

A GARCH Option Pricing Model in Incomplete Markets*

Giovanni Barone-Adesi^a

Robert F. Engle^b

Loriano Mancini^c

^aSwiss Finance Institute at the University of Lugano, Switzerland

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

^cSwiss Banking Institute, University of Zurich, Switzerland

First draft: October 2004

This version: February 2007

***Correspondence Information:** Giovanni Barone-Adesi, Swiss Finance Institute at the University of Lugano, Via Buffi 13, CH-6900 Lugano, Switzerland, E-mail address: baroneg@lu.unisi.ch. E-mail addresses for Robert Engle and Loriano Mancini: rengle@stern.nyu.edu, mancini@isb.unizh.ch. Giovanni Barone-Adesi and Loriano Mancini gratefully acknowledge the financial support of the Swiss National Science Foundation (NCCR FinRisk). We received helpful comments from Robert Elliott, Jens Jackwerth, Claudia Ravanelli, as well as seminar participants at University of Konstanz, the 2004 AFFI Meetings in Paris and AMASES in Modena, and the 2005 RMFE Conference at the University of Florida.

A GARCH Option Pricing Model in Incomplete Markets

Abstract

We propose a new method for pricing options based on GARCH models with filtered historical innovations. In an incomplete market framework we allow for different distributions of the historical and the pricing return dynamics enhancing the model flexibility to fit market option prices. An extensive empirical analysis based on S&P 500 index options shows that our model outperforms other competing GARCH pricing models and ad hoc Black-Scholes models. Using our GARCH model and a nonparametric approach we obtain decreasing state price densities per unit probability as suggested by economic theory, validating our GARCH pricing model. Implied volatility smiles appear to be explained by the negative asymmetry of the filtered historical innovations. A new simplified delta hedging scheme is presented based on conditions usually found in option markets, namely the local homogeneity of the pricing function. We provide empirical evidence and we quantify the deterioration of the delta hedging in the presence of large volatility shocks.

Keywords: Option pricing, GARCH model, state price density, delta hedging, Monte Carlo simulation.

JEL Classification: G13.

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—as well documented in the financial literature. However the success of GARCH in modeling historical return variances only partially extends to option pricing. Duan (1995), Heston (1993a) and Heston and Nandi (2000) among others assume normal return innovations and parametric risk premiums to derive GARCH pricing models. These assumptions allow to consider in a unified framework the historical and the pricing (or risk neutral) asset return dynamics. Unfortunately, they also imply that, up to the risk premium, the conditional volatility of the historical and the pricing distributions be governed by the same model parameters. Empirical studies for instance by Chernov and Ghysels (2000) and Christoffersen and Jacobs (2004) show that this restriction leads to rather poor pricing and hedging performances. The reason is that a risk neutral investor would price the option as if the distribution of its return had a different drift but unchanged volatility. This is certainly true in the Black and Scholes (1973) theory and the related option pricing models. However, Black and Scholes derived the above property under very special assumptions.¹ Changing volatility in real markets makes the perfect replication argument of Black and 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.²

In the markets we consider the volatility of the pricing process is different from the historical volatility of the asset process. This occurs because investors will set state prices to reflect their aggregate preferences. Our model differs from other models in the financial literature because we rely on market incompleteness to allow for a pricing distribution different in shape from the historical 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, the stability of the pricing process over time and across maturities imposes substantial parameter restrictions. Furthermore, economic theory imposes further restrictions on investor preferences for aggregate wealth in different states, such as decreasing intertemporal marginal rate of substitutions. Pricing models are economically validated when these restrictions are satisfied.

¹Frictionless markets, continuous trading, unlimited riskless borrowing and lending opportunities at the same constant rate, geometric Brownian motion with a known and constant diffusion coefficient for the underlying asset.

²For instance Jarrow and Madan (1995) investigate the hedging of systematic jumps in asset returns when additional assets are introduced in the market.

An interesting result concerns hedging performance. The delta hedging strategy based on the Black-Scholes model calibrated at the implied volatility cannot be substantially improved using any alternative pricing model. This result is due to the fact that deltas from the Black-Scholes model can be derived applying directly the (local first degree) homogeneity of option prices with respect to asset and strike prices, without using the Black-Scholes formula. Hence hedge ratios from the Black-Scholes model calibrated at the implied volatility are the “correct” hedge ratios unless a very strong departure from local homogeneity of the pricing function occurs. This is not the case for the continuous, almost linear volatility smiles commonly observed on option markets. 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 the Black-Scholes model at each implied volatility does not give a model of option pricing, the hedging performance of this common procedure—as often implemented by practitioners in the financial industry—can be very accurate. Barone-Adesi and Elliott (2006) further investigate the computation of the hedge ratios under similar assumptions.

This paper presents three main contributions that are the new GARCH pricing model, the analysis of aggregate intertemporal marginal rate of substitutions and the new simplified hedging scheme. Importantly, an in depth empirical study underlies all the previous analysis. Our GARCH pricing model relies on the Glosten, Jagannathan, and Runkle (1993) asymmetric volatility model driven by empirical GARCH innovations. We perform an extensive empirical analysis using European options on the S&P 500 Index from January 2002 to December 2004. We compare in- and out-of-sample pricing performances of our approach, the GARCH pricing models of Heston and Nandi (2000) and Christoffersen, Heston, and Jacobs (2006) and the benchmark model of Dumas, Fleming, and Whaley (1998). The calibration exercise is particularly extended as all the pricing models are calibrated each week in our sample. Interestingly, we find that our GARCH pricing model outperforms all the other pricing methods both in- and out-of-sample in almost all comparisons. To economically validate our approach we estimate the state price densities per unit probability (or the aggregate intertemporal marginal rate of substitutions) for all the available maturities in our sample. In contrast to previous studies, (such as Jackwerth (2000) and Rosenberg and Engle (2002)), we undertake a much more extensive empirical analysis and more importantly our estimates of the state price densities per unit probability tend to be monotonically decreasing as predicted by economic theory. Finally, we study the hedging performance of the new simplified scheme and we provide empirical evidence that a major deterioration of the delta hedging occurs in the presence of large volatility shocks.

The financial literature on GARCH option pricing is quite large and in the following we provide

only a partial overview. Heston and Nandi (2000) derive a closed-form pricing formula assuming normal return innovations, a linear risk premium and the same GARCH parameters for the historical and the pricing asset processes. In our pricing model we rely on Monte Carlo simulations and then we can relax their assumptions. Duan (1996) calibrates a GARCH model to the FTSE 100 index options assuming Gaussian innovations and the locally risk neutral valuation relationship, that is the daily conditional variances under the historical and the pricing measures are equal. Engle and Mustafa (1992) propose a similar method calibrating a GARCH model to S&P 500 index options to investigate the persistence of volatility shocks under the historical and the pricing distributions. Recent studies on GARCH pricing models include for instance Christoffersen, Heston, and Jacobs (2006) and Christoffersen, Jacobs, and Wang (2006).

The first section of this paper develops our theoretical framework, the second one presents our empirical findings on option pricing. The third section discusses hedging and the fourth one summarizes our conclusions.

1 Theoretical framework

1.1 State price densities and derivative prices

The relation between state price density (SPD), asset price dynamic and the representative agent's preferences is now well understood, but for completeness and to develop the notation we provide a brief summary.

In a dynamic equilibrium model (such as Rubinstein (1976b) or Lucas (1978)), the price of any asset equals the expected present value of its future payoff, where the discounting is calculated using the riskless rate r and the expectation is taken with respect to the marginal rate of substitution weighted probability density function (PDF) of the payoffs. This PDF is called the SPD or risk neutral PDF (Cox and Ross (1976)) or equivalent martingale measure (see Harrison and Kreps (1979)) or pricing density. Formally, the current price ψ_t of an asset with a single payoff ψ_T at date $T = t + \tau$ is

$$\psi_t = E_{\mathbb{Q}}[\psi_T e^{-r\tau} | \mathcal{F}_t] = e^{-r\tau} \int_0^{\infty} \psi_T(S_T) q_{t,T}(S_T) dS_T, \quad (1)$$

where S_T is a state variable of the economy (e.g., aggregate consumption), \mathcal{F}_t is the available information to the agent up to and including time t , $q_{t,T}$ is the SPD (the PDF under the risk neutral measure \mathbb{Q}) at time t for payoffs liquidated at date T , and $E_{\mathbb{Q}}$ is the expectation under \mathbb{Q} . The

security price ψ_t can be equivalently represented as

$$\psi_t = E_{\mathbb{P}}[\psi_T M_{t,T} | \mathcal{F}_t] = \int_0^\infty \psi_T(S_T) M_{t,T}(S_T) p_{t,T}(S_T) dS_T, \quad (2)$$

where $M_{t,T}$ is the SPD per unit probability³ and $p_{t,T}$ is the PDF under the historical or objective measure \mathbb{P} at time t for payoffs liquidated at date T . In a continuum of states, the SPD defines the Arrow-Debreu (Arrow (1964) and Debreu (1959)) security price, which gives for each state of the economy, s , the price $M_{t,T}(s) p_{t,T}(s)$ of a security paying one dollar at time T if the future state S_T falls between s and $s + ds$. The SPD per unit probability $M_{t,T}(s)$ is then the market price of a security paying $1/p_{t,T}(s)$ dollars in the previous scenario and the security has an expected rate of return equals to $1/M_{t,T}(s) - 1$ under the historical measure \mathbb{P} . These expected rate of returns depend on the current state of the economy summarized in the information set \mathcal{F}_t .

Equations (2) shows the high information content of the SPD per unit probability. This equation can be used to determine equilibrium asset prices given historical asset price dynamics and agent preferences, or to infer the agents characteristics given the observed market asset prices. For instance, in a Lucas (1978) economy, the state variable S_T is the aggregate consumption C_T and $M_{t,T} = U'(C_T)/U'(C_t)$, that is the intertemporal marginal rate of substitution. Using an unconditional version of equation (2) and the aggregate consumption C_T as state variable, Hansen and Singleton (1982) and Hansen and Singleton (1983) estimate the risk aversion and the time preference of the representative agent, hence identifying the SPD per unit probability.

In general the SPD per unit probability will depend not only on current and future consumptions but also on all variables that affect marginal utility, such as past consumptions or equity market returns; see for instance the discussion in Rosenberg and Engle (2002, Section 2.1) and the references therein. Furthermore, due to the well-known measurement problems and the low temporal frequency of aggregate consumption data,⁴ several researchers have proposed alternative methods to estimate $M_{t,T}$ substituting consumption data with market data. For instance, Rosenberg and Engle (2002) project $M_{t,T}$ onto the payoffs of traded assets ψ_T , avoiding the issue of specifying the state variables in the SPD per unit probability. Ait-Sahalia and Lo (2000) and Jackwerth (2000) adopt similar approaches projecting $M_{t,T}$ onto equity return states using S&P 500 index option prices. They assume that investors have a finite horizon and that the equity index level equals the aggregate

³The SPD per unit probability is also known as asset pricing kernel (Rosenberg and Engle (2002)) or stochastic discount factor; see Campbell, Lo, and MacKinlay (1997) and Cochrane (2001) for comprehensive surveys of its role in asset pricing, and Ross (1978), Harrison and Kreps (1979), Hansen and Richard (1987), Hansen and Jagannathan (1991) for related works.

⁴For discussions of these issues see for example Ferson and Harvey (1992), Wilcox (1992), and Slesnick (1998).

wealth. In this paper we also make the same assumptions.⁵ Our goal will be to develop a pricing model for the derivative securities, ψ_T , that accounts for the most important features of equity returns. The model will be evaluated on the basis of statistical measures (derived by the mispricing of existing securities) as well as economic measures (whether or not the corresponding SPD per unit probability satisfies minimum economic criteria).

1.2 Asset price dynamics

In this section we develop the stochastic volatility model that captures the most important features of the equity return process. This model will be used in the option pricing analysis.

1.2.1 Historical return dynamics

A substantial amount of empirical evidence suggest that equity return volatility is stochastic and mean reverting, return volatility responds asymmetrically to positive and negative returns, and return innovations are non-normal; see for instance Ghysels, Harvey, and Renault (1996). In a discrete time setting, the stochastic volatility is often modeled using extensions of the autoregressive conditional heteroscedasticity (ARCH) model proposed by Engle (1982). Comprehensive surveys of the ARCH and related models are Bollerslev, Chou, and Kroner (1992) and Bollerslev, Engle, and Nelson (1994). In a continuous time setting, stochastic volatility diffusion model are commonly used; surveys of this literature are for instance Ghysels, Harvey, and Renault (1996) and Shephard (1996).

To model the equity index return we use an asymmetric GARCH specification with an empirical innovation density. The GARCH model of Bollerslev (1986) accounts for stochastic, mean reverting volatility dynamics. The asymmetry term is based on Glosten, Jagannathan, and Runkle (1993) (GJR). The empirical innovation density captures potential non-normalities in the true innovation density and we refer to this approach as the filtering historical simulation (FHS) method. Barone-Adesi, Bourgoin, and Giannopoulos (1998) introduce the FHS method to compute portfolio risk measures and Engle and Gonzalez-Rivera (1991) investigate the theoretical properties of the GARCH model.

Under the objective or historical measure \mathbb{P} the asymmetric GJR GARCH model is

$$\begin{aligned} \log(S_t/S_{t-1}) &= \mu + \varepsilon_t \\ \sigma_t^2 &= \omega + \beta\sigma_{t-1}^2 + \alpha\varepsilon_{t-1}^2 + \gamma I_{t-1}\varepsilon_{t-1}^2, \end{aligned} \tag{3}$$

⁵See for example Rubinstein (1976a) and Brown and Gibbons (1985) for the conditions under which the SPD per unit probability with consumption growth rate as state variable is equivalent to a SPD per unit probability with equity index return as state variable.

where $\varepsilon_t = \sigma_t z_t$, $z_t \sim f(0, 1)$ and $I_{t-1} = 1$, when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$, otherwise. When $\gamma > 0$ the model accounts for the leverage effect,⁶ that is bad news ($\varepsilon_{t-1} < 0$) rises the future volatility more than good news ($\varepsilon_{t-1} \geq 0$) of the same absolute magnitude. The constant expected return μ of the log-return is not usually compatible with time varying state prices, but it is unlikely that it will strongly affect the estimation of state prices. Hence, equation (3) can be viewed as an approximation. The scaled return innovation z_t is drawn from the empirical density function f that is obtained dividing each estimated return innovations $\hat{\varepsilon}_t$ by its estimated conditional volatility $\hat{\sigma}_t$. This set of estimated scaled innovations gives an empirical density function that incorporates excess skewness, kurtosis and other extreme return behaviors that are not captured in a normal density. We estimate the model parameters using maximum likelihood with a normal innovation density. Bollerslev and Wooldridge (1992) show that under some conditions this technique provides consistent parameter estimates even when the true innovation density f is non-normal. In the last case the parameter estimates are pseudo maximum likelihood (PML) estimates.

The GARCH literature provides several specifications for the conditional volatility and we could adopt a different GARCH model in this study. We favor the asymmetric GARCH model (3) because of its flexibility for instance in capturing the leverage effect. Engle and Ng (1993) document that the asymmetric GARCH model provides an adequate modeling of the news impact curve and it outperforms for example the EGARCH model of Nelson (1991). Furthermore, Rosenberg and Engle (2002) also investigate the asymmetric GARCH model (3) using S&P 500 index returns as in our empirical application and they find that the model performs well in describing the returns dynamics. In Section 2 we will undertake an extensive empirical analysis involving several years of S&P 500 index returns and in some instances there might be other GARCH models outperforming the asymmetric GARCH model (3). However, for comparison purposes we always maintain the GJR model in our study.

