An Empirical Comparison of Affine and Non-Affine Models for Equity Index Options

Peter Christoffersen   Kris Jacobs   Karim Mimouni
Faculty of Management, McGill University*

October 21, 2005

Abstract

The existing literature on equity index option valuation largely focuses on affine models, because they lead to closed-form solutions for option prices. This paper investigates the empirical biases associated with affine models, using data on S&P500 call options. We find that the root mean squared dollar error for a simple non-affine continuous-time stochastic volatility model is 17.4% lower than that of the benchmark continuous-time affine stochastic volatility model in-sample, and almost 20% out-of-sample. The analytical convenience of affine option valuation models therefore comes at a price, and non-affine models ought to be investigated more extensively. We also compare the empirical performance of affine and non-affine discrete-time models. While the performance of the discrete-time non-affine model is similar to that of the continuous-time non-affine model, the discrete-time affine model outperforms the continuous-time affine model. These findings have some interesting implications for the interpretation of existing limit results. At the methodological level, our analysis uses a novel technique based on the Auxiliary Particle Filter. This technique allows for an analysis of option valuation models using options data that imposes consistency with underlying equity returns. It is more straightforward to implement than existing methods and can be used in a variety of applications and on various loss functions.

JEL Classification: G12

Keywords: stochastic volatility; GARCH; affine model; option valuation; filtering; out-of-sample; particle filter; continuous-time limit.

*Christoffersen and Jacobs are also affiliated with CIRANO and CIREQ and want to thank FQRSC, IFM^2 and SSHRC for financial support. Mimouni was supported by a grant from IFM^2. We also thank Jin Duan, Nour Meddahi and Marcel Rindisbacher for helpful comments. Any remaining inadequacies are ours alone. Correspondence to: Peter Christoffersen, Faculty of Management, McGill University, 1001 Sherbrooke Street West, Montreal, Quebec, Canada, H3A 1G5; Tel: (514) 398-2869; Fax: (514) 398-3876; E-mail: peter.christoffersen@mcgill.ca.

1 Introduction

Following the finding that Black-Scholes (1973) model prices systematically differ from market prices, the literature on option valuation has formulated a number of theoretical models designed to capture these empirical biases. One particularly popular modeling approach has attempted to correct the Black-Scholes biases by modifying the assumption that volatility is constant across maturity and moneyness. Estimates from returns data and options data indicate that return volatility is time-varying, and modeling volatility clustering leads to significant improvements in the performance of option pricing models. It has also been demonstrated that it is necessary to model a leverage effect. The leverage effect captures the negative correlation between returns and volatility, and thus generates negative skewness in the distribution of the underlying asset return.[1]

The existing literature has almost exclusively modeled volatility clustering and the leverage effect within an affine structure. Affine models of option valuation are convenient because they lead to closed-form solutions for prices of European equity options. In particular, the affine Heston (1993) model, which accounts for time-varying volatility and a leverage effect, has been implemented in a large number of empirical studies. In order to address the limitations of the affine structure, the Heston (1993) model is often combined with models of jumps in returns and/or volatility.[2] However, relatively little is known about the empirical biases that result from imposing the affine structure.[3]

The existing literature has also almost exclusively modeled volatility clustering and the leverage effect using continuous-time stochastic volatility models. There exists a discrete-time literature on option valuation using GARCH processes, and in this literature the distinction between affine and non-affine models is also relevant. The affine GARCH dynamic in Heston and Nandi (2000) yields a closed-form solution for option prices, but the extensive empirical literature on GARCH processes that concerns itself with fitting and predicting return volatility almost exclusively uses non-affine processes.[4] A small number of studies investigate option valuation assuming non-affine GARCH

[1] The leverage effect was first characterized in Black (1976). For empirical studies that emphasize the importance of volatility clustering and the leverage effect for option valuation see among others Benzoni (1998), Chernov and Ghysels (2000), Eraker (2004), Heston and Nandi (2000), Nandi (1998) and Pan (2002).

[2] For empirical studies that implement the Heston (1993) model by itself or in combination with different types of jump processes, see for example Andersen, Benzoni and Lund (2002), Bakshi, Cao and Chen (1997), Bates (1996, 2000), Benzoni (2002), Chernov and Ghysels (2000), Huang and Wu (2004), Pan (2002), Eraker (2004) and Eraker, Johannes and Polson (2003).

[3] Jones (2003) and Benzoni (2002) are notable exceptions that investigate non-affine option valuation models. The non-affine model in Benzoni (2002) does not improve on the performance of the Heston (1993) model, while a number of specification tests in Jones (2003) favor the non-affine constant elasticity of substitution model over the Heston (1993) model.

[4] The literature on GARCH processes is too voluminous to cite in full here. The classical references are Engle (1982) and Bollerslev (1986). See Bollerslev, Chou and Kroner (1992) and Diebold and Lopez (1995) for reviews.

processes for the underlying securities.[5]

This paper investigates the empirical implications of adopting an affine framework for option valuation. We compare the empirical performance of the affine Heston (1993) stochastic volatility model (AF-SV) with that of a simple non-affine stochastic volatility model (NA-SV). We estimate model parameters using a long time series of cross sections of options data in a framework that imposes consistency with the underlying equity returns. We conduct this empirical analysis using a novel setup that uses the Auxiliary Particle Filter algorithm. This methodology provides a convenient filtering algorithm for latent factor models such as stochastic volatility models. Our new methodology is relatively easy to implement compared with existing likelihood-based methods, and it can be adapted to provide the best possible fit to the objective function of interest.[6] We also compare the empirical performance of the affine Heston and Nandi (2000) model (AF-GARCH) with that of the non-affine GARCH model (NA-GARCH) of Engle and Ng (1993).

We find that the affine framework is very restrictive. We conduct six in-sample exercises and six out-of-sample exercises, and despite the fact that both the AF-SV and the NA-SV models show signs of misspecification in several samples, the NA-SV model outperforms the Heston (1993) AF-SV model in all of these exercises. The NA-SV model also outperforms the AF-SV model for all moneyness and maturity categories. On average, the dollar RMSE of the NA-SV we investigate is approximately 17.4% lower than that of the AF-SV model in-sample, and the out-of-sample dollar RMSE is almost 20% lower. We therefore conclude that while the closed-form solution provided by affine models is convenient, this analytical convenience comes at a price, and non-affine models need to be studied more extensively.

Interestingly, the differences in RMSE between the non-affine discrete-time model (NA-GARCH) and the affine discrete-time model (AF-GARCH) are smaller than in the continuous-time case, approximately 7.5% in-sample and 10% out-of-sample.[7] This finding naturally raises important questions about the relationship between the empirical performance of discrete-time and continuous-time models. Our findings on this issue are perhaps less straightforward to interpret than those on the implications of the affine structure, but we are able to draw a number of important conclusions.

The literature contains a number of limit results relating certain classes of discrete-time and continuous-time models (see for example Duan (1997), Heston and Nandi (2000) and Nelson (1990)). Some researchers have interpreted these limit results as evidence that the performance of discrete-time and continuous-time models ought to be very similar when the continuous-time dynamic is the

[5] See Amin and Ng (1993), Bollerslev and Mikkelsen (1996), Engle and Mustafa (1992), Heston and Nandi (2000), Christoffersen and Jacobs (2004) and Duan, Ritchken and Sun (2002).

[6] Likelihood-based approaches fully exploit the available information, but do so at the cost of greater complexity. See Eraker (2004) and Jones (2003) for Bayesian approaches. See Bates (2004) for an approximate maximum likelihood approach that is relatively easier to implement.

[7] See also Hsieh and Ritchken (2000) for a related comparison.

limit of the discrete-time dynamic. This interpretation is somewhat contentious, because a given discrete-time model can have several continuous-time limits, and a given continuous-time model can be the limit for more than one discrete-time model (see for instance Corradi (2000)). While limit results are therefore theoretically intriguing, in some cases their practical relevance may be limited.

We find some very interesting differences in the empirical performance of models that are often thought of as equivalent. Most importantly, the AF-SV Heston (1993) model, which is the benchmark model in the continuous-time stochastic volatility literature, significantly underperforms the AF-GARCH Heston-Nandi (2000) model. This result is perhaps surprising, because Heston and Nandi (2000) demonstrate that a restricted Heston (1993) model can be seen as the limit of the AF-GARCH model, but the key is of course that this particular limit is just one of the many available. Which limit obtains depends on the particular mathematical construction used.

This finding does not mean that discrete-time GARCH models outperform continuous-time stochastic volatility models. Indeed, the performance of the non-affine NA-GARCH discrete time model of Engle and Ng (1993) is very similar to the performance of the NA-SV model we investigate. These findings demonstrate that it is difficult to make general statements about discrete-time and continuous-time models. Certain stylized facts may be more conveniently captured by specific continuous-time models, while others are more easily modeled using a particular discrete-time model. This paper merely provides a start to that discussion by documenting the empirical performance of some important benchmark models. After documenting that the benchmark AF-SV Heston (1993) model violates our prior by underperforming a related discrete-time model, we proceed by suggesting an aspect of the model that may cause the underperformance.

One final remark is in order. Traditionally, the existence of multiple limits has not been the only issue that complicated a comparison of discrete-time and continuous-time models. The two classes of models are typically implemented using very different econometric methods, which renders fair comparisons difficult. Our use of the Auxiliary Particle Filter algorithm allows for straightforward comparisons of latent factor volatility continuous time models with discrete time GARCH models, because each model is implemented using the same objective function and the same information set. The novel estimation setup in this paper therefore facilitates comparisons between different classes of models.

The paper proceeds as follows. In Section 2 we introduce the discrete- and continuous-time volatility models, and we discuss their implementation. In Section 3 we present and discuss the empirical results. Section 4 concludes.

2 Volatility Dynamics and Model Implementation

We now turn to a description of the four volatility models we estimate and we discuss their implementation for the purpose of option valuation. We first discuss our implementation of the Heston (1993) model, which is different from the implementation available in the literature. Our implementation is designed to provide the best possible fit for the model, as well as to facilitate the comparison with the discrete-time models. Interestingly, this implementation is relatively straightforward and substantially simpler compared to other methods that analyze options data while imposing consistency with underlying returns data. Subsequently, we discuss the specification of a non-affine stochastic volatility model, and finally we discuss the two discrete-time GARCH models, where we consider the GARCH(1,1) representation because it is most closely related to the Heston (1993) continuous-time model. Following a discussion of the implementation of the discrete-time Heston and Nandi (2000) model, we conclude with the specification of a non-affine discrete-time GARCH model.

The Heston (1993) model is arguably the most popular model in the index option valuation literature. Its closed-form solution is due to the model's affine structure. [8] Because the existing literature does not contain strong evidence on modifications of the stochastic volatility dynamic that outperform the Heston (1993) model out-of-sample for index options, it constitutes a good benchmark for our study.[9] The underlying reason for its success is that the Heston (1993) model captures two important stylized facts that are needed to model option prices: volatility clustering and the leverage effect. After accounting for these two stylized facts, additional modifications of the return and volatility dynamic do not seem to result in significant out-of-sample improvements in the fit. We refer to the affine Heston (1993) model as AF-SV below.

The second model we investigate has not been extensively analyzed in the literature. It is a continuous time non-affine stochastic volatility model and we refer to it as NA-SV below. It contains a latent factor volatility which is correlated with returns, but it does not allow for closed-form option valuation. Consequently it is more complex to implement than the Heston (1993) model.

The discrete-time GARCH option valuation literature has resulted in relatively few empirical studies,[10] while discrete-time volatility modeling using only returns data has spawned a large num-

[8] Affine models are also very popular in the term structure literature for exactly the same reason. See for instance Duffie and Kan (1996) and Dai and Singleton (2000).

[9] There is an extensive and growing literature on the use of jumps in returns and volatility to improve the performance of the Heston model. See Andersen, Benzoni and Lund (2002), Bakshi, Cao and Chen (1997), Bates (1996, 2000), Chernov, Gallant, Ghysels and Tauchen (2003), Eraker, Johannes and Polson (2003), Eraker (2004), Pan (2002), Broadie, Chernov and Johannes (2004), Carr and Wu (2004) and Huang and Wu (2004). Extending our comparison to models of this type is interesting, but beyond the scope of this paper.

[10] Duan (1995) and Amin and Ng (1993) provide theoretical foundations for this literature. Bollerslev and Mikkelsen

ber of competing models following the work of Engle (1982) and Bollerslev (1986). Our benchmark discrete-time specification is the model of Heston and Nandi (2000), which was designed with option valuation in mind. Like the Heston (1993) model, it contains a leverage effect, it allows for volatility clustering, and it leads to a closed-form solution due to its affine structure. Heston and Nandi (2000) have demonstrated that this model performs satisfactorily vis-a-vis ad-hoc benchmarks for the purpose of option valuation. We refer to this model as AF-GARCH.

The other discrete-time model we investigate is the non-affine NGARCH model of Engle and Ng (1993), henceforth referred to as NA-GARCH. This is the simplest model in the GARCH literature that contains both volatility clustering and a leverage effect, and was first considered for option valuation by Duan (1995). Christoffersen and Jacobs (2004a) demonstrate that several richer GARCH parameterizations do not improve on the option valuation performance of the NA-GARCH model. This model does not lead to a closed-form solution for option prices, and has to be analyzed using numerical methods. Because the NA-GARCH dynamic has proven extremely valuable in modeling equity returns as well as other financial time series, it is of interest to verify whether the focus on closed-from valuation results in the options literature comes at the cost of a deterioration in the model's empirical performance.

We now give the specifics on each of the four models as well as a detailed description of their empirical implementation, which imposes consistency between option prices and underlying index returns.

2.1 The Affine Stochastic Volatility Model (AF-SV)

The Heston (1993) continuous-time stochastic volatility model (AF-SV) is defined by the following two equations

$$dS = \mu S dt + \sqrt{V} S dw^S \quad (1)$$

$$dV = \kappa(\theta - V)dt + \sigma\sqrt{V} dw^V \quad (2)$$

with $corr(dw^S, dw^V) = \rho$. This model allows for volatility clustering through the autoregressive component of volatility, and for a leverage effect through a negative correlation coefficient ρ, which translates into negative skewness of the return distribution.[11] Under the assumption that the volatility risk premium $\lambda(S, V, t)$ is equal to λV, the risk neutral dynamic expressed in terms of the

(1996), Engle and Mustafa (1992), and Duan, Ritchken and Sun (2002) estimate model parameters using the underlying asset returns and subsequently value options. Heston and Nandi (2000) and Christoffersen and Jacobs (2004a) estimate model parameters using equity option prices, and impose consistency with the underlying returns.

