

Discrete Time Hedging of OTC Options in a GARCH Environment: A Simulation Experiment*

Gustaf E. Hagerud
Department of Finance
Stockholm School of Economics

Working Paper Series in Economics and Finance No. 165
June, 1996

Abstract

This paper examines the effect of using Black and Scholes formula for pricing and hedging options in a discrete time heteroskedastic environment. This is done by a simulation procedure where asset returns are generated from a GARCH(1,1)-t model. In the simulation a hypothetical trader writes an option and then delta-hedges his position until the option expires. It is shown that the variance of the returns on the hedged position is considerably higher in a GARCH(1,1) environment than in a homoskedastic environment. The variance of returns depends greatly on the level of kurtosis in the returns process and on the first-order autocorrelation in centered and squared returns.

*Financial support from Bankforskingsinstitutet is greatly acknowledged. I am deeply indebted to Timo Teräsvirta for valuable comments. Sune Karlsson has given me good suggestion on the simulation procedure. I would also like to thank Magnus Dahlquist, Frank De Jong, Johan Molin, Lars Nordén, and Jens Nystedt for helpful suggestions. The author is solely responsible for errors. Correspondence to Gustaf E. Hagerud, Department of Finance, Stockholm School of Economics, P.O. Box 6501, S-113 83 Stockholm, Sweden. E-mail: fingh@hhs.se.

1 Introduction

It is well known that the price of an option depends on the expected volatility of the underlying asset during the life of the option. For participants in the option markets, it is therefore essential to have reliable volatility predictions, and a great deal of effort is put into forecasting. Many empirical investigations have shown that the variance of returns can be modeled with GARCH processes.¹ Furthermore, volatility forecasts performed with GARCH have proven to be superior to other forecasting methods in a number of studies.² GARCH models are therefore becoming popular among practitioners in the option markets.

When practitioners calculate prices of European options this, they often use formulas where the variance is assumed to be constant during the life of the option, e.g. the Black and Scholes [1973] formula. This practice is followed even though the agents believe that the return process is heteroskedastic. This is primarily done because of calculational convenience. The practitioners know that they are using the wrong model, but do so because they believe that the Black and Scholes formula gives them a good approximation to the option value calculated under an assumption of non-constant volatility.

This paper examines the effect of using the Black and Scholes formula for valuing options when the volatility plugged into the formula is forecasted with a GARCH(1,1) model and when the variance of the asset return in fact follows a GARCH process. It is assumed that there do not exist any other options with the same underlying security, or any other assets correlated with the volatility of the underlying security. This implies that it is impossible to vega-hedge the option, and that the risk-neutral valuation technique of Harrison and Kreps [1979] cannot be used. The problem considered can therefore be compared to the situation a financial intermediary faces when it has written an OTC option on a company's stock for which no exchange traded options exist, and when it wishes to keep the option position unexposed to any stock price movements. The intermediary is forced to continuously hedge its position and the only security it that can use for this purpose is the underlying security.³

The Black and Scholes formula is derived under the assumption that securities prices move continuously in time, and that hedging can be performed at each instant. The GARCH model, however, is defined in discrete time. The investigation perform here can therefore be regarded as a study on how well

¹A introduction to the GARCH literature and a summary of empirical investigations can be found in Bollerslev, Chou and Kroner [1992].

²The effectiveness of volatility forecasts performed with GARCH models has been tested before by e.g. Engle, Hong, Kane and Noh [1993], Engle, Kane and Noh [1993], and Noh, Engle and Kane [1994]. Engle, Hong, Kane and Noh [1993] compared the profitability of using GARCH(1,1) forecasts relative to some other forecasting methods. They find the GARCH(1,1) forecasts were significantly superior for valuing one-day options on the NYSE index during the period 1962 to 1989. Engle, Kane and Noh [1993] perform a similar study, with the difference that the options priced are allowed to have longer maturities than one day. In the test, NYSE-index data from 1968 to 1991 is used. They conclude that pricing of NYSE index options of up to 90 days maturity is more accurate when a GARCH model is used. Noh, Engle and Kane [1994] report similar results when a GARCH forecast is compared to a forecast made from implied volatilities. The data used is S&P 500 index from 1986 to 1991. The evidence in favour of GARCH, however, is not unanimous. Heynen and Kat [1994] show that for a number of stock index series, a *stochastic volatility* model outperforms two GARCH specifications in volatility prediction. The period investigated is 1980 to 1992. A survey of the stochastic volatility class of models is found in Ghysels, Harvey, and Renault [1996].

³The interest rate is assumed to be constant, i.e. no interest risk is considered.

a continuous time pricing formula works in a discrete time environment. This problem has previously been investigated by e.g. Bossaerts and Hillion [1995]. However, their approach is quite different to the one employed here.

This paper does not focus on the issue of the relative effectiveness of GARCH volatility predictions compared to other forecasting methods. Instead, the investigation is on how well the combination of the Black and Scholes formula and GARCH volatility forecasts performs compared to using the Black and Scholes formula when the volatility is constant. A simulation procedure is followed where the variance process, and therefore also the return process is created with a random number generator. The data is generated under the assumption that the variance follows a GARCH(1,1) process. To test the effectiveness of the combination of the Black and Scholes formula and GARCH volatility forecasts, the return on a complete option writing strategy is calculated, complete in the sense that both the initial pricing of the option and the delta-hedging strategy until maturity is included. The delta-hedging is performed on equally spaced time intervals, during the entire life of the option. Deltas are calculated with the Black and Scholes formula using a GARCH volatility forecast. Every day a new volatility forecast is made. Given this volatility forecast, a new delta is calculated and a new delta-hedging decision is made. Transaction costs are assumed to be zero. The simulation is repeated many times to give an estimate of expected return from the tested strategy.

The simulations show that the variance of profits in the heteroskedastic environment is larger than when the volatility is constant. The dispersion increases with the maturity of the option. The average profit is almost the same in both the homo- and heteroskedastic case. The effectiveness of the hedging is largely affected by, (1) the level of kurtosis in the return process, and (2) the first-order autocorrelation in centered and squared returns.

The GARCH(1,1) model is presented in Section 2 which also describes how volatility forecasts are performed. Section 3 discusses the problems concerning the pricing of options in a discrete time heteroskedastic world. Hedging techniques in discrete time under non-constant volatility are described in Section 4. In Section 5 the simulation procedure is presented. The results are given in Section 6 and conclusions in Section 7.

