

Center for Economic Institutions
Working Paper Series


CEI Working Paper Series, No. 2005-12

"GARCH Options in Incomplete Markets"

Giovanni Barone-Adesi

Robert Engle

Loriano Mancini



Institute of Economic Research

Hitotsubashi University

2-1 Naka, Kunitachi, Tokyo, 186-8603 JAPAN

Tel: +81-42-580-8405

Fax: +81-42-580-8333

e-mail: cei-info@ier.hit-u.ac.jp

GARCH Options in Incomplete Markets*

Giovanni Barone-Adesi^a, Robert Engle^b and Liorano Mancini^a

^aInstitute of Finance, University of Lugano, Switzerland

^bDept. of Finance, Leonard Stern School of Business, New York University

First Version: March 2004

Revised: October 2004

***Correspondence Information:** Giovanni Barone-Adesi, Institute of Finance, University of Lugano, Via Buffi 13, CH-6900 Lugano, Tel: +41 (0)91 912 47 53, Fax: +41 91 912 46 47, E-mail address: BaroneG@lu.unisi.ch. E-mail addresses for Robert Engle and Liorano Mancini: REngle@stern.nyu.edu, Liorano.Mancini@lu.unisi.ch. Giovanni Barone-Adesi and Liorano Mancini gratefully acknowledge the financial support of the Swiss National Science Foundation (NCCR FINRISK).

GARCH Options in Incomplete Markets

Abstract

We propose a new method to compute option prices based on GARCH models. In an incomplete market framework, we allow for the volatility of asset return to differ from the volatility of the pricing process and obtain adequate pricing results. We investigate the pricing performance of this approach over short and long time horizons by calibrating theoretical option prices under the Asymmetric GARCH model on S&P 500 market option prices. A new simplified scheme for delta hedging is proposed.

Introduction

There is a general consensus that asset returns exhibit variances that change through time. GARCH models are a popular choice to model these changing variances. However the success of GARCH in modelling return variance hardly extends to option pricing. Models by Duan (1995), Heston (1993) and Heston and Nandi (2000) impose that the conditional volatility of the risk-neutral and the objective distributions be the same. Total variance, (the expectation of the integral of return variance up to option maturity), is then the expected value under the GARCH process. Empirical tests by Chernov and Ghysels (2000), (see also references therein), find that the above models do not price options well and their hedging performance is worse than Black-Scholes calibrated at the implied volatility of each option.

A common feature of all the tests to date is the assumption that the volatility of asset return is equal to the volatility of the pricing process. In other words, a risk neutral investor prices the option as if the distribution of its return had a different drift but unchanged volatility. This is certainly a tribute to the pervasive intellectual influence of the Black and Scholes (1973) model on option pricing. However, Black and Scholes derived the above property under very special assumptions, (perfect complete markets, continuous time and price processes). Changing volatility in real markets makes the perfect replication argument of Black-Scholes invalid. Markets are then incomplete in the sense that perfect replication of contingent claims using only the underlying asset and a riskless bond is impossible. Of course markets become complete if a sufficient, (possibly infinite), number of contingent claims are available. In this case a well-defined pricing density exists.

In the markets we consider the volatility of the pricing process is different from the volatility of the asset process. This occurs because investors will set state prices to reflect their aggregate

preferences. The pricing distribution will then be different from the return distribution. It is possible then to calibrate the pricing process directly on option prices. Although this may appear to be a purely fitting exercise, involving no constraint beyond the absence of arbitrage, verification of the stability of the pricing process over time and across maturities imposes substantial parameter restrictions. Economic theory may impose further restrictions from investors preferences for aggregate wealth in different states.

Carr, Geman, Madan, and Yor (2003) propose a similar set-up for Lévy processes. They use a jump process in continuous time. We propose to use discrete time and a continuous distribution for prices. Moreover we use GARCH models to drive stochastic volatility.

Heston and Nandi (2000) derived a quasi-analytical pricing formula for European options assuming a parametric linear risk premium, Gaussian innovations and the same GARCH parameters for the pricing and the asset process. In our pricing model we relax their assumptions. We allow for different volatility processes and time-varying, nonparametric risk premia—set by aggregate investors’ risk preferences. We use not only Monte Carlo simulation, but also filtered GARCH innovations.

Our method is different from Duan (1996), where a GARCH model is calibrated to the FTSE 100 index options assuming Gaussian innovations and the locally risk neutral valuation relationship, which implies that the conditional variance returns are equal under the objective and the risk neutral measures. Engle and Mustafa (1992) proposed a similar method to calibrate a GARCH model to S&P 500 index options in order to investigate the persistence of volatility shocks.

The final target is the identification of a pricing process for options that provides an adequate pricing performance. A surprising result concerns hedging performance. Hedging per-

formance, contrary to what is commonly sought in the stochastic volatility literature, cannot be significantly better than the performance of the Black-Scholes model calibrated at the implied volatility for each option. This result stems from the fact that deltas, (hedge ratios), for Black-Scholes can be derived applying directly the (first degree) homogeneity of option prices with respect to asset and strike prices, without using the Black-Scholes formulas. Therefore, hedge ratios from Black-Scholes calibrated at the implied volatility are the “correct” hedge ratios unless a very strong departure from “local homogeneity” occurs. This is not the case for the continuous, almost linear volatility smiles commonly found. In practice, for regular calls and puts, this is the case only for the asset price being equal to the strike price one instant before maturity. In summary, although it may be argued that calibrating Black-Scholes at each implied volatility does not give a model of option pricing, the hedging performance of this common procedure is almost unbeatable. Barone-Adesi and Elliott (2004) further investigate the computation of the hedge ratios under similar assumptions.

Our tests use closing prices of European options on the S&P 500 Index over several months. After estimating a GARCH model from earlier S&P 500 index data we search in a neighborhood of this model for the best pricing performance. Care is taken to prevent that our results be driven by microstructure effects in illiquid options.

The structure of the paper is the following. Section 1 presents option and state prices under GARCH models when the pricing process is driven by simulated, Gaussian innovations. Section 2 investigates the pricing performance of the proposed method when the pricing process is driven by filtered, estimated GARCH innovations. Section 3 discusses hedging results and Section 4 concludes.

1 Option and State Prices under the GARCH Model

Consider a discrete-time economy. Let S_t denote the closing price of the S&P 500 index at day t and y_t the daily log-return, $y_t := \ln(S_t/S_{t-1})$. Suppose that under the objective or historical measure \mathbb{P} , y_t follows an Asymmetric GARCH(1,1) model; see Glosten, Jagannathan, and Runkle (1993),

$$\begin{aligned} y_t &= \mu + \varepsilon_t, \\ \sigma_t^2 &= \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 + \gamma I_{t-1} \varepsilon_{t-1}^2, \end{aligned} \tag{1}$$

where $\omega, \alpha, \beta > 0$, $\alpha + \beta + \gamma/2 < 1$, μ determines the constant return (continuously compounded) of S_t , $\varepsilon_t = \sigma_t z_t$, $z_t \sim i.i.d.(0, 1)$ and $I_{t-1} = 1$, when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$, otherwise. The parameter $\gamma > 0$ accounts for the “leverage effect”, that is the stronger impact of “bad news” ($\varepsilon_{t-1} < 0$) rather than “good news” ($\varepsilon_{t-1} \geq 0$) on the conditional variance σ_t^2 .