[11] Note that following Chernov and Ghysels (2000), Eraker, Johannes and Polson (2003) and Eraker (2004) we use a simple constant specification for the stock return drift.

physical parameters is

$$\begin{aligned} dS &= rSdt + \sqrt{V}Sdw^{*S} & (3)\\ dV &= (\kappa(\theta - V) + \lambda V)\,dt + \sigma\sqrt{V}dw^{*V} & (4)\\ &= (\kappa - \lambda)(\kappa\theta/(\kappa-\lambda) - V)dt + \sigma\sqrt{V}dw^{*V} & (5)\end{aligned}$$

with $corr(dw^{*S}, dw^{*V}) = \rho$. Heston (1993) demonstrates that this model admits a closed form solution for option prices, which is presented here in terms of the physical parameters $\kappa, \theta, \lambda, \rho$ and σ in order to facilitate the description of our estimation procedure below.

$$C(V) = SP_1 - Ke^{-r(T-t)}P_2 \tag{6}$$

where

$$P_j = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \text{Re}\left[\frac{\exp(-i\phi\ln(K))f_j(\ln(S), V, T; \phi))}{i\phi}\right]d\phi, \qquad j = 1, 2$$

and

$$\begin{aligned} f_j(\ln(S), V, T; \phi) &= \exp(C(T-t;\phi) + D(T-t;\phi)V + i\phi\ln(S))\\ C(\tau;\phi) &= r\phi i\tau + \frac{\kappa\theta}{\sigma^2}\left((b_j - \rho\sigma\phi i + d)\tau - 2\ln\left[\frac{1 - g\exp(d\tau)}{1-g}\right]\right)\\ D(\tau;\phi) &= \frac{b_j - \rho\sigma\phi i + d}{\sigma^2}\left(\frac{1-\exp(d\tau)}{1 - g\exp(d\tau)}\right)\end{aligned}$$

$$\begin{aligned} d &= \sqrt{(\rho\sigma\phi i - b_j)^2 - \sigma^2(2\mu_j\phi i - \phi^2)}\\ g &= \frac{b_j - \rho\sigma\phi i + d}{b_j - \rho\sigma\phi i - d}\\ \mu_1 &= \frac{1}{2}, \quad \mu_2 = -\frac{1}{2}, \quad b_1 = \kappa - \lambda - \rho\sigma, \quad b_2 = \kappa - \lambda\end{aligned}$$

The Heston model has been investigated empirically in a large number of studies. Often it is used as a building block together with models of jumps in return and volatility. For our purpose, it is important to note that the model can be estimated and investigated empirically using a number of different techniques. First, the model's parameters can be estimated using a single cross-section of option prices (for example see Bakshi, Cao and Chen (1997)). A second type of implementation of the Heston model uses multiple cross sections of option prices but does not use the information in the underlying asset returns. Instead, for every cross section a different initial volatility is estimated, leading to a highly parameterized problem (see for instance Bates (2000) and Huang and Wu (2004)). A third group of papers provide a likelihood-based analysis of the stochastic volatility model. For instance, Eraker (2004) provides a Markov Chain Monte Carlo analysis. Finally, Chernov and Ghysels (2000) use the efficient method of moments and Pan (2002) uses a method of moments technique as well. These methods can also combine information in the options data as well as the underlying returns.

In this paper we implement the Heston model in a novel way. This is mainly motivated by our desire to compare the performance of discrete-time and continuous-time methods in a meaningful way. Our implementation uses nonlinear least squares (NLS) estimation techniques to minimize

$$\$MSE = \frac{1}{N^T} \sum_{t,i} (C_{i,t} - C_i(V))^2 \tag{7}$$

with respect to the physical parameters $\mu, \kappa, \theta, \lambda, \rho$ and σ, where $N^T = \sum_{t=1}^{T} N_t$, T is the total number of days included in the options sample, N_t is the number of options with various strikes prices and maturities included in the sample at date t, $C_{i,t}$ is the market price of option i quoted on day t and $C_i(V)$ is the model price. In our opinion, this type of objective function guarantees the best possible performance for the model in- and out-of-sample. This is motivated by the insights of Granger (1969), Weiss (1996) and Weiss and Andersen (1984) who demonstrate that the choice of objective function (also labeled loss function) is an integral part of model specification. It follows that estimating a model using one objective function and evaluating it using another one amounts to suboptimal choice of objective function. Christoffersen and Jacobs (2004b) demonstrate that this issue is empirically relevant for the estimation of the deterministic volatility functions in Dumas, Fleming and Whaley (1998). We therefore implement the Heston model in a way that is consistent with these insights.

The problem with (7) is that V is unobservable, and that it has to be filtered from observed data. Filtering the latent volatility factor on observed index returns as we do below avoids overfitting and ensures that the option valuation model is consistent with both options and returns data. Our implementation uses the Auxiliary Particle Filter (APF) algorithm.[12] As shown by Pitt and Shephard (1999) the APF offers a convenient filtering algorithm for non-linear models such as the stochastic volatility model we consider here. Because the APF procedure is relatively new in finance, we now discuss the implementation of this method in more detail.[13]

2.1.1 Volatility Transformation and Discretization

To prevent V from becoming negative, we work with $f(V) = \ln(V)$. Using Ito's lemma, the dynamic of interest is therefore

$$d\ln(V) = \frac{1}{V}\left(\kappa(\theta - V) - \frac{1}{2}\sigma^2\right)dt + \sigma\frac{1}{\sqrt{V}}dw^V \tag{8}$$

[12] We have also implemented the Sampling-Importance-Resampling (SIR) particle filter as a robustness check, and this yields similar results.

[13] Johannes, Polson and Stroud (2002) discuss the use of the particle filter to estimate parameters for continuous-time jump-diffusion models on returns data.

Note that equations (1) and (2) specify how the unobserved state is linked to observed stock prices. This relationship allows us to infer the volatility path using the returns data. We first need to discretize equations (1) and (8). There are different discretization methods and every scheme has certain advantages and drawbacks. We use the Euler scheme which is easy to implement and has been found to work well for this type of applications. Discretizing equation (1) (using $\ln(S)$ instead of S and applying Ito's lemma again) and (8) gives

$$\ln(S_{t+1}) = \ln(S_t) + \left(\mu - \frac{1}{2}V_t\right) + \sqrt{V_t}\varepsilon_{t+1}^S \tag{9}$$

$$\ln(V_{t+1}) = \ln(V_t) + \frac{1}{V_t}\left(\kappa\left(\theta - V_t\right) - \frac{1}{2}\sigma^2\right) + \sigma\frac{1}{\sqrt{V_t}}\varepsilon_{t+1}^V \tag{10}$$

We implement the discretized model in (9) and (10) using daily returns, and all parameters will be expressed in daily units below. The model is characterized by six structural parameters: μ, κ, θ, σ, λ and ρ for which we have to choose a set of starting values. Subsequently, we have to choose an initial variance V_0 (the starting value for the variance path). We set the initial variance equal to the model-implied unconditional variance, $V_0 = \theta$.

Our optimization algorithm minimizes (7) using an iterative procedure. At each iteration, the volatility is filtered using the information embedded in observed returns. Since the minimization is performed relative to option prices, option data also indirectly contribute to the determination of the volatility path. Finally, using the filtered volatility and the structural parameters option prices are computed according to Heston's formula and the MSE is calculated. This procedure is repeated until the optimum is reached. We now describe volatility filtering in more detail.

2.1.2 Filtering the volatility path using the APF algorithm

The idea underlying the APF technique is to infer the volatility path from the observed returns data. V_t is propagated one day ahead using equation (8) into N possible states (or particles).[14] Subsequently, we use an auxiliary variable ι and the available data to decide which particles to keep in order to simulate the one day ahead volatility.

Assume that we are at date t and we have an initial set of particles $\left\{V_t^j, W_t^j\right\}_{j=1}^N$ with V_t^j the volatility at day t for the state j, W_t^j the weight associated with state j at date t and $j = 1,..,N$.[15] We want to propagate V_t one day ahead into $\left\{V_{t+1}^j, W_{t+1}^j\right\}_{j=1}^N$ using the APF. This task requires the following steps:

Step 1: Selecting the particles

[14] We set $N = 500$ in the initial search. Once a candidate optimum is identified we confirm it by increasing N to $5,000$. The results change very little when N is increased.

[15] At time 0, the initial set is constructed by setting each particle equal to the unconditional variance θ and giving all particles equal weight, $1/N$.

The weight W_t^j reflects information available at time t only and does not include our expectations about $(t+1)$. So even if state j is very likely according to the realization of the stock price at date t, it is possible that the realization of the stock at date $(t+1)$ suggests a certain re-adjustment of the probability that state j has occurred. By combining the information available at t and our expectations about $(t+1)$, we can eliminate many states with a low realization probability right before the propagation step. This is achieved by:

I) Computing a summary location statistic for $(t+1)$ that reflects the information at t. We use the mean μ_{t+1}^j given by:

$$\mu_{t+1}^j = E\left(\ln(V_{t+1})|V_t^j\right)$$

II) Simulating the auxiliary variable ι^j

$$\iota^j \propto W_t^j p\left(\ln(S_{t+1})|\mu_{t+1}^j\right)$$

where $p\left(\ln(S_{t+1})|\mu_{t+1}^j\right)$ is the conditional density of $\ln(S_{t+1})$ which can be easily inferred from (9). This auxiliary variable is simply an index that tells us which particle to keep and which particle to discard. After this selection exercise we obtain N new particles which are implicitly functions of the auxiliary variable ι,

$$\{V(\iota)_t, W(\iota)_t\}_{j=1}^N$$

In order to keep the notation simple we will omit ι below.

Step 2: Simulating the state forward (Sampling)

This is done by computing V_{t+1} using equation (10) and taking the correlation into account. We have

$$\ln\left(\frac{S_{t+1}}{S_t}\right) = \left(\mu - \frac{1}{2}V_t\right) + \sqrt{V_t}\varepsilon_{t+1}^S$$

which gives

$$\varepsilon_{t+1}^S = \frac{\ln\left(\frac{S_{t+1}}{S_t}\right) - \left(\mu - \frac{1}{2}V_t\right)}{\sqrt{V_t}}$$

Since

$$\varepsilon_{t+1}^V = \rho\varepsilon_{t+1}^S + \sqrt{1-\rho^2}\varepsilon_{t+1}$$

where $corr(\varepsilon_{t+1}^S, \varepsilon_{t+1}) = 0$, we get

$$\ln(V_{t+1}) =$$

$$\ln(V_t) + \frac{1}{V_t}\left(\kappa\left(\theta - V_t\right) - \frac{1}{2}\sigma^2\right) + \sigma\frac{1}{\sqrt{V_t}}\left(\rho\frac{\ln\left(\frac{S_{t+1}}{S_t}\right) - \left(\mu - \frac{1}{2}V_t\right)}{\sqrt{V_t}} + \sqrt{1-\rho^2}\varepsilon_{t+1}\right)$$

We simulate N states (N particles) which describe the set of possible values of V_{t+1}.

Step 3: Computing and normalizing the weights (Importance Sampling)

At this point, we have a vector of N possible values of V_{t+1} and we know according to equation (9) that given the other available information, V_{t+1} is sufficient to generate $\ln(S_{t+2})$. Therefore, equation (9) offers a simple way to evaluate the likelihood that the observation S_{t+2} has been generated by V_{t+1}. Hence, we have to compute the vector W whose elements represent the weight given to each particle (or the likelihood or probability that the particle j has generated S_{t+2}). The likelihood is computed as follows:

$$W_{t+1}^j = \frac{1}{\sqrt{V_{t+1}^j}} \exp\left(-\frac{1}{2}\frac{\left(\ln\left(\frac{S_{t+2}}{S_{t+1}}\right) - \left(\mu - \frac{1}{2}V_{t+1}^j\right)\right)^2}{V_{t+1}^j}\right)$$

This is to be repeated for $j = 1,..,N$. Finally, because nothing guarantees that $\sum_{j=1}^N W_{t+1}^j = 1$, we have to normalize and set $\overline{W}_{t+1}^j = \frac{W_{t+1}^j}{\sum_{j=1}^N W_{t+1}^j}$. In summary therefore, at the end of Step 2, we obtain a set of N particles describing the density of V_{t+1}. This procedure (Steps 1, 2 and 3) is repeated for $t = 1, ...T$. To obtain the filtered volatility path, we then compute

$$\bar{V}_{t+1} = \sum_{j=1}^{N} \overline{W}_{t+1}^j V_{t+1}^j$$

for each t.

2.1.3 Computing option prices and evaluating the loss function

We are now in a position to evaluate option prices $C_i\left(\bar{V}_t\right)$ based on the filtered volatility path using Heston's closed form solution according to equation (6). We subsequently evaluate the loss function (7), using $C_i\left(\bar{V}_t\right)$ for the model price $C_i(V)$. We use a standard numerical optimization routine to update the model parameters and iterate until convergence is achieved.[16]

Notice that the methodology we have suggested here for estimating the continuous time stochastic volatility model relies on the same information set and uses the same objective function as the two discrete time models, which will be discussed in detail below. This will allow for a fair empirical comparison between models.

2.2 The Non-Affine Stochastic Volatility Model (NA-SV)

While the empirical literature does not offer much guidance on the choice of a non-affine stochastic volatility models, and several sensible specifications are available, we focus on a process which has the same number of parameters as the affine SV model above. Our non-affine model is inspired by

[16] In order to mitigate the impact of the choice of the initial variance V_0 we start iterating on the volatility dynamic on January 2, 1989 in all the estimation samples and for all the models.

the properties of the NGARCH model of Engle and Ng (1993),[17] which we use later in this paper as our NA-GARCH specification. We use the following NA-SV dynamic

$$dS = \mu S dt + \sqrt{V} S dw^S \tag{11}$$

$$dV = \kappa(\theta - V)dt + \sigma V dw^V \tag{12}$$

with $corr(dw^S, dw^V) = \rho$. In this model the innovations are scaled by the conditional variance rather than by the square root of the conditional variance as is the case in the AF-SV model. It must also be noted that our NA-SV model is related to the Engle and Ng (1993) model in a more precise way, because it can be shown to be the limit of the Engle and Ng (1993) model, when the limit is defined according to Duan (1997).

Under the assumption that the volatility risk premium $\lambda(S, V, t)$ is equal to λV, the risk neutral dynamic expressed in terms of the physical parameters is

$$dS = rSdt + \sqrt{V} S dw^{*S} \tag{13}$$

$$\begin{aligned} dV &= (\kappa(\theta - V) + \lambda V)\, dt + \sigma V dw^{*V} \qquad (14) \\ &= (\kappa - \lambda)(\kappa\theta/(\kappa - \lambda) - V)dt + \sigma V dw^{*V} \qquad (15) \end{aligned}$$

with $corr(dw^{*S}, dw^{*V}) = \rho$.