2 The GARCH(1,1)-t Model and Volatility Forecasting

The ARCH family of models was first introduced by Engle [1982] and further developed into Generalized ARCH, GARCH, by Bollerslev [1986]. As mentioned by Bollerslev, Chou and Kroner [1992], one of the most commonly used GARCH models is the GARCH(1,1) model. In the GARCH(1,1) model the conditional variance of asset return at time t , h_t , obeys the process

$$h_t = \gamma + \alpha \varepsilon_{t-1}^2 + \beta h_{t-1}, \quad (1)$$

where ε_t is the error term in the return process that will be assumed to be⁴

$$r_t = \varphi + \varepsilon_t.$$

The error term is assumed to have the form

$$\varepsilon_t = z_t h_t^{1/2}, \quad (2)$$

where z_t is i.i.d. with expected value zero and variance, μ_2 , equal to unity. Most often in the literature z_t is assumed to be normally distributed.

For the variance process to be stationary it is required that $\alpha + \beta < 1$. Furthermore, it is required that $\gamma > 0$, $\alpha > 0$ and $\beta \geq 0$, for the variance process to be positive and the GARCH(1,1) process to be defined.⁵ Teräsvirta [1996] shows that for any $z_t \sim \text{i.i.d.}(0, \mu_2)$ for which $\mu_4 \equiv E[z_t^4] < \infty$ and $\gamma_2 \equiv E[(\beta + \alpha z_t^2)^2] < 1$, the kurtosis of ε_t is given by

$$\kappa_\varepsilon = \kappa_z \frac{1 - \gamma_1^2}{1 - \gamma_2}, \quad (3)$$

where κ_z is the kurtosis of the process z_t , and $\gamma_1 \equiv E[(\beta + \alpha z_t^2)]$. The autocorrelation function of squared ε_t for any $z_t \sim \text{i.i.d.}(0, \mu_2)$ for which $\mu_4 = E[z_t^4] < \infty$ and $\gamma_2 = E[(\beta + \alpha z_t^2)^2] < 1$, following Teräsvirta [1996], is equal to

$$\rho(\varepsilon_t^2, \varepsilon_{t-p}^2) = \frac{\mu_2 \gamma_1^{p-1} \alpha (1 - \beta^2 - \alpha \beta \mu_2)}{1 - \beta^2 - 2\alpha \beta \mu_2} \quad p = 1, 2, \dots \quad (4)$$

Below it is assumed that z_t is distributed Student-t, with ν degrees of freedom, denoted $z_t \sim t(0, 1; \nu)$.⁶ The error term, ε_t , will therefore be conditionally distributed Student-t with expected value zero and variance h_t . The reason for choosing t-distributed innovations, instead of Gaussian, is that a Student-t distribution is more likely to generate returns that resemble empirically observed high-frequency financial time series. Teräsvirta [1996] shows that the GARCH(1,1) model with normal errors is unable to generate data with the level of kurtosis and pattern of first-order autocorrelation in squared residuals empirically observed.

In this case, when $z_t \sim t(0, 1; \nu)$

$$\kappa_z = \mu_4 = \frac{3(\nu - 2)}{(\nu - 4)} \quad \nu > 4 \quad (5)$$

$$\gamma_1 = \beta + \alpha \quad (6)$$

$$\gamma_2 = \beta^2 + 2\alpha\beta + \alpha^2 \kappa_z. \quad (7)$$

⁴In a number of empirical investigations, the return process is estimated as being autoregressive. For example Noh, Engle, Kane [1994] specify the return process as being AR(1). Since an autoregressive return process further complicates option valuation, such return processes will not be considered here. Furthermore, no consideration is given to a conditional mean specification where the conditional variance enters as an explanatory variable.

⁵For a rigorous description of the GARCH(1,1) process see Nelson [1990a] and Teräsvirta [1996].

⁶The GARCH model with conditionally t-distributed errors, denoted the GARCH-t model, was introduced in Bollerslev [1987].

Formula (3) then implies that the kurtosis of ε_t is equal to

$$\kappa_\varepsilon = \kappa_z \frac{1 - (\beta + \alpha)^2}{1 - \beta^2 - 2\alpha\beta - \alpha^2\kappa_z}, \quad (8)$$

given that $\nu > 4$ and $\gamma_2 < 1$.

Furthermore, the autocorrelation of $\{\varepsilon_t^2\}$ when $z_t \sim t(0, 1; \nu)$ is, according to formula (4)

$$\rho(\varepsilon_t^2, \varepsilon_{t-p}^2) = \frac{(\alpha + \beta)^{p-1} \alpha (1 - \beta^2 - \alpha\beta)}{1 - \beta^2 - 2\alpha\beta} \quad p = 1, 2, \dots \quad (9)$$

From equation (9) it can be concluded that the autocorrelation function of squared residuals will be independent of the kurtosis of the i.i.d. process.

One important characteristic of the GARCH model is that the unconditional variance is constant. In the GARCH(1,1) the unconditional variance, σ^2 , is equal to

$$\sigma^2 = \frac{\gamma}{1 - \alpha - \beta}. \quad (10)$$

Given the distributional assumptions of z_t , the parameters of the model can be estimated by maximum likelihood. In the case z_t is distributed Student-t the parameters of the model will be: φ , α , β , γ , and ν . If the model is to be used for prediction of volatility for asset returns, the estimations are most often done with daily observations. Engle, Kane and Noh [1993] show that for the NYSE index, a suitable sample size consists of 1000 observations.

Given estimated parameters $\hat{\alpha}$, $\hat{\beta}$, and $\hat{\gamma}$, the one-day-ahead conditional variance forecast, $\hat{h}_{t+1|t}$, is calculated as

$$\hat{h}_{t+1|t} = \hat{\gamma} + \hat{\alpha}\varepsilon_t^2 + \hat{\beta}h_t,$$

and the s -day-ahead forecast is

$$\begin{aligned} \hat{h}_{t+s|t} &= \hat{\gamma} + \hat{\alpha}E[\varepsilon_{t+s-1}^2 | h_t] + \hat{\beta}\hat{h}_{t+s-1|t} = \\ &= \hat{\gamma} + (\hat{\alpha} + \hat{\beta})\hat{h}_{t+s-1|t}. \end{aligned} \quad (11)$$

In option valuation the average variance over a specific time interval is often of interest. The average predicted variance during the period from $t + 1$ until $t + \tau$ is calculated as

$$\hat{v}_{t,\tau}^2 = \frac{1}{\tau} \sum_{i=1}^{\tau} \hat{h}_{t+i|t}. \quad (12)$$

When this forecast is used in connection with option valuation, it is often rewritten in standard deviation p.a. using the formula

$$\sqrt{250\hat{v}_{t,\tau}^2}, \quad (13)$$

where 250 represents the number of trading days per year.