1.2.2 Pricing return dynamics

The asset return model (3) is specified under the historical measure \mathbb{P} and cannot be directly used to price options. One possibility is to specify the SPD per unit probability (i.e., the change of

⁶The name leverage effect was introduced by Black (1976) who suggested that a large negative return increases the financial and operating leverage, and rises equity return volatility; see also Christie (1982). Campbell and Hentschel (1992) suggested an alternative explanation based on the market risk premium and volatility feedback effects; see also the more recent discussion by Bekaert and Wu (2000). We shall use the name leverage effect as it is commonly used by researchers when referring to the asymmetric reaction of volatility to positive and negative return innovations.

measure from \mathbb{P} to \mathbb{Q}) and to recover the risk neutral asset dynamics.⁷ This approach is particularly attractive because the GARCH parameters $\theta = \{\omega, \beta, \alpha, \gamma\}$ can be easily estimated using historical asset returns and used for pricing purposes. However, several studies such as Chernov and Ghysels (2000) and Christoffersen and Jacobs (2004) show that this approach leads to a rather poor pricing performance and it is largely dominated by option pricing models estimated only using option prices.⁸ This negative result is mainly due to the difficulties of specifying a correct SPD per unit probability, that in general has time varying level and shape; see for instance Rosenberg and Engle (2002).

We do not attempt to directly specify the change of measure from \mathbb{P} to \mathbb{Q} , and we propose to approximate it by calibrating a new set of risk neutral GARCH parameter $\theta^* = \{\omega^*, \beta^*, \alpha^*, \gamma^*\}$ directly on market option prices. The parameter θ^* can possibly be different from the GARCH parameter θ under \mathbb{P} , and this difference characterizes the SPD per unit probability. Hence, under the risk neutral measure \mathbb{Q} the asset return dynamic is approximated by an asymmetric GARCH model with different parameters,

$$\begin{aligned} \log(S_t/S_{t-1}) &= \mu^* + \varepsilon_t \\ \sigma_t^2 &= \omega^* + \beta^* \sigma_{t-1}^2 + \alpha^* \varepsilon_{t-1}^2 + \gamma^* I_{t-1} \varepsilon_{t-1}^2, \end{aligned} \tag{4}$$

where the risk neutral drift μ^* ensures that the asset return equals the risk free rate, that is $E_{\mathbb{Q}}[S_t/S_{t-1}|\mathcal{F}_{t-1}] = e^r$. The distribution of the scaled innovation z could also be changed to better approximate the change of measure from \mathbb{P} to \mathbb{Q} . We retain the same historical distribution of the innovation process under the risk neutral measure \mathbb{Q} because in our empirical application the parameter change from θ to θ^* already provides a flexible change of measure to fit well market options prices. More importantly, we will investigate whether or not the induced SPD per unit probability is economically sustainable and satisfies the usual economic restrictions on the level and shape. In that case, it will be economically validated. The SPD per unit probability is estimated by discounting the ratio of the historical and pricing densities derived from the two different GARCH models, that is $M_{t,t+\tau} = e^{-r\tau} q_{t,t+\tau}/p_{t,t+\tau}$. No a priori model is imposed on $M_{t,t+\tau}$ and its estimate will provide a substantial amount of information on the aggregate investor preferences.

⁷For instance Rubinstein (1976b) and Brennan (1979) originally used log-normal distributions and power utilities to characterize the SPD per unit probability, while Heston (1993b) and Stutzer (1996) combine log-exponential distributions with power utility and exponential utility for the same purpose.

⁸Although Chernov and Ghysels (2000) find that the pricing of long term options in the Heston (1993a) model could be improved using security and option contracts jointly and Christoffersen and Jacobs (2004) find that the out-of-sample performances of certain GARCH models estimated on historical data are not so inferior to the ones of the same GARCH models calibrated only on option data.

If the return innovation z was to be normal, the SPD per unit probability would be restricted as in Duan (1995, Lemma A.1) because both historical and risk neutral distributions would be normal and the risk neutral asset returns could not follow a GARCH process with different parameters.⁹ We remark that our pricing model does not rely on normal innovations, but on non-normal empirical innovations. Hence in the spirit of Chernov and Ghysels (2000), we combine the historical asset returns given by the FHS innovations $\{z_t\}$ and the cross sectional market option prices to infer the option pricing model. We recall that the innovation process $\{z_t\}$ is not transformed and we exploit the flexibility of θ^* to compensate for that.

The previously mentioned calibration approach is usually undertaken in the option pricing literature; see for instance Engle and Mustafa (1992), Duan (1996), Heston and Nandi (2000), and Christoffersen, Jacobs, and Wang (2006).¹⁰ In contrast to the previous papers however we consider wider moneyness and/or maturity ranges for option prices and we undertake a more extensive calibration exercise, calibrating our GARCH pricing model and the competing models on each week for the three years in our sample. Furthermore, we estimate the SPD per unit probability for all the available maturities and we perform an hedging analysis based on our GARCH pricing model.

Under the GARCH model the distribution of temporally aggregated asset returns cannot be derived analytically. Hence, for any given horizon $T = t + \tau$, the SPD is estimated using the FHS method and the calibrated pricing GARCH parameters. The τ periods return density is estimated by simulating several τ periods return paths. A return path is simulated by drawing an estimated past innovation $z_{1,t+1}$, updating the conditional variance σ_{t+1}^2 , drawing a second innovation $z_{2,t+2}$, updating the conditional variance σ_{t+2}^2 , and so on up to day T . The τ periods simulated return is $S_T/S_t = \exp(\tau\mu^* + \sum_{i=1}^{\tau} \sigma_{t+i} z_{i,t+i})$. To reduce the Monte Carlo variance we use the empirical martingale simulation method proposed by Duan and Simonato (1998), where the simulated asset price paths are re-scaled to ensure that the risk neutral expectation of the underlying asset equals its forward price. Denoting by $S_T^{(l)}$ the simulated asset price at time T in the l -th sample path, the price of a call option at time t with strike K and maturity T is given by $e^{-r\tau} \sum_{l=1}^L \max(S_T^{(l)} - K, 0)$, where L is the total number of simulated sample paths. Put option prices are similarly computed.

⁹We are grateful to the referee for pointing this out.

¹⁰See also Amin and Ng (1997) for related work on interest rate models.

2 Empirical analysis

2.1 The data

To test our model we use European options on the S&P 500 index (symbol: SPX). The market for these options is one of the most active index options market in the world. Expiration months are the three near-term months and three additional months from the March, June, September, December, quarterly cycle. Strike price intervals are 5 and 25 points. The options are European and have no wild card features. SPX options can be hedged using the active market on the S&P 500 futures. Consequently, these options have been the focus of many empirical investigations, including Ait-Sahalia and Lo (1998), Chernov and Ghysels (2000), Heston and Nandi (2000), and Carr, Geman, Madan, and Yor (2003).

We consider closing prices of the out-of-the-money (OTM) put and call SPX options for each Wednesday¹¹ from January 2, 2002 to December 29, 2004. Option data and all the other necessary data are downloaded from OptionMetrics. The average of bid and ask prices are taken as option prices, and options with time to maturity less than 10 days or more than 360 days, implied volatility larger than 70%, or prices less than \$0.05 are discarded, which yields a sample of 29,211 observations. Put and call options are equally represented in the sample, that is 50.7% and 49.3%, respectively. In-the-money options are not actively traded compared to at-the-money and out-of-the-money options. For instance, the daily volume for out-of-the-money put options is usually several times as large as the volume for in-the-money puts. This phenomenon started after the October 1987 crash and reflects the strong demand by portfolio managers for protective puts—inducing the well-known implied volatility smile. Using only out-of-the-money options avoids potential issues associated to liquidity problems. Moreover, some empirical studies investigate model pricing performances separately using calls and puts, and the empirical findings based on calls and puts are rather similar; see, for instance, Bakshi, Cao, and Chen (1997) and Dumas, Fleming, and Whaley (1998).

Using the term structure of default-free interest rates, the riskless interest rate for each given maturity τ is obtained by linearly interpolating the two interest rates for maturities straddling τ . This procedure is repeated for each contract and each day in our sample.

We divide the option data into several categories according to either time to maturity or moneyness, m , defined as the ratio of the strike price over the asset price, K/S . A put option is said to be deep out-of-the-money if its moneyness $m < 0.85$; out-of-the-money if $0.85 \leq m < 1$; a call option

¹¹In our sample all but two days are Wednesdays. In those cases we take the subsequent trading day, but for simplicity we refer to all days as Wednesdays.

is said to be out-of-the-money if $1 \leq m < 1.15$; and deep out-of-the-money if $m \geq 1.15$. An option contract can be classified, by the time to maturity, as short maturity (< 60 days); medium maturity (60–160 days); and long maturity (> 160 days).

Table 1 describes the 29,211 option prices, the implied volatilities, and the bid-ask spreads used in our empirical analysis. The average put (call) prices range from \$0.77 (\$0.34) for short maturity, deep OTM options to \$38.80 (\$34.82) for long maturity, OTM options. OTM put and call options account for, respectively, 27% and 25% of the total sample. Short and long maturity options account for, respectively, 33% and 36% of the total sample. The table also shows the volatility smile and the corresponding term structure. For each given set of maturities the smile across moneyness is evident. When the time to maturity increases the smile tends to become flatter and the bid-ask spreads tend to narrow. The number of options on each Wednesday is on average 186.1, with a standard deviation of 22.3, a minimum of 142 and a maximum of 237 option contracts.

During the sample period, the S&P 500 index ranges from a minimum of \$776.8 to a maximum of \$1,213.5 with an average level of \$1,029.5. The average daily log-return is quite close to zero (6.6×10^{-5}), the standard deviation is 22.98% on an annual base,¹² and skewness and kurtosis are 0.25 and 4.98, respectively.

2.2 Calibration of the GARCH model with the FHS method

The pricing GARCH parameter $\theta^* = \{\omega^*, \beta^*, \alpha^*, \gamma^*\}$ are determined by calibrating GARCH option prices computed with the FHS method on the cross section of SPX out-of-the-money option prices. The procedure is in two steps. For each Wednesday t , in the first step the GARCH model (3) is estimated using n historical log-returns of the underlying asset up to time t , $\{\log(S_j/S_{j-1}), j = 1 - n + t, 2 - n + t, \dots, t\}$. The scaled innovations $\{z_j = \varepsilon_j/\sigma_j\}$ are also estimated. In the second step, the pricing GARCH parameters θ^* are obtained by minimizing the mean square error criterion function $\sum_{i=1}^{N_t} e(K_i, T_i)^2$, where $e(K_i, T_i)$ is the difference between the GARCH model price and the market price of the SPX option i with strike K_i and maturity T_i .¹³ N_t is the number of SPX out-of-the-money options on a given Wednesday t . The GARCH option price is computed using the FHS

¹²The standard deviation of the S&P 500 log-returns is approximately in line with the GARCH unconditional volatility estimates reported in Tables 2 and 3.

¹³To minimize the criterion function we use the Nelder-Mead simplex direct search method implemented in the Matlab function `fminsearch`, which does not require the computation of gradients. To ensure the convergence of the calibration algorithm, the FHS innovations used to simulate the GARCH sample paths are kept fix across all the iterations of the algorithm. Starting values for the pricing parameters θ^* are the GARCH parameters estimated under the historical measure \mathbb{P} and obtained in the first step of the FHS calibration method.

method as described in the previous section and the current conditional volatility σ_{t+1} is obtained using historical GARCH estimates and asset returns. Under a theoretical viewpoint this approach is motivated by a limiting argument, that is the SPD per unit probability $M_{t,t+\tau}$ tends to one, when τ goes to zero. Then the current historical and risk neutral volatilities coincide. Under a practical viewpoint, this approach reduces the calibration time because σ_{t+1} is kept fix at each iteration of the calibration algorithm.

In the GARCH framework the conditional variance σ_{t+1}^2 is readily computed using the historical asset returns and the historical GARCH parameter as in our approach or the risk neutral GARCH parameter as in Heston and Nandi (2000) or Christoffersen, Heston, and Jacobs (2006). This is an important advantage of the GARCH pricing models over continuous time stochastic volatility models where the current instantaneous volatility is not observable and is usually calibrated to option prices hence increasing the computational burden; see for instance Bakshi, Cao, and Chen (1997).

2.3 Benchmark model and alternative GARCH pricing models

In the empirical analysis we compare our option pricing model to three competing approaches, that are a benchmark model and the GARCH pricing models of Heston and Nandi (2000) and Christoffersen, Heston, and Jacobs (2006).

We follow Dumas, Fleming, and Whaley (1998) to define the benchmark model. For each Wednesday t the implied volatilities of the cross section of the SPX options are smoothed across strikes and time to maturities by fitting the following function

$$\sigma^{\text{bs}} = a_0 + a_1K + a_2K^2 + a_3\tau + a_4\tau^2 + a_5K\tau, \quad (5)$$

where σ^{bs} is the Black-Scholes implied volatility for an option with strike K and time to maturity τ . Option prices are then obtained by plugging in the Black-Scholes formula the fitted implied volatilities. The equation (5) is estimated every Wednesday using ordinary least squares. This approach is called the ad hoc Black-Scholes model. Although theoretically inconsistent, ad hoc Black-Scholes methods are routinely used in the option pricing industry and they represent a more challenging benchmark than simple Black-Scholes models (with a single implied volatility for the whole cross section of options), because they allow for different implied volatilities to price different options. Indeed, Dumas, Fleming, and Whaley (1998) show that this approach outperforms the deterministic volatility function option valuation model introduced by Derman and Kani (1994), Dupire (1994), and Rubinstein (1994).¹⁴

¹⁴See for instance Buraschi and Jackwerth (2001) for further discussions on this point.

We also consider the GARCH model with Gaussian innovations of Heston and Nandi (2000) (HN, in brief) and the GARCH model with Inverse Gaussian innovations of Christoffersen, Heston, and Jacobs (2006) (IG, in brief). Heston and Nandi (2000) apply the inversion of the characteristic function technique introduced by Heston (1993a) to derive a closed form expression for European option prices under GARCH models with Gaussian innovations. To capture the well documented negative asymmetry in GARCH innovations, Christoffersen, Heston, and Jacobs (2006) extend the previous technique to GARCH models with possibly negative skewed inverse Gaussian innovation. Here we only recall the pricing formulae and the specification of the two GARCH models. We refer the reader to the corresponding papers for an in depth discussion of the models.

In the HN model, the asset return dynamic under the risk neutral measure \mathbb{Q} is

$$\begin{aligned}\log(S_t/S_{t-1}) &= r - \sigma_t^2/2 + \sigma_t z_t \\ \sigma_t^2 &= \omega_{\text{hn}}^* + \beta_{\text{hn}}^* \sigma_{t-1}^2 + \alpha_{\text{hn}}^* (z_{t-1} - \gamma_{\text{hn}}^* \sigma_{t-1})^2,\end{aligned}$$

where z_t is a standard Gaussian innovation, and in the IG model

$$\begin{aligned}\log(S_t/S_{t-1}) &= r + \nu \sigma_t^2 + \eta^* y_t \\ \sigma_t^2 &= w^* + b^* \sigma_{t-1}^2 + c^* y_{t-1} + a^* \sigma_{t-1}^4 / y_{t-1},\end{aligned}$$

where y_t follows an inverse Gaussian distribution with parameter $\delta_t = \sigma_t^2 / \eta^{*2}$. In both models, at time t the call option C_t with strike price K and time to maturity τ is worth

$$C_t = e^{-r\tau} \zeta^*(1) \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re \left[\frac{K^{-i\phi} \zeta^*(i\phi + 1)}{i\phi \zeta^*(1)} \right] d\phi \right) - e^{-r\tau} K \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re \left[\frac{K^{-i\phi} \zeta^*(i\phi)}{i\phi} \right] d\phi \right), \quad (6)$$

where $\Re[\cdot]$ denotes the real part of a complex number, $\zeta^*(\phi)$ is the conditional moment generating function at time t of the log-price $X_T = \log(S_T)$,

$$\zeta^*(\phi) = E_{\mathbb{Q}}[e^{\phi X_T} | \mathcal{F}_t] = S_t^\phi e^{A_t + B_t \sigma_{t+1}^2}.$$

The coefficients A_t 's and B_t 's are computed backward starting from the terminal condition $A_T = B_T = 0$, and using the following recursive equations in the HN model

$$\begin{aligned}A_t &= A_{t+1} + \phi r + B_{t+1} \omega_{\text{hn}}^* - \frac{1}{2} \log(1 - 2\alpha_{\text{hn}}^* B_{t+1}) \\ B_t &= \phi \left(\gamma_{\text{hn}}^* - \frac{1}{2} \right) - \frac{\gamma_{\text{hn}}^{*2}}{2} + \beta_{\text{hn}}^* B_{t+1} + \frac{1/2(\phi - \gamma_{\text{hn}}^*)^2}{1 - 2\alpha_{\text{hn}}^* B_{t+1}}\end{aligned}$$

and in the IG model

$$\begin{aligned}A_t &= A_{t+1} + \phi r + B_{t+1} w^* - \frac{1}{2} \log(1 - 2a^* \eta^{*4} B_{t+1}) \\ B_t &= B_{t+1} b^* + \phi \nu^* + \eta^{*-2} - \eta^{*-2} \sqrt{(1 - 2a^* \eta^{*4} B_{t+1})(1 - 2c^* B_{t+1} - 2\eta^* \phi)}.\end{aligned}$$

Similarly to GJR model with FHS method, on each Wednesday the HN and IG models will be calibrated on the cross section of out-of-the-money SPX options by minimizing the dollar mean square error of the pricing errors.

