

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

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

January 13, 2006

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 25-27% lower than that of the benchmark continuous-time affine stochastic volatility model in and 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. We provide some intuition for these findings. 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 straightforward to implement and it can be used in a variety of applications and on various loss functions.

JEL Classification: G12

Keywords: Stochastic volatility; GARCH; option valuation; filtering; particle filter.

*Christoffersen and Jacobs are also affiliated with CIRANO and CIREQ and want to thank FQRSC, IFM² and SSHRC for financial support. Mimouni was supported by IFM². We thank seminar participants at Case Western Reserve, Northwestern, Rice University, SBFSIF, and University of Calgary, and in particular Torben Andersen, Alex David, Vadim di Pietro, Jin Duan, Jeff Fleming, Anurag Gupta, Michael Johannes, Nour Meddahi, Barbara Ostdiek, Brad Paye, Marcel Rindisbacher, Peter Ritchken, and Neal Stoughton 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 small 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 processes for the underlying securities.⁵

¹The 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 (2002), Chernov and Ghysels (2000), Eraker (2004), Heston and Nandi (2000), Nandi (1998) and Pan (2002).

²For 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).

³Jones (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.

⁴The 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.

⁵See Amin and Ng (1993), Bollen and Rasiel (2003), Bollerslev and Mikkelsen (1996), Christoffersen and Jacobs

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 three in-sample exercises and three out-of-sample exercises, and despite the fact that both the AF-SV and the NA-SV models show signs of misspecification, 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 root mean-squared error (RMSE) of the NA-SV we investigate is approximately 25.5% lower than that of the AF-SV model in-sample, and the out-of-sample dollar RMSE is 27.5% 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.⁸ 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 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

(2004a), Duan, Ritchken and Sun (2002), Engle and Mustafa (1992), and Heston and Nandi (2000),

⁶Likelihood-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.

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

⁸See for example Duan (1997), Heston and Nandi (2000), Nelson (1990), and Ritchken and Trevor (1999).

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 but slightly worse than 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 somewhat surprisingly underperforms relative to 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 algorithm⁹ 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 Model Specification

We now turn to a description of the four volatility models we investigate empirically below. We first discuss the Heston (1993) model, then 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

⁹See Pitt and Shephard (1999).

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.¹⁰ 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.¹¹ 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,¹² while discrete-time volatility modeling using only returns data has spawned a large number 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, which instead

¹⁰Affine 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).

¹¹There 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.

¹²Duan (1995) and Amin and Ng (1993) provide theoretical foundations for this literature. Bollerslev and Mikkelsen (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.

have to be computed using Monte Carlo 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-form valuation results in the options literature comes at the cost of a deterioration in the model’s empirical performance.

We now give the specifications of each of the four models and contrast them with recent model-free empirical findings in the realized volatility literature.¹³

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 for the stock price, S , and the variance, V ,

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

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

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.¹⁴ Note also that the AF-SV model implies that the instantaneous change in variance is heteroskedastic via the \sqrt{V} term in the volatility diffusion.

In order to explore the AF-SV model further consider the instantaneous volatility dynamic implied by the model. Using Ito’s lemma, we can write

$$d\sqrt{V} = \mu(V)dt + \frac{1}{2}\sigma dw^V \quad (3)$$

Note that the the AF-SV model implies that the instantaneous change in volatility should be Gaussian and homoskedastic: V does not show up in any way in the diffusion term for $d\sqrt{V}$. This implication is strong and can be assessed empirically quite easily.

Using daily realized volatilities from 1990 to 2002 the top panel of Figure 1 shows a quantile-quantile (QQ) plot of the daily realized volatility changes compared with the Gaussian distribution.¹⁵ The deviations of the data points from the straight line indicates that the Gaussian distribution is not a good assumption for daily changes in volatility. The observed tails (both left and right) are considerably fatter than the normal distribution would suggest. The middle panel in Figure 1 scatter plots the daily volatility changes against the daily volatility level. According to the AF-SV model this scatter plot should reveal no systematic patterns. However, as the volatility level increases on the horizontal axis a cone-shaped pattern in the daily volatility changes on the vertical

¹³See for example Andersen, Bollerslev, Diebold and Labys (2003), and Ait-Sahalia, Mykland and Zhang (2005).

¹⁴Note 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.

¹⁵The realized volatility data was graciously provided to us by and is documented in Andersen, Bollerslev and Diebold (2005).

axis is apparent. Confirming this pattern the bottom panel of Figure 1 scatter plots the absolute daily volatility changes against the daily volatility level. A simple OLS regression line is shown for reference. Notice a clearly positive relationship between the volatility level and the magnitude of volatility changes. This pattern is in conflict with the homoskedastic volatility implication of the AF-SV model.