Recursive substitution on equation (11) gives

$$\begin{aligned} \hat{h}_{t+s|t} &= \hat{\gamma} + (\hat{\alpha} + \hat{\beta}) (\hat{\gamma} + (\hat{\alpha} + \hat{\beta}) \hat{h}_{t+s-2|t}) \\ &= \hat{\gamma} \left[\sum_{i=0}^{s-2} (\hat{\alpha} + \hat{\beta})^i \right] + (\hat{\alpha} + \hat{\beta})^{s-1} \hat{h}_{t+1|t}. \end{aligned} \quad (14)$$

From equations (14) and (10) it can be concluded that

$$\widehat{h}_{t+s|t} \rightarrow \frac{\widehat{\gamma}}{1 - \widehat{\alpha} - \widehat{\beta}} = \widehat{\sigma}^2 \text{ as } s \rightarrow \infty. \quad (15)$$

Thus the variance forecast asymptotically approaches the estimated unconditional variance, $\widehat{\sigma}^2$. The predicted variance for a period far ahead of the current period will therefore be close to the unconditional variance. This is illustrated in Figure 1. Given equations (12) and (15), it is clear that the average volatility for long periods will be close to $\widehat{\sigma}^2$. The length of the forecast needed for the prediction to be close to the unconditional variance depends on the size of $(\widehat{\alpha} + \widehat{\beta})$. If $(\widehat{\alpha} + \widehat{\beta})$ is close to unity, the variance process moves slowly towards $\widehat{\sigma}^2$. If $(\widehat{\alpha} + \widehat{\beta})$ is low the variance process will be close to $\widehat{\sigma}^2$ most of the time. When the true values of α and β are low, the behavior of the process will be similar to that of a homoskedastic process.

In some studies (see e.g. Engle and Bollerslev [1986]), $(\alpha + \beta)$ has been estimated to values indistinguishable from unity. This causes the variance process to be non-stationary. This observation motivated the integrated GARCH (IGARCH) model of Engle and Bollerslev [1986]. In the simulations below, $(\alpha + \beta)$ will in some cases be close to one. Still, problems concerning non-stationarity in the variance process will not be dealt with. Furthermore, only processes where the kurtosis of the return process exists will be studied.

A number of empirical investigations have found a negative correlation between the variance and the return on some assets.⁷ This is one of the motivations for the exponential GARCH (EGARCH) model proposed by Nelson [1991]. From the specification of the GARCH(1,1) process (1) one can conclude that the conditional variance is independent of the path of the stock price. It is only the absolute value of the error term that affects the variance. The issue of correlation between the volatility and sign of the return is therefore not addressed in this study.

3 The Pricing of Options in a GARCH(1,1) Environment

When GARCH(1,1) volatility forecasts are used for options valuation, one common practice is that the predicted average volatility, calculated with equation (13), is plugged into the Black and Scholes [1973] formula, see e.g. Engle, Hong, Kane and Noh [1993].⁸ The purpose of this study is to evaluate how this practice affects the pricing and hedging of options. From the description in the previous section, it is evident that two of the crucial assumptions used in the Black and Scholes derivation are violated in the GARCH(1,1) world. First, in the GARCH(1,1) model, world time is discrete. Second, volatility is not constant. In this section, a few comments will be made regarding these two violations.

In this study it will be assumed that stock prices move at discrete moments in time. No assumptions will be made concerning the preferences of the agents, except that they prefer more to less. If the risk-

⁷See Nelson [1991] for a short survey.

⁸This practice can be motivated in some way by a result presented by Hull and White [1987]. They show that when the price of an asset and its instantaneous variance follow two independent geometrical Brownian motions, the price of a European option on the asset will be equal to the Black and Scholes price, with the variance set equal to the average variance during the life of the option.

neutral valuation techniques of Harrison and Kreps [1978] should be used, it is required that agents at all times can hedge themselves perfectly against all risk factors. In the GARCH(1,1) model, the conditional variance of the stock in the next period is deterministic, as can be seen from equation (1). Therefore the only risk factor to consider is the stock price movement, i.e. the delta of the option. However, since the stock price on the next discrete point in time can take an uncountable number of values, it is impossible to hedge the price risk perfectly. This is a general problem in all discrete time models which are not of the binomial type. Thus, even if the variance were constant, arbitrage arguments alone could not be used to price options. To price contingent claims in an environment where prices move discretely and where the stochastic process can give rise to more than the two possible prices in the next period, behavioral assumptions about the agents in the economy need to be made. More precisely, the agents attitude toward risk needs to be known. This approach is taken by Duan [1995] when he develops an option pricing model in a discrete time GARCH(1,1) environment. To obtain an option pricing expression, Duan has to make very strong assumptions about the utility functions of agents and /or about the distribution of aggregate consumption in the economy.

Instead of assuming that stock prices move discretely, the discrete events in the GARCH(1,1) model could be viewed as representing a countable sample from the uncountable event space generated by a continuous stochastic process. For the sake of argument, assume that this stochastic process is of conceptually the same form as the GARCH process, such that there is only one source of risk in the economy and that the non-constant variance is given by a deterministic function of the past price trajectory. If that were the case, and hedging could be performed continuously, the model would be complete, and risk-neutral valuation techniques could be used to price contingent claims. Generally, however, nothing is known of the form of the continuous process that, when sampled once a day, generates data with GARCH(1,1) characteristics. Moreover, since the focus here is on the GARCH(1,1) process in discrete time, this insight cannot help to solve the problem considered in this paper. A topic related to these issues is the estimation of the instantaneous conditional variance in a diffusion process by a GARCH model estimation. This problem is thoroughly investigated by Nelson [1992].

