table 6 gives a detailed analysis of all the models' out-of-sample performances by categorizing their valuation errors (both the RMSE and the MOE in SEK, and the percentage valuation error (% Error)) with respect to the options' moneyness and maturities.

Table 6

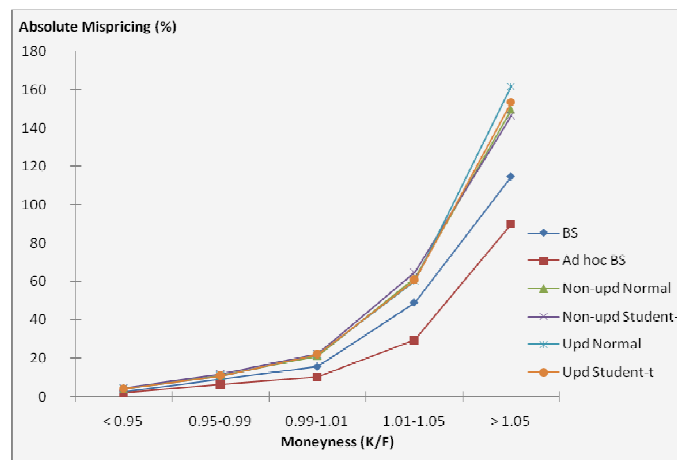
Out-of-sample valuation errors for call options

Model	Moneyness	Days to Expiration									
		< 40			40 - 70			>70			
		RMSE	% Error	MOE	RMSE	% Error	MOE	RMSE	% Error	MOE	
BS											
	< 0.95	1.55	2.35	0.01	2.18	3.57	-1.97	6.12	6.53	-0.50	
	[0.95-0.99)	2.99	8.99	-1.24	3.91	9.37	-1.00	3.36	7.22	-1.14	
	[0.99-1.01)	2.31	15.56	0.28	2.44	10.20	0.28	2.78	9.27	-0.06	
	[1.01-1.05]	2.47	48.96	0.94	3.51	29.98	1.75	4.31	25.28	1.40	
	> 1.05	1.20	114.58	0.71	2.43	58.94	2.00	3.26	48.33	1.47	
Ad hoc BS											
	< 0.95	1.29	1.96	0.75	1.46	2.40	-1.01	3.18	3.39	1.23	
	[0.95-0.99)	2.10	6.32	-0.56	2.41	5.77	-0.52	2.74	5.89	-1.22	
	[0.99-1.01)	1.50	10.12	0.10	1.73	7.24	-0.25	2.38	7.93	-0.28	
	[1.01-1.05]	1.49	29.53	0.22	2.15	18.36	0.17	3.26	19.13	0.78	
	> 1.05	0.94	89.96	0.19	1.90	45.97	0.58	2.57	38.14	0.90	
GARCH (non-updated)											
	Normal	< 0.95	2.51	4.14	0.44	2.48	3.62	-1.07	5.32	5.48	1.27
		[0.95-0.99)	3.74	11.02	0.87	4.25	9.81	1.48	5.59	11.49	1.06
		[0.99-1.01)	3.64	21.14	1.43	4.47	16.36	2.50	5.05	14.91	3.02
		[1.01-1.05]	4.01	61.62	1.28	6.47	42.82	2.86	7.86	38.70	3.28
		> 1.05	1.38	149.56	-0.10	3.50	59.71	1.42	4.67	48.85	2.36
	Student-t	< 0.95	2.73	4.50	0.59	2.75	4.02	0.70	6.18	6.32	1.67
		[0.95-0.99)	4.00	11.74	1.04	4.67	10.72	1.84	6.21	12.68	1.65
		[0.99-1.01)	3.90	22.32	1.74	4.88	17.61	2.72	5.54	16.14	3.04
		[1.01-1.05]	4.35	64.94	1.48	7.07	45.69	3.16	8.63	41.58	3.58
		> 1.05	1.51	146.48	0.01	4.05	65.32	1.71	5.30	53.02	2.79
GARCH (updated)											
	Normal	< 0.95	2.53	4.16	0.68	5.86	10.73	0.89	5.61	5.77	1.46
		[0.95-0.99)	3.66	10.82	0.73	4.35	11.90	1.83	5.73	11.77	1.15
		(0.99-1.01)	3.79	21.92	1.51	4.98	22.95	2.43	5.21	15.33	2.53
		[1.01-1.05]	3.87	60.25	1.13	6.23	41.40	1.87	7.61	36.29	3.27
		> 1.05	1.32	161.63	-0.21	3.32	60.16	1.19	4.58	49.11	2.26
	Student-t	< 0.95	2.53	4.16	0.71	6.10	11.15	1.17	5.56	5.72	1.43
		[0.95-0.99)	3.69	10.87	0.91	4.46	12.11	1.87	5.70	11.66	1.40
		[0.99-1.01)	3.86	22.13	1.67	5.18	23.64	2.42	5.33	15.51	2.82
		[1.01-1.05]	3.97	60.81	1.26	6.44	42.18	2.81	7.86	36.97	3.36
		> 1.05	1.31	153.50	-0.16	3.40	60.11	1.28	4.74	49.63	2.46