Note that the assumption on the volatility risk premium is the same as in the AF-SV model. In both cases the risk-neutralization can be obtained using a no-arbitrage argument, but in the AF-SV case the risk-neutral dynamic can also be obtained using a utility-based argument. For the NA-SV model, a utility-based risk-neutralization may be possible, but this cannot be done under the assumption that $\lambda(S, V, t)$ is equal to λV. If we change both the variance dynamic and the price of risk, the resulting change in model fit can be due to either assumption. In order to isolate the importance of the volatility dynamic, we keep the volatility risk premium specifications constant across the affine and non-affine SV models.[18]

The NA-SV model presented here has an unobserved variance factor and no closed form option valuation formula. Thus we need to implement it using the auxiliary particle filter to construct the variance path (as in AF-SV) and Monte Carlo simulation to calculate the option prices. We use 1000 simulated paths and a number of numerical techniques to increase numerical efficiency: the empirical martingale method of Duan and Simonato (1999), stratified random numbers, antithetic variates and a control variate technique. The model is estimated by minimizing (7) with respect to the physical parameters $\mu, \kappa, \theta, \lambda, \rho$ and σ.

[17] See also Duan (1997) and Ritchken and Trevor (1999).

[18] See Lewis (2000) for a thorough discussion of these issues.

2.3 The Affine GARCH(1,1) Model (AF-GARCH)

Heston and Nandi (2000) propose a class of affine GARCH models (AF-GARCH) that allow for a closed-form solution for the price of a European call option. We investigate the GARCH(1,1) version of this model, which is given by

$$\ln(S_{t+1}) = \ln(S_t) + r + ph_{t+1} + \sqrt{h_{t+1}} z_{t+1} \tag{16}$$

$$h_{t+1} = w + bh_t + a\left(z_t - c\sqrt{h_t}\right)^2 \tag{17}$$

where S_{t+1} denotes the underlying asset price, r the risk free rate, p the price of risk and h_{t+1} the daily variance on day $t+1$ which is known at the end of day t. The z_{t+1} shock is assumed to be i.i.d. $N(0,1)$. The Heston-Nandi model captures time variation in the conditional variance in ways similar to Engle (1982) and Bollerslev (1986). The parameter c represents the leverage effect, which captures the negative relationship between returns and volatility (Black (1976)) and results in a negatively skewed conditional distribution of multi-day returns. Note also that using the conventional GARCH notation, the conditional variance for day $t+1$, denoted h_{t+1}, is known at the end of day t.

Variance persistence can be computed via

$$b + ac^2 \equiv 1 - \kappa$$

and the unconditional variance can be computed via

$$(w+a)/\left(1 - b - ac^2\right) = (w+a)/\kappa \equiv \theta$$

Now we can rewrite the variance process as

$$h_{t+1} - h_t = \kappa(\theta - h_t) + a\left(\left(z_t^2 - 1\right) - 2cz_t\sqrt{h_t}\right) \tag{18}$$

which suggests the model's relationship with the diffusion volatility models considered above.

The risk-neutral dynamics for the GARCH(1,1) model (16)-(17) are given by[19]

$$\begin{aligned} \ln(S_{t+1}) &= \ln(S_t) + r - \tfrac{1}{2}h_{t+1} + \sqrt{h_{t+1}} z_{t+1}^* \\ h_{t+1} &= w + bh_t + a(z_t^* - (c + p + 0.5)\sqrt{h_t})^2 \end{aligned} \tag{19}$$

with $z_t^* \sim N(0,1)$ under the risk neutral measure.

We provide an analysis of this model using data on equity option prices as well as the time series of underlying equity returns. In order to value options at each date t, we need an estimate of the conditional volatility h_{t+1} on that particular date. One of the appealing aspects of discrete-time

[19] For the underlying theory on risk neutral distributions in discrete time option valuation see Rubinstein (1976), Brennan (1979), Amin and Ng (1993), Duan (1995), Camara (2003), Heston and Nandi (2000) and Schroder (2004).

GARCH models is that this filtering problem is extremely simple. Indeed, the filtering problem is solved by noting that from (16) we have

$$z_{t+1} = (R_{t+1} - r - ph_{t+1}) / \sqrt{h_{t+1}} \tag{20}$$

where $R_t = \ln(S_t/S_{t-1})$. Substituting (20) in (17), it can be seen that the updating from h_t to h_{t+1} is done using an updating function that exclusively involves observables

$$h_{t+1} = w + bh_t + a((R_t - r) / \sqrt{h_t} - (c+p)\sqrt{h_t})^2 \tag{21}$$

Model parameters are obtained by using the nonlinear least squares (NLS) estimation techniques to minimize

$$\$MSE = \frac{1}{NT} \sum_{t,i} (C_{i,t} - C_i(h_{t+1}))^2 \tag{22}$$

where $NT = \sum_{t=1}^{T} N_t$, T is the total number of days included in the options sample and N_t is the number of options included in the sample at date t, $C_{i,t}$ is the market price of option i quoted on day t and $C_i(h_{t+1})$ is the model price.

The implementation is therefore relatively simple: the NLS routine is called with a set of parameter starting values. The variance dynamic in (21) is then used to update the variance from day to day and the GARCH(1,1) option valuation formula from Heston and Nandi (2000) is used to compute the model prices. At time t, a European call option with strike price K that expires at time T can be calculated from

$$\begin{aligned} C(h_{t+1}) &= e^{-r(T-t)} E_t^*[Max(S_T - K, 0)] \\ &= S_t P_1 - K e^{-r(T-t)} P_2 \end{aligned} \tag{23}$$

where

$$\begin{aligned} P_1 &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re}\left[\frac{K^{-i\phi} f^*(t,T;i\phi+1)}{i\phi S_t e^{r(T-t)}}\right] d\phi \\ P_2 &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re}\left[\frac{K^{-i\phi} f^*(t,T;i\phi)}{i\phi}\right] d\phi \end{aligned} \tag{24}$$

and where $f^*(t,T;i\phi)$ is the conditional characteristic function of the logarithm of the spot price under the risk neutral measure, which is characterized by a set of difference equations with terminal conditions. See Heston and Nandi (2000) for these equations.

2.4 The Non-Affine GARCH Model (NA-GARCH)

We compare the AF-GARCH model with the non-affine NGARCH model of Engle and Ng (1993). We choose the NGARCH model because it is relatively easy to analyze, and because it has been shown to provide a good description of underlying equity returns. We will refer to this model as NA-GARCH. The model is given by

$$\ln(S_{t+1}) = \ln(S_t) + r + p\sqrt{h_{t+1}} - 0.5h_{t+1} + \sqrt{h_{t+1}}z_{t+1} \tag{25}$$

$$h_{t+1} = w + bh_t + ah_t(z_t - c)^2 \tag{26}$$

The variance persistence can be computed via

$$b + a\left(1 + c^2\right) \equiv 1 - \kappa$$

and the unconditional variance can be computed via

$$w/\left(1 - b - a\left(1 + c^2\right)\right) = w/\kappa \equiv \theta$$

Note we can rewrite the variance process as

$$h_{t+1} - h_t = \kappa(\theta - h_t) + ah_t\left(\left(z_t^2 - 1\right) - 2cz_t\right) \tag{27}$$

Notice that the variance of the shock term is $2 + 4c^2$ so that we have

$$h_{t+1} - h_t = \kappa(\theta - h_t) + \sigma h_t v_t$$

where $\sigma = a\sqrt{2 + 4c^2}$ and $v_t = \left(\left(z_t^2 - 1\right) - 2cz_t\right)/\sqrt{2 + 4c^2}$. Notice also that we have

$$Corr\left(z_t, v_t\right) = \frac{-2c}{\sqrt{2 + 4c^2}} \equiv \rho$$

all of which suggests the NA-GARCH model's close relationship with the NA-SV model.

Note that this model differs in some subtle ways from the Heston-Nandi model in (16)-(17). The Heston-Nandi model was engineered with the specific purpose of yielding closed-from option prices. The specification in (25)-(26) does not yield closed form option prices, but was designed to provide a good fit to the underlying equity returns. The question of interest is if the restrictions built into affine models such as (16)-(17) reduce the ability of the model to fit the data.

The risk-neutral dynamics for the NA-GARCH model (25)-(26) can be obtained using the same theoretical arguments underlying the Heston-Nandi model

$$\ln(S_{t+1}) = \ln(S_t) + r - 0.5h_{t+1} + \sqrt{h_{t+1}}z_{t+1}^* \tag{28}$$

$$h_{t+1} = w + bh_t + ah_t(z_t^* - (c + p))^2$$

with $z_t^* \sim N(0, 1)$. We can then estimate the model by minimizing (22), using the updating rule

$$h_{t+1} = w + bh_t + ah_t\left(\left[\left(R_t - r + 0.5h_{t-1}\right)/\sqrt{h_{t-1}}\right] - (c + p)\right)^2 \tag{29}$$

Option prices are computed numerically according to

$$C(h_{t+1}) = e^{-r(T-t)} E_t^*[Max(S_T - K, 0)]$$

where the expectation is calculated by Monte Carlo simulation of the daily returns from (28). We use the same Monte-Carlo setup as in the NA-SV case, with 1000 simulated paths and a number of numerical techniques to increase numerical efficiency: the empirical martingale method of Duan and Simonato (1999), stratified random numbers, antithetic variates and a control variate technique.

3 Empirical Results

This section presents the empirical results. We first discuss the data, followed by an empirical evaluation of the four models under investigation and a detailed discussion of the differences in performance of these models in- and out-of-sample.

3.1 Data

We conduct our empirical analysis using six years of data on S&P 500 call options, for the period 1990-1995. We apply standard filters to the data following Bakshi, Cao and Chen (1997). We only use Wednesday and Thursday options data. For the in-sample analysis, we use the Wednesday data. Wednesday is the day of the week least likely to be a holiday. It is also less likely than other days such as Monday and Friday to be affected by day-of-the-week effects. The decision to pick one day every week is to some extent motivated by computational constraints. The optimization problems are fairly time-intensive, and limiting the number of options reduces the computational burden. Using only Wednesday data allows us to study a fairly long time-series, which is useful considering the highly persistent volatility processes. An additional motivation for only using Wednesday data is that following the work of Dumas, Fleming and Whaley (1998), several studies have used this setup (see for instance Heston and Nandi (2000)).

We obtain six sets of parameter estimates in the in-sample analysis. We simply split the six years of data in six datasets, one for each calendar year, and perform annual estimation exercises. For each estimation sample, we use a volatility updating rule starting from the model implied unconditional variance on January 1, 1989.

Table 1 presents descriptive statistics for the options data for the 1990-1995 Wednesday in-sample data by moneyness and maturity. Panels A and B indicate that the data are standard. Panel C displays the volatility smirk in the data. The slope of the smirk clearly differs across maturities. We summarize the data for all six estimation samples in one set of tables to save space. Descriptive statistics for the yearly samples (not reported here) reveal similar stylized facts. The slope of the smirk changes over time, but the smirk is present throughout the sample. The top panel of Figure 1 gives some indication of the pattern of implied volatility over time. For the 313 Wednesdays of options data used in the empirical analysis, we present the average implied volatility of the options on each Wednesday. It is evident from Figure 1 that there is substantial clustering in implied volatilities. It can also be seen that volatility is higher in the early part of the sample. The bottom panel of Figure 1 presents a time series for the 30-day at-the-money volatility (VIX) index from the CBOE for our sample period. A comparison with the top panel clearly indicates that the options data in our sample are representative of market conditions, although the time series based on our sample is of course a bit more noisy due to the presence of options with different moneyness and maturities.

After conducting six in-sample estimations, we proceed to conduct separate out-of-sample analyses for each of the six sample years using the trading day following each in-sample Wednesday. We refer to these datasets as the Thursday data. Table 2 presents descriptive statistics for the out-of-sample data. The patterns in the data are clearly similar to those in the in-sample data in Table 1.

3.2 Parameter Estimates and Option Mean-Squared-Errors

Table 3 presents the parameter estimates for each of the four models and for each of the six annual estimation samples. The parameters for the AF-SV and NA-SV models are directly interpretable individually; for example, κ denotes the daily variance mean reversion and θ denotes the daily unconditional variance. These parameters appear to be relatively stable over time, even though the parameter estimates for 1990 and 1991 seem to be somewhat different from the estimates for the subsequent four years. The most striking aspect for the 1990 and 1991 estimates is that the correlation parameter ρ hits the prespecified boundary of -0.999 for the AF-SV and NA-SV models. This seems to indicate that it is difficult for the AF-SV and NA-SV models to match the negative correlation in the data for this period. The NA-SV model also hits the boundary for ρ in 1993 and 1994.

For the AF-GARCH and NA-GARCH models, we report the parameters from the specifications in (18) and (27) in order to facilitate comparision with the SV models. For the non-affine models the parameters are quite comparable across models. For the affine models, a comparison is less straightforward, partly because the conditional correlation is time-varying in the AF-GARCH model. Note that c, the parameter that determines the size of the leverage effect in the AF-GARCH model, is higher in 1990 and 1991, consistent with the ρ estimates for the AF-SV model.

To further facilitate the comparison between the different models, Table 4 complements Table 3 by focusing on variance persistence and the unconditional volatility reported under both the physical and risk neutral measures. The variance persistence is close to one for all models, which is consistent with other findings in the literature. It is higher under the risk neutral measure, which indicates a negative price of volatility risk. Often the persistence is largest for the AF-SV specification. The unconditional volatility displays considerable variation over time. In 1991 the persistence in the NA-GARCH model is very close to 1, leading to an unrealistically large estimate of the unconditional variance.[20] Keeping in mind the difficulties of the AF-SV and NA-SV models to capture the leverage effect mentioned above, we therefore conclude that it is challenging for most models to provide a satisfactory fit to the 1991 data.

The in- and out-of-sample RMSEs from the four models are reported in Table 5 for each of the six samples. First note that the NA-SV model is best overall and the AF-SV model is worst overall both in- and out-of-sample. The overall difference between these models is 17.4% in-sample and 19.8% out of sample. The performance of the NA-GARCH and NA-SV models is similar in- and out-of-sample with the NA-SV model performing somewhat better overall. This correspondence between the fit of the non-affine continuous-time and discrete-time models does not obtain for the affine models. The fit of the AF-GARCH model falls in between the AF-SV and the non-affine models. The AF-GARCH is approximately 7.3% better than the AF-SV in sample and approximately 8.1% better out of sample. Looking across the six samples, a very robust conclusion obtains: the non-affine models substantially outperform the affine models in every year, both in- and out-of-sample. Nevertheless, we note that the parameter estimates indicate potential misspecification of the NA-SV model. In particular, the estimate of ρ is equal to the pre-specified boundary of -0.999 in several samples. This observation confirms that further work on the non-affine model is needed. The objective of this paper is not to find the best possible non-affine model. We merely want to demonstrate that a simple (and possibly misspecified) non-affine model improves on the benchmark affine model.

Figure 2 provides further perspective on the similarities and differences between the four models by providing volatility sample paths for the models. The figure plots volatility paths for the four models using two sets of parameters estimates for each, the estimates for the 1990 sample (left column) and the estimates for the 1995 sample (right column). It can be seen that despite the fact that there are significant differences in the estimated long-run volatility between the model estimates for the 1990 sample, the sample paths for the four models are quite similar. For the 1995 estimates, the models all display increases in volatility around the same period, but the overall sample paths are rather different for the four models. When comparing the sample paths for the 1990 and 1995 estimates for a given model, it is clear that the 1995 estimates display lower persistence. The larger volatility of volatility estimates in the 1995 sample for all models except the AF-SV are also apparent.