Before ending this section, it might be appropriate to make some mention of the consequences of letting the length of the time steps in the GARCH model approach zero. This issue was first considered by Nelson [1990b]. In his article, Nelson presents tools for investigating the relationship between stochastic difference equations and Ito processes. Nelson then applies these techniques to two examples of ARCH models, the GARCH(1,1) model and the AR(1)-EGARCH model of Nelson [1991]. The continuous time limit of the GARCH(1,1) model is shown to have a process for the conditional variance with a stationary distribution that is the inverted gamma.

4 Hedging of Options in a GARCH(1,1) Environment

The previous section discussed pricing of options and it was noted, as done previously in the literature, that options cannot be priced using risk-neutral techniques when heteroskedasticity is present.⁹ In this section, a few words will be said about hedging in the GARCH(1,1) environment.

Hedging an option entails reducing exposure to different risk factors. The factors typically correspond to different random variables. A *hedge portfolio* has zero exposure to some or all of these factors. If the random variables move continuously, the hedging should ideally be performed continuously. If the random variables move discretely, hedging should be done at each time step.

The amount of hedging needed to reduce the risk of each factor can be expressed using *hedge parameters*. If the variance of the underlying asset is assumed to be constant, only one hedge parameter, the option delta, needs to be considered. The delta used in this context will henceforth be called the Black and Scholes delta, Δ_{BS} . Δ_{BS} can be calculated as the first derivative of the Black and Scholes formula with respect to the underlying security price¹⁰

$$\Delta_{BS} = \frac{\partial FBS}{\partial S}.$$

If it is assumed that the volatility may change once and then stay at the new level until the expiration of the option, the Black and Scholes vega can be used as a hedge parameter

$$\Lambda_{BS} = \frac{\partial FBS}{\partial v}.$$

Further, the Black and Scholes gamma can be written

$$\Gamma_{BS} = \frac{\partial^2 FBS}{\partial S^2}.$$

When the volatility is stochastic, extended hedge parameters have to be developed. These can be derived using Taylor series expansion. Following Engle and Rosenberg [1994], the discrete time delta is then given by

$$\Delta_{SV} = \Delta_{BS} + \Lambda_{BS} \frac{1}{2\hat{v}} \frac{\partial \hat{v}^2}{\partial S}, \quad (16)$$

where \hat{v} denotes the forecasted variance. In the GARCH(1,1) environment, when the first derivative of variance with respect to price is zero, the stochastic volatility delta will, according to equation (16), be equal to the Black and Scholes delta, $\Delta_{SV} = \Delta_{BS}$. This will also be true for gamma, i.e. $\Gamma_{SV} = \Gamma_{BS}$.

To create a delta-neutral hedge in the GARCH(1,1) environment, the Black and Scholes delta can be used. If an agent is short the option, the hedging is done by buying an amount delta of the underlying security. Since vega and gamma for the underlying security is equal to zero, the underlying security cannot be used for vega- and gamma-hedging. To hedge these two factors, other derivative securities with non-zero vega and gamma are needed. In this study, it is assumed that no other derivative securities exist, and consequently no vega and gamma hedging can be performed.

⁹One can use risk-neutral methods if there exists a traded security that is perfectly correlated with the volatility, but this is generally not the case.

¹⁰Note that in a discrete environment, the hedge parameters will not give a perfect hedge but rather one which on average reduces risk.

5 Simulation Experiment

As noted in the introduction, the object of this study is to try to evaluate the effectiveness of using GARCH(1,1) volatility prediction in combination with the Black and Scholes formula, when the return process is heteroskedastic. This is attempted by using a simulation procedure. It will be assumed that the variance of return follows a GARCH(1,1) process and that the agents know the true parameters of the process. The simulation procedure is performed in the following steps:

1. At time zero, a hypothetical trader writes a European call option. Different maturities and strike prices are considered. The premium is deposited at the risk free rate, assumed to be constant and equal to zero. The option is priced using the Black and Scholes formula. The volatility is set equal to a forecasted average volatility. The forecast is performed using a GARCH(1,1) model. The parameters of the model (1) will be

$$\begin{aligned}\alpha &= 0.0204 \\ \beta &= 0.9700 \\ \gamma &= 4.31 \cdot 10^{-7} \\ \nu &= 5\end{aligned}$$

The level of conditional volatility at date zero is set randomly. This is done by letting a random number generator create a return series for 250 days, assuming that the return follows the specified heteroskedastic process. At the beginning of this return series, the conditional variance is set to the unconditional variance. The last conditional variance in the series is the level assumed to prevail when the option is written.

2. After having written the option, the trader hedges the position by buying Δ_{BS} of the underlying security, where the delta is calculated with the Black and Scholes formula using the forecasted volatility.
3. Given the conditional variance for the next period, $h_{t+1|t}$, the return for the next period is generated given its distribution, $r_{t+1} \sim t(\varphi, h_{t+1|t}; 5)$.¹¹ φ is chosen to be zero for convenience.
4. From the new stock price calculated in step 3, a new delta is calculated.¹² If this delta is different from the previous one, the hedge will be adjusted by either selling or buying the underlying security.
5. Step 3 and 4 are repeated four times every day, i.e. the stock moves four times per day and hedging is performed each time the stock price moves. Note that the conditional variance and therefore also the volatility forecast is constant during each day.

¹¹Note that in the GARCH(1,1) model the forecasted conditional variance, $\hat{h}_{t+1|t}$, is equal to the true conditional variance, h_{t+1} .

¹²When the deltas are calculated during the day, the time to expiration is specified in days and fractions of a day. This proved to be important for obtaining correct deltas when the time to expiration was short.

6. When one day has passed the error term, ε_t , for that day is calculated. Given this error term and given the current conditional variance, the conditional variance for the next day is calculated using equation (1). From this new conditional variance, the stock movements for the following day can be generated. Furthermore, a new volatility forecast is calculated.
7. Steps 3 to 6 are repeated until the option expires when the position is closed.
8. All the cash flows generated in the steps above are summarized and recorded.
9. The simulation is repeated 1000 times to give an estimate of the expected value and the variance of the return from the investigated strategy.

It should be noted that even if the price of the underlying asset does not change over time, the option price will. It will therefore be necessary to hedge the position even when the price of the underlying asset is unchanged. However, this change will usually be minimal.