Note: Out-of-sample valuation errors categorized with respect to moneyness and maturity. Moneyness is defined as K/F. % Error is the ratio of the RMSE to the average option price for that option category.

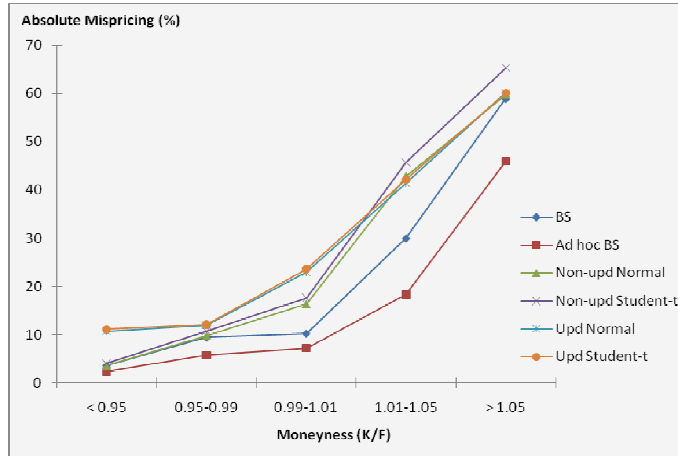
According to the RMSE and the % Error, the ad hoc BS model dominates the other models throughout all the moneyness- and the maturity categories. The three measures indicate that the performance of the non-updated HN GARCH(1,1) model is higher in all the categories under the Normal distribution, except for short-term (< 40 days) deep out-of-the-money ($K/F > 1.05$) options and medium-term (40-70 days to expire) deep in-the-money options where the results are mixed. The performance of the updated version of the model is according to the RMSE in general also higher under the Normal distribution, but with mixed results for the case of short term deep out-of-the-money options. This version of the HN GARCH(1,1) model is however indicated by all the measures to perform better under the Student-t distribution for long-term (>70 days) options in the 0.95-0.99 moneyness category. An interesting result is that according to the RMSE, the updated version performs in general better than the non-updated one under the Student-t distribution, whereas this is not the case under the Normal distribution. Figures 5A, 5B, and 5C, illustrate the percentage out-of-sample pricing errors for options of the three different maturities by moneyness.

Figure 5A Percentage out-of-sample pricing errors




Note: This figure presents the models' percentage out-of-sample pricing errors, defined as $100 \times \text{RMSE}/\text{Option Price}$, for options with a time to maturity of less than 40 days in the second half of each year.

Figure 5B Percentage out-of-sample pricing errors



Note: This figure presents the models' percentage out-of-sample pricing errors, defined as $100 \times \text{RMSE}/\text{Option Price}$, for options with a time to maturity of between 40 and 70 days in the second half of each year.

Figure 5C Percentage out-of-sample pricing errors



Note: This figure presents the models' percentage out-of-sample pricing errors, defined as $100 \times \text{RMSE}/\text{Option Price}$, for options with a time to maturity of more than 70 days in the second half of each year.

4. Conclusions and Discussion

This thesis is based on Heston and Nandi (2000), who present a closed-form discrete-time GARCH (HN GARCH(p,q)) model for pricing European options. The model values options by using the volatilities computed directly from the history of asset prices. It incorporates both the correlation between spot prices and their volatility, and the volatility smile.

The aim of this thesis is to check how the model performs on Swedish data and if there are any significant changes to its performance when estimating it using the Normal- and the Student-t distribution.