In the implementation below we will rely on the dynamics for log variance which using Ito's lemma can be written as

$$d\ln(V) = \mu(V)dt + \sigma \frac{1}{\sqrt{V}}dw^V \quad (4)$$

Notice that the log variance dynamics also imply heteroskedastic instantaneous changes and this time with a factor of $\frac{1}{\sqrt{V}}$.

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

When departing from the affine model framework many directions can be taken. We focus on a process which has the same number of parameters as the affine SV model above and which has appealing empirical implications as we shall see shortly. We will assume the following NA-SV dynamic

$$dS = \mu Sdt + \sqrt{V}Sdw^S \quad (5)$$

$$dV = \kappa(\theta - V)dt + \sigma Vdw^V \quad (6)$$

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

To explore this model further, consider now the implications of this model for instantaneous volatility dynamics. We can write

$$d\sqrt{V} = \mu(V)dt + \frac{1}{2}\sigma\sqrt{V}dw^V \quad (7)$$

so that volatility changes are heteroskedastic as Figure 1 suggests they should be.

Note next that the NA-SV model implies that the instantaneous changes in log variance are homoskedastic. Using Ito's lemma we can write

$$d\ln(V) = \mu(V)dt + \sigma dw^V$$

Figure 2 summarizes the empirics of realized log variances. The top panel shows that daily changes in the realized log variances follow the Gaussian distribution quite closely. This feature is implied by the NA-SV model and it clearly differs from the finding for the daily volatility changes in Figure 1. Furthermore, the middle panel of Figure 2 shows the scatter of daily log variance changes against the log variance level. Note that no cone-shaped pattern is apparent. The bottom panel in Figure 2 confirms this first impression and shows a virtually flat line when regressing absolute changes in log variance on log variance levels. Note that the absence of a clear relationship between

changes in the log variances and the log variance level in Figure 2 lends credence to the NA-SV specification but it also casts further doubt on the AF-SV specification where the instantaneous changes in log variance are heteroskedastic as shown in (4) above.

2.3 The Affine GARCH 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

$$\ln(S_{t+1}) = \ln(S_t) + r + pV_t + \sqrt{V_t}z_{t+1} \quad (8)$$

$$V_t = w + bV_{t-1} + a \left(z_t - c\sqrt{V_{t-1}} \right)^2 \quad (9)$$

where S_{t+1} denotes the underlying asset price, r the risk free rate, p the price of risk and V_t the daily variance for day $t + 1$ which is known at the end of day t .¹⁶ 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.

In order to explore the model further, notice that variance persistence can be computed via

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

and the unconditional variance can be computed via

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

Now we can rewrite the variance process as

$$V_t - V_{t-1} = \kappa(\theta - V_{t-1}) + av_t \quad (10)$$

where $v_t = ((z_t^2 - 1) - 2cz_t\sqrt{V_{t-1}})$, and where we have

$$StdDev_{t-1}(v_t) = \sqrt{2 + 4c^2V_{t-1}}$$

The fact that the variance appears in a square-root form here suggests the AF-GARCH model's relationship with the AF-SV model considered above.

¹⁶The timing convention we use here is slightly non-standard in the GARCH literature but it facilitates comparisons with the SV models.

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.¹⁷ We will refer to this model as NA-GARCH. The model is given by

$$\ln(S_{t+1}) = \ln(S_t) + r + p\sqrt{V_t} - 0.5V_t + \sqrt{V_t}z_{t+1} \quad (11)$$

$$V_t = w + bV_{t-1} + aV_{t-1}(z_t - c)^2 \quad (12)$$

As in the AF-GARCH model, we can define the variance persistence via

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

and the unconditional variance can be computed via

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

Now we can rewrite the variance process as

$$V_t - V_{t-1} = \kappa(\theta - V_{t-1}) + aV_{t-1}((z_t^2 - 1) - 2cz_t) \quad (13)$$

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

$$V_t - V_{t-1} = \kappa(\theta - V_{t-1}) + \sigma V_{t-1}v_t$$

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

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

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 (8)-(9). The Heston-Nandi model was engineered with the specific purpose of yielding closed-form option prices. The specification in (11)-(12) 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 (8)-(9) reduce the ability of the model to fit the options data.

3 Volatility Model Implementation

The Heston AF-SV 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

¹⁷See for example Duan (1997).

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 SV models 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. For all the four models our implementation uses the nonlinear least squares (NLS) estimation techniques to minimize

$$MSE = \frac{1}{NT} \sum_{t,i} (C_{i,t} - C_i(V_t))^2 \quad (14)$$

with respect to the structural parameters. $N^T = \sum_{t=1}^T N_t$, T denotes 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 $C_i(V_t)$ 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), 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 implement the SV models in a way that is consistent with these insights.

3.1 The AF-SV Model

The problem with (14) is that V_t is a latent factor in the two SV models, and that it has to be filtered from observed return 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 for the two SV models.¹⁸ As shown by Pitt and Shephard (1999) the APF offers a convenient

¹⁸We have also implemented the Sampling-Importance-Resampling (*SIR*) particle filter as a robustness check, and this yields similar results.