The representative agent in the economy is an expected utility maximizer and the utility function is time-separable and additive. At time $t = 0$, the following Euler equation from the standard expected utility maximization argument gives the price of a contingent T -claim ψ_T ,

$$\begin{aligned} \psi_0 &= E_{\mathbb{P}}[\psi_T U'(C_T)/U'(C_0)|\mathcal{F}_0] = E_{\mathbb{P}}[\psi_T Y_{0,T}|\mathcal{F}_0] \\ &= E_{\mathbb{Q}}[\psi_T e^{-rT}|\mathcal{F}_0], \end{aligned}$$

where $E_{\mathbb{G}}[\cdot]$ denotes the expectation under the measure \mathbb{G} , r is the risk-free rate, $U'(C_t)$ is the marginal utility of consumption at time t and \mathcal{F}_t is the information set available up to and including time t . The state price density per unit probability process Y is defined by $Y_{t,T} := e^{-r(T-t)} L_t$ and

$$L_t = \frac{d\mathbb{Q}_t}{d\mathbb{P}_t} = \frac{q dS}{p dS} = \frac{q}{p},$$

where \mathbb{Q} is the risk neutral measure absolutely continuous with respect to \mathbb{P} , the subindex t denotes the restriction to \mathcal{F}_t , q and p (time subscripts are omitted) are the corresponding density functions. When the financial market is incomplete, L_t is not unique and is determined by the representative agent's preferences. Intuitively, if $p(\bar{S}_T)$ was a discrete probability, the state price density evaluated at \bar{S}_T , $Y_{t,T}(\bar{S}_T)p(\bar{S}_T)$, gives at time t the price of \$1 to be received if state \bar{S}_T occurs. The state price per unit probability, $Y_{t,T}(\bar{S}_T)$, is then the market price of a state contingent claim that pays $1/p(\bar{S}_T)$ if state \bar{S}_T , which has probability $p(\bar{S}_T)$, occurs. The expected rate of return of such a claim under the physical measure \mathbb{P} is $1/Y_{t,T}(\bar{S}_T) - 1$. As marginal utilities of consumptions decrease when the states of the world "improve", $Y_{t,T}$ is expected to decrease in S_T .

1.1 Monte Carlo Option Prices

Monte Carlo simulation is used to compute the GARCH option prices, because the distribution of temporally aggregated asset returns cannot be derived analytically. We present the computation of a European call option price; other European claims can be priced similarly.

At time $t = 0$ the dollar price of a European call option with strike price $\$K$ and time to maturity T days is computed by simulating log-returns in model (1) under the risk neutral measure \mathbb{Q} . Specifically, we draw T independent standard normal random variables $(z_i^*)_{i=1,\dots,T}$, we simulate (y_i, σ_i^2) in model (1) under the risk neutral parameters $\omega^*, \alpha^*, \beta^*, \gamma^*, \mu = r - d - \sigma_i^2/2$, where r is the risk-free rate and d is the dividend yield on a daily basis, and we compute $S_T^{(n)} = S_0 \exp(\sum_{i=1}^T y_i)$. Then, we compute the discounted call option payoff $C^{(n)} = \exp(-rT) \max(0, S_T^{(n)} - K)$. Iterating the procedure N times gives the Monte Carlo estimate for the call option price, $C_{mc}(K, T) := N^{-1} \sum_{n=1}^N C^{(n)}$. To reduce the variance of the Monte

Carlo estimates we use the method of antithetic variates; cf., for instance, Boyle, Broadie, and Glassermann (1997). Specifically, $C^{(n)} = (C_a^{(n)} + C_b^{(n)})/2$, where $C_a^{(n)}$ is computed using $(z_i^*)_{i=1,\dots,T}$ and $C_b^{(n)}$ using $(-z_i^*)_{i=1,\dots,T}$. Each option price C_{mc} is computed simulating $2N$ sample paths for S . In our calibration exercises we set $N = 10,000$. To further reduce the variance of the Monte Carlo estimates we calibrate the mean as in the empirical martingale simulation method proposed by Duan and Simonato (1998). Scaling the simulated values $S_T^{(n)}$, $n = 1, \dots, N$, by a multiplicative factor, the method ensures that the risk neutral expectation of the underlying asset is equal to the forward price, i.e. $N^{-1} \sum_{n=1}^N \tilde{S}_T^{(n)} = S_0 \exp((r - d)T)$, where $\tilde{S}_T^{(n)} := S_T^{(n)} S_0 \exp((r - d)T) (N^{-1} \sum_{n=1}^N S_T^{(n)})^{-1}$. Then, option prices are computed using $\tilde{S}_T^{(n)}$. In our calibration exercises at least 100 simulated paths of the underlying asset end at maturity “in the money” for almost all the deepest out of the money options.

1.2 Calibration of the GARCH Model

The risk neutral parameters of the GARCH model, $\theta^* = (\omega^* \alpha^* \beta^* \gamma^*)$, are determined by calibrating GARCH option prices computed by Monte Carlo simulation on market option prices taken as averages of bid and ask prices at the end of one day.

Specifically, let $P^{mkt}(K, T)$ denote the market price in dollars at time $t = 0$ of a European option with strike price $\$K$ and time to maturity T days. The risk neutral parameters θ^* are determined by minimizing the mean squared error (mse) between model option prices and market prices. The mse is taken over all strikes and maturities,

$$\theta^* := \arg \min_{\theta} \sum_{i=1}^m \left(P^{garch}(K_i, T_i; \theta) - P^{mkt}(K_i, T_i) \right)^2, \quad (2)$$

where $P^{garch}(K, T; \theta)$ is the theoretical GARCH option price and m is the number of European options considered for the calibration at time $t = 0$.

As an overall measure of the quality of the calibration we compute the average absolute pricing error (ape) with respect to the mean price,

$$\text{ape} := \frac{\sum_{i=1}^m |Pgarch(K_i, T_i; \theta^*) - Pmkt(K_i, T_i)|}{\sum_{i=1}^m Pmkt(K_i, T_i)}. \quad (3)$$

1.3 Empirical Results

We calibrate the GARCH model to European options on the S&P 500 index observed on a random date $t :=$ August 29, 2003 and we set $t = 0$. Estimates of σ_0^2 and z_0 are necessary to simulate the risk neutral GARCH volatility and are obtained in the next section.

1.3.1 Estimation of the GARCH Model

Percentage daily log-returns, $y_t \times 100$, of the S&P 500 index are computed from December 11, 1987 to August 29, 2003 for a total of 4,100 observations. Model (1) is estimated using the Pseudo Maximum Likelihood (PML) estimator based on the nominal assumption of conditional normal innovations. The parameter estimates are reported in Table 1. The current August 29, 2003 estimates on a daily base of σ_0^2 and z_0 are 0.635 and 0.604, respectively, and will be used as starting values to simulate the risk neutral GARCH volatility in the calibration exercise.

1.3.2 Calibration of the GARCH Model with Gaussian Innovations

Initially we calibrate the GARCH model (1) to the closing prices (bid-ask averages) of out of the money European put and call options on the S&P 500 index observed on August 29, 2003. Precisely, we only consider option prices strictly larger than \$0.05—discarding 40 option prices to avoid that our results be driven by microstructure effects in very illiquid options—and maturities $T = 22, 50, 85, 113$ days for a total of $m = 118$ option prices. Strike prices range from \$550 to \$1,250, $r = 0.01127/365$, $d = 0.01634/365$ on a daily basis and $S_0 = \$1,008$.

To solve the minimization problem (2) we use the Nelder-Mead simplex direct search method implemented in the Matlab function `fminsearch`. This function does not require the computation of gradients. Starting values for the risk neutral parameters θ^* are the parameter estimates given in Table 1. Calibrated parameters, root mean squared error (rmse) and ape measure for the quality of the calibration are reported in the first row of Table 2. The “leverage effect” in the volatility process under the risk neutral measure \mathbb{Q} ($\gamma^* = 0.288$) is substantially larger than under the objective measure \mathbb{P} ($\gamma = 0.075$). The average pricing error is quite low and equals to 2.54%. Figure 1 shows the pricing performance of the GARCH model which seems to be satisfactory. Figure 2 shows the calibration errors defined as $P^{garch} - P^{mkt}$. Such errors tend to be larger for near at the money options (these options have the largest prices) and for deep out of the money put options.