There are two major differences concerning the methodology used in this thesis compared to the one Heston and Nandi (2000) use. These concern the data frequency and the estimation method; daily data is used instead of intra-daily, and the HN GARCH(1,1) model is estimated via maximum likelihood (ML) instead of non-linear least squares (NLLS).

The results show that when the HN GARCH(1,1) model is estimated using the Student-t distribution, its out-of-sample valuation performance increases in general when its parameters are updated. This is however not the case when it is estimated using the Normal distribution. The HN GARCH(1,1) model is still shown to suffer from significant mispricing errors as it is greatly outperformed by both the BS- and the ad hoc BS model. This meagre forecasting performance of the model is caused by poor estimates of two of its most important parameters, namely the volatility of volatility- and the skewness parameter. This naturally leads to the questioning of whether good estimates are difficult to obtain in general or not. If this would be the case, then only the ease of implementing the two other models should lead to them being preferred. However, it would be wrong to conclude that the HN GARCH(1,1) model is useless seeing that its estimates are obtained via MLE and not via NLLS as in Heston and Nandi (2000). Apart from this, the results show that the HN GARCH(1,1) model generally performs better when estimated using the Normal- than the Student-t distribution.

Further extensions to consider for future research could be to drop the restrictions made here. The most interesting one is perhaps the estimation method. Another extension could be, as Heston and Nandi (2000) suggest, to estimate the model with multiple lags.

5. References

- Black, F. and Scholes, M. (1973) "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-659.
- Black, F. (1976) "Studies of Stock Price Volatility Changes," *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistic Section*, 177-181.
- Bollerslev, T. (1986) "Generalized Autoregressive Conditional Heteroscedasticity," *Journal of Econometrics*, 31, 307-327.
- Bollerslev, T. Chou, R. Y. and Kroner, K. F. (1992) "ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5-59.
- Christoffersen, P. and Jacobs, K. (2004) "The Importance of the Loss Function in Option Pricing," *Journal of Financial Econometrics*, 72, 291-318.
- Cox, J. and Ross, S. (1976) "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3, 145-166.
- Duffie, D. (1996) "Dynamic Asset Pricing Theory," 2nd edition, Princeton, New Jersey: Princeton University Press.
- Dumas, B. Fleming, J. and Whaley, R. (1998) "Implementing Volatility Functions: Empirical Tests," *Journal of Finance*, 53, 2059-2106.
- Feller, W. (1971) *An Introduction to Probability Theory and Its Applications: Volume 2*, Wiley and Sons, New York.
- Heston, S. (1993) "A Closed-Form Solution for Options with Stochastic Volatility, with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327-343.
- Heston, S. and Nandi, S. (2000) "A Closed-Form GARCH Option Valuation Model," *Review of Financial Studies*, 13, 585-625.
- Heynen, R. (1993) "An Empirical Investigation of Observed Smile Patterns," *Review of Future Markets*, 13, 317-353.
- Hull, J. and White, A. (1987) "The Pricing of Options on Assets with Stochastic Volatility," *Journal of Finance*, 42, 281-300.
- Hull, J.C. (2006) "Options, Futures, and Other Derivatives," 6th edition, *Pearson Prentice Hall*, New Jersey.

Kendall, M. and Stuart, A. (1977) "The Advanced Theory of Statistics," 4th edition, New York: Macmillian.

Merton, R. C. (1973) "Theory of Rational Option Pricing," *The Bell Journal of Economics and Management Science*, 4, 141-183.

Ross, S. (1989) "Institutional Markets, Financial Markets and Financial Innovation," *Journal of Finance*, 44, 541-556.

Rubinstein, M. (1985) "Non-Parametric Tests of Alternative Option Pricing Models Using all Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 Trough August 31, 1978," *Journal of Finance*, 40, 455-480.

Rubinstein, M. (1994) "Implied Binomial Trees," *Journal of Finance*, 49, 771-818.

Rouah, F. D. and Vainberg, G. (2007) "Option Pricing Models and Volatility Using Excel-VBA," John Wiley and Sons, New Jersey.

Scott, L.O. (1987) "Option Pricing when the Variance Changes Randomly: Theory, Estimation and an Application," *The journal of Financial and Quantitative Analysis*, 4, 419-438.

Wiggins, J. B. (1987) "Option Values Under Stochastic Volatility: Theory and Empirical Tests," *Journal of Financial Economics*, 19, 351-372.

Wilmott, P. (1998) "Derivative: The Theory and Practice of Financial Engineering," West Sussex: Wiley.