2.4 Implementing the GARCH pricing models

On each Wednesday, the GARCH model (3) is estimated using the PML approach and $n = 3,500$ historical log-returns, that is a sample size of approximately fourteen years. As an example Figure 1 shows the S&P 500 log-returns, the estimated GARCH volatility $\{\sigma_t\}$ and the scaled innovations $\{z_t\}$ up to a randomly chosen date, July 9, 2003. A sufficiently long sample size ensures that the innovation distribution is adequately estimated. As a robustness check, we repeated all the estimations and calibration exercises using a sample size $n = 2,500$. This point will be further discussed in the next section. In the calibration of the GARCH pricing model (4), the SPD is estimated by simulating $L = 20,000$ sample paths of the underlying asset S with the FHS method and the empirical martingale simulation approach as already discussed.¹⁵ The sample paths are simulated in parallel running only one for-loop, on a standard Pentium IV processor with 1GB RAM. Hence our Monte Carlo approach could be potentially inaccurate—as any other simulation method—but the computation time is roughly the same as for the competing GARCH pricing models, where one for-loop is needed to compute the coefficients A_t 's and B_t 's.¹⁶ Whether or not the Monte Carlo option prices are sufficiently accurate will be reflected in the empirical pricing performance of our approach. However, for almost all the deepest out-of-the-money options, at least 100 simulated paths end “in-the-money”. Figure 2 shows the calibration of the GJR model with FHS method on the cross section of out-of-the-money SPX options on July 9, 2003. Given the wide range of moneyness and maturities considered, the calibration is quite satisfactory. An extensive empirical analysis will be undertaken in the next section.

When computing option prices using the HN and IG models, the integrals in equation (6) are computed numerically via discretization. We take as integration domain the interval $(0, 100)$ and we evaluate the integrand function on 5,000 equally spaced mid-points. Then, the integrals are computed by averaging the function values over such an interval.¹⁷ As the moment generating

¹⁵For a few selected days we repeated the calibration exercises using $L = 10,000$ sample paths and we obtained very similar pricing results. As the computation burden was still manageable we used $L = 20,000$ sample paths in our empirical study.

¹⁶If the number of simulated sample paths L has to be increased and the computer memory is not enough then the Monte Carlo approach is substantially slower than the competing GARCH approaches.

¹⁷The choice of the integration domain and the subinterval length are selected to match option prices on a few

function ζ^* does not depend on the strike K , for each given maturity τ , the coefficients A_t 's and B_t 's need to be computed only once. Hence, following the previous numerical procedure a whole cross section of option prices can be evaluated in a few seconds, largely reducing the calibration time. As in Christoffersen, Heston, and Jacobs (2006), we set the riskless rate $r = 0.05/365$ in the GARCH pricing formula. At each step of the calibration procedure, the conditional variances σ^2 are initialized at the unconditional variance levels and then updated using the corresponding pricing GARCH dynamics. The iterations are started 250 trading days before the first option date to allow for the models to find the right conditional variances. Hence, in contrast to other stochastic volatility models, such as the Heston (1993a) or the Bakshi, Cao, and Chen (1997) models, the pricing formula (6) does not require the calibration of σ_{t+1} , simplifying our calibration procedure.

Christoffersen, Heston, and Jacobs (2006) show that the ν^* parameter in the IG model is not a free parameter, but it is constrained to ensure that the underlying asset earns the risk free rate under the risk neutral measure \mathbb{Q} . This observation explains why ν^* does not appear in Table 3, which shows the calibrated risk neutral parameters.

Finally, when implementing all the GARCH pricing formulae the dividends paid by the stocks in the S&P 500 index have to be taken into account. The dividends are treated in different ways in the option pricing literature; see for instance Aït-Sahalia and Lo (1998) and Heston and Nandi (2000) for two alternative procedures. We use the dividend yields downloaded from OptionMetrics to compute an ex-dividend spot index level.

2.5 In-sample model comparison

For each Wednesday from January 2, 2002 to December 29, 2004 we estimate the GJR model using the PML method and $n = 3,500$ past S&P 500 log-returns; see the first panel in Table 2. Hence these estimates are under the historical measure \mathbb{P} and are used in the first step of the FHS calibration procedure. Then on each Wednesday we calibrate the GJR model with FHS innovations as well as the HN and the IG models to the cross section of out-of-the-money SPX options. Table 3 shows summary statistics for the GARCH parameters under the pricing measure \mathbb{Q} . It is known that option prices are more sensible for example to γ^* , γ_{hn}^* , c^* and η^* than ω^* , ω_{hn}^* and w^* . Indeed, the first set of parameters turn out to be more stable than the second one, confirming the finding for instance in Heston and Nandi (2000). As expected all the unconditional volatilities and the corresponding persistency measures are more stable than the single parameters. For each Wednesday, we also selected days computed using the Romberg's numerical integration method over the interval $(10^{-6}, 200)$ with tolerance 10^{-4} . Some differences between the two approaches are observed only for a few deepest out-of-the-money option prices.

estimate the ad hoc Black-Scholes model (5) (BS, in brief) and Table 4 shows summary statistics for the parameter estimates a_i , $i = 0, \dots, 5$. The modest standard deviations of the parameter estimates confirm that the implied volatility smile is a persistent phenomenon in SPX option market; see also Table 1.

Following Dumas, Fleming, and Whaley (1998) and Heston and Nandi (2000), to assess the quality of the fitted model, we report several measurements of fit. The dollar root mean square error (RMSE) is the square root of the averaged squared deviations between the model prices and the market prices; the mean absolute error (MAE) is the average of the absolute valuation error when the model price is outside the bid-ask spread; the mean outside error (MOE) is the average of the valuation error when the model price is outside the bid-ask spread; the minimum and maximum pricing errors; the percentage of positive pricing errors; the average pricing error as a percentage of the bid-ask spread.

For all pricing models Table 5 shows the in-sample pricing errors summarized by the previous measurements of fit. Table 6 disaggregates the in-sample pricing errors by moneyness/maturity categories; see also the left plots in Figure 3 for the in-sample absolute mispricing. The overall conclusion is that in terms of all measurements of fit our GJR model with FHS innovations outperforms all the other pricing models. Occasionally, the GJR is somewhat outperformed by other models—for instance in terms of RMSE by the HN model during 2003 in Table 5, or by the IG model for deep out-of-the-money put options in Table 6—but its good pricing performance is remarkably stable. Although the percentage of positive pricing errors in Table 5 for the HN and the IG models seem to be too low, no systematic biases seem to emerge for the MOE's in Table 6. Hence the previous finding is due to slight underestimations of option prices. Finally, all the GARCH pricing models always outperform the benchmark ad hoc Black-Scholes model, showing also some difficulties of the BS model in fitting the whole cross section of option prices.

2.5.1 News impact curves

It might be surprising that the GJR with FHS innovations outperforms all the other GARCH pricing models, sometimes even by a large extent. After all, also the HN and IG models can account for the same stylized facts about the asset returns and volatility dynamics. To shed light on the differences among the GARCH models we compute the news impact curve (NIC) introduced by Engle and Ng (1993). The news impact curve measures the impact on the conditional variance σ^2 of a shock in the scaled innovation z , when the volatility is at its long run level. Figure 4 shows the NIC for the GJR, HN and IG models using the pricing GARCH parameters in Table 3. The plots have to

be carefully interpreted because the NIC is estimated using the average pricing parameter values and to simplify the interpretation the vertical axis shows the conditional volatility rather than the conditional variance. However, the findings are remarkably stable across the three years 2002, 2003 and 2004. In all GARCH models a negative shock z raises the volatility more than a positive shock of equal magnitude, with some exceptions for the IG model. This effect is much stronger in GJR than in the HN model. Moreover, the Gaussian innovation in the HN model very rarely will be below, say, -3 , while in the GJR model FHS innovations are often below this threshold; see for instance Figure 1. The IG model can potentially react very strongly to negative shocks because the inverse Gaussian innovation, y_t , enters as $1/y_t$ in the GARCH dynamics. Hence NIC tends to explode when y_t goes to zero, although the probability of this event depends on the degrees of freedom parameter $\delta_t = \sigma_t^2/\eta^{*2}$.

2.5.2 Comparison with the CGMY models

We also compare the pricing performance of the GJR model with FHS innovations to the CGMYSA model proposed by Carr, Geman, Madan, and Yor (2003). Under their model the underlying asset follows a mean corrected, exponential Lévy process time changed with a Cox, Ingersoll and Ross process. Carr, Geman, Madan, and Yor (2003) calibrate the CGMYSA model to the out-of-the-money SPX 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. As an overall measure of the quality of the calibration they propose the average absolute pricing error (APE) with respect to the mean price,

$$\text{APE} = \frac{\sum_{i=1}^{N_t} |P^{\text{model}}(K_i, T_i) - P^{\text{market}}(K_i, T_i)|}{\sum_{i=1}^{N_t} P^{\text{market}}(K_i, T_i)}, \quad (7)$$

where P^{model} is either the GJR or the CGMYSA model price, and P^{market} is the SPX market option price. The pricing performances of the GJR and CGMYSA models are reported in Table 7. The average absolute pricing errors are somewhat in favor of the CGMYSA model, but this model has nine parameters while the GJR model has four parameters. Carr, Geman, Madan, and Yor (2003) also proposed 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, in terms of the APE measure, the GJR model outperforms the VGSA and NIGSA models in five and four out of six cases, respectively. Finally, there is evidence that the GJR parameters tend to change from month to month, but the pricing performance is quite stable especially in terms of the APE measure.

2.6 Out-of-sample model comparison

Out-of-sample forecasting of option prices is an interesting challenge for any pricing method. It tests not only the goodness-of-fit of the pricing formulae, but also whether or not the pricing methods overfit the option prices in the in-sample period. For each Wednesday in our sample, in-sample model estimates are used to price SPX options one week later (hence out-of-sample) using asset prices, time to maturities and interest rates relevant on the next Wednesday.¹⁸ The conditional variances are updated in our GARCH model with FHS innovations using the current historical parameter estimates in Table 2, while in the other GARCH pricing models using the calibrated pricing parameters in Table 3. The actual S&P 500 daily log-returns are also used to update all conditional variances.

Table 8 shows the out-of-sample pricing errors summarized by the different measurements of fit and Table 9 disaggregates the out-of-sample pricing errors by moneyness/maturity categories; see also the right plots in Figure 3 for the out-of-sample absolute mispricing. Interestingly, the out-of-sample results largely confirm the in-sample ones and overall our GJR model with FHS innovations outperforms all the competing pricing models. Most of the out-of-sample measurements of fit share the same patterns as the in-sample ones. Hence our GARCH pricing model is flexible enough to achieve a good pricing performance, capturing the pricing mechanisms, without overfitting the data.

In all the previous in- and out-of-sample pricing exercises the GJR is estimated and calibrated using $n = 3,500$ past S&P 500 log-returns. As a robustness check we repeat all the previous analysis using $n = 2,500$. We obtain very similar results in terms of parameter estimations and calibrations as well as pricing performances. Hence the corresponding tables are omitted, but are available from the authors upon request.

The simulated sample paths used to calibrate the GJR GARCH pricing model contain a large amount of information on the underlying and derivative securities. For instance, storing the simulated sample paths a sensitivity analysis (such as computing delta and gamma) can be easily obtained following the direct approach in Duan (1995) or the regression methodology approach in Longstaff and Schwartz (2001). Hence our Monte Carlo calibration procedure can be computationally intensive—like any other calibration approach—but the resulting amount of information could balance the computational costs. Moreover, the previous out-of-sample results show that for pricing purposes our GARCH pricing model does not need to be recalibrated every day. Then the whole cross section

¹⁸The one week ahead forecast horizon is also adopted by Dumas, Fleming, and Whaley (1998) and Heston and Nandi (2000), among others, and for comparative purposes we adopt it here as well.

of option contracts can be priced in about one second running the parallel simulation.

2.7 State Price Density per unit probability

In this section we estimate the SPD per unit of probability, $M_{t,t+\tau}$, using historical and pricing densities, $p_{t,t+\tau}$ and $q_{t,t+\tau}$, respectively. The historical density $p_{t,t+\tau}$ is estimated by simulating with the FHS method the GJR model fitted to historical S&P 500 log-returns up to time t .¹⁹ The density $p_{t,t+\tau}$ depends on the expected return, which is notoriously difficult to estimate. Indeed, Merton (1980) argues that positive risk premia should be explicitly modeled. Hence, as in Jackwerth (2000), we set the risk premium at 8% per year. The pricing density $q_{t,t+\tau}$ is estimated by simulating with the FHS method the GJR model calibrated on the cross section of out-of-the-money SPX options observed on date t . Then, $M_{t,t+\tau}$ is estimated by taking the discounted ratio of the pricing over the historical densities, $M_{t,t+\tau} = e^{-r\tau} q_{t,t+\tau} / p_{t,t+\tau}$. This procedure gives semiparametric estimates of $M_{t,t+\tau}$, because the densities $p_{t,t+\tau}$ and $q_{t,t+\tau}$ are based on parametric GARCH models, but no a priori functional form is imposed on $M_{t,t+\tau}$. For a randomly chosen date, namely July 9, 2003, Figure 5 shows the historical and pricing densities and the SPD per unit probability for different times to maturities estimated using the GJR models with the FHS method; see also Table 10 for the numerical values of $M_{t,t+\tau}$. The state price densities per unit probability are rather stable across the different maturities and fairly monotonically decreasing in S_T , with some exceptions for the longest maturity and the lowest states. The state price densities per unit probability outside the plotted values of S_T tend to be unstable, as the density estimates are based on very few observations. The documented stability of SPD per unit probability is consistent with aggregate risk aversion changing slowly through time at a given wealth level.

To assess the economic relevance of the FHS approach in estimating SPD per unit probability, we also estimate $M_{t,t+\tau}$ using the historical and pricing GJR models driven by Gaussian innovations. The pricing GJR model with Gaussian innovations is also calibrated on the cross section of out-of-the-money SPX options, obtaining the corresponding pricing GARCH parameters.²⁰ Then, the pricing density $q_{t,t+\tau}$ is obtained by simulating the GJR model with the pricing GARCH parameters and the randomly drawn Gaussian innovations. The historical density $p_{t,t+\tau}$ is obtained by simulating the

¹⁹All the density functions are estimated using the Matlab function `ksdensity` with the Gaussian kernel and the optimal default bandwidth for estimating Gaussian densities.

²⁰The calibration procedure is as the same as for the GJR model with the FHS method, but where the FHS innovations are replaced by randomly drawn Gaussian innovations. The same Gaussian innovations are also used to estimate the historical and pricing GARCH densities for the estimation of $M_{t,t+\tau}$'s.