1.3.3 State Price Density Estimates with Gaussian Innovations

For the maturities $T = 22, 50, 85, 113$ days we compute the state price densities per unit probability of $S_T, Y_{0,T}$, as the discounted ratio of the risk neutral density over the objective density. Under the objective measure \mathbb{P} , the asset prices S are simulated assuming the drift $\mu = r + 0.08/365 - \sigma_t^2/2$ in equation (1) and the parameter estimates in Table 1. Under the risk neutral measure \mathbb{Q} , $\mu = r - d - \sigma_t^2/2$ and the calibrated GARCH parameters are given in the first row of Table 2. The density functions are estimated by the Matlab function `ksdensity` using the Gaussian kernel and the optimal default bandwidth for estimating Gaussian densities.

Figure 3 shows the estimated risk neutral and objective densities and the corresponding state price densities per unit probability; see also Table 3. As expected the state price densities are quite stable across maturities and monotonic, decreasing in S_T . However, the high values

on the left imply very negative expected rate of return for out of the money puts, that appear intuitively “overpriced”. As an example, a state price per unit probability of \$6 corresponds to an expected rate of return of $1/6 - 1 = -0.833$ for a simple state contingent claim. State price densities outside the reported values for S_T tend to be unstable, as the density estimates are based on very few observations.

2 GARCH Option Prices with Filtering Historical Simulations

In this section we investigate the pricing performance of the GARCH model when the simulated, Gaussian innovations—used to drive the GARCH process under the risk neutral measure—are replaced by historical, estimated GARCH innovations. We refer to this approach as the Filtering Historical Simulation (FHS) method. Barone-Adesi, Bourgoin, and Giannopoulos (1998) introduced the FHS method to estimate portfolio risk measures.

This procedure is in two steps. Suppose we aim at calibrating the GARCH model on market option prices $P^{mkt}(K_i, T_i)$, $i = 1, \dots, m$ observed on day $t := 0$. In the first step, the GARCH model is estimated on the historical log-returns of the underlying asset $y_{-n+1}, y_{-n+2}, \dots, y_0$ up to time $t = 0$. The scaled innovations of the GARCH process $\hat{z}_t = \hat{\varepsilon}_t \hat{\sigma}_t^{-1}$, for $t = -n+1, \dots, 0$, are also estimated.

In the second step, the GARCH model is calibrated to the market option prices by solving the minimization problem (2). The theoretical GARCH option prices, $P^{garch}(K, T; \theta^*)$, are computed by Monte Carlo simulations as in Section 1.1, but the Gaussian innovations are replaced by innovations \hat{z}_t 's estimated in the first step, randomly drawn with uniform probabilities. To preserve the negative skewness of the estimated innovations the method of the antithetic variates is not used.

2.1 Calibration of the GARCH model with FHS Innovations

We apply this two steps procedure to the option prices on the S&P 500 observed on a random date July 9, 2003. Specifically, in the first step we estimate the GARCH model (1) on $n = 3,800$ historical returns of the S&P 500 index from December 14, 1988 to July 9, 2003 and we estimate the corresponding innovations \hat{z} . In the second step, we calibrate the GARCH model to the out of the money put and call options with maturities $T = 10, 38, 73, 164, 255, 346$ days for a total of $m = 151$ option prices; 45 options with bid price lower than \$0.05 are discarded. The PML estimates of model (1) are reported in Table 4. The last panel in Figure 4 shows the estimated scaled innovations, \hat{z}_t 's, used to drive the GARCH process under the risk neutral measure. The skewness and the kurtosis of the empirical distribution of \hat{z} are -0.6 and 7.4 , respectively. Calibration results are reported in the first row of Table 5 and Figure 5. The average pricing error is 3.5% and the overall pricing performance is quite satisfactory given the wide range of strikes and maturities of the options used for the calibration.

We calibrate the GARCH model using the FHS method also on the same options considered in the calibration for August 29, 2003. The results are reported in the second row of Table 2. Given the limited number of options used in this calibration, the GARCH pricing model with Gaussian innovation has already a very low pricing error. However, using the FHS method both the rmse and the ape measure are reduced by about 10%. The asymmetry parameter γ^* decreases from 0.288 to 0.201 when filtered, estimated innovations rather than Gaussian innovations are used, because of the negative skewness, -0.61 , of the filtered innovations.

2.2 State Price Density Estimates with FHS Innovations

The state price densities per unit probability on July 9, 2003, computed similarly as in Section 1.3.3, are shown in Figure 6. Using FHS innovations, the asymmetry parameter γ^* is now very close to γ (cf. Tables 4–5) and state prices per unit probability are still monotone, but much closer to each other. In particular the state prices per unit probability on the left are now in line with the remaining ones. This implies that “excess” out of the money put prices can be explained by the skewness of FHS innovations. The volatility smile—computed using out of the money European put and call options—for 38 days to maturity on this date is reported in Figure 7. Notice that the sample period to estimate the GARCH model (1) starts after the October 1987 crash. Such a large negative return would inflate the variance estimates and this tends to produce non monotone state price densities per unit probability.

The state price densities per unit probability on August 29, 2003 using the FHS method are quite close to those on July 9, 2003 and are omitted.

2.3 Short Run Stability of the GARCH Pricing Model

To investigate the stability of the pricing performance for the GARCH model over a “short” time horizon, i.e. one month, we calibrate the model for several dates from July 9 to August 8, 2003 on out of the money European option prices with maturities less than a year. The calibration results are reported in Table 5. The GARCH parameters tend to change over time, but the pricing performances are quite stable in terms of rmse and ape measures. Moreover, the estimates of the long run level of the risk neutral variance $E_{\mathbb{Q}}[\sigma^2]$ are quite stable and about 1% on a daily base.

To check for the stability of the GARCH parameters we calibrate one GARCH model to

the option prices on July 9, 10, 11, and 14, 2003. The initial variances and innovations, σ_0^2 's and z_0 's, for the dates July 10, 11, 14 are computed updating the corresponding estimates for July 9, i.e. 0.793, -0.667 , and using the objective GARCH estimates in Table 4. This procedure ensures that future, not yet available information is not used for the fitting of earlier option prices. The GARCH parameter of the “pooled” calibration are $\omega_{pool}^* = 0.016$, $\alpha_{pool}^* = 0.000$, $\beta_{pool}^* = 0.924$, $\gamma_{pool}^* = 0.121$, which imply a long run level of the risk neutral variance $E_{\mathbb{Q}}[\sigma_{pool}^2] = 0.99$. Table 6 compares the pricing errors—the differences between theoretical and observed option prices—of the pooled calibration with the corresponding errors for the single day calibration given in Table 5. As expected the rmse's for the pool calibration are larger than the corresponding rmse's for the single day calibrations. However, differences are small and the correlation between the two pricing errors is on average 0.92, meaning that the two pricing performances are quite close.

2.4 Long Run Stability of the GARCH Pricing Model and Comparison with CGMYSA Model

To investigate the pricing performance of the GARCH model over a “long” time horizon, i.e. one year, we calibrate the model on out of the money European option prices with maturities between a month and a year for the dates January 12, March 8, May 10, July 12, September 13 and November 8 for the year 2000. For each calibration we use about the last seven years of S&P 500 daily log-returns to implement the FHS method. We also compare the pricing performance of the GARCH model with the CGMYSA model proposed by Carr, Geman, Madan, and Yor (2003) for the dynamic of the underlying asset, which is a mean corrected, exponential Lévy process time changed with a Cox, Ingersoll and Ross process. Average absolute pric-

ing errors are somewhat in favour of the CGMYSA model as this model has nine parameters while the GARCH model has four parameters. The results are reported in Table 7. There is evidence that the GARCH parameters tend to change from month to month, but the pricing performance is quite stable especially in terms of the ape measure. Moreover, the mean and the standard deviation of the ape measures for the GARCH model are 4.07, 1.03 and for the CGMYSA model are 3.91, 1.17, respectively. Hence, the pricing performance of the GARCH model is more stable than the pricing performance of the CGMYSA model, but the last model is superior in terms of average ape measure. Carr, Geman, Madan, and Yor (2003) proposed also more parsimonious (six parameters) models, namely the VGSA and NIGSA models, which are, respectively, finite variation and infinite variation mean corrected, exponential Lévy processes with infinite activity for the underlying asset. For the previous dates, the GARCH model outperforms the VGSA and NIGSA models in five and four out of six cases, respectively.