The parameter values used in the simulation have been obtained from an estimation on 900 daily returns for the S&P 500 index from November 6, 1991 to July 11, 1995.¹³ In the estimation, the degrees of freedom were estimated to 4.6, but for simplicity ν is set to an integer value. The parameters give the unconditional variance $4.49 \cdot 10^{-5}$, corresponding to a standard deviation p.a. equal to 10.6 percent. The kurtosis implied by the parameters is 10.90, which is calculated using formulas (5) and (8). Furthermore, the first autocorrelation of squared residuals will, according to formula (9), be 0.04. Even though $\rho(\varepsilon_t^2, \varepsilon_{t-1}^2)$ is relatively low, the level of kurtosis and the level of autocorrelation are similar to values obtained in empirical investigations, as can be seen in Teräsvirta [1996].

The results obtained in the simulation procedure are compared to results obtained in simulations where the return process is homoskedastic. These simulations also follow the eight steps specified above, but with constant volatility. Returns are generated from a model where the error term is distributed Student-t with the same degrees of freedom as in the heteroskedastic case. The variance is held constant and equal to the unconditional variance for the GARCH(1,1) process. When the options are priced and when the deltas are calculated, the true volatility is used in the Black and Scholes formula. The level of kurtosis in these returns is equal to that of z_t , which can be calculated using formula (5). Thus, when ν is equal to five, the kurtosis will be equal to nine.

6 Results

Simulations with options of six different maturities have been performed: 1M, 2M, 3M, 4M, 5M, and 6M. The results from simulations with at-the-money call options are presented in Table 1.¹⁴ In the table, rows labelled GARCH(1,1) give the results from simulations with heteroskedastic return processes, and rows labelled Homoskedastic give the results from simulations where the variance of returns is constant. The first two lines report average profits of the hedging procedure as percentage of the initial stock price.

¹³I am deeply indebted to Tobias Rydén and Stefan E. Åsbrink for providing us with the estimated parameters.

¹⁴Since the interest rate is equal to zero, the at-the-money options will also be at-the-money-forward.

The reported values indicate that the average profit is almost the same for the homo- and heteroskedastic cases, and close to zero. This agrees with the finding of Engle and Rosenberg [1994] that $\Delta_{SV} = \Delta_{BS}$.¹⁵

The rows labelled Standard deviation in Table 1 report the dispersion of profits as percentage of initial stock price. In the constant volatility case, the standard deviation is at the same level for all maturities, whereas the standard deviation in the GARCH(1,1) cases increases with the time to maturity. For the 1M option, the standard deviation is 10 percent higher in the GARCH(1,1) case than in the constant volatility case. This difference increases to 32 percent for the 2M option, and further to 95 percent for the 6M option.

Note that the figures reported in Table 1 are based on profits in percent of initial stock prices.. It might appear to be more natural to report the figures in percent of the initial option prices. However, since the initial option prices will vary with each iteration in the GARCH(1,1) case, such a procedure was found to be less suitable. To give an indication of how the figures compare to the option prices, the initial option prices in the homoskedastic case are given on the last row of Table 1. These prices are also the median prices in the heteroskedastic case. The figures show that the standard deviations reported are relatively large compared to the option prices. For example, the standard deviation of 0.20 percent reported for the 3M option in the homoskedastic case, constitutes 9 percent of the option price.

Figures 2 to 4 plot the profits from the simulations performed on options with maturities of one, three and six months, respectively. The graphs clearly show that the variance of return increases with the maturity of the option.

Since the returns generated in the constant variance case are distributed Student-t, the standard deviation would be expected to be larger in these simulation than in simulations done with normally distributed returns. To test this, a simulation was performed for a three month option with normal errors. The variance of returns was also in this case set equal to the unconditional variance of the GARCH(1,1) process. The kurtosis of this return series, like the kurtosis of any normally distributed variable will be equal to three. The standard deviation in this case was equal to 0.11, which should be compared to 0.20 for the 3M option in the Student-t case. It can therefore be concluded that with Student-t distributed errors, the standard deviation in profits is almost double that in the case with normal errors.

Table 2 contain the results from four simulations done with call options that are in-the-money and out-of-the-money. The maturity of the options is three months. The in-the-money options are 10 percent in the money and the out-of-the-money options are 10 percent out of the money. The results have a similar pattern as in the at-the-money case. For the in-the-money options, the standard deviation is 100 percent higher in the GARCH(1,1) simulation than in the constant volatility simulation. The difference for the out-of-the-money option is 171 percent. The difference in standard deviation is therefore larger in

¹⁵Engle and Rosenberg [1995] examine whether $\Delta_{SV} = \Delta_{BS}$ is also true in practice. They estimate risk factors using Monte Carlo techniques when the variance is assumed to follow a GARCH process. Their study is performed on data from four different markets: S&P 500 index, bond index futures, weighted foreign exchange rate index, and oil futures. They assume that the error term is distributed student-t, and allows the variance to be negatively correlated with the level of return. They find that GARCH deltas are similar to those calculated with the Black and Scholes formula. GARCH gammas are found to be significantly higher than the Black and Scholes gammas.

these cases than in the at-the-money case, where the difference is 50 percent for the three month options. However, it should be noted that the *level* of standard deviation is lower when the option is not written at-the-money.

The test of this section, will present results from simulations done for testing how the standard deviation of profits is dependent on the parameter values in the GARCH(1,1)-t model. The simulations are done with European at-the-money call options with an initial maturity of three months.

The first test determines how sensitive the results are to the level of the parameter γ . Table 3 contains these results. The first two lines simply restate the results given in Table 1, whereas the last two lines give the results from two simulations done where γ is twice as high as before. Since α and β are unchanged, the kurtosis and the autocorrelation in squared residuals will be unchanged. However, the increase in γ will double the unconditional variance. As expected, the increase in the unconditional variance increases the standard deviation both in the stochastic case and in the constant volatility case. However, the relative difference is almost constant, at 50 percent when γ is equal to $4.31 \cdot 10^{-07}$, and at 53 percent when γ is equal to $8.62 \cdot 10^{-07}$. This finding is of particular interest to agents in markets where volatility is much higher than for the S&P 500 index.