GJR model using the historical parameter and the previous Gaussian innovations used to simulate the pricing density. To reduce the variance of the Monte Carlo estimates we use the method of antithetic variates; see, for instance, Boyle, Broadie, and Glassermann (1997).²¹ Figure 5 also shows the SPD per unit probability estimated with the GJR model with Gaussian innovations, that are clearly higher than $M_{t,t+\tau}$ with the FHS method for the low states, but very close for medium and high states of the S&P 500. This difference is economically significant. High values of $M_{t,t+\tau}$ on the left imply that simple state contingent claims are “overpriced” and have very negative expected rate of returns. As an example, a SPD per unit probability of \$3.8 estimated using the GJR model with Gaussian innovations implies an expected rate of returns of $1/(3.8) - 1 = -73.7\%$ over an horizon of 73 days, that might be “unreasonably” low. In contrast, for the same contingent claim a SPD per unit probability of \$1.4 estimated using the GJR model with FHS innovation implies an expected rate of returns of $1/(1.4) - 1 = -28.6\%$, that is certainly a more “reasonable” return. As a consequence, according to $M_{t,t+\tau}$ based on Gaussian innovations, out-of-the-money puts appear to be “overpriced”, while using FHS innovations they are not. Therefore, the negative skewness in the FHS innovation can account for the excess out-of-the-money put prices and provides an adequate pricing of the downside market risk.

We support the previous analysis by estimating the SPD per unit probability for each Wednesday t and each time to maturity τ in our sample, using both GJR GARCH models with FHS and Gaussian innovations. We obtain 862 density estimates of $M_{t,t+\tau}$ for each model. The pricing GARCH parameters for the GJR with Gaussian innovations are reported in the second panel of Table 2. To visually represent the SPD per unit probability we follow a similar approach as in Jackwerth (2000) showing average estimates of $M_{t,t+\tau}$. For each date t and maturity τ , the $M_{t,t+\tau}$ ’s are first estimated on a grid of 100 points of the gross return $\{S_{t+\tau}/S_t\}$, and then averaged across each year and each short, medium, long maturity categories. The averaged SPD per unit probability, and ± 0.5 times the empirical standard deviations, are plotted in Figure 6. The findings are aggregated, but still indicative of the respective periods and time to maturities covered and the results largely confirm the previous discussion based on $M_{t,t+\tau}$ estimated on July 9, 2003. In particular, the SPD per unit probability with FHS method are quite stable across time and maturities, and for low states they are significantly lower than the SPD per unit probability with Gaussian innovations. Moreover, in contrast to previous studies (for instance Jackwerth (2000) and Rosenberg and Engle (2002)) almost

²¹The method of antithetic variates is not used in our FHS approach to preserve the negative asymmetry of the historical GARCH innovations.

all the SPD per unit probability are monotonically decreasing as expected.²²

Comparing Tables 2 and 3 shows that the asymmetry parameter γ^* tend to be larger when the GJR model is calibrated using Gaussian innovation rather than FHS innovations. Indeed, a similar phenomenon occurs in the Black-Scholes model. The Gaussian distribution of the log-returns cannot account for the market expectations on the frequency of large negative returns and to recover observed out-of-the-money put option prices the volatility parameter has to be increased. This phenomenon induces the well-known volatility smile; see Table 1 for the volatility smile in our database. In the GARCH framework, despite the stochastic volatility, Gaussian innovations cannot account for the previous market expectations and to fit out-of-the-money put prices a large asymmetry parameter γ^* is needed. Using FHS innovations the GJR model can more easily price the out-of-the-money put prices giving SPD per unit probability in line with economic reasonings. When GARCH models are driven by Gaussian innovations, Duan (1995) shows that the historical and the pricing parameters are equal. Table 2 provides empirical evidence that in the GJR model the two sets of parameters are quite different. This finding points out the theoretical inconsistency of the GJR model with Gaussian innovations and can support the GJR model with FHS innovations—given the documented good pricing performance.

Several other studies provide estimates of SPD per unit probability. Perhaps the most closely related study is provided by Rosenberg and Engle (2002). The main similarity is that in both cases, conditional SPD per unit probability are investigated, hence reflecting current market conditions such as high or low volatility. The main difference is that Rosenberg and Engle (2002) consider a flexible but parametric functional form for $M_{t,t+\tau}$, while here we provide nonparametric estimates of the ratio $q_{t,t+\tau}/p_{t,t+\tau}$. Aït-Sahalia and Lo (2000) and Jackwerth (2000) provide nonparametric estimates of SPD per unit probability (or closely related quantities such as risk aversion functions), but they exploit the time continuity of $M_{t,t+\tau}$ to average state prices and probabilities across time. Hence, their result can be interpreted as estimates of the average SPD per unit probability over the sample period, instead of the conditional estimate; for a recent extension see Constantinides, Jackwerth, and Perrakis (2006). Finally, compared to all the previous studies we provide estimates of SPD per unit probability for many more time horizons and much larger ranges of $S_{t+\tau}/S_t$.

²²In this paper we use a different methodology, sample period and moneyness/maturity ranges for the SPX options compared to the other studies.

3 Hedging

In this section we present a new approach to compute hedge ratios for implementing classical delta hedging strategies and we study hedging performances in the presence of volatility shocks.

3.1 Numerical example

Our main result is that the delta hedging strategy based on the Black-Scholes model calibrated on the implied volatility surface cannot be largely enhanced using any alternative pricing model. To understand why this is the case we provide a simple numerical example in Table 11. The three rows in the middle are call option market prices.²³ The first row is obtained multiplying the middle row by 0.9 and the last row is obtained multiplying the middle row by 1.1, that is assuming a first degree homogeneous pricing model, $C(\kappa S, \kappa K) = \kappa C(S, K)$, where C is the call option pricing function and κ is close to one. In general, market option prices might not be homogenous functions of the underlying asset and strike prices. However, for the computation of hedge ratios we will not impose the homogeneity property everywhere, but only as a local approximation. The last assumption is weaker and based on the local continuity of the pricing function C as a function of the stock price S and the strike price K .

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 $\Delta_{55} := (2.64 - 1.00)/(53.9 - 49)$. Notice that the homogeneity property was used to determine option prices associated to values of the underlying asset that are different from the current value, $S = 49$. Taking the average of these two ratios, for the strike price $K = 50$ gives an estimate of delta equals to 0.518, which is almost identical to the delta 0.522 from the Black-Scholes model calibrated at the implied volatility 20.0% for the middle row. Therefore, the “almost” model free hedge ratio (based on market prices and the local homogeneity of the pricing function) is very close to the Black-Scholes hedge ratio commonly used by practitioners. This result implies that pricing models alternative to Black-Scholes calibrated at the implied volatility will generally lead to very similar hedge ratios, if they fit well market prices. It appears therefore that deltas are to a large degree determined by market option prices, independently of the chosen model.

In the above computation of the hedge ratio 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

²³The implied volatilities are 21.2%, 20.0% and 22.1%, respectively, the time to maturity is 20/52 years and the risk free rate r is 5%. These options are used for illustrative purposes and as an example of volatility smile.

be homogeneous functions of the stock and the strike price and the hedge ratio computed using the previous two approaches would be very close; see Barone-Adesi and Elliott (2006) for a numerical experiment. 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} . To investigate the discretization error let $\Delta(K)$ denote the delta as a function of the strike price K . For small intervals of the strike price the delta hedge ratio is approximated by

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

where $\Delta'(K)$ is the first derivative of $\Delta(K)$. 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 small if the strike price increment is small relative to the asset price and its volatility. This is generally the case for the strike price increments of the SPX options, that are about 1% of the strike price. The reader may verify this simple result on the options of his choice.

3.2 Impact of volatility shocks on delta hedging

The only significant deterioration of hedging occurs in the presence of large volatility shocks, which diminish the effectiveness of delta hedging. To document this phenomenon we compare the out-of-sample forecasts of option prices for two different days affected by volatility shocks of different magnitudes. We consider the Thursday $(t_1 + 1) =$ September 19, 2002, where a relatively large increase in the volatility was generated by the preceding Wednesday negative log-return $\log(S_{t_1}/S_{t_1-1}) = -3.01\%$, and the Thursday $(t_2 + 1) =$ July 10, 2003, where a modest change in volatility was induced by the preceding Wednesday log-return $\log(S_{t_2}/S_{t_2-1}) = -1.35\%$.²⁴ Then, to evaluate the out-of-sample forecasts we run the following Mincer-Zarnowitz regressions (see Mincer and Zarnowitz (1969) and Chong and Hendry (1986)) for each date t_i , $i = 1, 2$,

- 1) $P_{t+1}^{\text{market}} = \eta_0 + \eta_1 P_{t,t+1}^{\text{bs}} + \text{error}_{t+1}$
- 2) $P_{t+1}^{\text{market}} = \eta_0 + \eta_1 P_{t,t+1}^{\text{bs}} + \eta_2 P_{t,t+1}^{\text{garch}} + \text{error}_{t+1}$
- 3) $P_{t+1}^{\text{market}} = \eta_0 + \eta_2 P_{t,t+1}^{\text{garch}} + \text{error}_{t+1},$

²⁴Recall that in the GARCH model, the volatility of the log-return $\log(S_{t+1}/S_t)$ is already determined at date t and depends on the return innovation ε_t driving the log-return $\log(S_t/S_{t-1})$ at time t .

where P_{t+1}^{market} are the out-of-the-money SPX option prices observed on date $t + 1$, $P_{t,t+1}^{\text{bs}}$ are the one day ahead out-of-sample Black-Scholes forecasts of option prices at time $t + 1$ computed by plugging in the Black-Scholes formula the stock price S_{t+1} , the risk free rate r , the dividend yield d at time $t + 1$ and the implied volatility observed at time t , i.e. on the Wednesdays September 18, 2002 and July 9, 2003. $P_{t,t+1}^{\text{garch}}$ are the one day ahead out-of-sample GARCH forecasts obtained using S_{t+1} , r , d at time $t + 1$, the GARCH parameter calibrated at time t and σ_{t+1} updated according to the historical GARCH model estimated at time t . We apply the previous filtering criteria to the SPX option prices on the Thursdays $t_1 + 1$ and $t_2 + 1$. If the out-of-sample forecasts are accurate, η_0 should be close to zero, and η_1 or η_2 close to one. The ordinary least square estimates, t-statistics and the variance of the forecast errors of the previous regressions are given in the first two panels of Table 12. All R^2 are above 0.99 and the F-statistics are highly significant, meaning that overall the forecasts are quite accurate. The variance of the forecast errors using only the Black-Scholes forecasts in regressions 1) are smaller than using only the GARCH forecasts in regressions 3) for both Wednesdays 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. These results are due to the advantage of the Black-Scholes model calibrated at each point of the volatility surface at time t , and therefore having zero pricing errors initially. However, for Wednesday $t_1 = \text{September 18, 2002}$, where a large volatility shock occurs, from regression 1) to regression 2) the variance of the prediction error is reduced by 50% adding the GARCH forecast as a regressor. The η_2 parameter is highly significant and the likelihood ratio (LR) test of the regression model 1) versus the regression model 2) strongly reject the null hypothesis, i.e. the regression model 1). Hence, the GARCH model carries on a large amount of information on option price dynamics. The main reason is that it provides a dynamic model for the volatility, while the Black-Scholes model does not. Interestingly, the large positive value of η_0 in regression 1) for the Wednesday t_1 means that the Black-Scholes forecasts heavily underestimate P_{t+1}^{market} . Adding the GARCH forecasts accounts for the aforementioned increase in volatility, reducing η_0 and the bias in the forecast regression 2). In contrast to the Black-Scholes model, the GARCH model is able to detect the increase in volatility due to the large negative returns, and this phenomenon is reflected in its forecasts of volatilities and option prices. For the Wednesday $t_2 = \text{July 9, 2003}$, the reduction in the variance of the prediction errors is more modest and equals to 19%, although the η_2 parameter is highly significant and the LR test strongly reject the null hypothesis given by the regression model 1). Hence, the GARCH forecasts have more explanatory power when the change in volatility is more pronounced. This point will be deeply investigated in a while.

To support the previous findings, we repeat the regression analysis for each Wednesday t from January 2, 2002 to December 29, 2004. All the regressions have R^2 above 0.98 and F-statistics highly significant. In all cases but one, the coefficients η_2 are statistically significant at 1% confidence level. The p-values for LR tests of the regression model 1) versus model 2) are below 10% in 79.7% of the occasions, and below 1% in 70.6% of the occasions. These findings strongly confirm the explanatory power of GARCH forecasts in predicting volatilities and option prices.

To deepen the analysis we investigate whether the explanatory power of the GARCH model is stronger when the Wednesday log-return is negative rather than positive—for the leverage effect negative returns are expected to increase the volatility more than positive returns. The last two panels in Table 12 show two subsets of results conditioning on the Wednesday log-return $\log(S_t/S_{t-1})$ being rather negative, i.e. below -0.01 , or rather positive, i.e. above 0.01 on a daily base.²⁵ Table 12 presents some empirical evidence to support the hypothesis that the explanatory power of GARCH forecasts is stronger when the Wednesdays log-return is rather negative. For instance, when $\log(S_t/S_{t-1}) < -0.01$, the average η_2 parameter in regressions 2) is 0.24, which is larger than 0.17 in the last panel where $\log(S_t/S_{t-1}) > 0.01$. For the two subsets of negative and positive Wednesdays log-returns, Table 12 shows the means and the standard deviations of the LR p-values of the regression model 1) versus model 2). The skewness of the LR p-values distributions are 2.38 and 1.57, and 83.3% and 54.2% of p-values are below 10%, respectively. This result confirms that GARCH model has higher explanatory power when the Wednesday log-return is rather negative. Figure 7 shows on the vertical axis the reduction in the variance of the forecast errors from the regression model 1) to the regression model 2) and on the horizontal axis the log-return $\log(S_t/S_{t-1})$ for each Wednesday t . We also plot the regression line of the variance reductions on the S&P 500 Wednesdays log-returns. Most of the daily log-returns are close to zero and the R^2 is only 3.8%, but the p-value for the F-test is 2% and the slope is statistically away from zero at 5% confidence level. Also this plot confirms that the GARCH forecasts have stronger impacts in reducing the variance of the forecast errors when the Wednesday log-returns are negative.

Unfortunately, our GARCH forecasts are conditioned on the current index level and cannot be used to improve significantly delta hedging. Their explanatory powers simply indicate that delta hedging is less effective in the presence of large volatility shocks.

²⁵In this analysis we discard the daily log-returns are around zero to evaluate the impact of rather large log-returns on volatility. When also the log-returns around zero are included in this analysis, the results for the two subsets of negative and positive log-returns in Table 12 tend to be rather close.

4 Conclusions

We propose a new method for pricing options based on GARCH models with filtered historical innovations. In an incomplete market framework we allow for the objective and the pricing distributions to have different shapes enhancing the model flexibility to fit market option prices. An extensive empirical analysis based on S&P 500 index options shows that our model outperforms other competing GARCH pricing models and ad hoc Black-Scholes models. In contrast to previous financial literature, using our GARCH model and a nonparametric approach we obtain decreasing state price densities per unit probability as suggested by economic theory and validating our GARCH pricing model. Furthermore, implied volatility smiles appear to be explained by the negative asymmetry of the filtered historical innovations, with no need of “unreasonably” high state prices for out-of-the-money puts. A new simplified delta hedging scheme is presented based on conditions usually found in index option markets, namely the local homogeneity of the pricing function. We provide empirical evidence and we quantify the deterioration of the delta hedging in the presence of large volatility shocks. Further refinements of pricing and stability issues are left to future research.