[20] The high unconditional variance does not signal model misspecification because the conditional variance is a function of the product of the unconditional variance and one minus the persistence.

3.3 Pricing Errors Over Time and Across Moneyness and Maturity

We now analyze the performance of the four models in more detail. Figure 3 addresses the performance of the models over time. The top four panels present the RMSE on a week-by-week basis. It can clearly be seen that the four models display important similarities in terms of the pricing patterns that they can and cannot explain. This observation is confirmed by inspecting the week-by-week bias in Figure 4.

What is even more striking in Figure 3 are the similarities in RMSE over time for the two affine models on the one hand and the two non-affine models on the other hand. Whether the model is affine or not is clearly much more important for its performance than whether it is formulated in continuous time or not. This confirms the message from the overall RMSEs in Table 5, but the message is much more striking when delivered visually on a week-by-week basis as in Figure 3.

It is also interesting to visually inspect the relationship between the level of the volatility in the bottom panel of Figures 3 and 4 and the RMSE and bias for the different models. It is clear that in periods of high volatility, the RMSE for all models increases.

Tables 6 and 7 present an analysis of the in- and out-of-sample RMSE by moneyness and maturity. The most important conclusion from Tables 6 and 7 is that the two non-affine models outperform the corresponding affine models for almost every cell in the moneyness-maturity matrix, in-sample as well as out-of-sample. Consider in particular the All strike prices rows and the All maturities columns in each panel of Tables 6 and 7. It is striking that the NA-SV model performs the best in every single case and the NA-GARCH model performs second best also in every single case.

These tables also allow for some other important conclusions. For example, consider the difference between the AF-SV and AF-GARCH models. While the overall RMSE difference between the two models in Table 5 is approximately 8% in- and out-of-sample, there are important variations in the relative performance of the models across maturity. The overall RMSE of the AF-SV model is larger than that of the AF-GARCH model, but for short-maturity options the AF-SV model performs significantly better. This finding is perhaps somewhat surprising. While we believe that the limit arguments have some empirical value, our prior was that continuous-time models would prove to be somewhat restrictive because of the assumption of a continuous sample path. This restriction is well recognized in the continuous-time option valuation literature, and stochastic volatility models are augmented with jump models to improve their performance. However, jump models are believed to help the performance of stochastic volatility models mainly for short-maturity options. We therefore expected that if the Heston (1993) AF-SV model would underperform the AF-GARCH model, it would be for short-maturity options. Instead, the AF-SV model seems to underperform the AF-GARCH model mainly for longer maturities.

3.4 Conditional Moment Dynamics

Ultimately, the option prices for the different models are determined by the model-implied conditional density dynamics. Therefore, we now discuss model differences by focusing on various aspects of the conditional density.

In order to asses the different models' ability to generate time-variation in the asymmetry of the return distribution, Figure 5 plots the conditional covariance between returns and variances for each model. Obviously, for the two stochastic volatility models, we use the discretized version of the models. We refer to this conditional covariance as the conditional leverage path, which for the four models is given by

$$\begin{aligned} \text{AF-GARCH} &: \quad cov_t\left(R_{t+1}, h_{t+2}\right) = -2ach_{t+1} \\ \text{AF-SV} &: \quad cov_t\left(R_{t+1}, V_{t+1}\right) = \sigma\rho V_t \\ \text{NA-GARCH} &: \quad cov_t\left(R_{t+1}, h_{t+2}\right) = -2ach_{t+1}^{3/2} \\ \text{NA-SV} &: \quad cov_t\left(R_{t+1}, V_{t+1}\right) = \sigma\rho V_t^{3/2} \end{aligned} \tag{30}$$

Notice that critical differences between the affine and non-affine models show up very prominently in these conditional moments. The 3/2 term on the volatility of the non-affine models suggests that these models may be able to exhibit more variation in the conditional leverage paths. Figure 5 confirms this intuition. For each model we plot the daily conditional leverage path during 1990-1995, annualized by multiplying by 252. The left column uses the 1990 estimates from Table 3 and the right column uses the 1995 estimates. The four rows of panels correspond to the AF-GARCH, AF-SV, NA-GARCH and NA-SV models respectively. Notice that the scaling is different between the two columns, because the 1995 estimates imply a much larger level of (and variation in) the leverage effect.

The main conclusion is that for both sets of estimates the non-affine models imply more substantial leverage, as well as more substantial variation over time in the leverage effect. Given the importance of the leverage effect for option valuation, this may be a very important factor in explaining the differences in the fit between affine and non-affine models documented in Table 5.

Option prices are a function of the conditional variance, and therefore the variation in option prices over time is related to the conditional variance of variance. The conditional variance of variance of returns for the four models are given by

$$\begin{aligned} \text{AF-GARCH} &: \quad Var_t(h_{t+2}) = 2a^2 + 4a^2c^2h_{t+1} \\ \text{AF-SV} &: \quad Var_t(V_{t+1}) = \sigma^2 V_t \\ \text{NA-GARCH} &: \quad Var_t(h_{t+2}) = a^2(2+4c^2)h_{t+1}^2 \\ \text{NA-SV} &: \quad Var_t(V_{t+1}) = \sigma^2 V_t^2 \end{aligned} \tag{31}$$

Notice again that these conditional moments indicate important differences between affine and non-affine models. The conditional variance shows up in levels in the affine models and in squared form in the non-affine models, which again suggests that non-affine models will display more variation in the conditional variance of variance.[21] Figure 6 reports the empirical results for the volatility of variance. For each model, we plot the daily conditional volatility of variance path during 1990-1995 annualized by multiplying by 252. The left column uses the 1990 estimates from Table 3 and the right column uses the 1995 estimates. The four rows of panels correspond to the AF-GARCH, AF-SV, NA-GARCH and NA-SV models respectively, and again the scaling differs between the columns because the 1995 estimates imply higher conditional volatility of variance. Figure 6 indicates that for both sets of estimates the non-affine models also display much more time-variation in the volatility of variance. These differences between the models further help us understand the superior fit of the non-affine models.

3.5 State Price Densities

Figures 7 and 8 provide additional intuition for the differences in performance between the four models. Using estimates for the 1990 and 1995 samples respectively, these figures depict the simulated state price densities for a one-month, three-month and one-year horizon. Each row of panels reports the risk neutral distribution of the index return according to the AF-GARCH, AF-SV, NA-GARCH and NA-SV models respectively. The normal distribution corresponding to the Black-Scholes model is superimposed for reference. The left column reports the 1-month horizon, the center column the 3-month horizon distribution and the right column shows the 1-year distribution. The distributions are constructed by simulating daily returns from each model setting the initial spot variance equal to the unconditional variance. Kernel density estimates are then constructed from the standardized simulated returns.

It can clearly be seen that deviations from normality are critical, and that the estimated parameters for the four models imply different deviations from normality. It is interesting to note that especially for the affine models, the leverage effects generate a substantial amount of skewness in the risk-neutral return distributions, even at the 1-year horizon. This finding is consistent with the nonparametric evidence in Ait-Sahalia and Lo (1998) that skewness persists at long horizons.

3.6 Discussion

The main objective of the empirical comparison between the four option valuation models is to investigate the implications of assuming an affine model structure. While this is a simple and

[21] Notice also that in the affine GARCH model the variance of variance will be constant when $c = 0$ whereas this is not the case in the non-affine models nor in the AF-SV model.

important question, the available literature does not contain a conclusive answer, although specification tests in Jones (2003) indicate that generalizations of the affine framework might be useful. Our use of the Auxiliary Particle Filter, which is new in the option valuation literature, allows us to compare the benchmark AF-SV Heston (1993) model with a simple NA-SV model along a simple dimension such as the dollar RMSE. Our results are very clear: while the affine model is analytically convenient, it substantially underperforms the non-affine model. Interestingly, the discrete-time NA-GARCH model also outperforms the AF-GARCH model, but the differences in fit are much smaller.

These empirical results also allow us to comment on the relationship between the empirical performance of discrete-time and continuous-time models. Existing theoretical limit results have sometimes been interpreted as suggesting that the performance of these models ought to be similar, and indeed as suggesting that a study of the empirical differences is not worthwhile. However, this interpretation is contradicted at a theoretical level by the fact that a single discrete-time model can be linked with several continuous-time limits and vice versa. Our empirical results clearly confirm that the relationship between discrete-time and continuous-time models is far from obvious.

Our most important result in this respect is that the AF-SV Heston (1993) model, which is one of the most popular option valuation models, performs rather poorly when compared to the AF-GARCH Heston-Nandi (2000) model. This may seem surprising, because Heston and Nandi (2000) demonstrate that a restricted version of the Heston (1993) model can be obtained as a limit of their model. However, close inspection of the proof in Heston and Nandi (2000) reveals that this limit result is a very special case, and the restricted Heston (1993) limit may not be the limit that is most relevant from an empirical perspective.

In order to explore this issue further, consider the conditional correlation between returns and volatility for the four models, which can be computed from (30) and (31)

$$
\begin{aligned}
\text{AF-GARCH} &: \quad Corr_t\left(R_{t+1}, h_{t+2}\right) = \frac{-2ch_{t+1}}{\sqrt{(2+4c^2h_{t+1})h_{t+1}}} \\
\text{AF-SV} &: \quad Corr_t\left(R_{t+1}, V_{t+1}\right) = \frac{\sigma\rho V_t}{\sqrt{V_t\sigma^2 V_t}} = \rho \\
\text{NA-GARCH} &: \quad Corr_t\left(R_{t+1}, h_{t+2}\right) = \frac{-2c}{\sqrt{2+4c^2}} = \rho \\
\text{NA-SV} &: \quad Corr_t\left(R_{t+1}, V_{t+1}\right) = \frac{\sigma\rho V_t^{3/2}}{\sqrt{V_t\sigma^2 V_t^2}} = \rho
\end{aligned} \tag{32}
$$

Clearly an important difference is that for the AF-GARCH model the conditional correlation is time-varying whereas it is constant for the other three models. Thus, while the AF-GARCH and AF-SV models are similar in many ways, they differ along this important dimension. To investigate the empirical significance of this restriction, we implemented an affine continuous-time stochastic volatility model that contains a time-varying conditional correlation (not reported).

Option valuation in this model has to be done using Monte Carlo simulation. This model improves on the empirical performance of the Heston (1993) model, and in fact its performance is not very different from that of the AF-GARCH model, suggesting that time-variation in the conditional correlation improves empirical performance.

4 Conclusions and Directions for Future Work

This paper provides an empirical comparison of affine and non-affine option valuation models. The in-sample RMSE of a non-affine stochastic volatility model (NA-SV) is approximately 17.4% lower than that of the affine Heston (1993) model (AF-SV), and the out-of-sample RMSE is approximately 20% lower. The non-affine model outperforms the affine model in all six in-sample exercises as well as in all six out-of-sample exercises. The non-affine model also outperforms the affine model for all moneyness and maturity categories.

We also study the differences between the affine discrete-time GARCH option model of Heston and Nandi (2000) (AF-GARCH) and a non-affine GARCH model (NA-GARCH). Interestingly, the differences in RMSE between these two models are smaller, approximately 7.5% in-sample and 10% out-of-sample. While the performance of the NA-GARCH model is very similar to that of the NA-SV model, the AF-GARCH model performs significantly better than the AF-SV model, in-sample as well as out-of-sample. The AF-GARCH does not lead to a uniformly better fit than the AF-SV model, though: the AF-SV model outperforms the AF-GARCH model for short maturities. This is somewhat surprising because the literature focuses on improving the AF-SV model's performance at short maturities.

These empirical results allow us to draw two conclusions. First, regarding the distinction between affine and non-affine models: while the focus of the option valuation literature on affine models is well motivated, because the resulting closed-form solutions are extremely convenient, our results suggest that this analytical convenience comes at a price, and non-affine models need to be studied more extensively. Second, our results are relevant for the relationship between continuous-time and discrete-time option valuation models. Our empirical results on this issue have to be interpreted cautiously and do not allow for general conclusions. In our opinion, the most important aspect of our empirical comparison between the empirical performance of discrete-time and continuous-time models is what it says about the performance of some important benchmarks in the literature.

Our most interesting conclusion in this respect is that the most important model in the literature on option valuation under stochastic volatility, the AF-SV Heston (1993) model, performs rather poorly when compared to the AF-GARCH Heston-Nandi (2000) model. This may seem surprising, because Heston and Nandi (2000) demonstrate that a restricted version of the Heston (1993) model can be obtained as a limit of their model. However, close inspection of the proof in Heston and Nandi (2000) reveals that this limit result is a very special case, and the restricted Heston (1993) limit may not be the limit that is most relevant from an empirical perspective. Note also that the AF-SV Heston (1993) model is effectively a model with two stochastic shocks, while the AF-GARCH Heston and Nandi (2000) model, like other GARCH models, only contains one stochastic shock. It is safe to say that the relationship between discrete-time models and continuous-time models is complex, and our paper certainly does not provide the final answer to the empirical relationship between discrete-time and continuous-time models. For instance, it is a puzzle why the AF-SV model outperforms the AF-GARCH model for short maturities, even though it underperforms on average. One possible explanantion is that the objection function we use (dollar RMSE) puts more weight on the longer term and more expensive options.

Finally, at the methodological level, this paper presents a new method to estimate continuous-time option valuation models. While this method may seem complex, it is rather flexible and it is straightforward to implement. Its is substantially easier than existing methods used to investigate options data and underlying equity returns jointly. We plan to study the performance of this method in more detail in future work. Another interesting avenue for future work is the search for non-affine models that better fit the data, perhaps following Jones (2003). Moreover, it will also be interesting to investigate whether jump processes that are successful in the affine literature can be used to improve the fit of non-affine models.

References

[1] Ait-Sahalia, Y. and A. Lo (1998), Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices, Journal of Finance, 53, 499-547.

[2] Amin, K. and V. Ng (1993), ARCH Processes and Option Valuation, Manuscript, University of Michigan.

[3] Andersen, T.G., L. Benzoni, and J. Lund (2002), Estimating Jump-Diffusions for Equity Returns, Journal of Finance, 57, 1239-1284.

[4] Bakshi, G., C. Cao, and Z. Chen (1997), Empirical Performance of Alternative Option Pricing Models, Journal of Finance, 52, 2003-2049.

[5] Bates, D. (1996), Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options, Review of Financial Studies, 9, 69-107.

[6] Bates, D. (2000), Post-'87 Crash Fears in the S&P 500 Futures Option Market, Journal of Econometrics, 94, 181-238.

[7] Bates, D. (2004), Maximum Likelihood Estimation of Latent Affine Processes, Working Paper, University of Iowa.