Second, tests how sensitive the results are to the level of kurtosis in the return process are performed. The results from these simulations are shown in Table 4. Row one restates the results from Table 1. The following two rows show the results when the level of kurtosis is lowered by increasing the degree of freedom in the Student-t distribution. In the last row, the degree of freedom is equal to infinity, which implies that the innovations are in fact normally distributed. Since α and β are unchanged, the autocorrelation in squared residuals will also be unchanged. As can be seen in the last column, the standard deviation in profits decreases as the level of kurtosis falls.

Third, it is tested how the level of first-order autocorrelation in squared residuals influences the distribution of profits. This has been done in a number of simulations where the parameter values have been changed such that the first-order autocorrelation in squared residuals has increased, but leaving the unconditional variance and kurtosis unchanged. The results from these simulations are shown in Table 5. The values in the last column indicate that the dispersion of profits increases when the level of autocorrelation in squared residuals increases, though not monotonically. This result is of major importance since, as noted above, the level of first-order autocorrelation in the original simulations, 0.04, is relatively low.

Finally, it is investigated how the level of persistence in the conditional variance, measured by $\alpha + \beta$, affects the distribution of profits. In the simulations above, $\alpha + \beta$ has always been close to 0.99. That will give a half-life for the conditional variance of approximately 70 trading days. Table 6 shows the results from the original simulation and three other simulations where the parameters values are changed so that the level of persistence has decreased. The half-life on the last three rows are approximately 15, 11, and 10 days, respectively. As can be seen in columns six and seven, the level of kurtosis and autocorrelation in squared residuals are kept almost constant. The results presented in Table 6 indicate that the level of persistence in the conditional variance have no major effect on the standard deviation in profits.

7 Conclusion

The objective of this simulation study has been to investigate if it is recommendable to follow the practice of calculating option prices and deltas with the Black and Scholes formula when the return process follows a GARCH(1,1) process. In Section 6 it was shown that the variance of return was higher when the return followed a GARCH(1,1) process rather than a constant volatility process. The average profit, however, was almost the same in both cases. The results for in-the-money options and out-of-the-money options pointed in the same direction. Furthermore, it is shown that the dispersion of profits increases with the level of kurtosis and the level of first-order autocorrelation in squared residuals.

The findings mentioned above imply that an agent who is not risk-neutral and believes that the return follows a GARCH(1,1) process must compensate by charging a higher premium when writing options. Two good indicators of the size the risk in the hedging are, (i) the kurtosis and (ii) the level of first-order autocorrelation in squared residuals.

References

- Black, Fischer, and Myron Scholes [1973], "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, 637-654.
- Bollerslev, Tim [1986], "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307-327.
- Bollerslev, Tim [1987], "A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *The Review of Economics and Statistics*, 9, 542-547.
- Bollerslev, Tim, Ray Chou, and Kenneth Kroner [1992], "ARCH Modeling in Finance," *Journal of Econometrics*, 52, 5-59.
- Bossaerts, Peter, and Pierre Hillion [1995], "Local Parametric Analysis of Hedging in Discrete Time," Working Paper at INSEAD.
- Duan, Jin-Chuan [1995], "The GARCH Option Pricing Model," *Mathematical Finance*, 5, 13-32.
- Engle, Robert F. [1982], "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987-1007.
- Engle, Robert F., and Tim Bollerslev [1986], "Modelling the Persistence of Conditional Variances," *Econometric Review*, 5, 1-50.
- Engle, Robert F., Che-Hsiung Hong, Alex Kane, and Jaesun Noh [1993], "Arbitrage Valuation of Variance Forecasts with Simulated Options," in *Advances in Futures and Options Research*, 6, 393-415.
- Engle, Robert F., Alex Kane, and Jaesun Noh, [1993], "Index-Option Pricing with Stochastic Volatility and the Value of Accurate Variance Forecasts," Discussion Paper, UC-San Diego 93-43.
- Engle, Robert F., and Joshua Rosenberg [1994], "Hedging Options in a GARCH Environment: Testing the Term Structure of Stochastic Volatility Models," NBER Working Paper Series no. 4958.
- Engle, Robert F., and Joshua Rosenberg [1995], "GARCH Gamma," *Journal of Derivatives*, Summer 1995, 47-59.
- Ghysels, Eric, Andrew Harvey, and Eric Renault [1996], "Stochastic Volatility," in G.S. Maddala and C.R. Rao (eds.) *Handbook of Statistics*, Vol. 14: Statistical Methods in Finance, 119-182, Amsterdam: Elsevier.
- Harrison, J.M., and D. Kreps [1979], "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory* 20, 381-408.
- Heynen, Ronald and Harry M. Kat [1994], "Volatility Prediction: A Comparison of the Stochastic Volatility, GARCH(1,1), and EGARCH(1,1) Models," *Journal of Derivatives* Winter 1994, 50-65.

- Hull, John, and Alan White, [1987], "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 281-300.
- Nelson, Daniel B. [1990a], "Stationarity and Persistence in the GARCH(1,1) Model," *Econometric Theory*, 6, 318-334.
- Nelson, Daniel B. [1990b], "ARCH Models as Diffusion Approximations," *Journal of Econometrics* 45, 7-38.
- Nelson, Daniel B. [1991], "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 347-370.
- Nelson, Daniel B. [1992], "Filtering and Forecasting with Misspecified ARCH Models I," *Journal of Econometrics* 52, 61-90.
- Noh, Jaesun, Robert F. Engle, and Alex Kane [1994], "Forecasting Volatility and Option Prices of the S&P 500 Index," in *Journal of Derivatives*, Fall 1994, 17-30.
- Teräsvirta, Timo [1996], "Two Stylized Facts and the GARCH(1,1) Model," Working Paper in Economics and Finance No. 96, Stockholm School of Economics.

Table 1. Results for at-the-money call options

The table shows results from simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. The rows labeled Mean return report average profit as percentage of the initial stock price, and the rows labeled Standard deviation report the dispersion of profits as percentage of initial stock price. The rows labeled GARCH(1,1) give the results from simulations with heteroskedastic return processes, with parameter values $\alpha = 0.0204$, $\beta = 0.970$, $\gamma = 4.31 \cdot 10^{-7}$ and $\nu = 5$. The rows labeled Homoskedastic report the results from simulations with a constant volatility return process. The last row reports initial option prices calculated with the Black and Scholes formula in the homoskedastic case. The number of trading days for the different maturities is: 21, 42, 63, 83, 104, and 125.