3 Hedging

Extension to the GARCH setting of the delta hedging, Engle and Rosenberg (2002), does not show an improvement on the delta hedging strategy based on the Black-Scholes model calibrated at the implied volatility. To understand why this is the case consider the example presented in Table 8. The three rows in the middle are market option prices from Hull's book. The first row is obtained multiplying the middle row times 0.9 and the last row is obtained multiplying the middle row by 1.1, that is assuming an homogeneous pricing model.

Incremental ratios, that is change in option price over change in stock price, can be computed between the first two and then again the last two rows, i.e. $\Delta_{45} := (5.60 - 2.16)/(49 - 44.1)$ and $\Delta_{55} := (2.64 - 1.00)/(53.9 - 49)$. Taking the average of these two ratios, for the strike

price $K = 50$ we obtain an estimate of delta equals to 0.518, which is almost identical to the delta from the Black-Scholes model calibrated at the implied volatility for the middle row, i.e. 0.522—the implied volatility is equal to 0.2 when $r = 0.05$ and $T = 20/52$ years. Hence, the application of first-degree homogeneity to non-homogeneous prices has led to an essentially correct hedge ratio! To understand this paradoxical result consider the sources of errors in the above computations. There is a discretization error and an error due to the volatility smile. In fact, in the absence of a volatility smile, Black-Scholes option prices would be homogeneous functions of the stock and the strike price. The discretization error leads to a discrete delta which is approximately the average of the Black-Scholes deltas computed at the two extremes of each interval and approximated by Δ_{45} and Δ_{55} . Formally, denote by $\Delta(K)$ the delta as a function of the strike price K . For small intervals the delta hedge is approximated by

$$\Delta(50) \approx \frac{\Delta(50) + \overbrace{\Delta'(50)(45 - 50)}^{>0} + \Delta(50) + \overbrace{\Delta'(50)(55 - 50)}^{<0}}{2} \approx \frac{\Delta_{45} + \Delta_{55}}{2}.$$

Therefore, the two discrete ratios considered, Δ_{45} and Δ_{55} , are affected by opposite errors up to the first order. Taking their average eliminates these errors. The only error left is due to the smile effect. However, this error is very small if the strike price increment is small relative to the asset price and its volatility. See Barone-Adesi and Elliott (2004) for further discussion. The reader may verify this simple result on the options of his choice. It appears therefore that deltas are to a large degree determined by market option prices, independently of the chosen model. Therefore, models alternative to Black-Scholes calibrated at the implied volatility will generally lead to very similar hedge ratios, if they fit well market prices. The only significant deterioration of hedging occurs in the presence of large volatility shocks, which diminish the effectiveness of delta hedging. To observe this compare a day with a modest change in volatility,

e.g. $t_2 :=$ July 10, 2003, with a day in which a large negative index return led to a large increase in volatility, e.g. $t_1 :=$ January 24, 2003. Specifically, for the day t_1 we consider out of the money put and call options with maturities equal to 30, 58, 86, 149, 240, 331 days for a total of 160 option prices and for the day t_2 we consider the same options as in Section 2.3. Then, we run the following set of regression for $t + 1 = t_1, t_2$

$$1) P_{t+1}^{mkt} = \eta_0 + \eta_1 P_{t,t+1}^{bs} + error,$$

$$2) P_{t+1}^{mkt} = \eta_0 + \eta_1 P_{t,t+1}^{bs} + \eta_2 P_{t,t+1}^{garch} + error,$$

$$3) P_{t+1}^{mkt} = \eta_0 + \eta_2 P_{t,t+1}^{garch} + error,$$

where P_{t+1}^{mkt} are the option prices observed on time $t + 1$, $P_{t,t+1}^{bs}$ are the Black-Scholes forecasts of option prices for $t + 1$ computed by plugging in the Black-Scholes formula S_{t+1} , r , d at time $t + 1$ and the implied volatility observed on time t (i.e. January 23 and July 9, 2003, respectively). $P_{t,t+1}^{garch}$ are the GARCH forecasts obtained using S_{t+1} , the GARCH parameter calibrated at time t and σ_{t+1} updated according to the objective estimates at time t .

The ordinary least square (OLS) estimates of the previous regressions are given in Table 9. In terms of the error variance the Black-Scholes forecasts in regressions 1) are superior to the GARCH forecasts in regressions 3) for both days t_1 and t_2 . Moreover, in the regressions 2) the weights η_1 of the Black-Scholes forecasts are larger than the weights η_2 for the GARCH forecasts. This is due to the “initial advantage” of the Black-Scholes forecasts, i.e. the zero pricing error at time t . However, for the day January 24, 2003, from regression 1) to regression 2) the variance of the prediction error is reduced about 60% adding the GARCH forecast as a regressor. Hence, the GARCH model carries on large amount of information on option price dynamics. Moreover, the GARCH model provides a dynamic model for the risk neutral

volatility, while the Black-Scholes model does not.

Interestingly, the Black-Scholes forecasts tend to underestimate option prices observed on January 24, 2003 (while the GARCH forecasts tend to overestimate option prices). An explanation is the following. The daily log-return of the S&P 500 for January 24, 2003 is -2.97% , which induces an increase in the volatility of the underlying asset. Such an increase in the volatility can not be detected by the Black-Scholes model with constant implied volatility, but it is reflected in the GARCH forecasts of volatilities and option prices. This effect is stronger in days with large returns. For the day July 10, 2003 the reduction in the variance of the prediction error is only 11%, as the return of the S&P 500 is -1.36% only.

Unfortunately, our GARCH price forecast is conditioned on the current index and it cannot be used to improve significantly delta hedging. Its explanatory power simply indicates that delta hedging is less effective in the presence of large volatility shocks. They are linked to the index return in a nonlinear fashion in the GARCH model.

4 Conclusion

Casting the option pricing problem in incomplete markets allows for more flexibility in the calibration of market prices. Investors' preferences can be inferred comparing the physical and the pricing distributions. Using filtered historical simulation the volatility smile appears to be explained by innovation skewness, with no need of much higher state prices for out of the money puts. Delta hedging does not require a large computational effort under conditions usually found in index option markets, removing a major drawback of simulation-based option pricing. Further refinements of pricing and stability issues are left to future research.

References

- Barone-Adesi, G., F. Bourgoïn, and K. Giannopoulos, 1998, “Don’t look back,” *Risk*, 11, 100–103.
- Barone-Adesi, G., and R. J. Elliott, 2004, “Cutting the hedge,” Working paper.
- Black, F., and M. Scholes, 1973, “The valuation of options and corporate liabilities,” *Journal of Political Economy*, 81, 637–654.
- Boyle, P., M. Broadie, and P. Glassermann, 1997, “Monte Carlo Methods for Security Pricing,” *Journal of Economic Dynamics and Control*, 21, 1267–1321.
- Carr, P., H. Geman, D. B. Madan, and M. Yor, 2003, “Stochastic Volatility for Lévy Processes,” *Mathematical Finance*, 13, 345–382.
- Chernov, M., and E. Ghysels, 2000, “A Study towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation,” *Journal of Financial Economics*, 56, 407–458.
- Duan, J.-C., 1995, “The GARCH Option Pricing Model,” *Mathematical Finance*, 5, 13–32.
- , 1996, “Cracking the Smile,” *Risk*, 9, 55–59.
- Duan, J.-C., and J.-G. Simonato, 1998, “Empirical martingale simulation for asset prices,” *Management Science*, 44, 1218–1233.
- Engle, R. F., and C. Mustafa, 1992, “Implied ARCH Models from Options Prices,” *Journal of Econometrics*, 52, 289–311.