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.¹⁹

3.1.1 Volatility Transformation and Discretization

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

$$\begin{aligned} d\ln(S) &= \left(\mu - \frac{1}{2}V\right) dt + \sqrt{V}dw^S \\ d\ln(V) &= \frac{1}{V} \left(\kappa(\theta - V) - \frac{1}{2}\sigma^2\right) dt + \sigma \frac{1}{\sqrt{V}}dw^V \end{aligned} \quad (15)$$

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 (15). 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.²⁰ Discretizing (15) gives

$$\ln(S_{t+1}) = \ln(S_t) + \left(\mu - \frac{1}{2}V_t\right) + \sqrt{V_t}\varepsilon_{t+1}^S \quad (16)$$

$$\ln(V_{t+1}) = \ln(V_t) + \frac{1}{V_t} \left(\kappa(\theta - V_t) - \frac{1}{2}\sigma^2\right) + \sigma \frac{1}{\sqrt{V_t}}\varepsilon_{t+1}^V \quad (17)$$

We implement the discretized model in (16) and (17) 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 (14) using an iterative procedure on the structural parameters. At each iteration, the volatility is filtered using the information embedded in observed returns and the structural parameters. Using the filtered volatility and the structural parameters option prices are computed according to Heston's formula and the *MSE* is calculated. This procedure searches in the structural parameter space until the optimum is reached.

We now describe the volatility filtering step in more detail.

¹⁹Johannes, Polson and Stroud (2002) discuss the use of the particle filter to estimate parameters for continuous-time jump-diffusion models on returns data.

²⁰See for example Eraker (2001).

3.1.2 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 (17) into N possible particles.²¹ 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 $\{V_t^j, W_t^j\}_{j=1}^N$ with V_t^j the volatility at day t for particle j , W_t^j the weight associated with particle j at date t and $j = 1, \dots, N$.²² We want to propagate V_t^j one day ahead into $\{V_{t+1}^j, W_{t+1}^j\}_{j=1}^N$ using the *APF*. This task requires the following steps:

Step 1: Selecting the particles

Because of the Euler discretization in (16) and (17) the weight W_t^j will reflect information available at time $t + 1$. Furthermore, even if particle j is very likely according to the realization of the stock price at date $t + 1$, the conditional density of S_{t+2} may suggest a certain re-adjustment of the probability that state j has occurred. By combining the information available at $t + 1$ and the conditional density of S_{t+2} , we can eliminate many particles with a low probability before the propagation step. This is achieved by:

I) Computing a summary location statistic for the unobserved V_{t+1} that reflects the information at time 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

$$\iota^j \propto W_t^j p\left(\ln(S_{t+2})|\mu_{t+1}^j\right)$$

where $p\left(\ln(S_{t+2})|\mu_{t+1}^j\right)$ is the conditional density of $\ln(S_{t+2})$ which can be easily inferred from (16). This auxiliary variable tells us in which proportion to resample the various particles in the original set.²³

After this resampling step 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, and j will refer to the resampled particles from now on.

Step 2: Simulating the state forward (Sampling)

²¹We 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.

²²At time 0, the initial set is constructed by setting each particle equal to the unconditional variance θ and giving all particles equal weight, $1/N$.

²³See Bolic (2004) for details on the resampling step.

This is done by computing V_{t+1}^j for the resampled particles using equation (17) and taking the correlation into account. We have

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

which gives

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

Since

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

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

$$\begin{aligned} \ln(V_{t+1}^j) = & \ln(V_t^j) + \frac{1}{V_t^j} \left(\kappa \left(\theta - V_t^j \right) - \frac{1}{2}\sigma^2 \right) + \sigma \frac{1}{\sqrt{V_t^j}} \left(\rho \frac{\ln\left(\frac{S_{t+1}}{S_t}\right) - \left(\mu - \frac{1}{2}V_t^j\right)}{\sqrt{V_t^j}} + \sqrt{1 - \rho^2}\varepsilon_{t+1}^j \right) \end{aligned}$$

We simulate 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 (16) that given the other available information, V_{t+1} is sufficient to generate $\ln(S_{t+2})$. Therefore, equation (16) 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 has generated S_{t+2}). The likelihood is computed as follows:

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

This is to be repeated for $j = 1, \dots, N$. Finally, because nothing guarantees that $\sum_{j=1}^N W_{t+1}^j = 1$, we have to normalize and set $\bar{W}_{t+1}^j = \frac{W_{t+1}^j}{\sum_{j=1}^N W_{t+1}^j}$. 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 $t = 1, \dots, T$. To obtain the filtered volatility path, we then compute

$$\bar{V}_{t+1} = \sum_{j=1}^N \bar{W}_{t+1}^j V_{t+1}^j \tag{18}$$

for each t .