	Variance	Maturity of option					
	Procees	1M	2M	3M	4M	5M	6M
Mean return	GARCH(1,1)	0.01	0.01	0.03	0.02	0.02	0.01
	Homoskedastic	0.01	0.00	0.01	0.00	0.00	0.01
Standard deviation	GARCH(1,1)	0.21	0.25	0.30	0.32	0.35	0.39
	Homoskedastic	0.19	0.19	0.20	0.22	0.21	0.20
Initial opt. price	Homoskedastic	1.22	1.73	2.12	2.43	2.73	2.99

Table 2. Results for in-the-money and out-of-the-money call options

The table shows results from simulations done with options that initially were 10 percent in-the-money, and 10 percent out-of-the-money. 1000 simulations are performed for each option investigated. The rows labeled Mean return report average profit as percentage of the initial stock price, and the rows labeled Standard deviation report the dispersion of profits as percentage of the initial stock price. The rows labeled GARCH(1,1) give the results from simulations with heteroskedastic return processes, with parameter values $\alpha = 0.0204$, $\beta = 0.970$, $\gamma = 4.31 \cdot 10^{-7}$ and $\nu = 5$. The rows labeled Homoskedastic report the results from simulations with a constant volatility return process. The maturity of the options is three months, which corresponds to 63 trading days.

	Variance Procees	In-the-money	Out-of-the-money
Mean return	GARCH(1,1)	-0.01	0.01
	Homoskedastic	0.00	0.01
Standard deviation	GARCH(1,1)	0.08	0.12
	Homoskedastic	0.04	0.07

Table 3. Response to a higher γ

The table shows results from simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. The rows labeled GARCH(1,1) report the results from simulations with heteroskedastic return processes, and the rows labeled Homoskedastic report the results from simulations with a constant volatility return process. σ^2 is calculated with formula (10). The standard deviation figures represent the dispersion of profits as percentage of the initial stock price. The maturity of the options is three months, which corresponds to 63 trading days.

Variance	Parameter values					St.
Process	α	β	γ	ν	σ^2	dev.
GARCH(1,1)	0.0204	0.970	$4.31 \cdot 10^{-7}$	5	$4.49 \cdot 10^{-5}$	0.30
Homoskedastic				5	$4.49 \cdot 10^{-5}$	0.20
GARCH(1,1)	0.0204	0.970	$8.62 \cdot 10^{-7}$	5	$8.98 \cdot 10^{-5}$	0.43
Homoskedastic				5	$8.98 \cdot 10^{-5}$	0.28

Table 4. Response to the level of kurtosis

The table shows results from simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. σ^2 is calculated with formula (10), κ_ε with formulas (5) and (8), and $\rho(\varepsilon_t^2, \varepsilon_{t-1}^2)$ with formula (9). The standard deviation figures represent the dispersion of profits as percentage of the initial stock price. The maturity of the options is three months, which corresponds to 63 trading days.

Parameter values							St.
α	β	γ	ν	σ^2	κ_ε	$\rho(\varepsilon_t^2, \varepsilon_{t-1}^2)$	dev.
0.0204	0.970	$4.31 \cdot 10^{-7}$	5	$4.49 \cdot 10^{-5}$	10.9	0.04	0.30
0.0204	0.970	$4.31 \cdot 10^{-7}$	6	$4.49 \cdot 10^{-5}$	6.7	0.04	0.27
0.0204	0.970	$4.31 \cdot 10^{-7}$	∞	$4.49 \cdot 10^{-5}$	3.1	0.04	0.19

Table 5. Response to the level of first-order autocorrelation in squared residuals

The table shows results from simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. The rows labeled GARCH(1,1) report the results from simulations with heteroskedastic return processes, and the rows labeled Constant Volatility report the results from simulations with homoskedastic return processes. σ^2 is calculated with formula (10), κ_ε with formulas (5) and (8), and $\rho(\varepsilon_t^2, \varepsilon_{t-1}^2)$ with formula (9). The standard deviation figures represent the dispersion of profits as percentage of the initial stock price. The maturity of the options is three months, which corresponds to 63 trading days.

Parameter values							St.
α	β	γ	ν	σ	κ_ε	$\rho(\varepsilon_t^2, \varepsilon_{t-1}^2)$	dev.
0.0204	0.970	$4.31 \cdot 10^{-7}$	5	$4.49 \cdot 10^{-5}$	10.9	0.04	0.30
0.0407	0.950	$4.16 \cdot 10^{-7}$	6	$4.49 \cdot 10^{-5}$	10.9	0.12	0.35
0.0505	0.940	$4.26 \cdot 10^{-7}$	7	$4.49 \cdot 10^{-5}$	10.9	0.16	0.47
0.0594	0.930	$4.75 \cdot 10^{-7}$	8	$4.49 \cdot 10^{-5}$	10.9	0.19	0.43

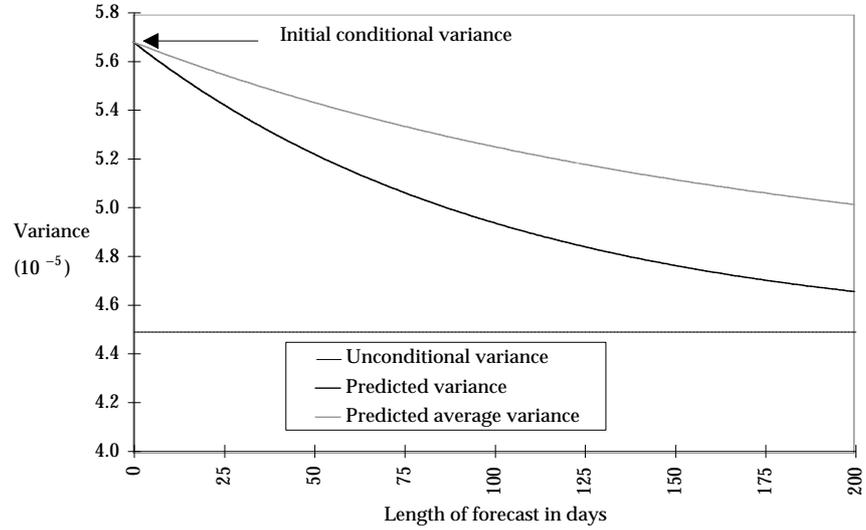
Table 6. Response to the level of persistence in the conditional variance

The table shows results from simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. σ^2 is calculated with formula (10), κ_ε with formulas (5) and (8), and $\rho(\varepsilon_t^2, \varepsilon_{t-1}^2)$ with formula (9). The standard deviation figures represent the dispersion of profits as percentage of the initial stock price. The maturity of the options is three months, which corresponds to 63 trading days.