Engle, R. F., and J. V. Rosenberg, 2002, “Empirical Pricing Kernels,” *Journal of Financial Economics*, 64, 341–372.

Glosten, L. R., R. Jagannathan, and D. E. Runkle, 1993, “On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks,” *Journal of Finance*, 48, 1779–1801.

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 S. Nandi, 2000, “A Closed-Form GARCH Option Valuation Model,” *Review of Financial Studies*, 13, 585–625.

Table 1: PML estimates of the GARCH model (1), $y_t \times 100 = \mu + \varepsilon_t$, $\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 + \gamma I_{t-1}\varepsilon_{t-1}^2$, $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$ otherwise, $\varepsilon_t = \sigma_t z_t$, $z_t \sim i.i.d.(0, 1)$, (p -values in parenthesis) for the S&P 500 index daily log-returns y_t in percentage from December 11, 1987 to August 29, 2003.

μ	ω	α	β	γ
0.033	0.009	0.006	0.946	0.075
(0.008)	(0.000)	(0.416)	(0.000)	(0.000)

Table 2: Calibrated parameters of the GARCH model (1), $\sigma_t^2 = \omega^* + \alpha^* \varepsilon_{t-1}^2 + \beta^* \sigma_{t-1}^2 + \gamma^* I_{t-1} \varepsilon_{t-1}^2$, $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$ otherwise, $\varepsilon_t = \sigma_t z_t$, $z_t \sim i.i.d.(0, 1)$, using Gaussian innovations (first row) and FHS method (second row) on August 29, 2003 out of the money European put and call options ($m = 118$) and time to maturities $T = 22, 50, 85, 113$ days. The root mean squared error (rmse) is in \$, the ape measure is defined in equation (3).

	ω^*	α^*	β^*	γ^*	rmse	ape%
Gauss. z	0.037	0.000	0.833	0.288	0.27	2.54
FHS	0.037	0.000	0.870	0.201	0.24	2.29

Table 3: State price densities estimates per unit of probability, $Y_{0,T}$, time to maturities $T = 22, 50, 85, 113$ days for August 29, 2003. $Y_{0,T} := e^{-rT} L_0$ and $L_0 = d\mathbb{Q}_0/d\mathbb{P}_0$, where \mathbb{Q} is the risk neutral measure absolutely continuous with respect to the objective measure \mathbb{P} and the subindex $t = 0$ denotes the restriction to \mathcal{F}_0 .

S_T	900	1,000	1,100	1,200
$Y_{0,22}$	1.882	1.001	0.437	—
$Y_{0,50}$	1.284	1.011	0.773	0.254
$Y_{0,85}$	1.197	1.003	0.844	0.597
$Y_{0,113}$	1.281	1.028	0.834	0.641

Table 4: PML estimates of the GARCH model (1), $y_t \times 100 = \mu + \varepsilon_t$, $\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 + \gamma I_{t-1}\varepsilon_{t-1}^2$, $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$ otherwise, $\varepsilon_t = \sigma_t z_t$, $z_t \sim i.i.d.(0, 1)$, (p -values in parenthesis) for the S&P 500 index daily log-returns y_t in percentage from December 14, 1988 to July 9, 2003.

μ	ω	α	β	γ
0.033	0.012	0.005	0.936	0.093
(0.008)	(0.000)	(0.547)	(0.000)	(0.000)

Table 5: Calibrated parameters of the GARCH model (1), $\sigma_t^2 = \omega^* + \alpha^* \varepsilon_{t-1}^2 + \beta^* \sigma_{t-1}^2 + \gamma^* I_{t-1} \varepsilon_{t-1}^2$, $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$ otherwise, $\varepsilon_t = \sigma_t z_t$, $z_t \sim i.i.d.(0, 1)$, under the risk neutral measure \mathbb{Q} , using FHS on several days and m out of the money European put and call options. T is the time to maturity in days. The root mean squared error (rmse) is in \$, the ape measure is defined in equation (3).

date	ω^*	α^*	β^*	γ^*	$E_{\mathbb{Q}}[\sigma^2]$	m	$\min(T)$	$\max(T)$	rmse	ape%
Jul 9	0.019	0.000	0.912	0.138	1.00	151	10	346	0.64	3.50
Jul 10	0.008	0.000	0.953	0.078	1.00	148	9	345	0.49	2.75
Jul 11	0.016	0.000	0.921	0.125	0.98	146	8	344	0.64	3.64
Jul 14	0.009	0.000	0.949	0.083	0.96	146	5	341	0.43	2.33
Jul 16	0.011	0.000	0.946	0.086	1.00	141	3	339	0.67	3.59
Jul 21	0.005	0.000	0.964	0.061	0.86	156	26	334	0.94	3.61
Jul 25	0.054	0.000	0.787	0.319	1.03	165	22	330	0.69	4.24
Jul 30	0.010	0.000	0.943	0.092	0.97	161	17	325	0.40	2.26
Aug 1	0.022	0.000	0.912	0.137	1.12	163	15	323	0.59	3.38
Aug 4	0.016	0.000	0.928	0.117	1.21	163	12	320	1.02	5.64
Aug 8	0.017	0.000	0.925	0.119	1.10	159	8	316	0.65	3.69

Table 6: Comparison between pricing errors, i.e. the differences between theoretical and observed option prices, of the calibration pool for July 9, 10, 11, 14, and the single day calibrations. The root mean squared error (rmse) is in \$, $\text{corr}(\text{err single day}, \text{err pool})$ denotes the correlation between the pricing errors for the single day calibration and the corresponding pricing errors for the pooled calibration.

	Jul 9	Jul 10	Jul 11	Jul 14	average
rmse single day	0.639	0.487	0.636	0.434	0.549
rmse pool	0.725	0.584	0.686	0.481	0.619
$\text{corr}(\text{err single day}, \text{err pool})$	0.935	0.877	0.943	0.895	0.915

Table 7: Calibrated parameters of the GARCH model (1), $\sigma_t^2 = \omega^* + \alpha^* \varepsilon_{t-1}^2 + \beta^* \sigma_{t-1}^2 + \gamma^* I_{t-1} \varepsilon_{t-1}^2$, $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$ otherwise, $\varepsilon_t = \sigma_t z_t$, $z_t \sim i.i.d.(0, 1)$, under the risk neutral measure \mathbb{Q} , using FHS on m out of the money European put and call options for the year 2000 and comparison with the CGMYSA model. The root mean squared error (rmse) is in \$, the ape measure is defined in equation (3).

date	ω^*	α^*	β^*	γ^*	$E_{\mathbb{Q}}[\sigma^2]$	m	rmse	ape%	ape% CGMYSA
Jan	0.016	0.000	0.914	0.155	1.80	177	1.62	4.78	3.78
Mar	0.118	0.000	0.635	0.600	1.82	143	1.61	5.13	5.23
May	0.158	0.000	0.526	0.839	2.90	155	1.93	4.74	5.48
Jul	0.006	0.000	0.963	0.065	1.38	159	0.91	2.34	3.26
Sep	0.041	0.000	0.866	0.189	1.04	151	1.08	3.67	2.87
Nov	0.017	0.000	0.903	0.159	0.97	169	1.22	3.74	2.85

Table 8: “Homogeneous hedging of the smile”. The three rows in the middle are market option prices from Hull’s book. The first row is obtained multiplying the middle row times 0.9 and the last row is obtained multiplying the middle row by 1.1.