3.1.3 Computing option prices and evaluating the loss function

In order to price options with the AF-SV model we need the risk neutral dynamics. 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 and λ is

$$dS = rSdt + \sqrt{V}Sdw^{*S} \quad (19)$$

$$dV = (\kappa(\theta - V) - \lambda V)dt + \sigma\sqrt{V}dw^{*V} \quad (20)$$

$$= (\kappa + \lambda)(\kappa\theta/(\kappa + \lambda) - V)dt + \sigma\sqrt{V}dw^{*V} \quad (21)$$

with $\text{corr}(dw^{*S}, dw^{*V}) = \rho$. Heston (1993) demonstrates that this model admits a closed form solution for option prices, which can be written as

$$C(V_t) = SP_1 - Ke^{-r(T-t)}P_2 \quad (22)$$

where the equations for P_1 and P_2 are given in the appendix.

We are now in a position to evaluate option prices $C_i(\bar{V}_t)$ based on the filtered volatility path using Heston's closed form solution according to equation (22). We subsequently evaluate the loss function (14), using $C_i(\bar{V}_t)$ for the model price $C_i(V_t)$. We use a standard numerical optimization routine to update the model parameters and iterate until convergence is achieved.²⁴

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.

3.2 The NA-SV Model

In order to compute option prices under the NA-SV dynamics we assume that the volatility risk premium $\lambda(S, V, t)$ is equal to λV , so that the risk neutral dynamic expressed in terms of the physical parameters and λ is

$$dS = rSdt + \sqrt{V}Sdw^{*S} \quad (23)$$

$$dV = (\kappa(\theta - V) - \lambda V)dt + \sigma Vdw^{*V} \quad (24)$$

$$= (\kappa + \lambda)(\kappa\theta/(\kappa + \lambda) - V)dt + \sigma Vdw^{*V} \quad (25)$$

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

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. In order to

²⁴In 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.

isolate the importance of the volatility dynamic, we keep the volatility risk premium specifications constant across the affine and non-affine SV models.²⁵ While clearly empirically appealing the NA-SV model presented here has no known closed form option valuation formula and so we will be computing option prices below using Monte Carlo simulation.

The NA-SV model 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 (1998), stratified random numbers, antithetic variates and a control variate technique. The model is estimated by minimizing (14) with respect to the structural parameters $\mu, \kappa, \theta, \rho, \sigma$, and λ .

3.3 The AF-GARCH Model

The risk-neutral dynamics for the GARCH(1,1) model (8)-(9) are given by²⁶

$$\begin{aligned}\ln(S_{t+1}) &= \ln(S_t) + r - \frac{1}{2}V_t + \sqrt{V_t}z_{t+1}^* \\ V_t &= w + bV_{t-1} + a(z_t^* - (c + p + 0.5)\sqrt{V_{t-1}})^2\end{aligned}\tag{26}$$

with $z_t^* \sim N(0, 1)$ under the risk neutral measure. At time t , a European call option with strike price K that expires at time T can be calculated from

$$C(V_t) = S_t P_1 - K e^{-r(T-t)} P_2$$

where the formulas for P_1 and P_2 are provided in the appendix.

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 V_t on that particular date. One of the appealing aspects of discrete-time GARCH models is that this filtering problem is extremely simple. Indeed, the filtering problem is solved by noting that from (8) we have

$$z_{t+1} = (R_{t+1} - r - pV_t) / \sqrt{V_t}\tag{27}$$

where $R_t = \ln(S_t/S_{t-1})$. Substituting (27) in (9), it can be seen that the updating from V_{t-1} to V_t is done exclusively using observables

$$V_t = w + bV_{t-1} + a((R_t - r) / \sqrt{V_{t-1}} - (c + p)\sqrt{V_{t-1}})^2\tag{28}$$

Model parameters are again obtained by using the nonlinear least squares (NLS) estimation techniques to minimize (14).

²⁵See Lewis (2000) for a thorough discussion of these issues.

²⁶For 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).

The implementation is therefore relatively simple: the NLS routine is called with a set of parameter starting values. The variance dynamic in (28) 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.

3.4 The NA-GARCH Model

The risk-neutral dynamics required for option valuation with the NA-GARCH model (11)-(12) can be obtained as

$$\begin{aligned}\ln(S_{t+1}) &= \ln(S_t) + r - 0.5V_t + \sqrt{V_t}z_{t+1}^* \\ V_t &= w + bV_{t-1} + aV_{t-1}(z_t^* - (c + p))^2\end{aligned}\tag{29}$$

with $z_t^* \sim N(0, 1)$. As in the case of the NA-SV model, no closed-form solution exists for option valuation in the NA-GARCH model instead Monte Carlo simulation is required.

We can then estimate the model by minimizing (14), using the updating rule

$$V_t = w + bV_{t-1} + aV_{t-1} \left(\left[(R_t - r + 0.5V_{t-1}) / \sqrt{V_{t-1}} \right] - (c + p) \right)^2\tag{30}$$

Option prices are computed numerically according to

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

where the expectation is calculated by Monte Carlo simulation of the daily returns from (29). 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 (1998), stratified random numbers, antithetic variates and a control variate technique.