References

- Aït-Sahalia, Y., and A. W. Lo, 1998, “Nonparametric estimation of state-price densities implicit in financial assets prices,” *Journal of Finance*, 53, 499–548.
- , 2000, “Nonparametric risk management and implied risk aversion,” *Journal of Econometrics*, 94, 9–51.
- Amin, K. I., and V. K. Ng, 1997, “Inferring future volatility from the information in implied volatility in Eurodollar options: A new approach,” *Review of Financial Studies*, 10, 333–367.
- Arrow, K., 1964, “The role of securities in the optimal allocation of risk bearing,” *Review of Economic Studies*, 31, 91–96.
- Bakshi, G., C. Cao, and Z. Chen, 1997, “Empirical performance of alternative option pricing models,” *Journal of Finance*, 52, 2003–2049.
- Barone-Adesi, G., F. Bourgoïn, and K. Giannopoulos, 1998, “Don’t look back,” *Risk*, 11, 100–103.
- Barone-Adesi, G., and R. J. Elliott, 2006, “Cutting the hedge,” *Computational Economics*, forthcoming.
- Bekaert, G., and G. Wu, 2000, “Asymmetric volatility and risk in equity markets,” *Review of Financial Studies*, 13, 1–42.
- Black, F., 1976, “Studies of stock market volatility changes,” in *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistic Section*, pp. 177–181.
- Black, F., and M. Scholes, 1973, “The valuation of options and corporate liabilities,” *Journal of Political Economy*, 81, 637–654.
- Bollerslev, T., 1986, “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 31, 307–327.
- Bollerslev, T., R. Y. Chou, and K. F. Kroner, 1992, “ARCH modeling in finance—a review of the theory and empirical evidence,” *Journal of Econometrics*, 52, 5–59.
- Bollerslev, T., R. F. Engle, and D. B. Nelson, 1994, “ARCH models,” in *Handbook of Econometrics*, ed. by R. F. Engle, and D. L. McFadden. North-Holland, Amsterdam, pp. 2959–3038.

- Bollerslev, T., and J. M. Wooldridge, 1992, “Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances,” *Economics Letters*, 12, 143–172.
- Boyle, P., M. Broadie, and P. Glassermann, 1997, “Monte Carlo methods for security pricing,” *Journal of Economic Dynamics and Control*, 21, 1267–1321.
- Brennan, M. J., 1979, “The pricing of contingent claims in discrete time models,” *Journal of Finance*, 34, 53–68.
- Brown, D. P., and M. R. Gibbons, 1985, “A simple econometric approach for utility-based asset pricing models,” *Journal of Finance*, 40, 359–381.
- Buraschi, A., and J. Jackwerth, 2001, “The price of a smile: Hedging and spanning in option markets,” *Review of Financial Studies*, 14, 495–527.
- Campbell, J. Y., and L. Hentschel, 1992, “No news is good news: An asymmetric model of changing volatility in stock returns,” *Journal of Financial Economics*, 31, 281–318.
- Campbell, J. Y., A. W. Lo, and C. MacKinlay, 1997, *The econometrics of financial markets*. Princeton University Press, Princeton, NY.
- 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.
- Chong, Y. Y., and D. F. Hendry, 1986, “Econometric evaluation of linear macroeconomic models,” *Review of Economic Studies*, 53, 671–690.
- Christie, A., 1982, “The stochastic behavior of common stock variances: Value, leverage and interest rate effects,” *Journal of Financial Economics*, 10, 407–432.
- Christoffersen, P., S. Heston, and K. Jacobs, 2006, “Option valuation with conditional skewness,” *Journal of Econometrics*, 131, 253–284.
- Christoffersen, P., and K. Jacobs, 2004, “Which GARCH model for option valuation?,” *Management Science*, 50, 1204–1221.

- Christoffersen, P., K. Jacobs, and Y. Wang, 2006, "Option valuation with long-run and short-run volatility components," working paper, McGill University, Canada.
- Cochrane, J. H., 2001, *Asset pricing*. Princeton University Press, Princeton.
- Constantinides, G. M., J. C. Jackwerth, and S. Perrakis, 2006, "Mispricing of S&P 500 index options," working paper, University of Chicago.
- Cox, J. C., and S. A. Ross, 1976, "The valuation of options for alternative stochastic processes," *Journal of Financial Economics*, 3, 145–166.
- Debreu, G., 1959, *Theory of value*. Wiley, New York.
- Derman, E., and I. Kani, 1994, "Riding on the smile," *Risk*, 7, 32–39.
- 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.
- Dumas, B., J. Fleming, and R. E. Whaley, 1998, "Implied volatility functions: Empirical tests," *Journal of Finance*, 53, 2059–2106.
- Dupire, B., 1994, "Pricing with a smile," *Risk*, 7, 18–20.
- Engle, R., and G. Gonzalez-Rivera, 1991, "Semiparametric ARCH models," *Journal of Business and Economic Statistics*, 9, 345–359.
- Engle, R. F., 1982, "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation," *Econometrica*, 50, 987–1007.
- Engle, R. F., and C. Mustafa, 1992, "Implied ARCH Models from Options Prices," *Journal of Econometrics*, 52, 289–311.
- Engle, R. F., and V. Ng, 1993, "Measuring and testing the impact of news on volatility," *Journal of Finance*, 48, 1749–1778.
- Ferson, W. E., and C. R. Harvey, 1992, "Seasonality and consumption-based asset pricing," *Journal of Finance*, 47, 511–552.

- Ghysels, E., A. C. Harvey, and E. Renault, 1996, “Stochastic volatility,” in *Handbook of Statistics*, ed. by G. S. Maddala, and C. R. Rao. North-Holland, Amsterdam, pp. 119–191.
- 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.
- Hansen, L. P., and R. Jagannathan, 1991, “Implications of security market data for models of dynamic economies,” *Journal of Political Economy*, 99, 225–262.
- Hansen, L. P., and S. Richard, 1987, “The role of conditioning information in deducing testable restrictions implied by asset pricing models,” *Econometrica*, 55, 587–613.
- Hansen, L. P., and K. J. Singleton, 1982, “Generalized instrumental variables estimation of nonlinear rational expectations models,” *Econometrica*, 50, 1269–1286.
- , 1983, “Stochastic consumption, risk aversion, and the temporal behavior of asset returns,” *Journal of Political Economy*, 91, 249–265.
- Harrison, M., and D. Kreps, 1979, “Martingales and arbitrage in multiperiod securities markets,” *Journal of Economic Theory*, 20, 381–408.
- Heston, S., 1993a, “A closed-form solution for options with stochastic volatility, with applications to bond and currency options,” *Review of Financial Studies*, 6, 327–343.
- , 1993b, “Invisible parameters in option prices,” *Journal of Finance*, 48, 933–947.
- Heston, S., and S. Nandi, 2000, “A closed-form GARCH option valuation model,” *Review of Financial Studies*, 13, 585–625.
- Jackwerth, J. C., 2000, “Recovering risk aversion from option prices and realized returns,” *Review of Financial Studies*, 13, 433–451.
- Jarrow, R., and D. B. Madan, 1995, “Option pricing using the term structure of interest rates to hedge systematic discontinuities in asset returns,” *Mathematical Finance*, 5, 311–336.
- Longstaff, F., and E. S. Schwartz, 2001, “Valuing American options by simulations: A simple least-squares approach,” *Review of Financial Studies*, 14, 113–147.
- Lucas, R., 1978, “Asset prices in an exchange economy,” *Econometrica*, 46, 1429–1446.

- Merton, R. C., 1980, "On estimating the expected return on the market: An exploratory investigation," *Journal of Financial Economics*, 8, 323–361.
- Mincer, J., and V. Zarnowitz, 1969, "The Evaluation of Economic Forecasts," in *Economic Forecasts and Expectations*, ed. by J. Mincer. National Bureau of Economic Research, New York.
- Nelson, D. B., 1991, "Conditional heteroskedasticity in asset returns: A new approach," *Econometrica*, 59, 347–370.
- Rosenberg, J. V., and R. F. Engle, 2002, "Empirical pricing kernels," *Journal of Financial Economics*, 64, 341–372.
- Ross, S. A., 1978, "A simple approach to the valuation of risky streams," *Journal of Business*, 51, 453–475.
- Rubinstein, M., 1976a, "The strong case for the generalized logarithmic utility model as the premier model of financial markets," *Journal of Finance*, 31, 551–571.
- , 1976b, "The valuation of uncertain income streams and the pricing of options," *Bell Journal of Economics*, 7, 407–425.
- , 1994, "Implied binomial trees," *Journal of Finance*, 49, 771–818.
- Shephard, N., 1996, "Statistical aspects of ARCH and stochastic volatility," in *Time Series Models in Econometrics, Finance, and other Fields*, ed. by D. R. Cox, O. E. Barndorff-Nielsen, and D. V. Hinkley. Chapman & Hall, London, pp. 1–67.
- Slesnick, D. T., 1998, "Are our data relevant to the theory? The case of aggregate consumption," *Journal of the American Statistical Association*, 16, 52–61.
- Stutzer, M., 1996, "A simple nonparametric approach to derivative security valuation," *Journal of Finance*, 51, 1633–1652.
- Wilcox, D., 1992, "The construction of U.S. consumption data: Some facts and their implications for empirical work," *American Economic Review*, 82, 922–941.

Moneyness K/S		Maturity					
		Less than 60		60 to 160		More than 160	
		Mean	Std.	Mean	Std.	Mean	Std.
< 0.85	Put price \$	0.77	1.20	2.54	3.20	8.69	7.63
	$\sigma^{\text{bs}}\%$	38.97	9.59	33.36	7.75	28.37	5.21
	Bid-Ask%	0.99	0.69	0.69	0.68	0.27	0.36
	Observations	1,798		2,357		2,845	
0.85–1.00	Put price \$	8.43	7.60	19.12	11.63	38.80	15.67
	$\sigma^{\text{bs}}\%$	22.38	7.05	21.88	5.35	21.28	3.91
	Bid-Ask%	0.19	0.21	0.10	0.05	0.06	0.03
	Observations	3,356		2,136		2,314	
1.00–1.15	Call price \$	7.45	7.81	15.79	12.60	34.82	18.19
	$\sigma^{\text{bs}}\%$	17.55	5.80	17.31	5.02	17.61	4.00
	Bid-Ask%	0.34	0.46	0.17	0.23	0.07	0.04
	Observations	2,955		2,128		2,247	
> 1.15	Call price \$	0.34	0.43	0.85	1.61	3.96	5.60
	$\sigma^{\text{bs}}\%$	34.87	12.35	24.34	8.27	18.87	4.48
	Bid-Ask%	1.81	0.47	1.49	0.71	0.83	0.78
	Observations	1,633		2,288		3,154	

Table 1: Database description. The table shows mean, standard deviation (Std.) and number of observations for each moneyness/maturity category of out-of-the-money SPX options observed on Wednesdays from January 2, 2002 to December 29, 2004, after applying the filtering criteria described in the text. σ^{bs} is the Black-Scholes implied volatility. Bid-Ask% is $100 \times (\text{ask price} - \text{bid price}) / \text{market price}$, where the market price is the average of the bid and ask prices. Moneyness is the strike price divided by the spot asset price, K/S . Maturity is measured in calendar days.

GJR	$\omega \times 10^6$		β		$\alpha \times 10^3$		γ		Persistence		Ann. Vol.	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Year												
2002	1.39	0.20	0.93	0.01	7.44	2.95	0.10	0.01	0.986	0.003	0.19	0.01
2003	1.28	0.30	0.93	0.01	8.31	2.79	0.11	0.01	0.988	0.003	0.21	0.02
2004	1.05	0.06	0.93	0.00	8.08	2.41	0.11	0.00	0.990	0.001	0.20	0.01

GAUSS	$\omega^* \times 10^6$		β^*		$\alpha^* \times 10^3$		γ^*		Persistence		Ann. Vol.	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Year												
2002	3.40	2.42	0.83	0.10	3.64	10.28	0.29	0.16	0.980	0.017	0.27	0.04
2003	2.66	3.05	0.85	0.17	2.61	6.75	0.27	0.31	0.985	0.017	0.26	0.05
2004	1.40	1.25	0.85	0.09	3.68	6.99	0.28	0.16	0.992	0.011	0.28	0.06

Table 2: The first panel shows the parameter estimates of the GJR model estimated for each Wednesday from January 2, 2002 to December 29, 2004 using the PML approach and 3,500 log-returns. The GJR model under the historical measure \mathbb{P} is $\log(S_t/S_{t-1}) = \mu + \varepsilon_t$, $\sigma_t^2 = \omega + \beta\sigma_{t-1}^2 + \alpha\varepsilon_{t-1}^2 + \gamma I_{t-1}\varepsilon_{t-1}^2$, where $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$, and $I_{t-1} = 0$ otherwise. The second panel shows the pricing parameters of the GJR model with Gaussian innovation calibrated on the cross section of out-of-the-money SPX options on each Wednesday from January 2, 2002 to December 29, 2004. The persistence is given by $\beta + \alpha + \gamma/2$ under the historical measure \mathbb{P} and $\beta^* + \alpha^* + \gamma^*/2$ under the risk neutral measure \mathbb{Q} . Ann. Vol. is the annualized volatility.

FHS		$\omega^* \times 10^6$		β^*		$\alpha^* \times 10^3$		γ^*		Persistence		Ann. Vol.			
Year	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	
2002	3.56	2.98	0.86	0.09	1.62	6.65	0.21	0.14	0.97	0.02	0.21	0.02	0.21	0.02	
2003	3.42	4.91	0.85	0.20	2.17	7.25	0.25	0.36	0.97	0.03	0.21	0.03	0.21	0.03	
2004	1.80	2.03	0.85	0.13	2.39	7.34	0.26	0.22	0.98	0.02	0.19	0.02	0.19	0.02	
HN		$\omega_{\text{hn}}^* \times 10^{12}$		β_{hn}^*		$\alpha_{\text{hn}}^* \times 10^6$		γ_{hn}^*		Persistence		Ann. Vol.			
Year	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	
2002	0.02	0.15	0.53	0.23	4.99	3.33	315.74	60.08	0.96	0.03	0.23	0.03	0.23	0.03	
2003	0.05	0.23	0.66	0.15	2.94	1.94	355.37	83.76	0.98	0.02	0.22	0.03	0.22	0.03	
2004	0.08	0.52	0.42	0.22	2.20	1.71	597.84	206.58	0.97	0.02	0.19	0.02	0.19	0.02	
IG		$w^* \times 10^6$		b^*		$a^* \times 10^{-3}$		$c^* \times 10^6$		$\eta^* \times 10^3$		Persistence		Ann. Vol.	
Year	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	
2002	0.14	0.47	0.14	0.20	15.82	14.54	10.77	8.25	-3.88	1.51	0.98	0.02	0.22	0.05	
2003	0.08	0.27	0.17	0.22	23.90	16.79	6.64	2.89	-3.33	1.17	0.99	0.00	0.22	0.03	
2004	0.04	0.10	0.09	0.14	26.29	14.65	5.18	1.64	-2.65	0.44	0.99	0.00	0.18	0.02	

Table 3: The table shows the pricing parameters of the GJR model with the FHS method, $\sigma_t^2 = \omega^* + \beta^* \sigma_{t-1}^2 + \alpha^* \varepsilon_{t-1}^2 + \gamma^* I_{t-1} \varepsilon_{t-1}^2$, where $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$, and $I_{t-1} = 0$ otherwise; the HN model, $\sigma_t^2 = \omega_{\text{hn}}^* + \beta_{\text{hn}}^* \sigma_{t-1}^2 + \alpha_{\text{hn}}^* (z_{t-1} - \gamma_{\text{hn}}^* \sigma_{t-1})^2$, where z_t is a standard Gaussian innovation; the IG model, $\sigma_t^2 = w^* + b^* \sigma_{t-1}^2 + c^* y_{t-1} + a^* \sigma_{t-1}^4 / y_{t-1}$, where y_t has an inverse Gaussian distribution with parameter $\delta_t = \sigma_t^2 / \eta^{*2}$. The GARCH models are calibrated on the cross section of out-of-the-money SPX options for each Wednesday from January 2002 to December 2004. The persistence is given by $\beta^* + \alpha^* + \gamma^*/2$ for FHS, $\beta_{\text{hn}}^* + \alpha_{\text{hn}}^* \gamma_{\text{hn}}^{*2}$ for HN, $b^* + a^* \eta^{*2} + c^* / \eta^{*2}$ for IG, and Ann. Vol. is the annualized volatility.

	a_0	$a_1 \times 10^3$	$a_2 \times 10^7$	$a_3 \times 10$	$a_4 \times 10$	$a_5 \times 10^4$
2002						
Mean	1.11	-1.41	6.23	-1.42	2.57	-2.02
Std.	0.12	0.22	1.09	1.37	1.41	1.83
2003						
Mean	1.07	-1.45	6.48	-1.88	1.64	-0.57
Std.	0.21	0.31	1.32	2.52	1.19	3.12
2004						
Mean	1.35	-1.68	5.68	-5.45	0.93	3.87
Std.	0.11	0.22	1.32	1.75	0.73	1.48

Table 4: The table shows the estimated coefficients of the ad hoc Black-Scholes model, $\sigma^{\text{bs}} = a_0 + a_1K + a_2K^2 + a_3\tau + a_4\tau^2 + a_5K\tau$, where σ^{bs} is the Black-Scholes implied volatility for an option with strike K and time to maturity τ . The model is estimated with ordinary least square on the out-of-sample SPX options for each Wednesday from January 2, 2002 to December 29, 2004.