[8] Benzoni, L. (2002), Pricing Options Under Stochastic Volatility: An Empirical Investigation, Working Paper, University of Minnesota.

[9] Black, F. (1976), Studies of Stock Price Volatility Changes, in: Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association, 177-181.

[10] Black, F., and M. Scholes (1973), The Pricing of Options and Corporate Liabilities, Journal of Political Economy, 81, 637-659.

[11] Bollen, N., and E. Rasiel (2003), The Performance of Alternative Valuation Models in the OTC Currency Options Market, Journal of International Money and Finance, 22, 33–64.

[12] Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, Journal of Econometrics, 31, 307-327.

[13] Bollerslev, T., Chou, R. and K. Kroner (1992), ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence, Journal of Econometrics, 52, 5-59.

[14] Bollerslev, T. and H.O. Mikkelsen (1996), Long-Term Equity AnticiPation Securities and Stock Market Volatility Dynamics, Journal of Econometrics, 92, 75-99.

[15] Brennan, M. (1979), The Pricing of Contingent Claims in Discrete-Time Models, Journal of Finance, 34, 53-68.

[16] Broadie, M., Chernov, M. and M. Johannes (2004), Model Specification and Risk Premiums: The Evidence from the Futures Options, manuscript, Graduate School of Business, Columbia University.

[17] Camara, A. (2003), A Generalization of the Brennan-Rubinstein Approach for the Pricing of Derivatives, Journal of Finance, 58, 805-819.

[18] Carr, P. and L. Wu (2004), Time-Changed Levy Processes and Option Pricing, Journal of Financial Economics, 17, 113–141.

[19] Chernov, M., A.R. Gallant, E. Ghysels, and G. Tauchen, (2003), Alternative Models for Stock Price Dynamics, Journal of Econometrics, 116, 225-257.

[20] 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 Option Valuation, Journal of Financial Economics, 56, 407-458.

[21] Christoffersen, P. and K. Jacobs (2004a), Which GARCH Model for Option Valuation?, Management Science, 50, 1204-1221.

[22] Christoffersen, P. and K. Jacobs (2004b), The Importance of the Loss Function in Option Valuation, Journal of Financial Economics, 72, 291-318.

[23] Corradi, V. (2000), "Reconsidering the Continuous Time Limit of the GARCH(1,1) Process," Journal of Econometrics, 96, 145-153.

[24] Dai, Q. and K.J. Singleton (2000), Specification Analysis of Affine Term Structure Models, Journal of Finance, 55, 1943-1978.

[25] Davidson, J. (1994), Stochastic Limt Theory: An Introduction for Econometricians, Oxford University Press.

[26] Diebold, F. and J. Lopez (1995), Modeling Volatility Dynamics, in Macroeconometrics: Developments, Tensions and Prospects, Kevin Hoover (ed.). Boston: Kluwer Academic Press.

[27] Duan, J.-C. (1995), The GARCH Option Pricing Model, Mathematical Finance, 5, 13-32.

[28] Duan, J.-C., (1997), Augmented GARCH(p,q) Process and its Diffusion limit, Journal of Econometrics, 79, 97-127.

[29] Duan, J.-C. and J.-G. Simonato (1998), Empirical Martingale Simulation for Asset Prices, Management Science, 44, 1218-1233.

[30] Duffie, D., and R. Kan (1996), A Yield-Factor Model of Interest Rates, Mathematical Finance, 6, 379-406.

[31] Dumas, B., Fleming, F. and R. Whaley (1998), Implied Volatility Functions: Empirical Tests, Journal of Finance, 53, 2059-2106.

[32] Engle, R. (1982), Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation, Econometrica, 50, 987-1008.

[33] Engle, R. and C. Mustafa (1992), Implied ARCH Models from Options Prices, Journal of Econometrics, 52, 289-311.

[34] Engle, R. and V. Ng (1993), Measuring and Testing the Impact of News on Volatility, Journal of Finance, 48, 1749-1778.

[35] Eraker, B. (2004), Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, Journal of Finance, 59, 1367-1403.

[36] Eraker, B., M. Johannes, and N. Polson (2003) The Role of Jumps in Returns and Volatility, Journal of Finance, 58, 1269-1300.

[37] Gallant, R., Chernov, M., Ghysels, E. and G. Tauchen (2003), Alternative Models for Stock Price Dynamics, Journal of Econometrics, 116, 225-257.

[38] Heston, S. (1993), A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, Review of Financial Studies, 6, 327-343.

[39] Heston, S. and S. Nandi (2000), A Closed-Form GARCH Option Pricing Model, Review of Financial Studies, 13, 585-626.

[40] Hsieh, K. and P. Ritchken (2000), An Empirical Comparison of GARCH Option Pricing Models, manuscript, Case Western Reserve University

[41] Huang, J.-Z. and L. Wu (2004), Specification Analysis of Option Pricing Models Based on Time-Changed Levy Processes, Journal of Finance, 59, 1405–1439.

[42] Johannes, M., Polson, N. and J. Stroud (2002), Nonlinear Filtering of Stochastic Differential Equations with Jumps, Working Paper, University of Chicago.

[43] Jones, C. (2003), The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets, Journal of Econometrics 116, 181-224.

[44] Lewis, A. (2000), Option Valuation Under Stochastic Volatility, Finance Press, Newport Beach, CA.

[45] Nandi, S. (1998), How Important is the Correlation Between Returns and Volatility in a Stochastic Volatility Model? Empirical Evidence from Pricing and Hedging in the S&P 500 Index Options Market, Journal of Banking and Finance 22, 589-610.

[46] Pan, J. (2002), The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, Journal of Financial Economics, 63, 3-50.

[47] Pitt, M., and N. Shephard (1999), Filtering via Simulation: Auxiliary Particle Filters, Journal of the American Statistical Association, 94, 590-599.

[48] Ritchken, P. and R. Trevor (1999), Pricing Options under Generalized GARCH and Stochastic Volatility Processes, Journal of Finance, 54, 377-402.

[49] Rubinstein, M. (1976), The Valuation of Uncertain Income Streams and the Pricing of Options, Bell Journal of Economics, 7, 407-425.

[50] Schroder, M. (2004), Risk Neutral Parameter Shifts and Derivatives Pricing in Discrete Time, Journal of Finance, 59, 2375-2402.

[51] White, H. (2001), Asymptotic Theory for Econometricians, Revised Edition, Academic Press.

Figure 1: Average Implied Volatility in S&P500 Option Data and the VIX

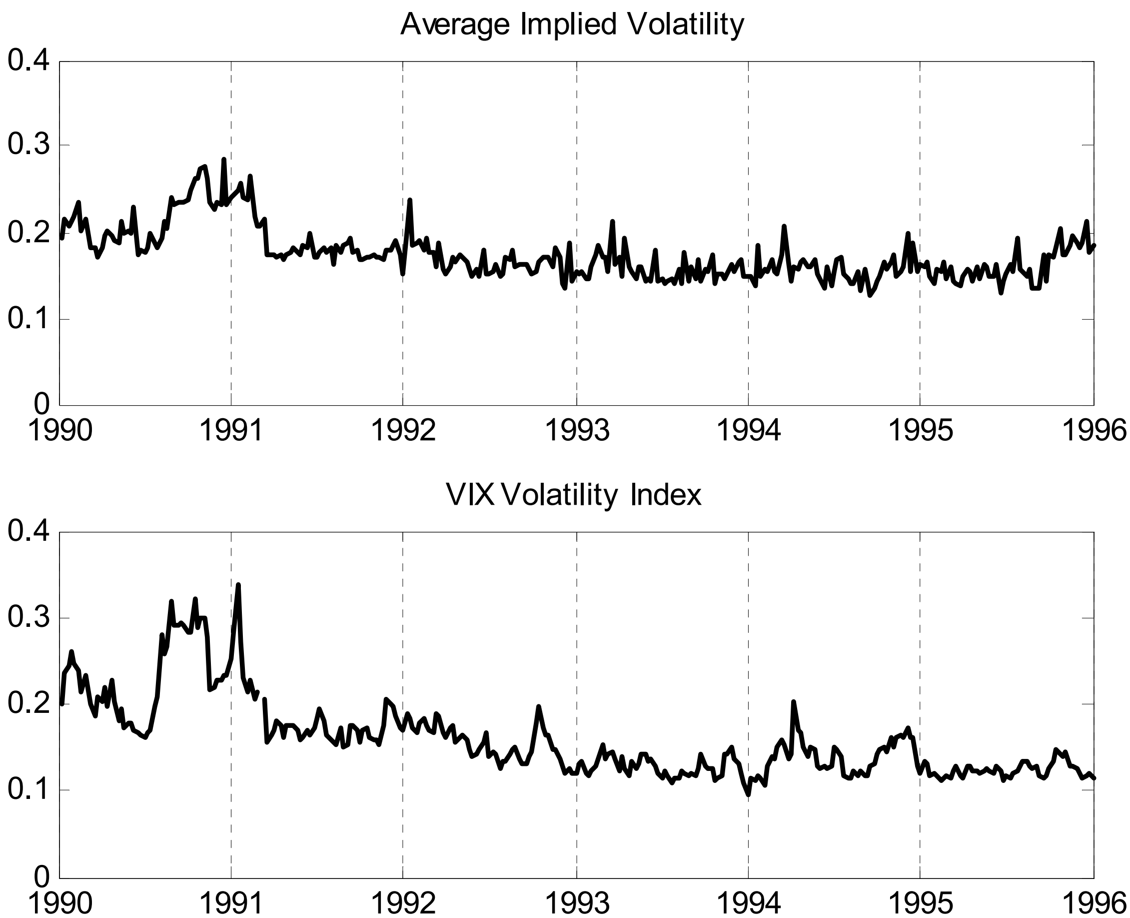

Notes to figure: The top panel plots the average implied Black-Scholes volatility each Wednesday during 1990-1995. The average is taken across maturities and strike prices using the call options in our data set. For comparison, the bottom panel shows the one-month, at-the-money VIX volatility index retrieved from the CBOE website.

Figure 2: Spot Volatility Paths from Each Model. 1990-1995.
Using 1990 and 1995 Estimates

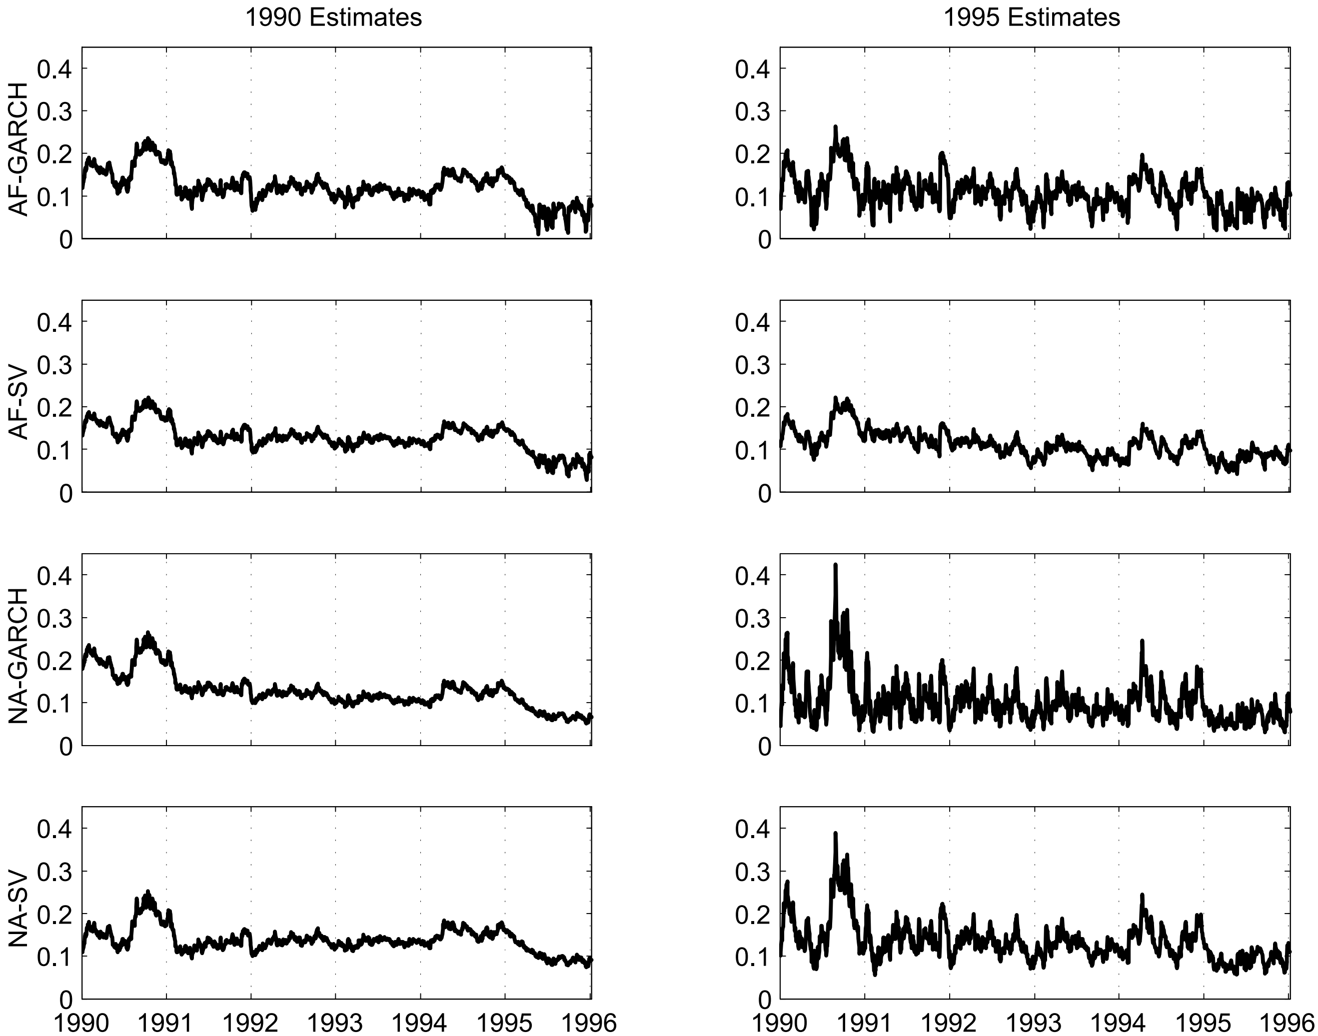

Notes to figure: For each model we plot the daily spot volatility path (annualized) during 1990-1995. The left column uses the 1990 estimates from Table 3 and the right column uses the 1995 estimates. The top row shows the volatility paths from the Heston-Nandi Affine GARCH model, the second row shows the Heston affine stochastic volatility model, the third row shows the non-affine GARCH model and the bottom row shows the non-affine stochastic volatility model.