4 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.

4.1 Data

We conduct our empirical analysis using three years of data on S&P 500 call options, for the period 1993-1995. 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.²⁷

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

Table 1 presents descriptive statistics for the options data for the 1993-1995 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 three 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 3 gives some indication of the pattern of implied volatility over time. For the 156 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 3 that there is substantial clustering in implied volatilities. The bottom panel of Figure 3 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 three in-sample estimations, we proceed to conducting separate out-of-sample analyses for each of the three 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 Thursday data. The patterns in the data are clearly similar to those in the in-sample data in Table 1.

Note that we are deliberately working with a relatively tranquil time-period for option prices. We are estimating four models that are simple in terms of their volatility memory properties (single factor models) and simple in terms of shock innovations (Gaussian), but which vary in their specifications of how shocks to volatility change the level of volatility. Increasing the sample period length to including more interesting episodes such as the 1990-91 recession / first Gulf War or the 1998 LTCM / Russian default / liquidity crisis would be useful but would require more sophisticated models potentially including multiple volatility components and non-Gaussian innovations. We leave such analysis for future work.

²⁷See for instance Heston and Nandi (2000).

4.2 Parameter Estimates and Option Mean-Squared-Errors

Table 3 presents the parameter estimates for each of the four models and for each of the three annual estimation samples. The parameters for the SV models are directly interpretable individually; for example, κ denotes the daily variance mean reversion and θ denotes the unconditional variance which we report in annual standard deviation terms. The most striking aspect of the estimates is that the correlation parameter ρ hits the prespecified boundary of -0.999 for the NA-SV model in 1993 and 1994 and gets very close to it in 1995. The correlation in the AF-SV model is also large in magnitude in all the years.

These estimates, which are driven by the objective to minimize the option price squared errors, indicate that the options want correlations which are much larger in magnitude than the estimates typically found in papers estimating the models on returns only. Below we will estimate the four models on returns only which yields correlation estimates that are indeed very close to those reported in the literature. We readily acknowledge that the parameter estimates indicate potential misspecification of the model and that further work on the models are needed. The objective of this paper is not to find the best possible affine or non-affine model. We merely want to demonstrate that a simple (and possibly misspecified) non-affine model improves significantly on the benchmark affine model.

For the AF-GARCH and NA-GARCH models, we report the parameters from the specifications in (10) and (13) in order to facilitate comparison 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 as we will discuss further below.

To further facilitate the comparison between the different models, the last two columns in Table 3 report the risk neutral variance mean-reversion and unconditional volatility. The variance mean-reversion is close to zero for all models, which is consistent with other findings in the literature. Due to the negative price of volatility risk in the SV models and the positive price of equity risk in the GARCH models, the risk neutral mean-reversion is lower than the physical mean-reversion in all models. For the same reason, the unconditional volatility is larger under the risk-neutral measures. Mean-reversion is always smallest in the AF-SV specification. The physical unconditional volatility displays some variation over time, but the risk-neutral unconditional volatility is actually quite stable.

Figure 4 provides further perspective on the similarities and differences between the four models by reporting volatility sample paths for the models. The figure plots volatility paths for 1993-1995 from the four models using the three sets of parameters estimates from 1993, 1994, and 1995 respectively. It can be seen that the sample paths are quite similar both across models and across estimates. However, it does appear that volatility itself is less volatile in the AF-SV model than in the other three models. We will investigate this issue further below.

The in- and out-of-sample RMSEs from the four models are reported in Table 4 for each of

the three 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 25.5% in-sample and 27.5% 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 hold 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 10.2% better than the AF-SV in sample and approximately 7.6% better out of sample. Looking across the three samples, a robust conclusion obtains: the non-affine models substantially outperform the affine models in every year, both in- and out-of-sample.

Table 5 tests the significance of the difference in weekly RMSE. In order to circumvent having to specify the variance matrix of the options errors which are likely to have strong maturity, moneyness and dynamic dependencies we simply test the difference in the average weekly RMSE across models. When doing so we use the Diebold and Mariano (1995) test allowing for autocorrelation in the weekly RMSEs. The Diebold-Mariano test has an asymptotic normal distribution and the results in Table 5 show that when the AF-SV model is taken as a benchmark, the two non-affine models are significantly better in and out of sample, and the AF-GARCH model is significantly better in sample. When the NA-SV model is used as a benchmark the two affine models are significantly worse, but the NA-GARCH model is not significantly different from the NA-SV. We thus conclude that the differences between affine and non-affine models are significant but the difference between the NA-GARCH and NA-SV models is not significant.