Panel A: Aggregate valuation errors across all years

	RMSE	MAE	MOE	Min	Max	Err>0%	ErrBD%	MAE%	MOE%
FHS	0.87	0.44	0.08	-6.02	4.64	50.50	23.69	21.16	10.75
HN	1.19	0.74	-0.20	-5.66	6.47	30.65	-44.20	23.93	-25.74
IG	1.27	0.77	-0.10	-5.41	6.34	32.06	-32.31	22.49	-25.82
BS	3.39	2.15	0.06	-15.88	28.66	50.82	69.95	57.91	39.46

Panel B: Valuation errors by years

	RMSE	MAE	MOE	Min	Max	Err>0%	ErrBD%	MAE%	MOE%
2002									
FHS	0.88	0.43	0.06	-5.63	4.64	47.57	14.67	14.80	4.82
HN	1.16	0.72	-0.21	-4.64	5.55	27.59	-48.00	19.55	-33.06
IG	1.26	0.76	-0.15	-5.14	5.89	27.45	-44.15	19.33	-33.87
BS	3.69	2.36	0.03	-15.25	28.66	48.38	51.20	58.95	39.79
2003									
FHS	0.93	0.47	-0.06	-6.02	3.68	44.63	0.95	20.90	6.44
HN	0.92	0.57	-0.28	-3.82	4.92	25.73	-64.94	22.10	-32.13
IG	0.93	0.56	-0.19	-3.72	6.09	27.75	-49.26	20.67	-32.09
BS	3.32	2.12	0.03	-15.88	15.35	49.86	68.37	58.59	38.69
2004									
FHS	0.79	0.43	0.23	-3.41	3.90	58.86	53.82	27.61	20.58
HN	1.41	0.91	-0.12	-5.66	6.47	38.24	-21.06	29.92	-12.60
IG	1.53	0.98	0.02	-5.41	6.34	40.62	-4.86	27.27	-12.08
BS	3.15	1.98	0.12	-13.98	15.18	54.11	89.74	56.26	39.85

Table 5: In-sample pricing errors. RMSE is the root mean square error of the dollar pricing error (model price – market price); MAE is the dollar average absolute pricing error outside the bid-ask spread (MAE% is in relative terms and in percentage); MOE is the dollar average pricing error outside the bid-ask spread (MOE% is in relative terms and in percentage); Min (Max) is the minimum (maximum) pricing error; Err>0% is the percentage of positive pricing errors; ErrBD% is the average of $100 \times (\text{model price} - \text{market price}) / (\text{bid-ask spread})$.

Moneyiness	Maturity														
	Less than 60						60 to 160						More than 160		
	RMSE	MOE	MAE%	RMSE	MOE	MAE%	RMSE	MOE	MAE%	RMSE	MOE	MAE%	RMSE	MOE	MAE%
< 0.85	FHS	0.47	0.09	59.46	0.69	0.28	58.26	0.90	0.50	37.39					
	HN	0.52	-0.19	55.97	0.78	-0.48	28.42	1.56	-1.28	26.95					
	IG	0.45	-0.21	43.56	0.76	-0.48	28.61	1.47	-1.15	23.50					
	BS	0.50	-0.25	38.07	0.86	-0.32	31.74	3.90	0.36	71.41					
0.85-1.00	FHS	1.09	0.09	14.58	1.18	-0.48	3.59	0.95	-0.27	1.16					
	HN	1.40	0.60	23.19	1.01	0.21	5.14	1.21	-0.22	1.89					
	IG	1.68	0.93	24.31	1.26	0.51	6.30	1.29	-0.23	1.92					
	BS	2.55	1.66	26.63	2.14	-0.33	9.44	6.30	-4.72	15.68					
1.00-1.15	FHS	1.10	0.18	19.22	0.96	-0.20	5.77	0.92	0.19	1.82					
	HN	1.53	0.14	35.69	1.43	-0.65	22.03	1.42	0.41	4.70					
	IG	1.69	0.61	30.79	1.47	-0.33	20.92	1.44	-0.03	5.63					
	BS	4.00	3.42	111.73	2.61	1.75	30.90	5.72	-3.00	15.33					
> 1.15	FHS	0.24	-0.08	31.82	0.27	-0.03	12.69	0.68	0.33	16.98					
	HN	0.31	-0.06	17.00	0.69	-0.34	28.71	0.98	-0.57	33.36					
	IG	0.29	-0.06	13.13	0.65	-0.31	27.30	1.06	-0.67	35.80					
	BS	0.67	0.24	49.47	0.77	0.37	60.25	3.03	-0.10	175.45					

Table 6: In-sample pricing errors disaggregated by moneyiness and maturities. For the legend see Table 5.

date	$\omega^* \times 10^6$	β^*	$\alpha^* \times 10^3$	γ^*	Persist.	Ann. Vol.	N_t	Min(τ)	Max(τ)	Avg.	RMSE	FHS APE%	CGMYSA APE%
Jan	1.56	0.91	0.00	0.15	0.991	0.26	177	38	339	28.20	1.62	4.78	3.78
Mar	11.79	0.64	0.00	0.60	0.935	0.26	143	45	283	26.31	1.61	5.13	5.23
May	15.80	0.53	0.00	0.84	0.946	0.33	155	38	311	31.28	1.93	4.74	5.48
Jul	0.55	0.96	0.00	0.07	0.996	0.22	159	38	339	28.33	0.91	2.34	3.26
Sep	4.11	0.87	0.00	0.19	0.960	0.19	151	38	276	20.61	1.08	3.67	2.87
Nov	1.65	0.90	0.00	0.16	0.982	0.18	169	38	318	25.32	1.22	3.74	2.85

Table 7: The table shows the pricing parameters of the GJR model with the FHS method, $\sigma_t^2 = \omega^* + \beta^* \sigma_{t-1}^2 + \alpha^* \varepsilon_{t-1}^2 + \gamma^* I_{t-1} \varepsilon_{t-1}^2$, $I_{t-1} = 1$ when $\varepsilon_{t-1} < 0$ and $I_{t-1} = 0$ otherwise. The GARCH model is calibrated on the cross section of out-of-the-money SPX options for each month in the year 2000. Persist. is the persistence given by $\beta^* + \alpha^* + \gamma^*/2$. Ann. Vol. is the the annualized volatility. N_t is the number of options on a given date. Min(τ) (Max(τ)) is the minimum (maximum) time to maturity. Avg. is the average out-of-the-money SPX price. RMSE is the dollar root mean square error. APE measure is defined in equation (7).

Panel A: Aggregate valuation errors across all years

	RMSE	MAE	MOE	Min	Max	Err>0%	ErrBD%	MAE%	MOE%
FHS	1.48	0.86	0.09	-9.20	7.84	49.92	31.80	26.52	13.09
HN	1.55	0.98	-0.20	-7.69	9.52	30.27	-44.90	25.36	-25.38
IG	2.36	1.33	0.19	-11.12	20.91	35.08	2.48	27.90	-19.61
BS	3.82	2.32	0.18	-18.19	63.31	50.97	88.89	66.38	47.03

Panel B: Valuation errors by years

	RMSE	MAE	MOE	Min	Max	Err>0%	ErrBD%	MAE%	MOE%
2002									
FHS	1.72	1.00	0.03	-9.20	7.10	43.61	18.01	19.64	5.57
HN	1.70	1.07	-0.27	-7.69	6.76	26.33	-53.00	21.16	-33.26
IG	2.22	1.36	-0.09	-11.12	12.65	29.11	-35.65	22.68	-31.56
BS	4.41	2.59	0.27	-16.03	63.31	49.59	80.71	70.94	51.32
2003									
FHS	1.36	0.78	-0.14	-8.16	6.06	41.12	-10.62	23.28	4.81
HN	1.29	0.81	-0.31	-6.99	5.59	24.51	-70.93	23.87	-32.44
IG	1.58	0.96	0.01	-6.52	9.50	30.02	-31.94	22.81	-29.44
BS	3.62	2.27	0.11	-18.19	26.76	49.32	78.23	68.51	47.51
2004									
FHS	1.35	0.81	0.37	-7.12	7.84	63.86	83.99	35.76	27.62
HN	1.63	1.07	-0.03	-5.99	9.52	39.21	-13.22	30.54	-11.67
IG	3.00	1.65	0.63	-7.08	20.91	45.19	69.07	37.38	0.37
BS	3.40	2.10	0.16	-13.46	32.05	53.75	106.26	60.27	42.72

Table 8: Out-of-sample pricing errors. For the legend see Table 5.

Moneyiness	Maturity														
	Less than 60						60 to 160						More than 160		
	RMSE	MOE	MAE%	RMSE	MOE	MAE%	RMSE	MOE	MAE%	RMSE	MOE	MAE%	RMSE	MOE	MAE%
< 0.85	FHS	0.54	0.12	69.43	0.92	0.29	65.21	1.38	0.51	43.09					
	HN	0.56	-0.19	54.61	0.91	-0.50	28.25	1.74	-1.26	26.55					
	IG	0.55	-0.18	45.05	1.04	-0.38	32.24	1.80	-0.98	23.69					
	BS	0.50	-0.23	38.34	0.90	-0.31	31.40	4.63	0.37	74.04					
0.85-1.00	FHS	1.52	0.21	19.35	1.95	-0.49	7.85	2.23	-0.36	3.98					
	HN	1.74	0.58	26.69	1.67	0.23	8.31	2.02	-0.19	3.72					
	IG	3.03	1.41	41.28	3.28	1.11	14.86	3.06	0.28	6.02					
	BS	2.87	1.77	30.10	2.36	-0.19	10.78	6.94	-4.63	16.46					
1.00-1.15	FHS	1.53	0.21	27.19	1.73	-0.19	12.44	2.21	0.24	5.52					
	HN	1.87	0.15	39.55	1.88	-0.65	24.00	2.06	0.46	6.28					
	IG	2.77	1.06	44.00	2.95	0.15	26.27	3.39	0.55	9.86					
	BS	4.24	3.47	115.43	3.03	1.95	36.25	6.28	-2.69	16.54					
> 1.15	FHS	0.26	-0.08	34.71	0.46	-0.05	15.92	1.04	0.33	21.34					
	HN	0.32	-0.06	16.80	0.76	-0.35	28.88	1.14	-0.61	34.71					
	IG	0.30	-0.06	12.22	0.70	-0.31	27.09	1.37	-0.65	36.85					
	BS	0.75	0.24	61.31	1.03	0.46	88.13	3.83	0.22	217.52					

Table 9: Out-of-sample pricing errors disaggregated by moneyiness and maturities. For the legend see Table 5.

S&P 500	800	900	1,000	1,100	1,200
$M_{t,t+38}$	1.72	1.26	1.08	0.64	0.20
$M_{t,t+73}$	1.37	1.27	1.15	0.71	0.36
$M_{t,t+164}$	1.56	1.48	1.22	0.84	0.48
$M_{t,t+255}$	1.69	1.64	1.26	0.90	0.59

Table 10: The table shows the state price densities per unit probability $M_{t,t+\tau}$ for different time to maturities $\tau = 38, 73, 164, 255$ days, estimated for July 9, 2003; see also Figure 5. $M_{t,t+\tau}$ is the discounted ratio of the pricing over the historical densities of the GJR GARCH model with FHS method.

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 11: “Homogeneous hedging of the smile”. The three rows in the middle are market option prices. 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.

	η_0		η_1		η_2		$Var[\text{error}_{t+1}]$		LR p-value	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
<hr/>										
$t_1 = 09/18/'02$										
1)	0.49		1.00				1.23			
t-stat.	4.82		221.78							
2)	0.08		0.58		0.42		0.61		0.00	
t-stat.	0.97		18.09		13.54					
3)	-0.40				0.98		1.72			
t-stat.	-3.24				187.60					
<hr/>										
$t_2 = 07/09/'03$										
1)	-0.03		1.00				0.16			
t-stat.	-0.70		499.38							
2)	-0.09		0.74		0.26		0.13		0.00	
t-stat.	-2.29		16.06		5.52					
3)	-0.26				1.00		0.37			
t-stat.	-4.00				330.01					
<hr/>										
Return ⁽⁻⁾										
1)	0.17	0.25	1.01	0.02			0.56	0.57		
t-stat.	3.49	5.10	443.74	218.23						
2)	0.02	0.18	0.76	0.30	0.24	0.30	0.43	0.46	0.11	0.27
t-stat.	1.09	3.90	22.64	17.43	5.68	5.12				
3)	-0.29	0.53			1.00	0.03	1.08	0.79		
t-stat.	-2.81	5.57			259.13	92.80				
<hr/>										
Return ⁽⁺⁾										
1)	-0.14	0.21	0.99	0.01			0.45	0.40		
t-stat.	-3.05	4.67	424.58	204.32						
2)	-0.16	0.18	0.82	0.20	0.17	0.20	0.42	0.39	0.16	0.23
t-stat.	-3.31	4.03	21.81	18.88	3.39	3.76				
3)	-0.21	0.44			0.99	0.02	1.12	0.92		
t-stat.	-3.31	4.84			240.46	67.51				

Table 12: Forecast regressions: 1) $P_{t+1}^{\text{market}} = \eta_0 + \eta_1 P_{t,t+1}^{\text{bs}} + \text{error}_{t+1}$, 2) $P_{t+1}^{\text{market}} = \eta_0 + \eta_1 P_{t,t+1}^{\text{bs}} + \eta_2 P_{t,t+1}^{\text{garch}} + \text{error}_{t+1}$, 3) $P_{t+1}^{\text{market}} = \eta_0 + \eta_2 P_{t,t+1}^{\text{garch}} + \text{error}_{t+1}$. P_{t+1}^{market} are the option prices observed on time $t+1$, $P_{t,t+1}^{\text{bs}}$ 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 . $P_{t,t+1}^{\text{garch}}$ 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 . $Var[\text{error}_{t+1}]$ is the variance of the regression errors, Std. is the standard deviation, LR is the p-value for the likelihood ratio test of the null hypothesis 1) versus the alternative hypothesis 2), and Return⁽⁻⁾ (Return⁽⁺⁾) are daily log-returns of the S&P 500 from t to $t+1$ smaller than -0.01 (larger than 0.01).