Figure 3: Weekly In Sample Root Mean Squared Error (RMSE) for each Model. 1990-1995. Average Weekly Implied Volatility (IV) is Shown for Reference

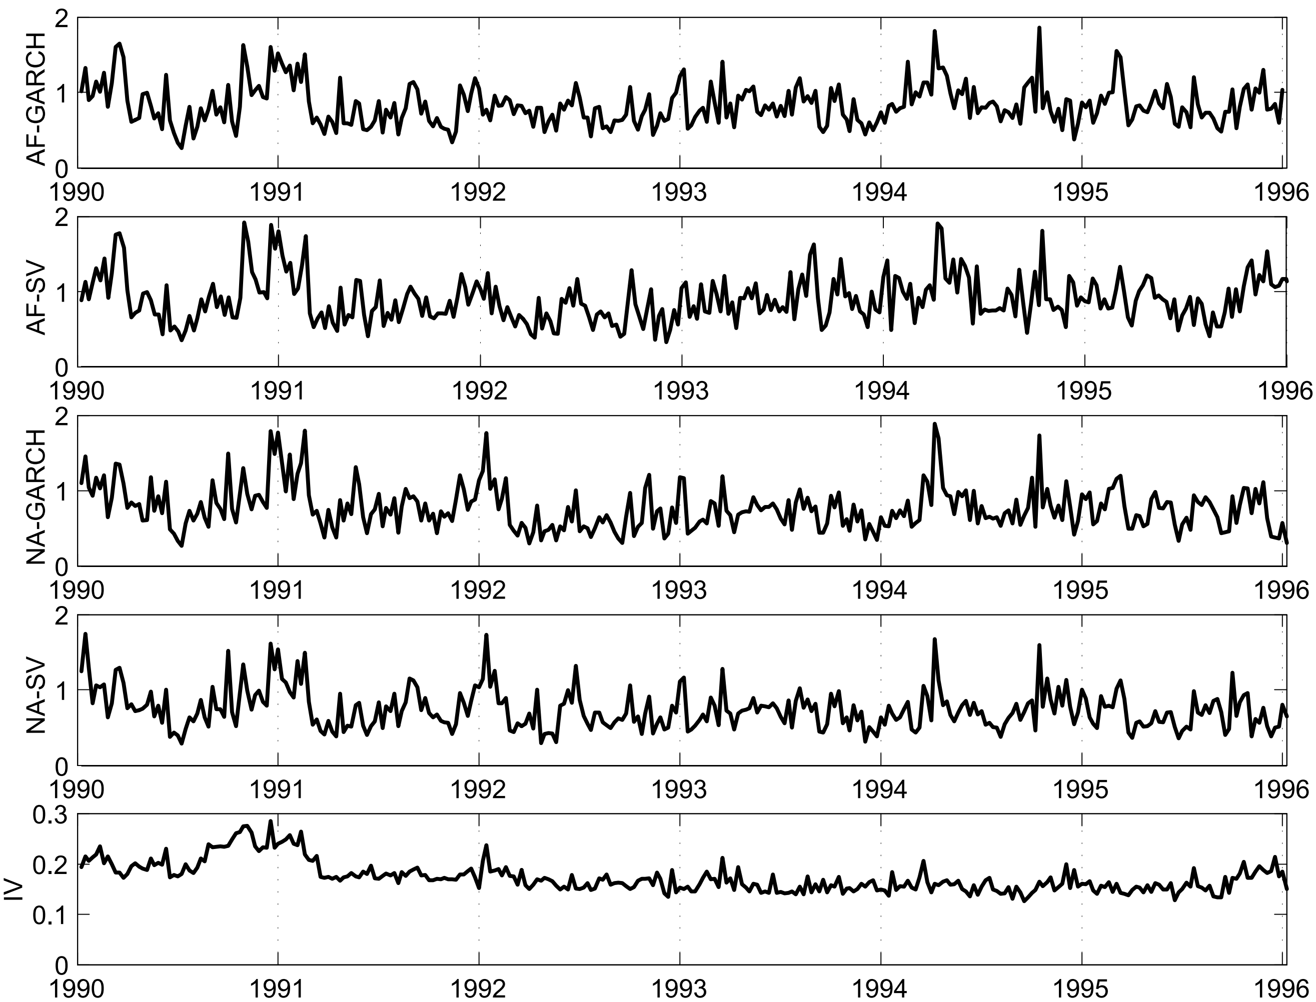

Notes to figure: The four top panels show the weekly root mean squared error (RMSE) for the AF-GARCH, AF-SV, NA-GARCH and NA-SV models respectively. The bottom panel shows the weekly average implied Black-Scholes volatility from Figure 1 for reference.

Figure 4: Weekly In Sample Bias for each Model. 1990-1995.
Average Weekly Implied Volatility (IV) is Shown for Reference

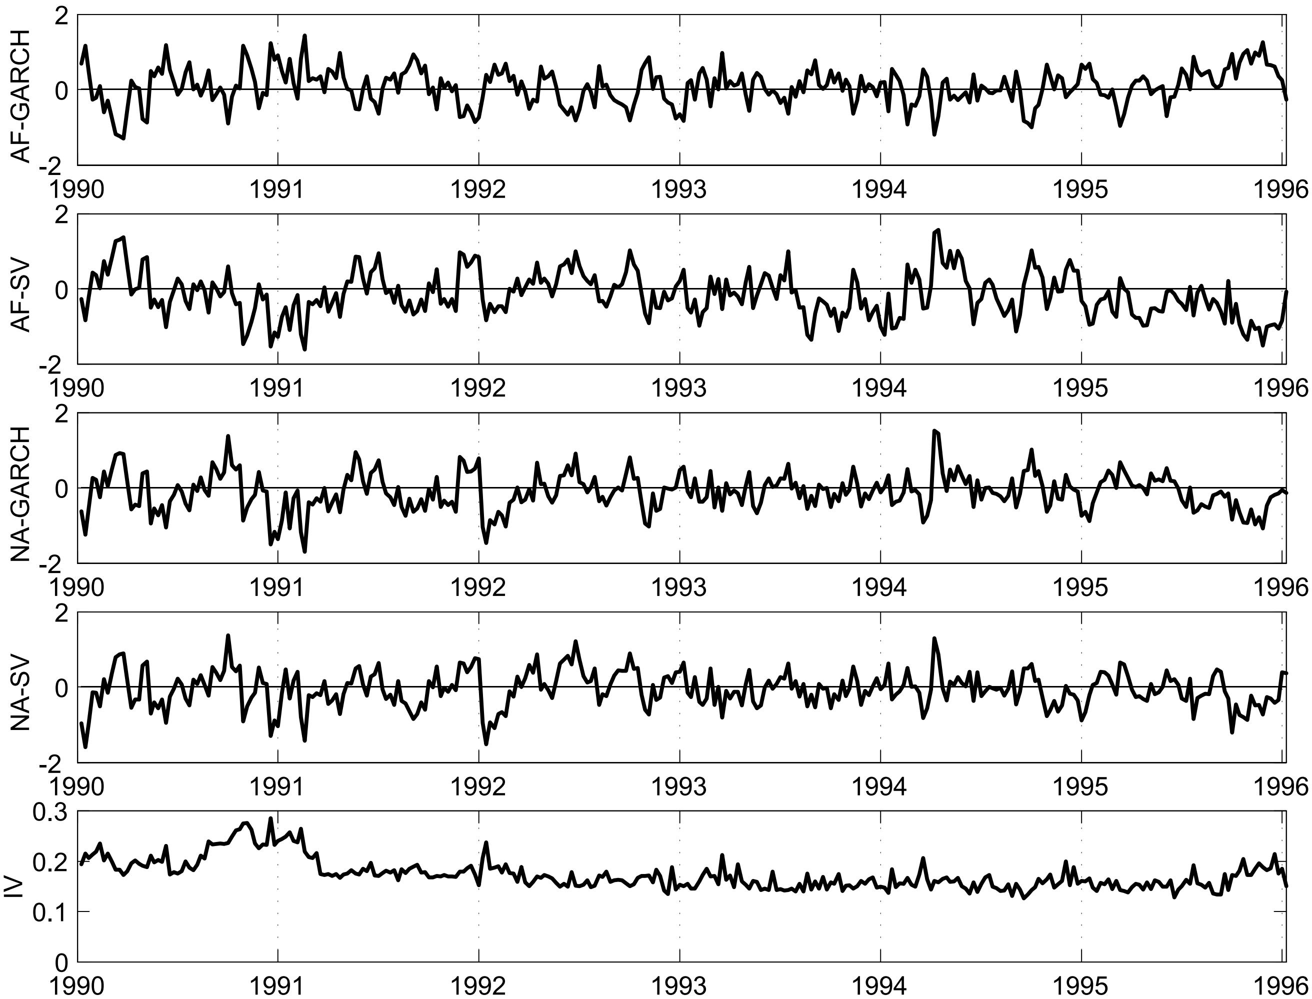

Notes to figure: The top four panels show the weekly bias for the AF-GARCH, AF-SV, NA-GARCH and NA-SV models respectively. The bottom panel shows the weekly average implied Black-Scholes volatility from Figure 1 for reference.

Figure 5: Conditional Leverage Paths from Each Model. 1990-1995.
Using 1990 and 1995 Estimates

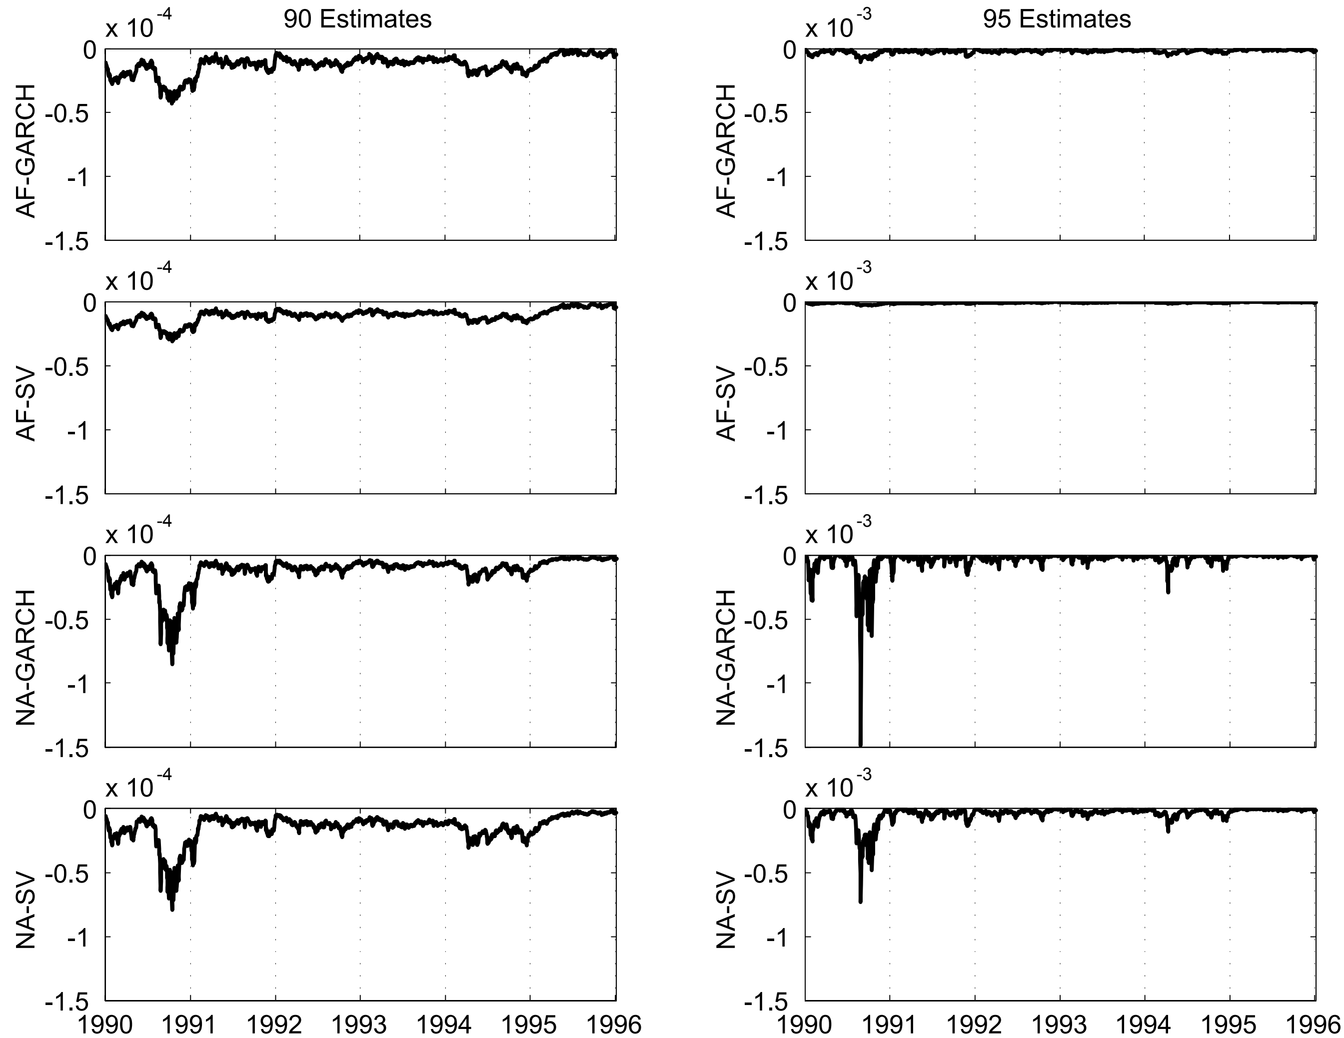

Notes to figure: For each model we plot the daily conditional leverage path defined as the conditional covariance between shocks to returns and shocks to variance. The paths are annualized by multiplying by 252 and plotted during 1990-1995. The left column uses the 1990 estimates from Table 3 and the right column uses the 1995 estimates. Note that the scales are different in the two columns. The top row shows the volatility paths from the Heston-Nandi Affine GARCH model, the second row shows the Heston affine stochastic volatility model, the third row shows the non-affine GARCH model and the bottom row shows the non-affine stochastic volatility model.