Strike price	Asset price	Option price
45	44.1	2.16
45	49	5.60
50	49	2.40
55	49	1.00
55	53.9	2.64

Table 9: OLS regression estimates and variance of forecast errors for time $t + 1$, i.e. January 24, 2003 (first panel) and July 10, 2003 (second panel): 1) $P_{t+1}^{mkt} = \eta_0 + \eta_1 P_{t,t+1}^{bs} + error$; 2) $P_{t+1}^{mkt} = \eta_0 + \eta_1 P_{t,t+1}^{bs} + \eta_2 P_{t,t+1}^{garch} + error$, 3) $P_{t+1}^{mkt} = \eta_0 + \eta_2 P_{t,t+1}^{garch} + error$, where P_{t+1}^{mkt} are the option prices observed on time $t + 1$, $P_{t,t+1}^{bs}$ are the Black-Scholes forecasts of option prices for $t + 1$ computed by plugging in the Black-Scholes formula S_{t+1} , r , d at time $t + 1$ and the implied volatility observed on time t (i.e. January 23 and July 9, respectively). $P_{t,t+1}^{garch}$ are the GARCH forecasts obtained using S_{t+1} , the GARCH parameter calibrated at time t and σ_{t+1} updated according to the estimates at time t .

	η_0	η_1	η_2	$Var[error]$
1)	0.823	0.996	—	0.761
2)	-0.037	0.558	0.436	0.316
3)	-1.073	—	0.988	1.035
<hr/>				
1)	-0.118	0.997	—	0.188
2)	-0.213	0.293	0.704	0.161
3)	-0.429	—	0.997	0.315

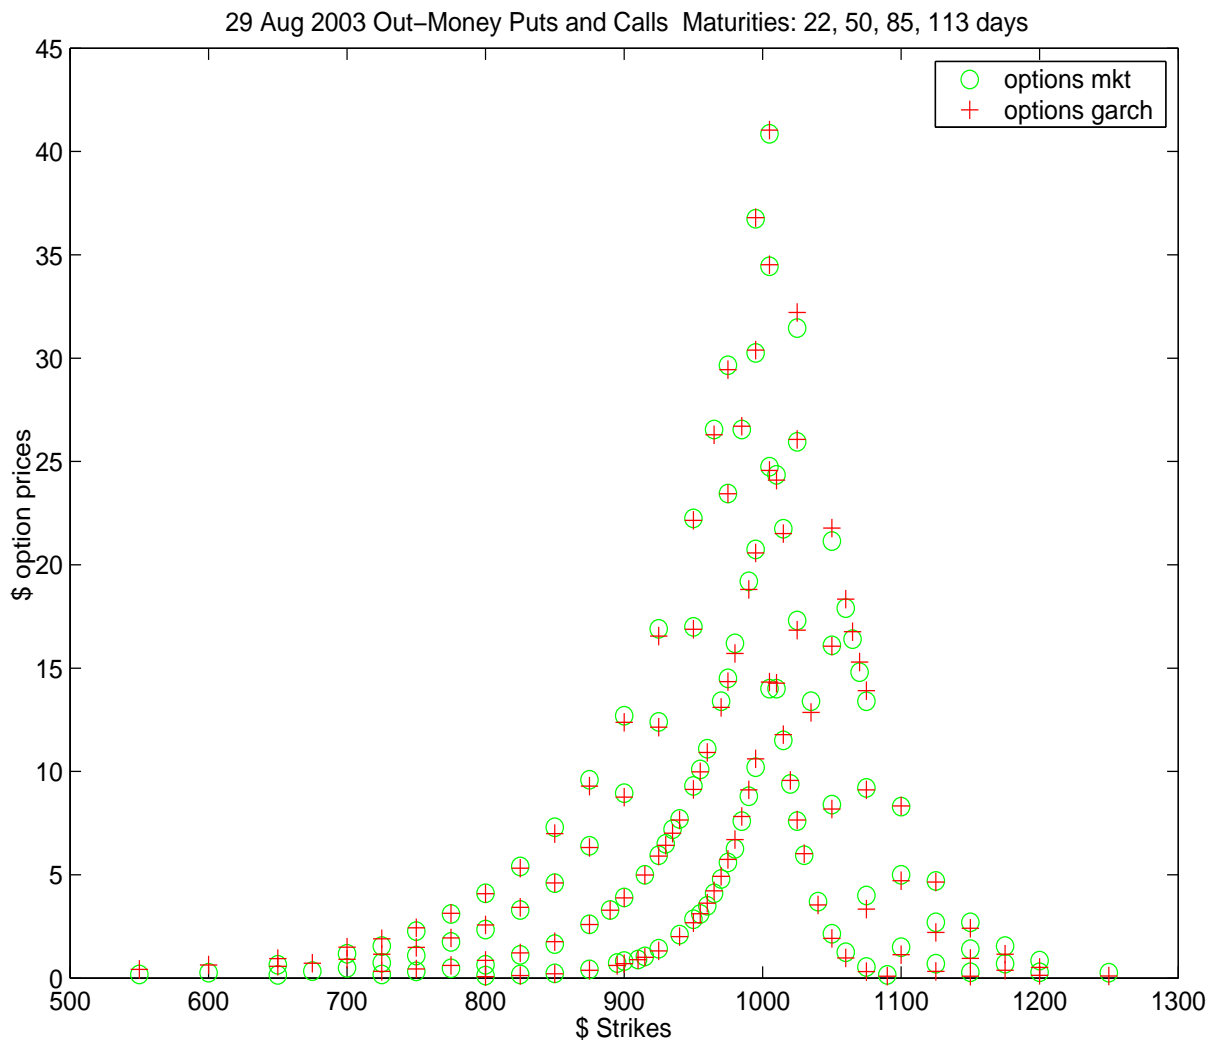


Figure 1: Monte Carlo calibration results of the GARCH model to $m = 118$ out of the money European put and call option prices observed on August 29, 2003.