Parameter values							St.
α	β	γ	ν	σ^2	κ_ε	$\rho(\varepsilon_t^2, \varepsilon_{t-1}^2)$	dev.
0.0204	0.970	$4.31 \cdot 10^{-7}$	5	$4.49 \cdot 10^{-5}$	10.9	0.04	0.30
0.0442	0.910	$20.56 \cdot 10^{-7}$	5	$4.49 \cdot 10^{-5}$	10.9	0.06	0.34
0.0503	0.890	$26.80 \cdot 10^{-7}$	5	$4.49 \cdot 10^{-5}$	10.9	0.07	0.32
0.0531	0.880	$29.85 \cdot 10^{-7}$	5	$4.49 \cdot 10^{-5}$	10.9	0.07	0.30

Figure 1. Predicted variance and predicted average variance as a function of the forecasting horizon

The forecasts are made with a GARCH(1,1) model. Parameter values are $\alpha = 0.0204$, $\beta = 0.970$ and $\gamma = 4.31 \cdot 10^{-7}$. The unconditional variance is equal to $4.49 \cdot 10^{-5}$, and the initial conditional variance is equal



to $5.70 \cdot 10^{-5}$.

Figure 2. Distribution of profits from writing an at-the-money option with a maturity of one month and delta hedging it until expiry

The figure shows results from two simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. The curve labeled GARCH(1,1) gives the results from simulations with heteroskedastic return processes, with parameter values $\alpha = 0.0204$, $\beta = 0.970$, $\gamma = 4.31 \cdot 10^{-7}$ and $\nu = 5$. The curve labeled Constant Volatility gives the results from simulations with homoskedastic return processes. The maturity of the options is one month, which corresponds to 21 trading days.

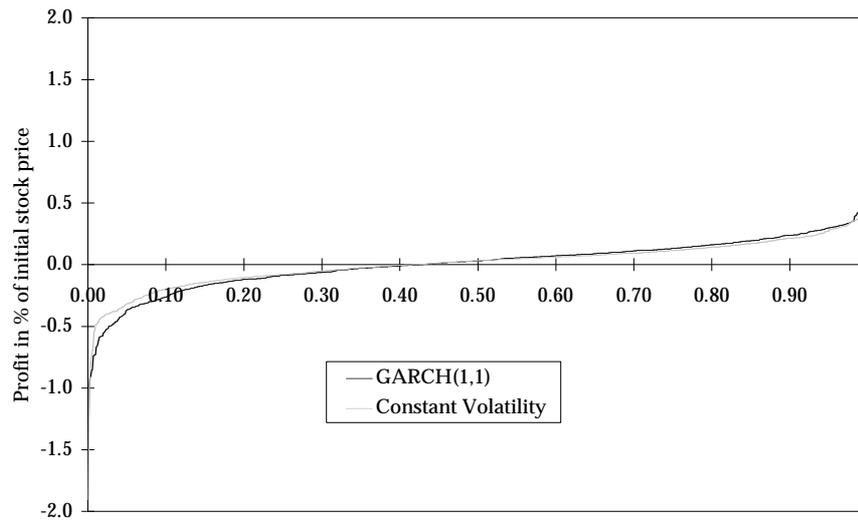


Figure 3. Distribution of profits from writing an at-the-money option with a maturity of three months and delta hedging it until expiry

The figure shows results from two simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. The curve labeled GARCH(1,1) gives the results from simulations with heteroskedastic return processes, with parameter values $\alpha = 0.0204$, $\beta = 0.970$, $\gamma = 4.31 \cdot 10^{-7}$ and $\nu = 5$. The curve labeled Constant Volatility gives the results from simulations with homoskedastic return processes. The maturity of the options is three months, which corresponds to 63 trading days.

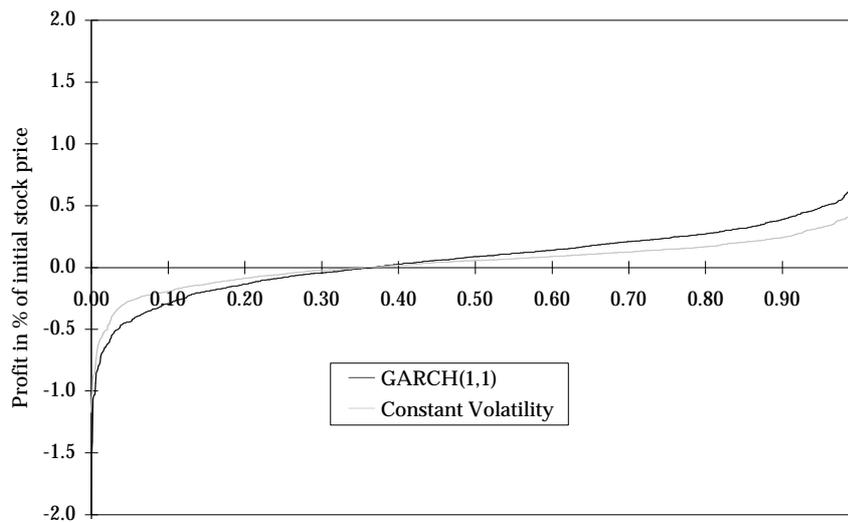


Figure 4. Distribution of profits from writing an at-the-money option with a maturity of six months and delta hedging it until expiry

The figure shows results from two simulations done with options that initially were at-the-money. 1000 simulations are performed for each option investigated. The curve labeled GARCH(1,1) gives the results from simulations with heteroskedastic return processes, with parameter values $\alpha = 0.0204$, $\beta = 0.970$, $\gamma = 4.31 \cdot 10^{-7}$ and $\nu = 5$. The curve labeled Constant Volatility gives the results from simulations with homoskedastic return processes. The maturity of the options is six months, which corresponds to 125 trading days.