We now analyze the dynamic performance of the four models in more detail. Figures 5 and 6 address the performance of the models over time. The four panels in Figure 5 present the RMSE on a week-by-week basis for each of the four models. It can clearly be seen that the four models display important similarities in terms of the dynamic pricing errors, however, the AF-SV model appears to be somewhat different from the other three models. Notice in particular the larger weekly RMSE in the first half of 1994 which corresponds to the period when the VIX in Figure 3 was high. This observation is confirmed by inspecting the week-by-week bias in Figure 6 where the AF-SV model underprices options around the peak of the VIX. Notice also how in 1995 the AF-SV model consistently overprices the options in almost every week of the year. This is less true in the other three models.

4.3 Pricing Errors Across Moneyness and Maturity

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 moneyness 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

virtually every single case and the NA-GARCH model performs second best also in virtually 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 8-10% 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 actually performs better. While perhaps somewhat surprising, this finding could be due to the fact that our MSE objective function puts relatively little weight on the inexpensive short-term options.

Tables 8 and 9 report the in- and out-of-sample bias across moneyness and maturity. In these tables we are looking for systematic over or underpricing of options with particular moneyness and maturity characteristics. Consider for example the AF-SV model in Panel A of Table 8. It has an overall bias of 24 cents but more importantly it systematically overprices out-of-the-money calls and underprices in-the-money calls. Such systematic biases are much less apparent for the other three models in Table 8. Table 9 shows the same pattern as Table 8 for the AF-SV model and now also for the NA-GARCH model.

Figures 7.A-C provide additional intuition for the differences in the cross-sectional performance between the four models. Using each of the three sets of estimates 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 standardized index return according to the AF-SV, NA-SV, AF-GARCH and NA-GARCH models respectively. The normal distribution corresponding to the Black-Scholes model is shown in dots 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 large, 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. It is also consistent with the biases found in Table 8 and 9: Only the NA-SV and the AF-GARCH models provide enough skewness in the state-price density to avoid systematically overpricing the out-of-the-money calls and underpricing the in-the-money calls.

4.4 Conditional Moment Dynamics

While the state price densities shown in Figures 7.A-C illustrate the option pricing models when fixing the conditional volatility at its unconditional level, ultimately, the option prices from the different models are determined by the dynamics of the conditional density. Therefore, we now discuss model differences by focusing on various conditional moments of the density.

In order to assess the different models' ability to generate time-variation in the asymmetry of the return distribution, Figure 8 plots the conditional covariance between returns and variances for each model. We refer to this conditional covariance as the conditional leverage path, which for the four models is given by

$$\begin{aligned}
 \text{AF-SV} & : \text{cov}_t(R_{t+1}, V_{t+1}) = \sigma\rho V_t \\
 \text{NA-SV} & : \text{cov}_t(R_{t+1}, V_{t+1}) = \sigma\rho V_t^{3/2} \\
 \text{AF-GARCH} & : \text{cov}_t(R_{t+1}, V_{t+1}) = -2acV_t \\
 \text{NA-GARCH} & : \text{cov}_t(R_{t+1}, V_{t+1}) = \sigma\rho V_t^{3/2}
 \end{aligned} \tag{31}$$

Notice that critical differences between the affine and non-affine models show up very prominently in these conditional moments where the power on the variance differs between the affine and the non-affine models.

For each model we plot in Figure 8 the daily conditional leverage path during 1993-1995, annualized by multiplying by 252. The left column uses the 1993 estimates from Table 3, the middle column uses 1994 estimates, and the right column uses the 1995 estimates. The four rows of panels correspond to the AF-SV, NA-SV, AF-GARCH and NA-GARCH models respectively.

The main conclusion is that for all three sets of estimates the non-affine models imply more variation over time in the leverage effect. Notice that particularly the AF-SV model is very different from the other three models. 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 Tables 4 and 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 is given by

$$\begin{aligned}
 \text{AF-SV} & : \text{Var}_t(V_{t+1}) = \sigma^2 V_t \\
 \text{NA-SV} & : \text{Var}_t(V_{t+1}) = \sigma^2 V_t^2 \\
 \text{AF-GARCH} & : \text{Var}_t(V_{t+1}) = 2a^2 + 4a^2 c^2 V_t \\
 \text{NA-GARCH} & : \text{Var}_t(V_{t+1}) = \sigma^2 V_t^2
 \end{aligned} \tag{32}$$

Notice again that the power on V_t in 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.²⁸ Figure 9 reports the empirical results for the volatility of variance. For each model, we plot the daily conditional volatility of variance path during 1993-1995 annualized by multiplying by 252. The left column uses the 1993 estimates from Table 3, the middle column uses 1994 estimates, and the right column uses the 1995 estimates. The four rows of panels correspond to the AF-SV, NA-SV, AF-GARCH and NA-GARCH models respectively. Figure 9 indicates that for all sets of estimates the non-affine models display much more time-variation in the volatility of variance. This is particularly true versus the AF-SV model and it is particularly true during the first half of 1994 when the level of volatility peaks (see Figure 3). These differences between the models further help us understand the superior fit of the non-affine models.