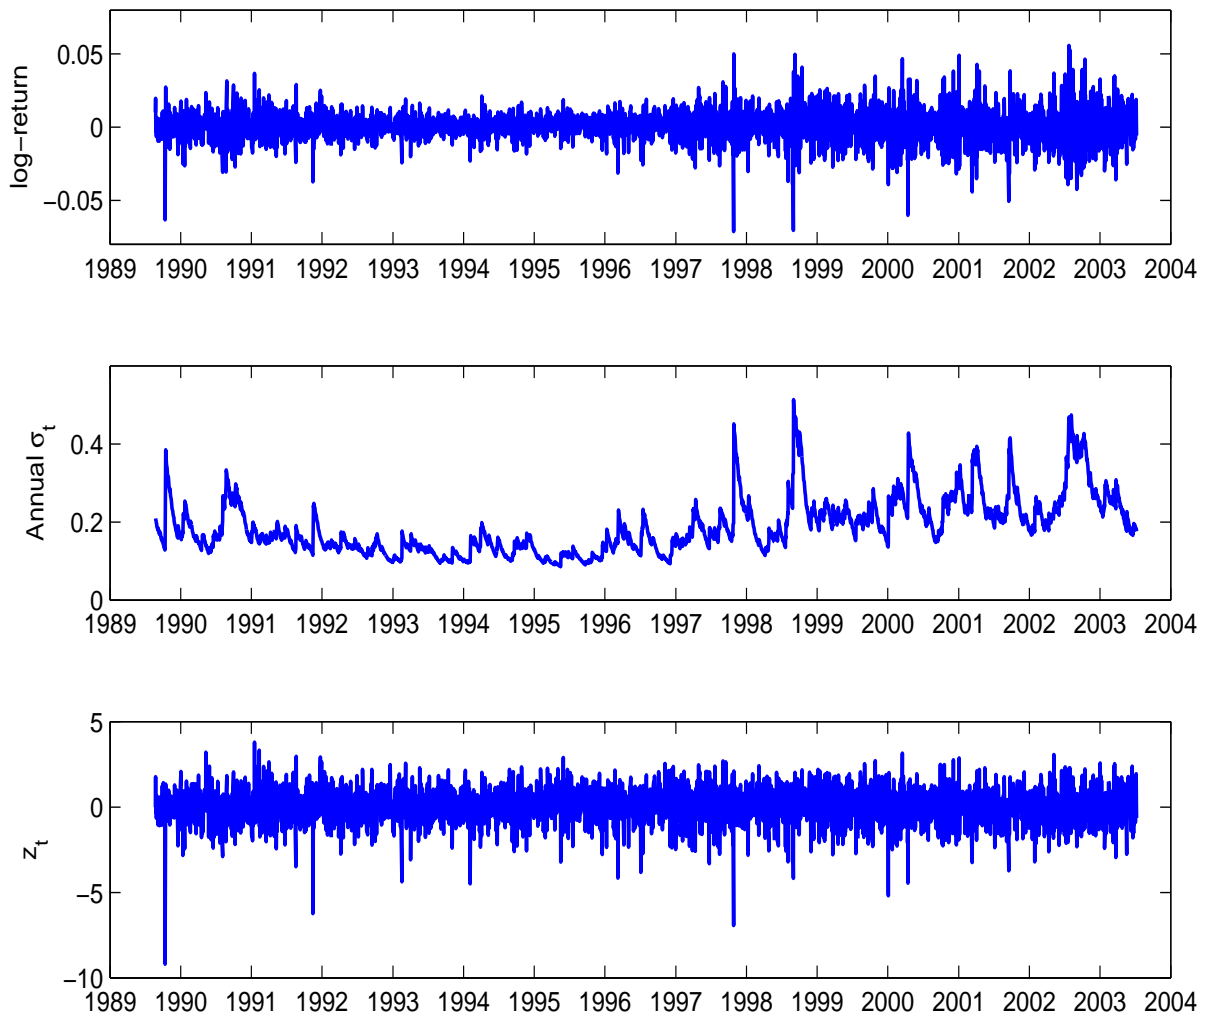


Figure 1: Daily log-return of the S&P 500 index from August 23, 1989 to July 9, 2003, (i.e. 3,500 log-returns), the conditional GARCH volatility σ_t (on an annual base) and the filtered innovations z_t . The GARCH model is estimated using the PML approach with the nominal assumption of Gaussian innovations

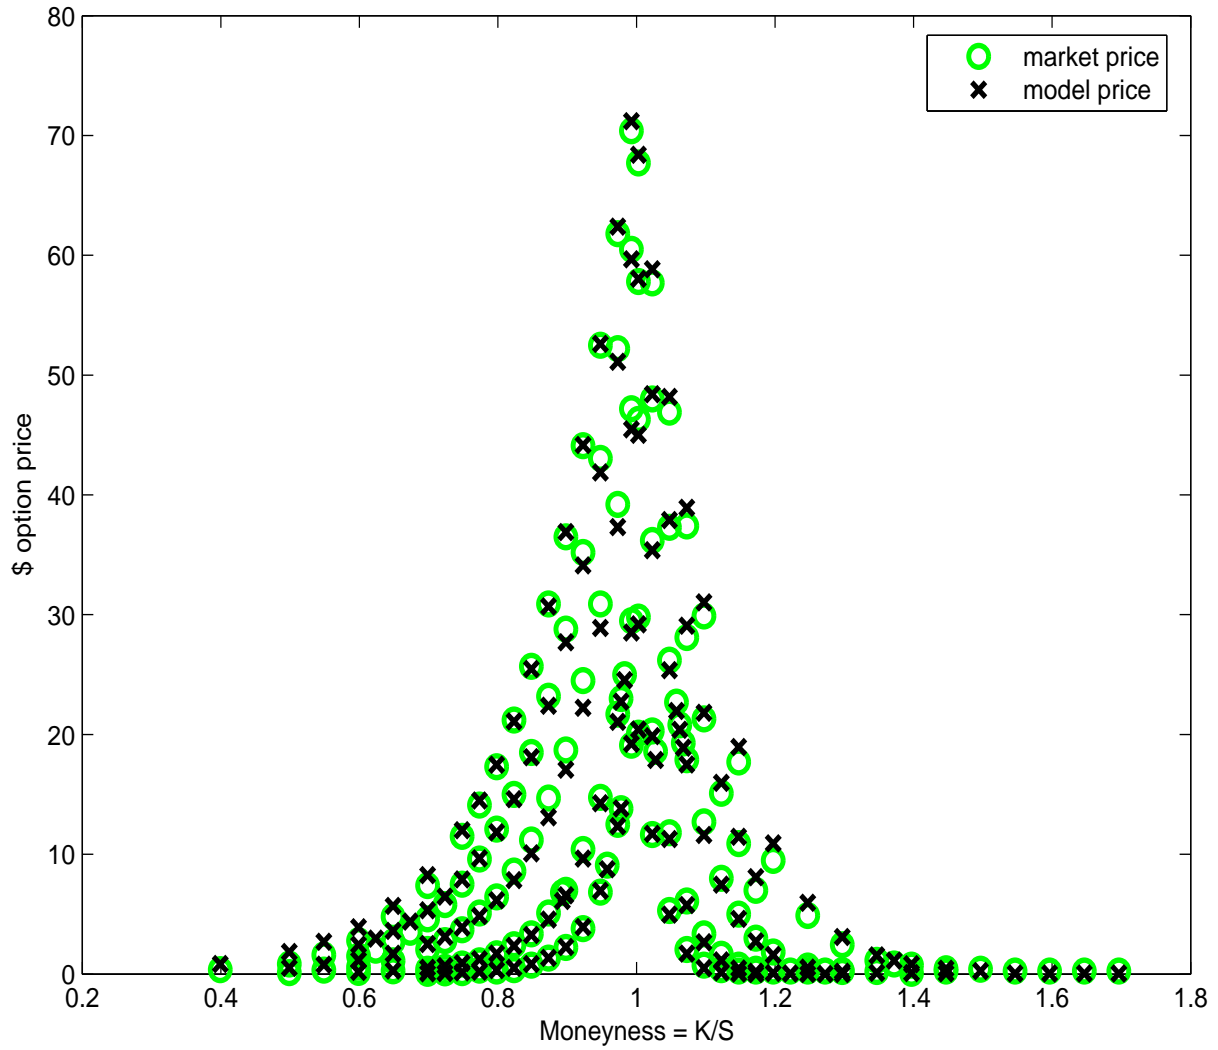


Figure 2: Calibration of the GJR GARCH model with the FHS method on the cross section of out-of-the-money SPX options on July 9, 2003. Moneyness is the strike price divided by the spot asset price.

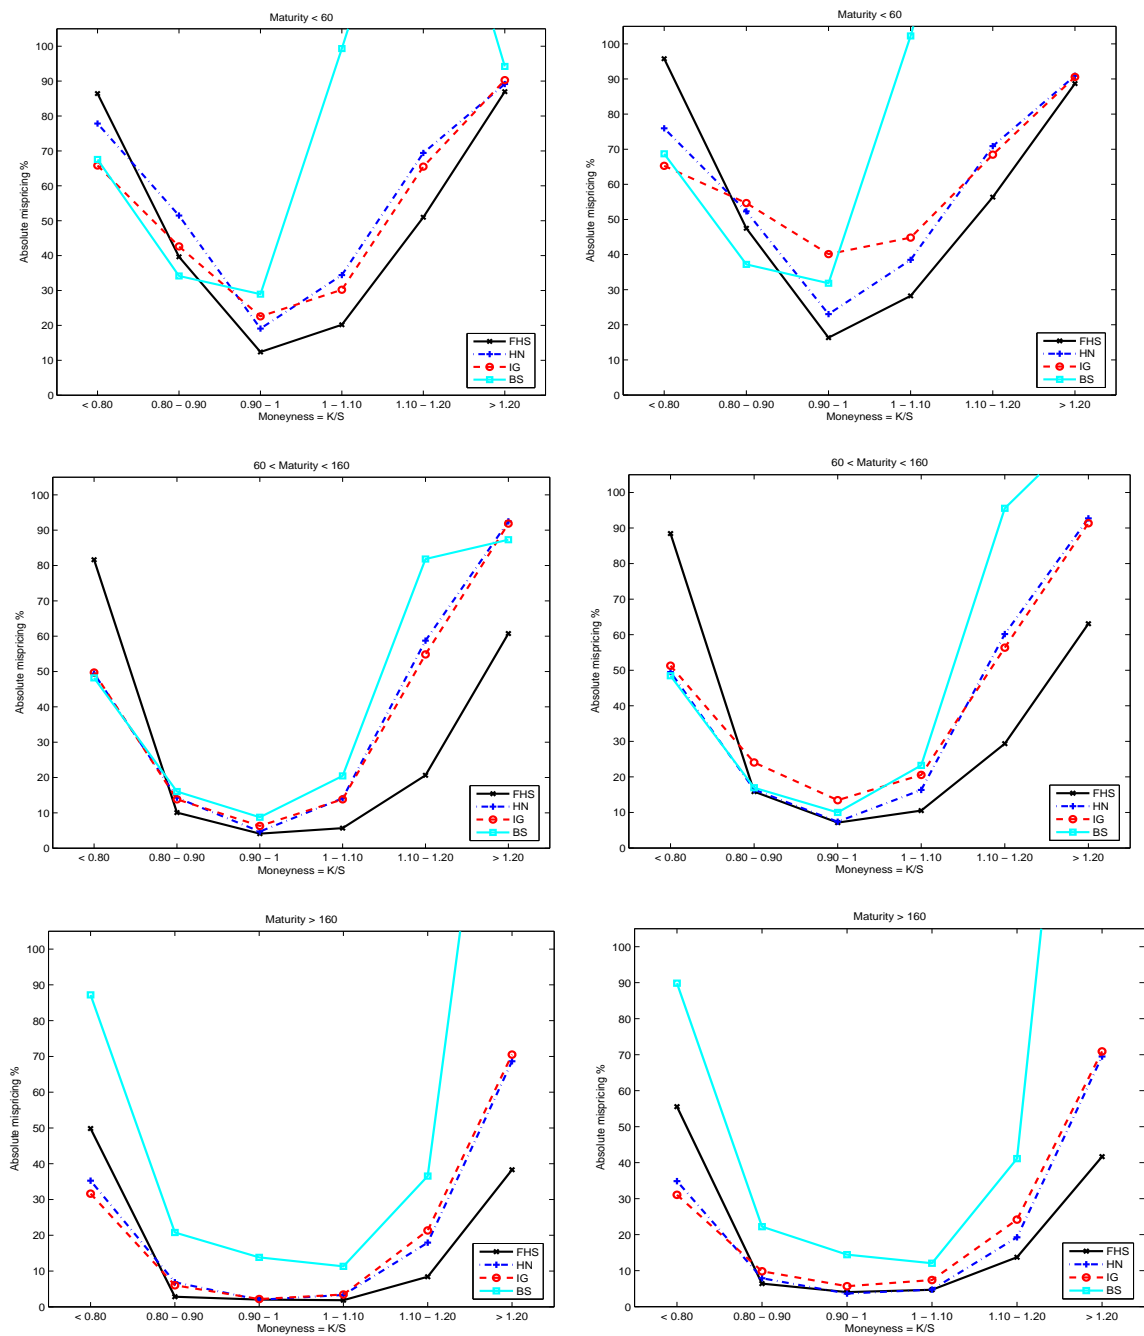


Figure 3: Absolute mispricing in percentage, i.e. $100 \times |\text{model price} - \text{market price}| / \text{market price}$, by the different pricing models, averaged across the Wednesdays from January 2002 to December 2004 for out-of-the-money SPX options. Left plots are in-sample comparisons and right plots are out-of-sample comparisons. Moneyness is the strike price divided by the spot asset price.

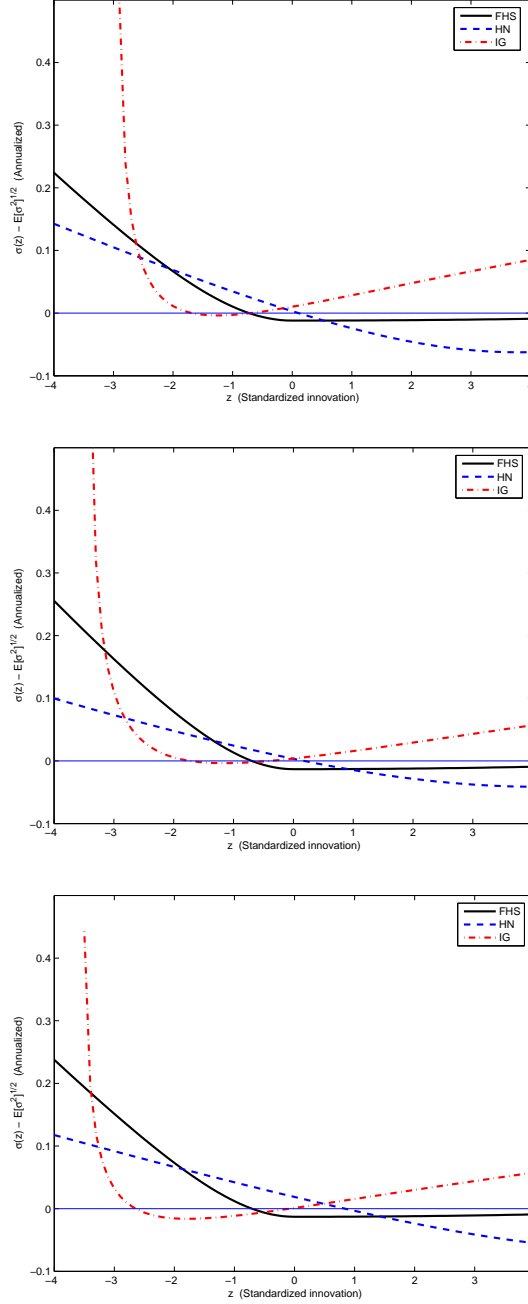


Figure 4: News impact curves the FHS, HN and IG GARCH models based on the average pricing parameters reported in Table 3 for the year 2003, 2003 and 2004 from up to bottom. The horizontal axis is the standardized GARCH innovations z . The vertical axis is the change in volatility from the long-run level, $\sigma(z) - \sqrt{E_{\mathbb{Q}}[\sigma^2]}$ on an annual base, as function of the standardized innovation. For GJR model $\sigma^2(z) = \omega^* + \beta^* E_{\mathbb{Q}}[\sigma^2] + E_{\mathbb{Q}}[\sigma^2](\alpha^* + \gamma^* I_{\{z < 0\}})z^2$, for the HN model $\sigma^2(z) = \omega_{\text{hn}}^* + \beta_{\text{hn}}^* E_{\mathbb{Q}}[\sigma^2] + \alpha_{\text{hn}}^* (z - \gamma_{\text{hn}}^* \sqrt{E_{\mathbb{Q}}[\sigma^2]})^2$, for the IG model $\sigma^2(z) = w^* + b^* E_{\mathbb{Q}}[\sigma^2] + c^* \tilde{y}(z) + a^* E_{\mathbb{Q}}[\sigma^2]^2 / \tilde{y}(z)$, where $\tilde{y}(z) = (z + \sqrt{\delta^*})\sqrt{\delta^*}$, and $\delta^* = E_{\mathbb{Q}}[\sigma^2] / \eta^{*2}$. Notice that $\tilde{y}(z)$ is the unscaled inverse Gaussian innovation and has to be positive.