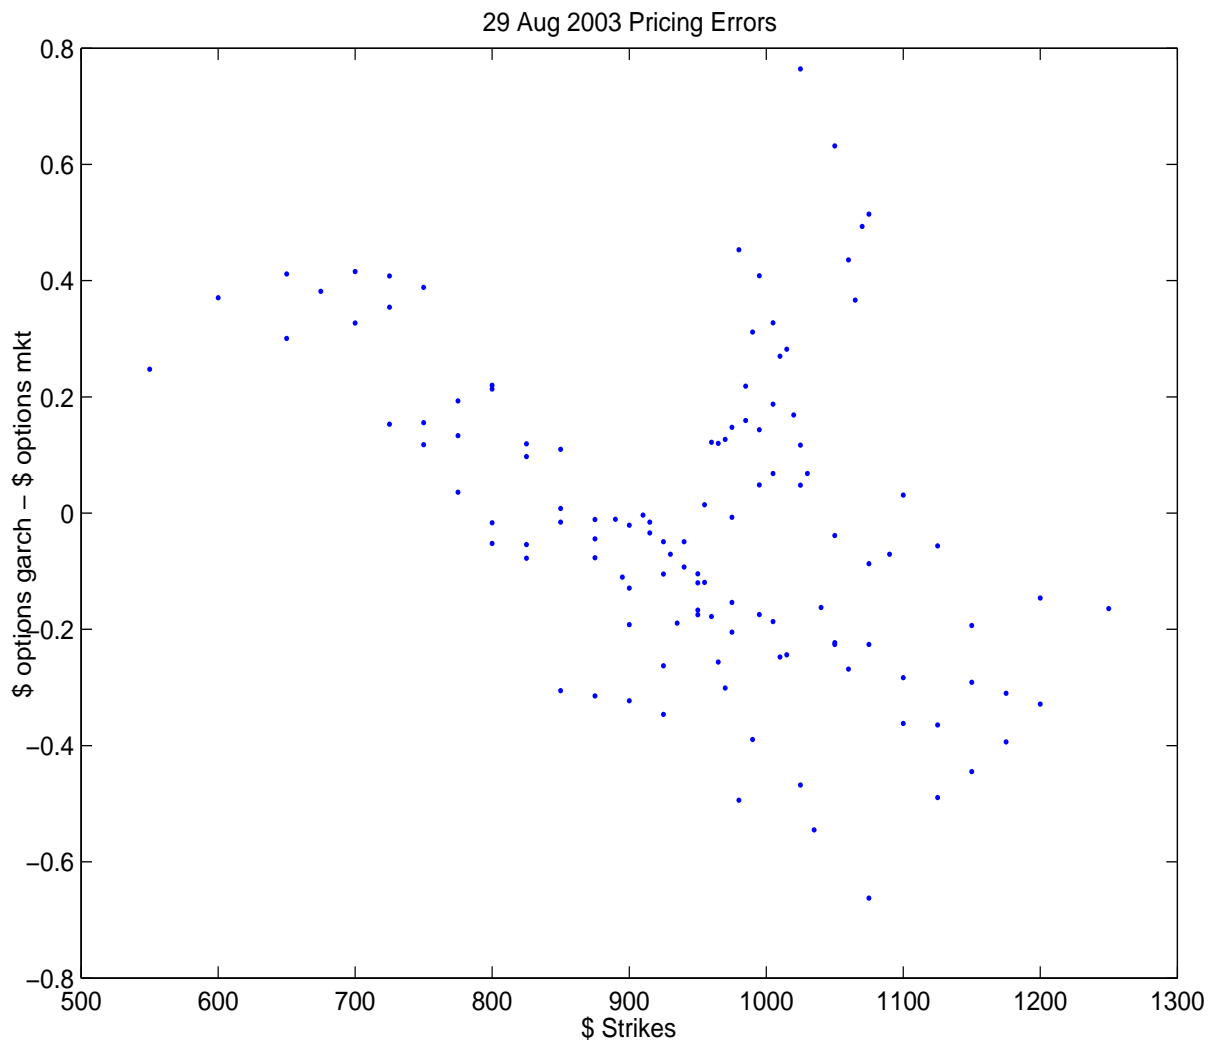


Figure 2: Pricing errors of the GARCH model for $m = 118$ out of the money European put and call option prices observed on August 29, 2003.




Figure 3: Risk neutral and objective density estimates (left plots) and state price density estimates per unit of probability (right plots) for August 29, 2003.




Figure 4: Daily log-return in percentage of the S&P 500 index from December, 14 1988 to July 9, 2003 (first panel), estimated conditional variances (second panel) and scaled innovations (third panel).

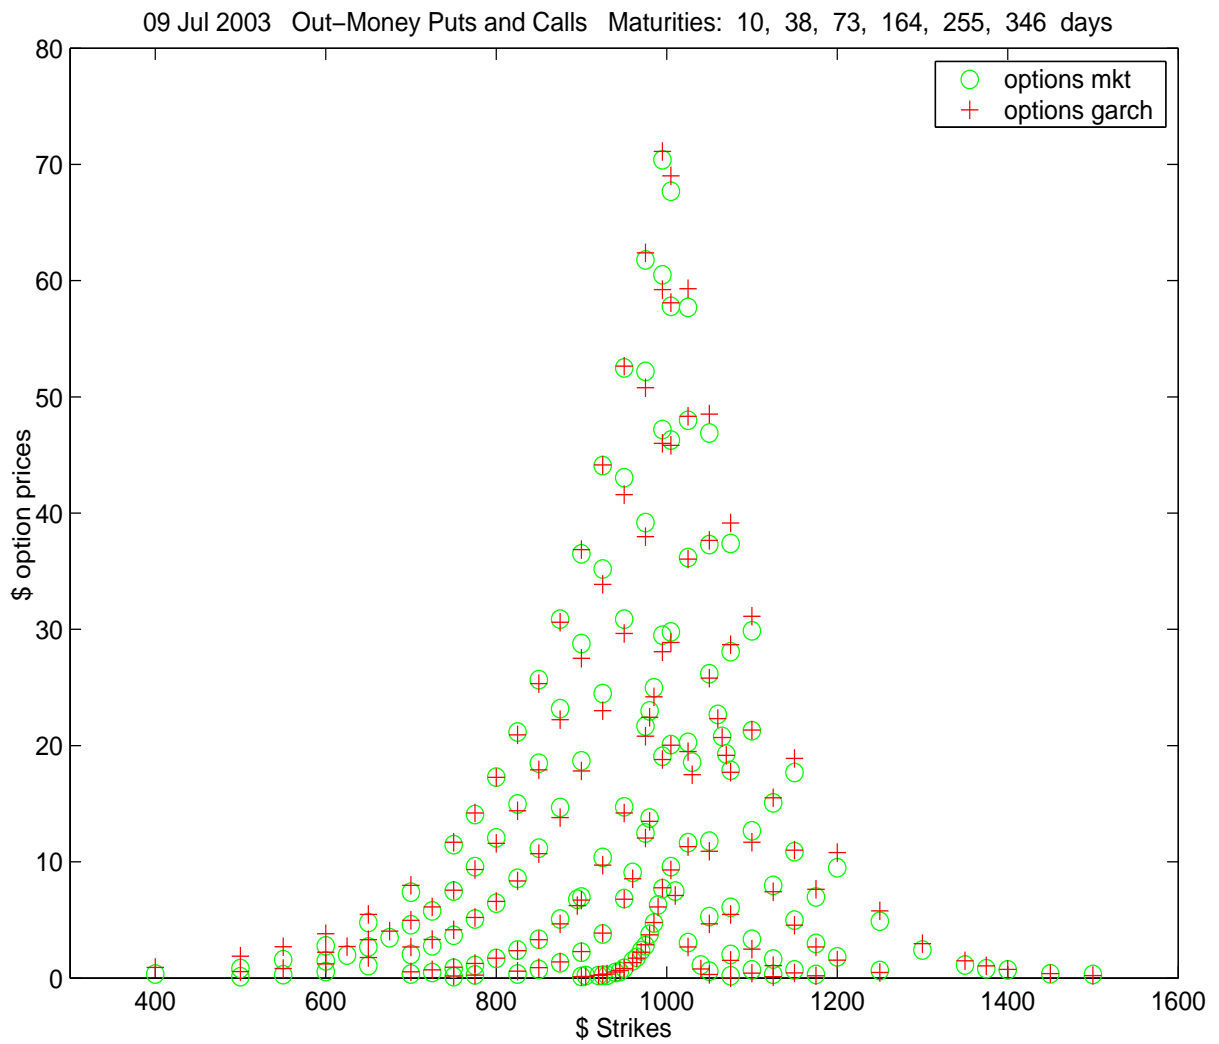


Figure 5: FHS calibration results of the GARCH model to $m = 151$ out of the money European put and call option prices observed on July 9, 2003.