4.5 Comparing the SV and GARCH Results

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 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 dimension such as the dollar RMSE. Our results are very clear: while the affine SV model is analytically convenient, it substantially underperforms the non-affine SV 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 special case, and the restricted Heston (1993) limit may not be the limit that is most relevant from an empirical perspective.

²⁸Notice 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.

In order to explore this issue further, consider the conditional correlation between returns and volatility for the four models, which can be computed from (31) and (32)

$$\begin{aligned}
\text{AF-SV} & : \text{Corr}_t(R_{t+1}, V_{t+1}) = \frac{\sigma \rho V_t}{\sqrt{V_t \sigma^2 V_t}} = \rho \\
\text{NA-SV} & : \text{Corr}_t(R_{t+1}, V_{t+1}) = \frac{\sigma \rho V_t^{3/2}}{\sqrt{V_t \sigma^2 V_t^2}} = \rho \\
\text{AF-GARCH} & : \text{Corr}_t(R_{t+1}, V_{t+1}) = \frac{-2cV_t}{\sqrt{(2 + 4c^2 V_t)V_t}} \\
\text{NA-GARCH} & : \text{Corr}_t(R_{t+1}, V_{t+1}) = \frac{\sigma \rho V_t^{3/2}}{\sqrt{V_t \sigma^2 V_t^2}} = \rho
\end{aligned} \tag{33}$$

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. Figure 10 plot the time-varying correlation in the AF-GARCH during the period 1993-1995 for the three sets of estimates obtained in Table 3. The horizontal line in each panel shows the estimated constant correlation in the AF-SV model for the relevant year. Notice that the time-varying correlation in the AF-GARCH model typically varies substantially from the constant correlation in the AF-SV model. This added flexibility in the AF-GARCH model can potentially explain some of the improvement in option RMSE over the AF-SV model.

4.6 Estimation on Returns Only

The estimates in Table 3 were intentionally obtained by minimizing the object of interest which in this paper was taken to be the option price mean squared error. The advantage of the estimation methodology we have suggested is that it can be tailored to any objection function of interest. Instead of option pricing MSE the objective of interest could relative price MSE or implied volatility MSE. More economically based objectives such as hedging error variance or trading profits could be entertained as well. Consider the special case where the researcher is interested not in option price fitting but instead in a return-based likelihood objective. In the two discrete time GARCH models the optimization problem is straightforward as the log-likelihood is simply given by

$$\ln(L) = -\frac{1}{2} \sum_{t=2}^T (\ln(2\pi) + \ln(V_t) + z_t^2/V_t)$$

where z_t is obtained from the relevant return equation and V_t from the variance dynamics. The initial observation can be used as a starting condition setting V_1 to θ .

Using the discretization in equations (16) and (17) and using the auxiliary particle filter we can similarly estimate the two SV models on returns by maximizing the function analog to the log-likelihood

$$\ln(L) = -\frac{1}{2} \sum_{t=1}^T (\ln(2\pi) + \ln(\bar{V}_t) + z_t^2/\bar{V}_t)$$

where \bar{V}_t is now computed from the particle filter as in equation (18).

Table 10 shows the results from this estimation exercise when using daily returns on the S&P500 from 1980 to 1999. Note that the objective function is much larger for the NA-SV than the AF-SV and similarly the NA-GARCH objective is much larger than the AF-GARCH objective. Thus, returns themselves favor the non-affine specifications in both the SV and the GARCH case. Notice also that the ρ and c parameters which capture the shock correlations in the four models are much smaller in magnitude in Table 10 than those reported in Table 3. The estimates in Table 10 are very close to the return-driven estimates found for example in Eraker, Johannes and Polson (2003) who also use the 1980-1999 sample period.

5 Conclusions and Directions for Future Work

This paper has provided 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 25.5% lower than that of the affine Heston (1993) model (AF-SV), and the out-of-sample RMSE is approximately 27.5% lower. The non-affine model outperforms the affine model in all three in-sample exercises as well as in all three out-of-sample exercises. The non-affine model also outperforms the affine model for virtually 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 10% in-sample and 15% 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: The AF-SV model outperforms the AF-GARCH model for 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 model in Heston and Nandi (2000). 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 somewhat surprising that AF-SV model outperforms the AF-GARCH model for short maturities, even though it underperforms on average. One possible explanation 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 easily 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 further improve the fit of non-affine models.