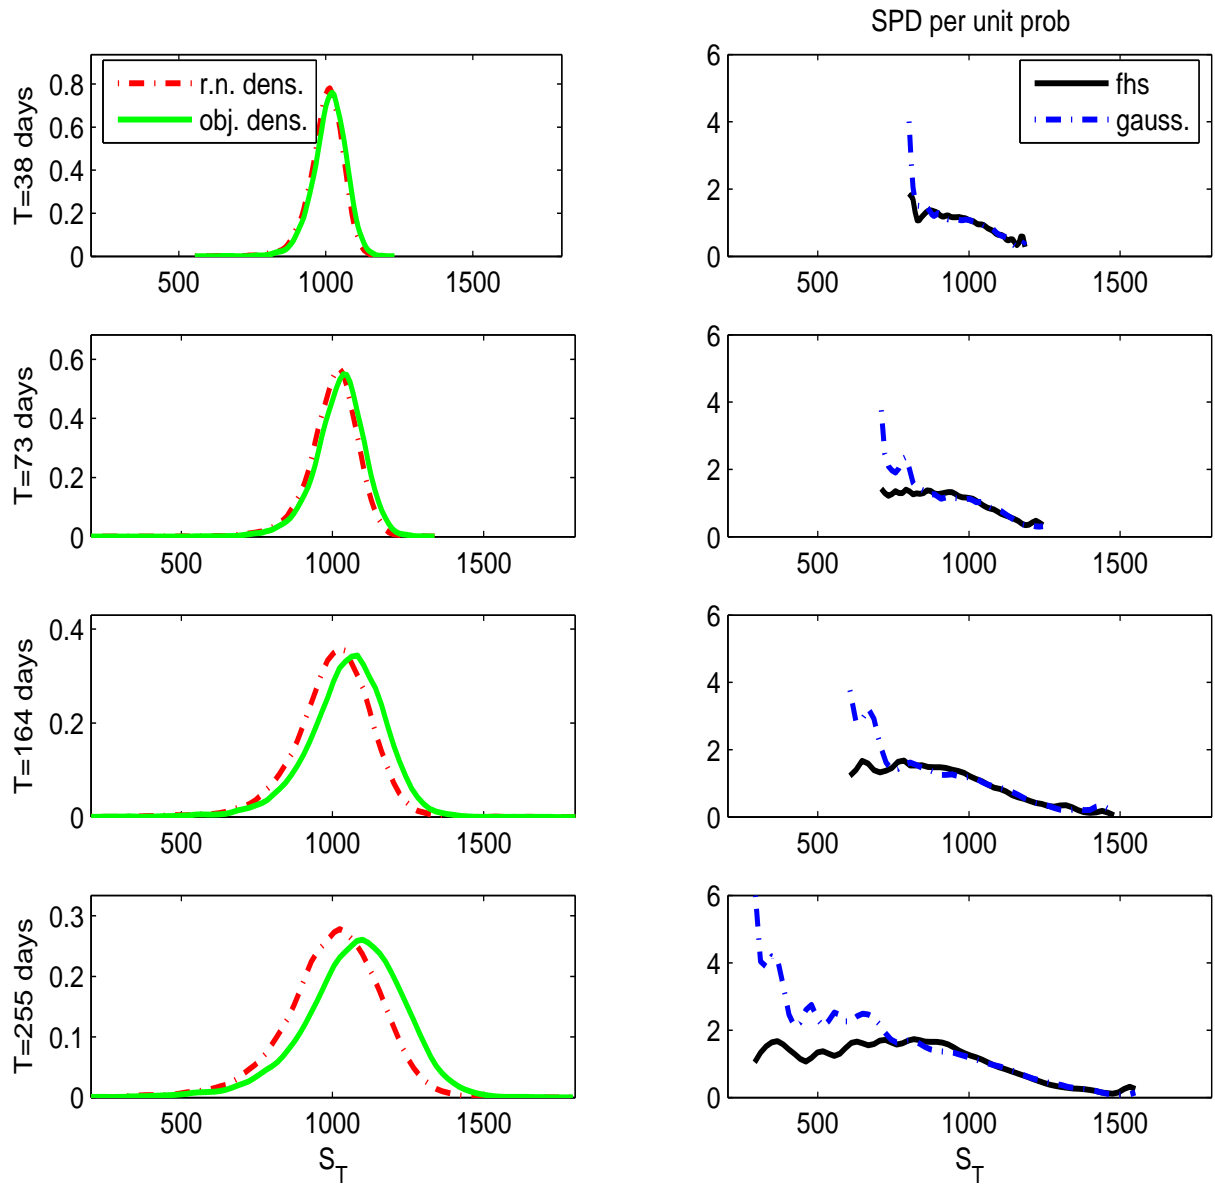


Figure 5: The left plots show the pricing and the historical density estimates for different time horizons on July 9, 2003. The pricing density is obtained by calibrating the GJR model with FHS method calibrated on the cross section of out-of-the-money SPX options. The historical density is obtained by fitting the GJR GARCH model to 3,500 historical S&P 500 log-returns and using the PML method. The right plots show the SPD per unit probability given by the ratio of the pricing over the historical density discounted by risk free rate. The Gaussian SPD are based on the GJR GARCH model calibrated with the Gaussian innovations.

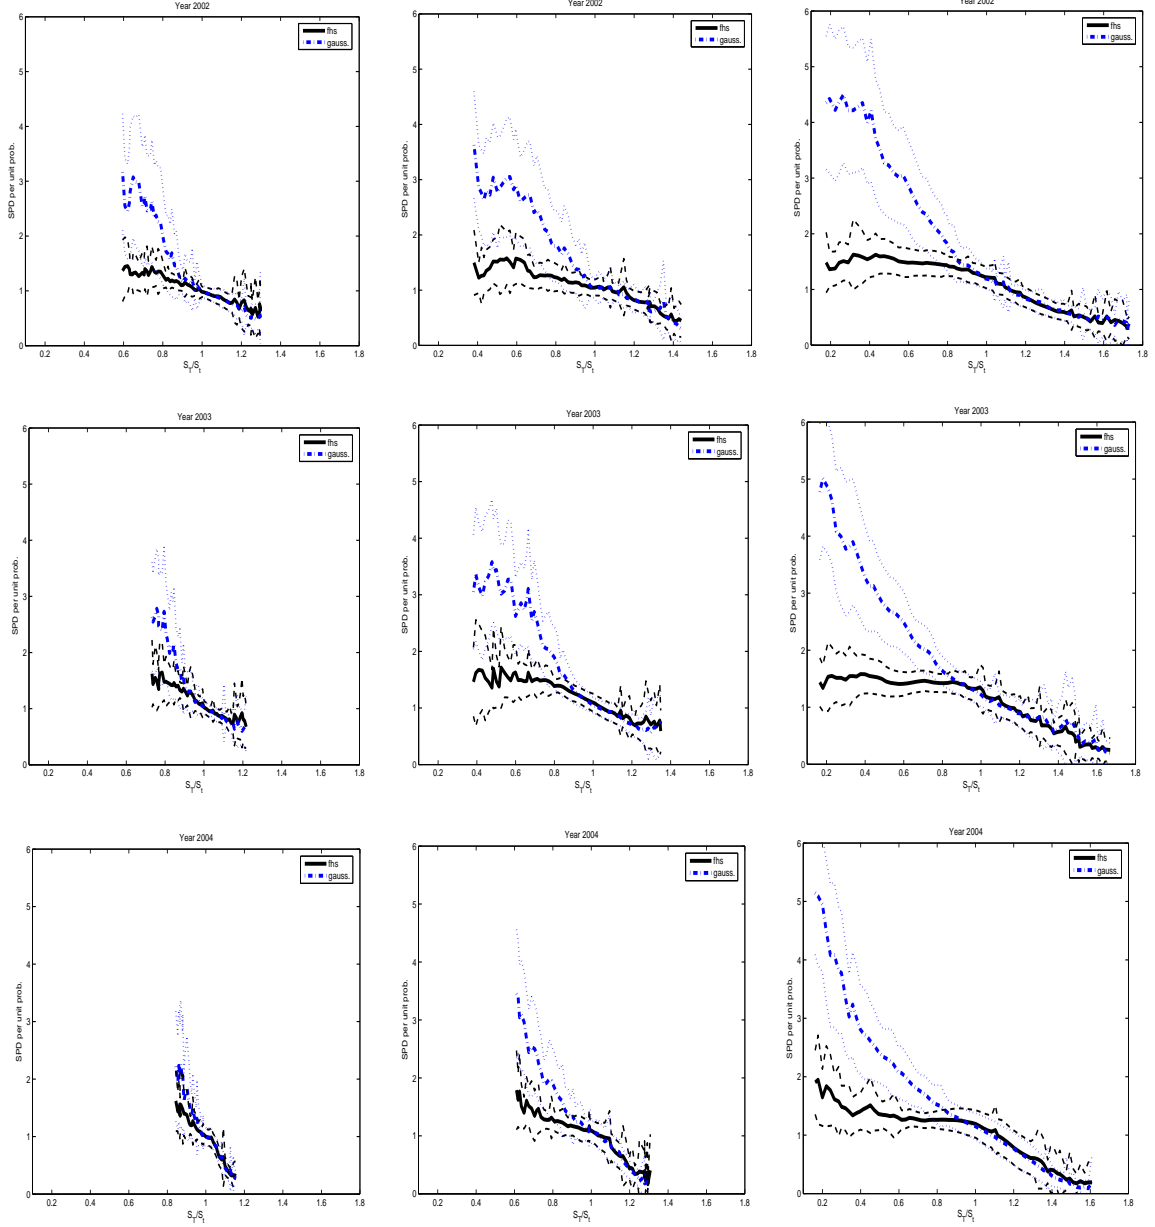


Figure 6: Average SPD per unit probability. For each Wednesdays t from January 2, 2002 to December 29, 2004 and for each time to maturity τ between 10 and 360 days, the SPD per unit probability $M_{t,t+\tau}$ are estimated by the discounted ratio of the pricing GARCH density $q_{t,t+\tau}$ over the historical GARCH density $p_{t,t+\tau}$. For each t , the pricing GJR model is calibrated on the cross section of out-of-the-money SPX options, and the historical GJR model is estimated using 3,500 historical S&P 500 log-returns via the PML method. Two estimates of $M_{t,t+\tau}$ are provided using (i) the FHS method and (ii) the Gaussian innovations. In case (i), $q_{t,t+\tau}$ and $p_{t,t+\tau}$ are obtained by simulating the pricing and historical GJR model using estimated GARCH innovations. In case (ii), GARCH innovations are replaced by Gaussian innovations. For each year and for each maturity category, the plots show the average of $M_{t,T}$ computed for a set of gross returns S_T/S_t , and also the average of $M_{t,T} \pm 0.5$ times the empirical standard deviation of $M_{t,T}$.

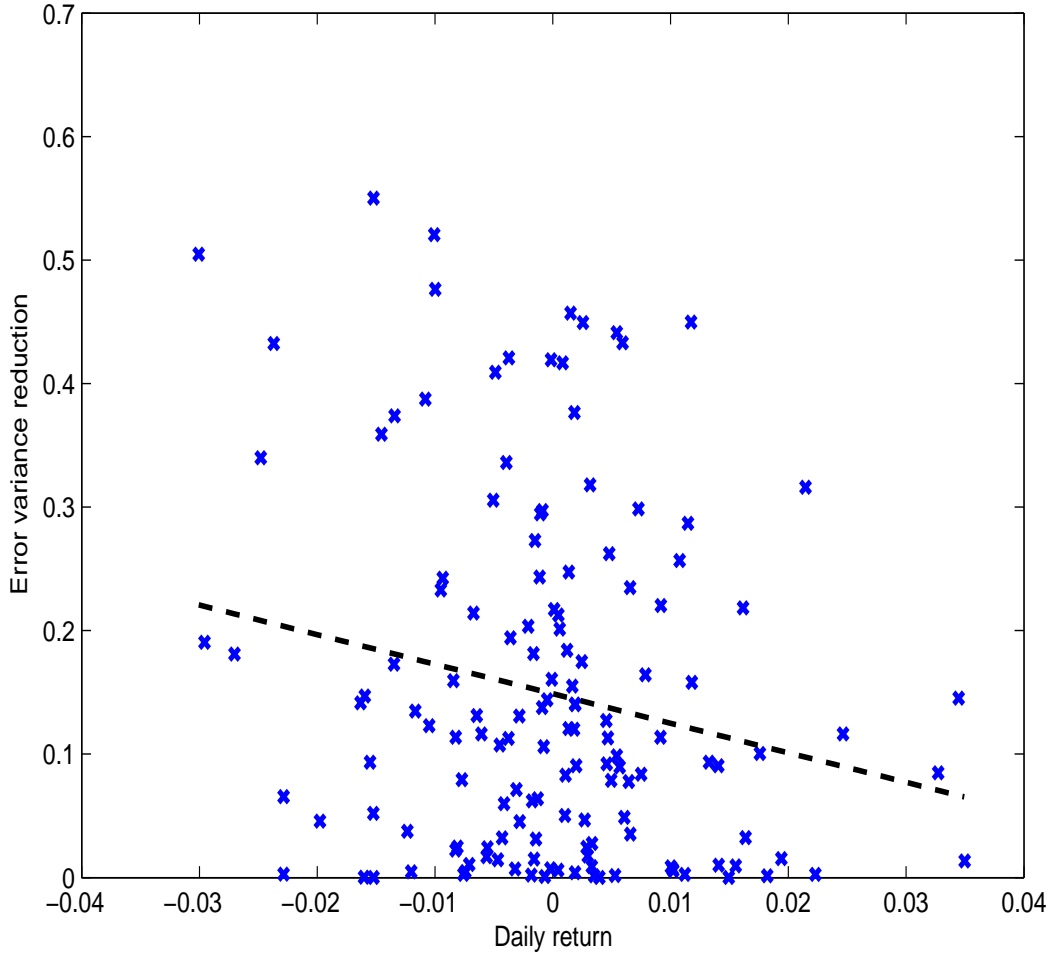


Figure 7: Reduction in the variance of forecast error when the FHS GARCH forecasts are added to the Black-Scholes forecasts for each Wednesday t from January 2, 2002 to December 29, 2004. The vertical axis displays the variance reduction, that is $(Var[error_{t,1}] - Var[error_{t,2}])/Var[error_{t,1}]$, where the errors are from the forecast regression 1) $P_{t+1}^{\text{market}} = \eta_0 + \eta_1 P_{t,t+1}^{\text{bs}} + error_{t,1}$, and the forecast regression 2) $P_{t+1}^{\text{market}} = \eta_0 + \eta_1 P_{t,t+1}^{\text{bs}} + \eta_2 P_{t,t+1}^{\text{garch}} + error_{t,2}$. See Table 12 for the definition of P_{t+1}^{market} , $P_{t,t+1}^{\text{bs}}$ and $P_{t,t+1}^{\text{garch}}$. The horizontal axis displays the daily S&P 500 log-return $\log(S_t/S_{t-1})$ for each Wednesday t .