Figure 6: Conditional Volatility of Variance Paths from Each Model. 1990-1995.
Using 1990 and 1995 Estimates

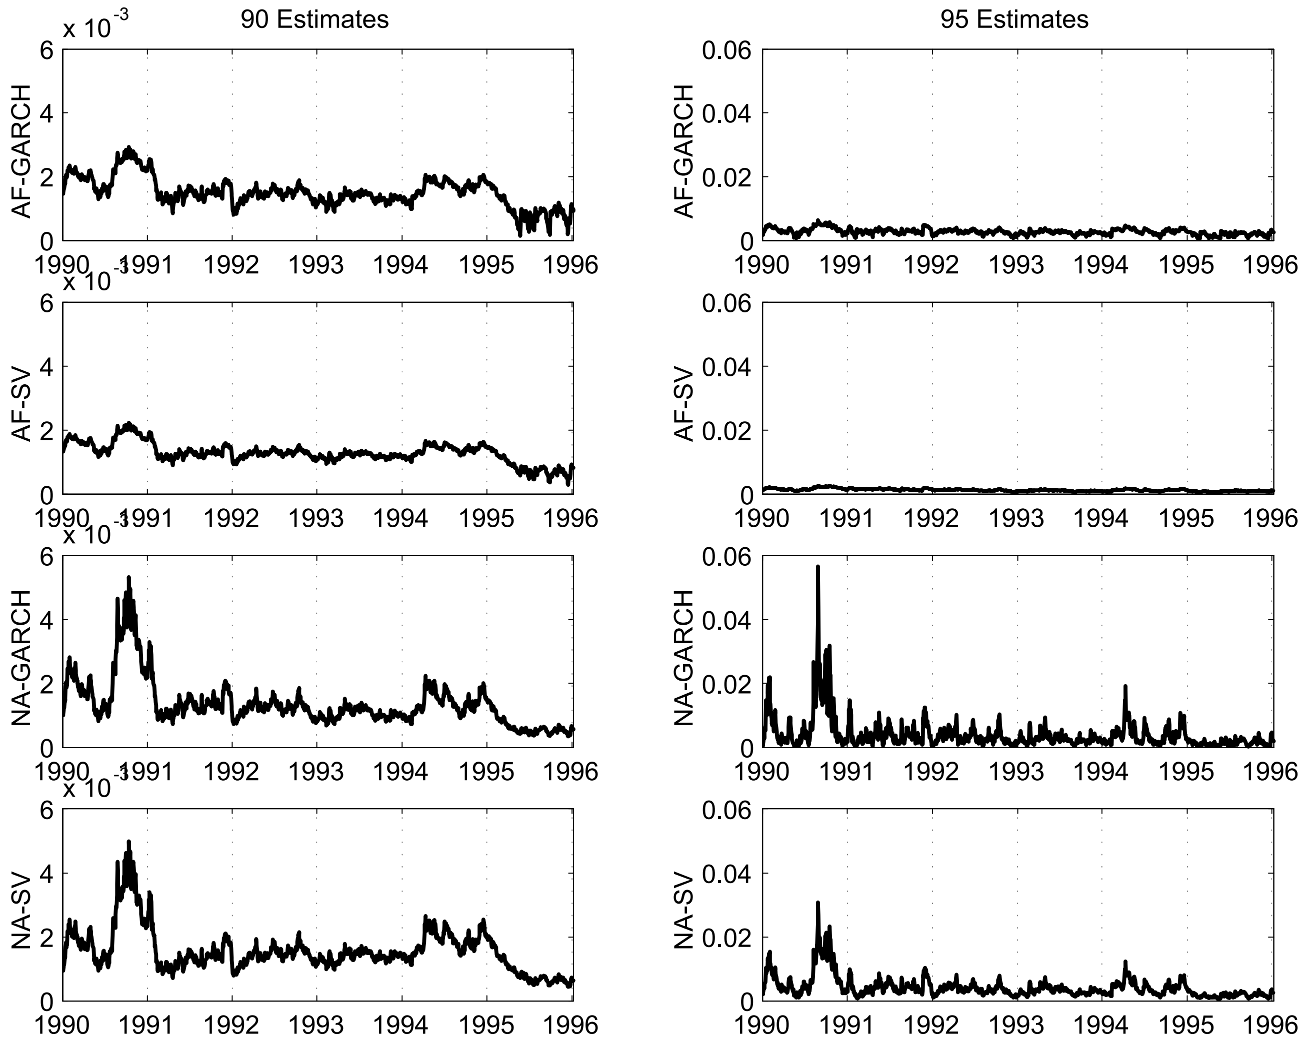

Notes to figure: For each model we plot the daily conditional volatility of variance path defined as the square root of the conditional variance of the conditional variance of returns. The paths are annualized by multiplying by 252 and plotted during 1990-1995. The left column uses the 1990 estimates from Table 3 and the right column uses the 1995 estimates. Note that the scales are different in the two columns. The top row of panels shows the volatility paths from the Heston-Nandi affine GARCH model (AF-GARCH), the second row shows the Heston affine stochastic volatility model (AF-SV), the third row shows the non-affine GARCH model (NA-GARCH) and the bottom row shows the non-affine stochastic volatility model (NA-SV).

Figure 7: Model Implied State Price Densities. 1990 Estimates
1-Month, 3-Month and 1-Year Horizon

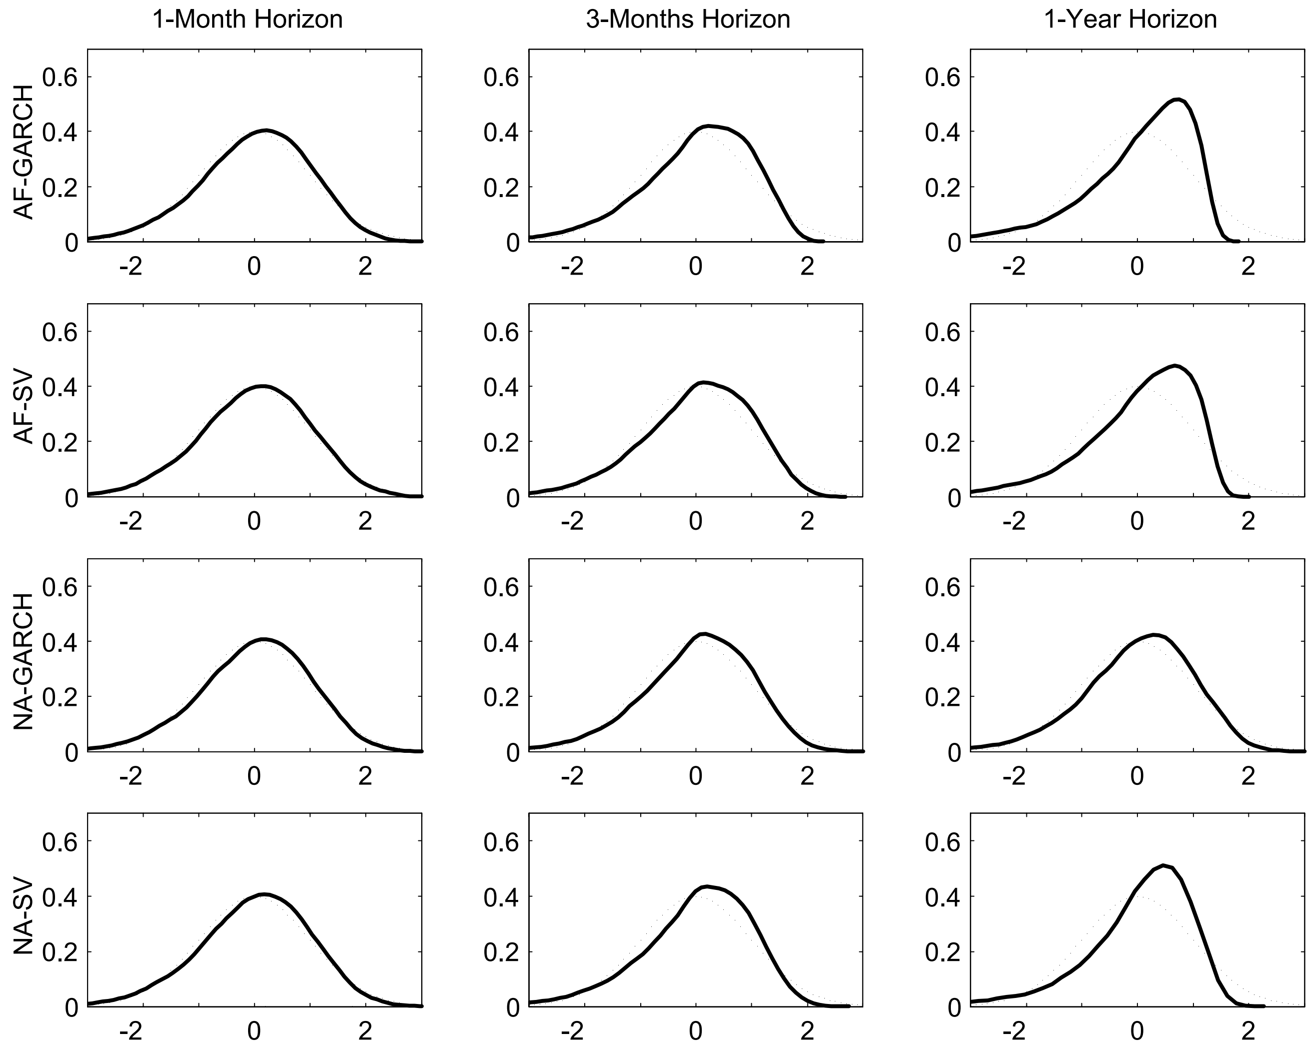

Notes to figure: Each row of panels reports the risk neutral distribution of the index return according to the AF-GARCH, AF-SV, NA-GARCH and NA-SV models respectively. The normal distribution corresponding to the Black-Scholes model is superimposed for reference. The left column reports the 1-month horizon, the center column the 3-month horizon distribution and the right column shows the 1-year distribution. The distributions are constructed by simulating daily returns from each model setting the initial spot variance equal to the unconditional variance. Kernel density estimates are then constructed from the standardized simulated returns. 1990 parameter estimates from Table 3 converted to the risk neutral measure are used to simulate the risk neutral returns.

Figure 8: Model Implied State Price Densities. 1995 Estimates
1-Month, 3-Month and 1-Year Horizon

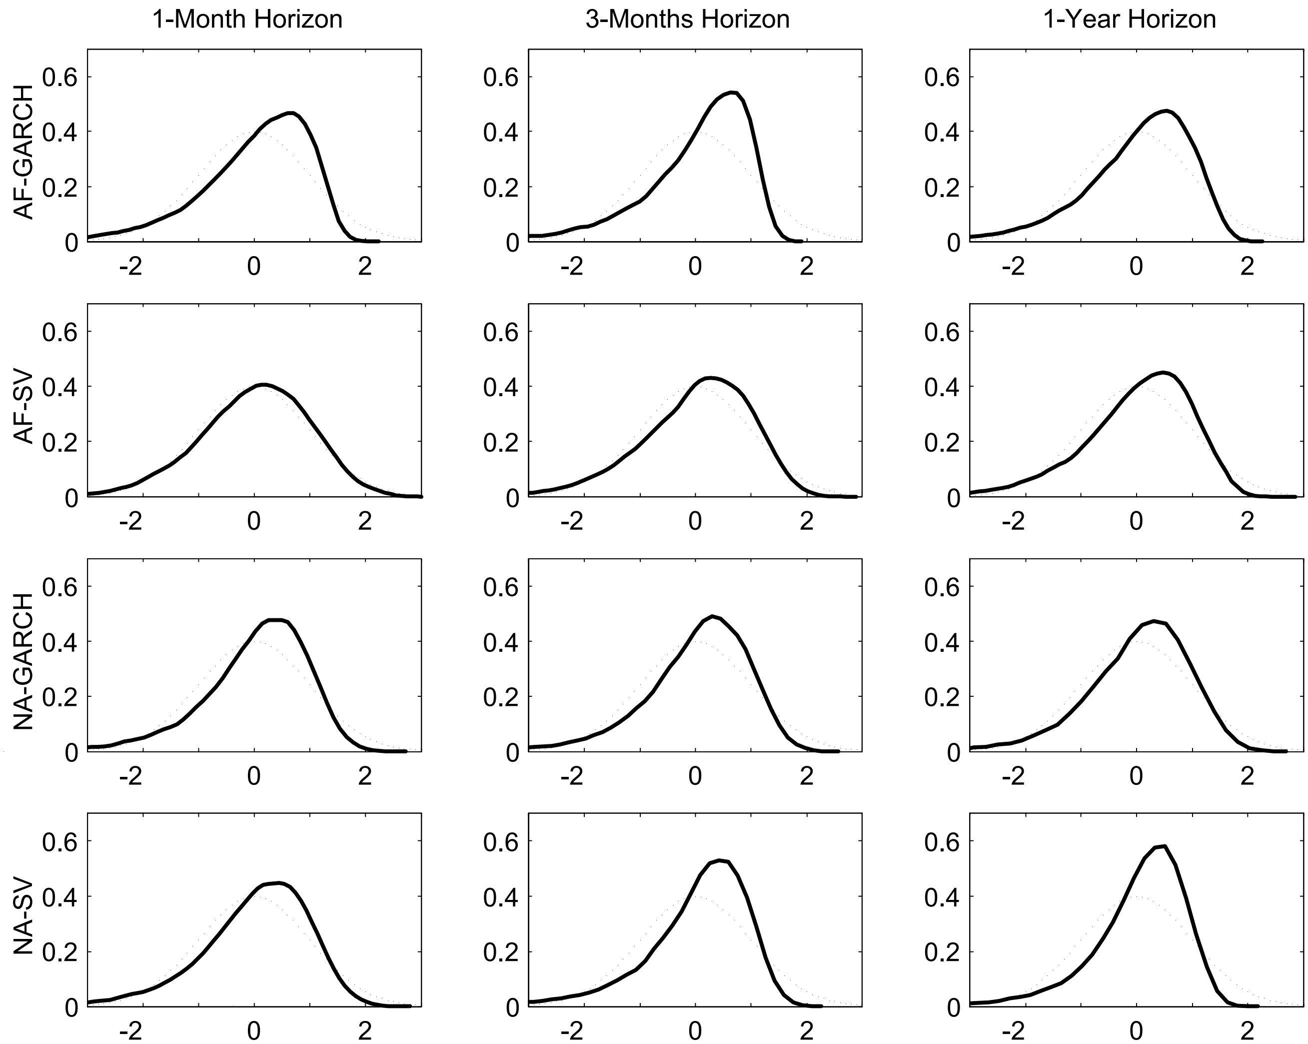

Notes to figure: Each row of panels reports the risk neutral distribution of the index return according to the AF-GARCH, AF-SV, NA-GARCH and NA-SV models respectively. The normal distribution corresponding to the Black-Scholes model is superimposed for reference. The left column reports the 1-month horizon, the center column the 3-month horizon distribution and the right column shows the 1-year distribution. The distributions are constructed by simulating daily returns from each model setting the initial spot variance equal to the unconditional variance. Kernel density estimates are then constructed from the standardized simulated returns. 1995 parameter estimates from Table 3 converted to the risk neutral measure are used to simulate the risk neutral returns.