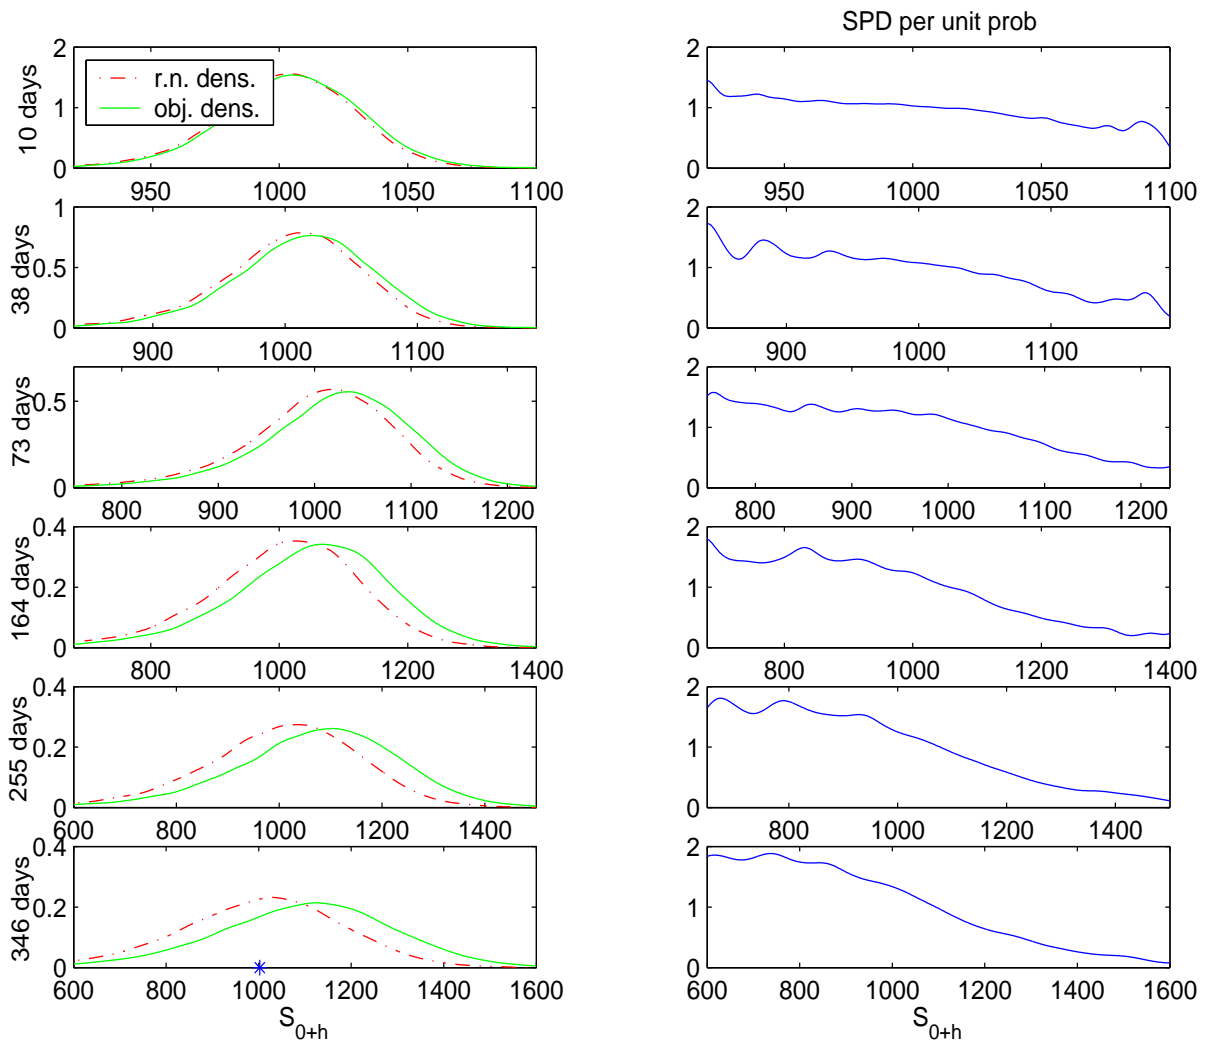


Figure 6: Risk neutral and objective density estimates (left plots) and state price density estimates per unit of probability (right plots) for July 09, 2003.

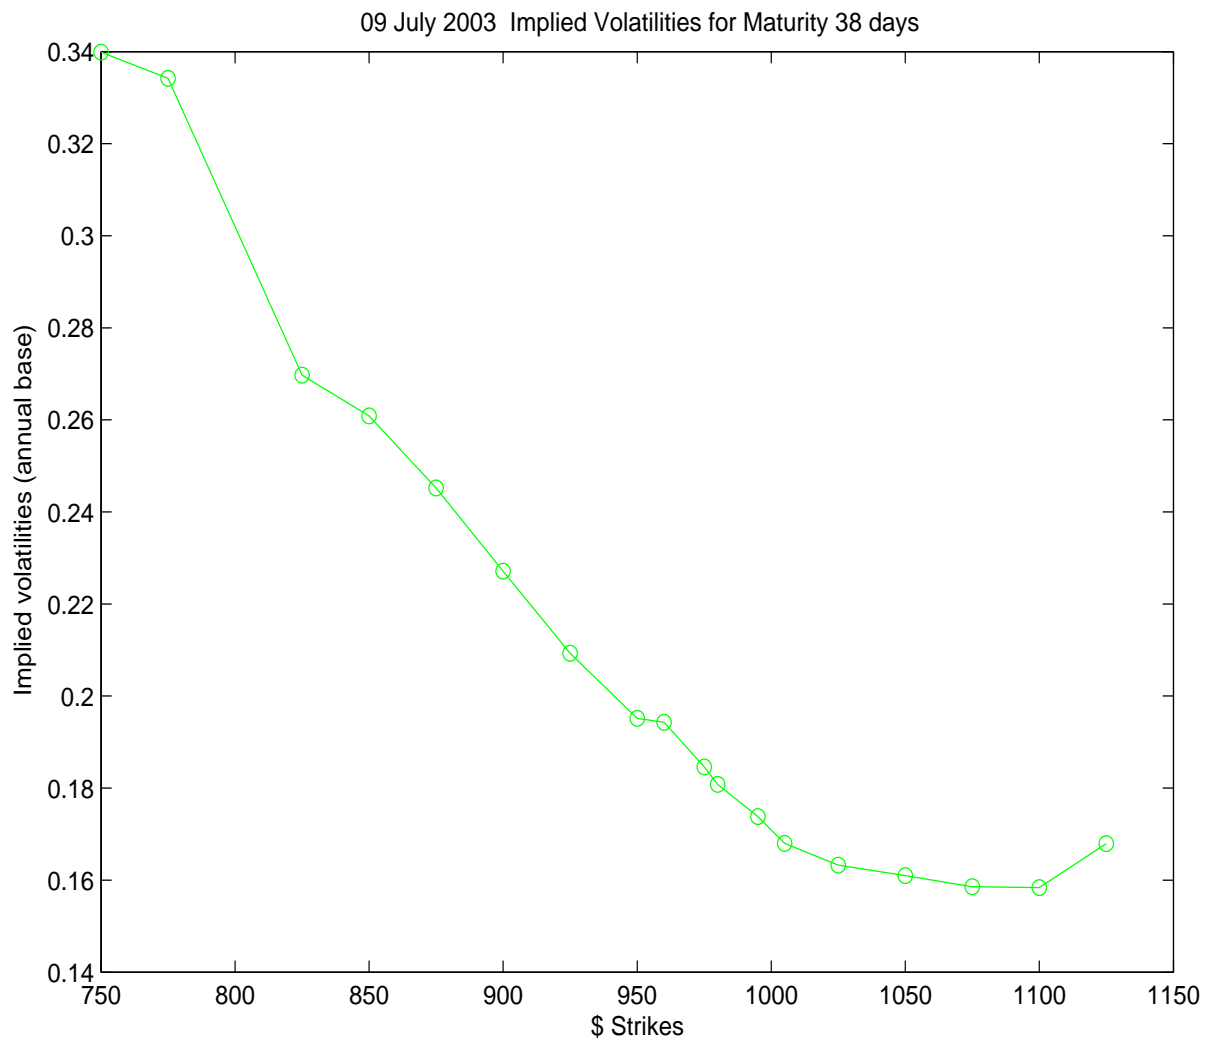


Figure 7: Implied volatilities observed on July 9, 2003 from out of the money European put and call options with maturity $T = 38$ days.