Table 1: S&P500 Index Call Option Data. 1990-1995. In Sample

Panel A. Number of Call Option Contracts

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	101	1,884	1,931	1,765	5,681
0.975<S/X<1	283	1,272	706	477	2,738
1<S/X<1.025	300	1,212	726	523	2,761
1.025<S/X<1.05	261	1,167	654	406	2,488
1.05<S/X<1.075	245	1,039	582	390	2,256
S/X>1.075	549	2,345	1,679	1,242	5,815
All	1,739	8,919	6,278	4,803	21,739

Panel B. Average Call Price

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.88	2.30	6.25	11.92	6.61
0.975<S/X<1	2.29	6.83	15.19	27.50	12.12
1<S/X<1.025	8.35	13.60	22.48	34.34	19.29
1.025<S/X<1.05	17.57	22.00	30.11	42.03	26.94
1.05<S/X<1.075	27.11	30.84	38.14	48.83	35.43
S/X>1.075	50.67	52.78	58.98	68.30	57.69
All	24.32	23.66	28.68	36.03	27.89

Panel C. Average Implied Volatility from Call Options

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.1625	0.1266	0.1348	0.1393	0.1341
0.975<S/X<1	0.1308	0.1296	0.1448	0.1562	0.1387
1<S/X<1.025	0.1527	0.1459	0.1558	0.1607	0.1522
1.025<S/X<1.05	0.1914	0.1647	0.1665	0.1657	0.1683
1.05<S/X<1.075	0.2429	0.1828	0.1775	0.1739	0.1875
S/X>1.075	0.3878	0.2353	0.1961	0.1868	0.2347
All	0.2639	0.1751	0.1639	0.1618	0.1780

Notes: The sample contains European call options on the S&P500 index. We use quotes within 30 minutes from closing on every Wednesday during the January 1, 1990 to December 31, 1995 period. We apply the moneyness and maturity filters used by Bakshi, Cao and Chen (1997). Implied volatilities are computed using the Black-Scholes formula.

Table 2: S&P500 Index Call Option Data. 1990-1995. Out of Sample

Panel A. Number of Call Option Contracts

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	91	1,797	1,870	1,705	5,463
0.975<S/X<1	274	1,231	723	452	2,680
1<S/X<1.025	279	1,183	679	487	2,628
1.025<S/X<1.05	264	1,093	614	350	2,321
1.05<S/X<1.075	197	941	519	258	1,915
S/X>1.075	357	1,407	1,003	707	3,474
All	1,462	7,652	5,408	3,959	18,481

Panel B. Average Call Price

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.86	2.28	6.18	11.98	6.62
0.975<S/X<1	2.21	6.65	15.09	27.46	11.98
1<S/X<1.025	8.17	13.42	22.45	33.85	18.98
1.025<S/X<1.05	17.07	21.80	29.82	41.64	26.38
1.05<S/X<1.075	26.67	30.34	37.53	47.75	34.26
S/X>1.075	47.60	49.90	56.25	69.09	55.40
All	20.33	19.70	24.39	31.59	23.67

Panel C. Average Implied Volatility from Call Options

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.1636	0.1260	0.1353	0.1400	0.1344
0.975<S/X<1	0.1293	0.1282	0.1444	0.1566	0.1379
1<S/X<1.025	0.1531	0.1442	0.1554	0.1608	0.1513
1.025<S/X<1.05	0.1943	0.1609	0.1651	0.1649	0.1667
1.05<S/X<1.075	0.2477	0.1815	0.1769	0.1763	0.1875
S/X>1.075	0.3903	0.2396	0.2116	0.2074	0.2461
All	0.2481	0.1670	0.1629	0.1629	0.1728

Notes: The sample contains European call options on the S&P500 index. We use quotes within 30 minutes from closing on every Thursday during the January 1, 1990 to December 31, 1995 period. We apply the moneyness and maturity filters used by Bakshi, Cao and Chen (1997). Implied volatilities are computed using the Black-Scholes formula.

Table 3: Parameter Estimates

AF-GARCH

	p	κ	θ	a	c
1990	3.5327	0.0062	5.62E-05	1.98E-07	1959.9
1991	1.3819	0.0028	9.22E-05	2.02E-07	1553.7
1992	15.2797	0.0790	3.79E-05	1.98E-06	599.8
1993	1.1586	0.0399	5.05E-05	1.65E-06	652.1
1994	0.9639	0.0359	5.30E-05	1.43E-06	761.7
1995	4.6551	0.0363	4.19E-05	9.43E-07	822.0

AF-SV

	μ	κ	θ	σ	ρ	λ
1990	6.30E-04	0.0047	3.44E-05	6.31E-04	-0.9990	3.38E-03
1991	4.32E-04	0.0033	8.67E-05	7.27E-04	-0.9900	5.34E-04
1992	5.34E-04	0.0162	6.92E-05	1.05E-03	-0.8629	3.24E-03
1993	6.54E-04	0.0122	4.51E-05	1.01E-03	-0.8301	4.16E-03
1994	6.57E-04	0.0150	3.64E-05	1.05E-03	-0.9359	6.85E-03
1995	6.16E-04	0.0132	4.32E-05	7.46E-04	-0.8616	3.92E-03

NA-GARCH

	p	κ	θ	σ	ρ
1990	0.1727	0.0166	2.81E-05	0.0814	-0.9898
1991	0.0017	0.0002	9.47E-04	0.0717	-0.9934
1992	0.0723	0.0208	5.46E-05	0.2114	-0.8286
1993	0.0467	0.0356	4.74E-05	0.3404	-0.9433
1994	0.1395	0.0341	2.54E-05	0.1982	-0.9176
1995	0.0896	0.0505	3.30E-05	0.3187	-0.9788

NA-SV

	μ	κ	θ	σ	ρ	λ
1990	6.26E-04	0.0052	6.72E-05	0.0783	-0.9990	3.35E-03
1991	2.21E-04	0.0007	3.01E-04	0.0683	-0.9990	1.27E-04
1992	7.39E-04	0.0171	6.49E-05	0.1615	-0.9409	6.80E-03
1993	2.49E-06	0.0223	7.53E-05	0.3102	-0.9990	3.35E-04
1994	5.96E-04	0.0412	2.36E-05	0.2032	-0.9990	3.15E-02
1995	1.27E-03	0.0142	5.16E-05	0.2038	-0.9730	5.51E-03

Notes: For each model, we perform six estimation exercises using Nonlinear Least Squares on the valuation errors. We use Wednesday option prices in each of the years 1990, 1991, 1992, 1993, 1994 and 1995 to conduct separate estimation exercises.

Table 4: Persistence and Annual Volatility

Panel A. Physical Persistence

	AF-GARCH	AF-SV	NA-GARCH	NA-SV
1990	0.9938	0.9953	0.9834	0.9948
1991	0.9972	0.9967	0.9998	0.9993
1992	0.9210	0.9838	0.9792	0.9829
1993	0.9601	0.9878	0.9644	0.9777
1994	0.9641	0.9850	0.9659	0.9588
1995	0.9637	0.9868	0.9495	0.9858
Average	0.9666	0.9892	0.9737	0.9832

Panel B. Risk Neutral Persistence

	AF-GARCH	AF-SV	NA-GARCH	NA-SV
1990	0.9969	0.9987	0.9974	0.9982
1991	0.9984	0.9973	0.9999	0.9995
1992	0.9590	0.9871	0.9906	0.9897
1993	0.9637	0.9919	0.9788	0.9781
1994	0.9673	0.9919	0.9905	0.9903
1995	0.9717	0.9907	0.9773	0.9913
Average	0.9762	0.9929	0.9891	0.9912

Panel C. Unconditional Physical Volatility (Annualized)

	AF-GARCH	AF-SV	NA-GARCH	NA-SV
1990	0.1190	0.0931	0.0842	0.1301
1991	0.1524	0.1478	0.4886	0.2754
1992	0.0978	0.1321	0.1173	0.1278
1993	0.1128	0.1066	0.1093	0.1378
1994	0.1155	0.0958	0.0801	0.0771
1995	0.1027	0.1044	0.0912	0.1140
Average	0.1167	0.1133	0.1618	0.1437

Panel C. Unconditional Risk Neutral Volatility (Annualized)

	AF-GARCH	AF-SV	NA-GARCH	NA-SV
1990	0.1694	0.1761	0.2143	0.2193
1991	0.1998	0.1616	0.7247	0.3070
1992	0.1356	0.1477	0.1748	0.1646
1993	0.1183	0.1313	0.1416	0.1388
1994	0.1211	0.1301	0.1516	0.1592
1995	0.1164	0.1244	0.1359	0.1457
Average	0.1434	0.1452	0.2572	0.1891

Notes: Using the parameter estimates reported in Table 3, we compute unconditional volatility and persistence for each of the four models using the formulas given in the paper. We compute physical as well as risk-neutral estimates for each of the six estimation exercises.

Table 5: RMSE In and Out of Sample

Panel A. RMSE In Sample

	AF-GARCH	AF-SV	NA-GARCH	NA-SV
1990	0.9711	1.0581	0.9667	0.9371
1991	0.8455	0.9034	0.7956	0.7960
1992	0.7758	0.7671	0.7602	0.7897
1993	0.8118	0.9395	0.7015	0.7087
1994	0.9591	1.0339	0.8761	0.7993
1995	0.8728	0.9371	0.7654	0.6933
Overall	0.8768	0.9461	0.8108	0.7814
Normalized	0.9267	1.0000	0.8569	0.8259

Panel B. RMSE Out of Sample

	AF-GARCH	AF-SV	NA-GARCH	NA-SV
1990	0.9784	1.0692	0.9730	0.9420
1991	0.9005	1.1470	0.8370	0.8389
1992	0.7806	0.7589	0.7370	0.7478
1993	0.6925	0.8275	0.5728	0.5818
1994	0.9577	1.0451	0.7686	0.6805
1995	0.9449	0.8971	0.8214	0.7735
Overall	0.8837	0.9621	0.7966	0.7718
Normalized	0.9185	1.0000	0.8280	0.8022

Notes: Using the parameter estimates from Table 3, we compute dollar RMSEs for all fou models, for each of the six in-sample and out-of-sample periods. RMSE refers to the square root of the mean-squared valuation errors.

Table 6: RMSE by Moneyness and Maturity. 1990-1995. In Sample

Panel A. AF-GARCH

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4760	0.7316	0.7652	1.0177	0.8380
0.975<S/X<1	0.8628	0.9601	0.8416	1.0448	0.9370
1<S/X<1.025	0.7844	0.9141	0.8056	0.9276	0.8760
1.025<S/X<1.05	0.5865	0.8202	0.7347	0.9607	0.8026
1.05<S/X<1.075	0.7884	0.7565	0.7618	1.1451	0.8411
S/X>1.075	0.8022	0.8393	0.8197	1.2265	0.9274
All	0.7634	0.8354	0.7903	1.0753	0.8768

Panel B. AF-SV

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4093	0.7753	0.8612	1.1353	0.9246
0.975<S/X<1	0.6455	0.9611	0.9314	1.0355	0.9395
1<S/X<1.025	0.5987	0.8760	0.9059	0.9657	0.8765
1.025<S/X<1.05	0.5835	0.8431	0.8756	1.0075	0.8581
1.05<S/X<1.075	0.7787	0.8597	0.9674	1.1373	0.9334
S/X>1.075	0.7740	0.9701	1.0266	1.2573	1.0383
All	0.6819	0.8886	0.9322	1.1319	0.9461

Panel C. NA-GARCH

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4581	0.6600	0.7317	0.9872	0.7956
0.975<S/X<1	0.6150	0.7800	0.7657	1.0183	0.8084
1<S/X<1.025	0.6508	0.7603	0.7831	0.9111	0.7863
1.025<S/X<1.05	0.6031	0.7306	0.7490	0.9153	0.7568
1.05<S/X<1.075	0.6001	0.7303	0.8082	1.0359	0.7992
S/X>1.075	0.7439	0.7693	0.8070	1.1160	0.8631
All	0.6534	0.7382	0.7712	1.0156	0.8108

Panel D. NA-SV

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4399	0.6265	0.7098	0.8984	0.7454
0.975<S/X<1	0.5607	0.7519	0.7476	0.9407	0.7703
1<S/X<1.025	0.6001	0.7488	0.7597	0.8515	0.7578
1.025<S/X<1.05	0.5754	0.7238	0.7249	0.8507	0.7326
1.05<S/X<1.075	0.5980	0.7304	0.7898	0.9542	0.7764
S/X>1.075	0.7442	0.7762	0.7997	1.0693	0.8510
All	0.6320	0.7272	0.7538	0.9456	0.7814

Notes: We use the NLS estimates from Table 3 to compute the dollar root mean squared option valuation error (RMSE) for various moneyness and maturity bins for the four models. The option prices used in the table are for the 1990-1995 in-sample period, which consists of Wednesday option prices.

Table 7: RMSE by Moneyness and Maturity. 1990-1995. Out of Sample

Panel A. AF-GARCH

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4946	0.7190	0.8009	1.0651	0.8647
0.975<S/X<1	0.8428	1.0015	0.9155	1.0993	0.9815
1<S/X<1.025	0.8022	0.9587	0.8885	0.9792	0.9294
1.025<S/X<1.05	0.5880	0.8434	0.7947	0.9241	0.8187
1.05<S/X<1.075	0.7508	0.7510	0.7959	1.0395	0.8077
S/X>1.075	0.5763	0.8319	0.9121	1.0392	0.8800
All	0.7005	0.8504	0.8484	1.0408	0.8837

Panel B. AF-SV

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4269	0.7645	0.8972	1.1529	0.9411
0.975<S/X<1	0.6152	1.0150	0.9954	1.0824	0.9883
1<S/X<1.025	0.5892	0.8896	0.9556	0.9684	0.8960
1.025<S/X<1.05	0.5579	0.8131	0.8740	0.9072	0.8203
1.05<S/X<1.075	0.7584	0.8172	0.9334	1.0908	0.8853
S/X>1.075	0.5743	1.3139	1.0634	1.0605	1.1351
All	0.6023	0.9588	0.9516	1.0829	0.9621

Panel C. NA-GARCH

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4398	0.6308	0.7394	1.0048	0.7977
0.975<S/X<1	0.5818	0.7935	0.8196	1.0346	0.8280
1<S/X<1.025	0.6390	0.7495	0.8198	0.9083	0.7891
1.025<S/X<1.05	0.5843	0.7086	0.7852	0.8298	0.7363
1.05<S/X<1.075	0.5547	0.7269	0.8008	1.0061	0.7755
S/X>1.075	0.5509	0.7774	0.8953	0.9248	0.8255
All	0.5750	0.7276	0.8022	0.9685	0.7966

Panel D. NA-SV

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.4199	0.6053	0.7251	0.9164	0.7520
0.975<S/X<1	0.5537	0.7754	0.8003	0.9609	0.7975
1<S/X<1.025	0.5927	0.7526	0.8082	0.8705	0.7755
1.025<S/X<1.05	0.5662	0.7010	0.7739	0.7740	0.7188
1.05<S/X<1.075	0.5500	0.7234	0.7990	0.9410	0.7619
S/X>1.075	0.5499	0.7835	0.8876	0.8909	0.8177
All	0.5549	0.7195	0.7905	0.9015	0.7718

Notes: We use the NLS estimates from Table 3 to compute dollar root mean squared option valuation errors (RMSEs) for various moneyness and maturity bins for the four models. The option prices used are for the 1990-1995 out-of-sample period, which consists of Thursday option prices.