6 Appendix: Option Pricing Formulae

6.1 AF-SV

Heston (1993) demonstrates that the AF-SV model admits a closed form solution for option prices, which is presented here in terms of the physical parameters $\kappa, \theta, \rho, \sigma$ as well as λ . We have

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

where

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

and

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

$$d = \sqrt{(\rho\sigma\phi i - b_j)^2 - \sigma^2(2\mu_j\phi i - \phi^2)} \quad g = \frac{b_j - \rho\sigma\phi i + d}{b_j - \rho\sigma\phi i - d}$$

$$\mu_1 = \frac{1}{2}, \quad \mu_2 = -\frac{1}{2}, \quad b_1 = \kappa - \lambda - \rho\sigma, \quad b_2 = \kappa - \lambda$$

6.2 AF-GARCH

Heston and Nandi (2000) show that the AF-GARCH option price is given by

$$C = SP_1 - Ke^{-r(T-t)}P_2$$

where again using Fourier inversion we have

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

with the characteristic function given by

$$\begin{aligned} f(\phi) &= S^\phi \exp(A(t, T, \phi) + B(t, T, \phi)V) \\ A(t, T, \phi) &= A(t+1, T, \phi) + \phi r + B(t, T, \phi)w - \frac{1}{2} \log(1 - 2aB(t+1, T, \phi)) \\ B(t, T, \phi) &= (p+c) - \frac{1}{2}c^2 + bB(t+1, T, \phi) + \frac{\frac{1}{2(\phi-c)^2}}{1 - 2aB(t+1, T, \phi)} \end{aligned}$$

The coefficients A and B are calculated by recursion from the following terminal conditions

$$A(T, T, \phi) = 0 \quad B(T, T, \phi) = 0$$

References


- [1] Ait-Sahalia, Y., P. Mykland, and L. Zhang (2005), How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise, *Review of Financial Studies*, 18, 351-416.
- [2] Ait-Sahalia, Y. and A. Lo (1998), Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices, *Journal of Finance*, 53, 499-547.
- [3] Amin, K. and V. Ng (1993), ARCH Processes and Option Valuation, Manuscript, University of Michigan.
- [4] Andersen, T.G., L. Benzoni, and J. Lund (2002), Estimating Jump-Diffusions for Equity Returns, *Journal of Finance*, 57, 1239-1284.
- [5] Andersen, T.G., T. Bollerslev, and F.X. Diebold (2005), Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility, Manuscript, University of Pennsylvania.
- [6] Andersen, T.G., T. Bollerslev, F.X. Diebold, P. Labys (2003), Modeling and Forecasting Realized Volatility, *Econometrica*, 71, 529-626.
- [7] Bakshi, G., C. Cao, and Z. Chen (1997), Empirical Performance of Alternative Option Pricing Models, *Journal of Finance*, 52, 2003-2049.
- [8] Bates, D. (1996), Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options, *Review of Financial Studies*, 9, 69-107.
- [9] Bates, D. (2000), Post-'87 Crash Fears in the S&P 500 Futures Option Market, *Journal of Econometrics*, 94, 181-238.
- [10] Bates, D. (2004), Maximum Likelihood Estimation of Latent Affine Processes, Working Paper, University of Iowa.
- [11] Benzoni, L. (2002), Pricing Options Under Stochastic Volatility: An Empirical Investigation, Working Paper, University of Minnesota.
- [12] 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.
- [13] Black, F., and M. Scholes (1973), The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, 81, 637-659.

- [14] Bolic M. (2004), Architectures for Efficient Implementation of Particle Filters, Ph.D. Thesis, State University of New York at Stony Brook, available from [www-sigproc.eng.cam.ac.uk/smc/index.html](http://www.sigproc.eng.cam.ac.uk/smc/index.html).
- [15] 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.
- [16] Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics*, 31, 307-327.
- [17] 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.
- [18] Bollerslev, T. and H.O. Mikkelsen (1996), Long-Term Equity Anticipation Securities and Stock Market Volatility Dynamics, *Journal of Econometrics*, 92, 75-99.
- [19] Brennan, M. (1979), The Pricing of Contingent Claims in Discrete-Time Models, *Journal of Finance*, 34, 53-68.
- [20] 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.
- [21] Camara, A. (2003), A Generalization of the Brennan-Rubinstein Approach for the Pricing of Derivatives, *Journal of Finance*, 58, 805-819.
- [22] Carr, P. and L. Wu (2004), Time-Changed Levy Processes and Option Pricing, *Journal of Financial Economics*, 17, 113–141.
- [23] Chernov, M., A.R. Gallant, E. Ghysels, and G. Tauchen, (2003), Alternative Models for Stock Price Dynamics, *Journal of Econometrics*, 116, 225-257.
- [24] 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.
- [25] Christoffersen, P. and K. Jacobs (2004a), Which GARCH Model for Option Valuation?, *Management Science*, 50, 1204-1221.
- [26] Christoffersen, P. and K. Jacobs (2004b), The Importance of the Loss Function in Option Valuation, *Journal of Financial Economics*, 72, 291-318.
- [27] Corradi, V. (2000), Reconsidering the Continuous Time Limit of the GARCH(1,1) Process, *Journal of Econometrics*, 96, 145-153.

- [28] Dai, Q. and K.J. Singleton (2000), Specification Analysis of Affine Term Structure Models, *Journal of Finance*, 55, 1943-1978.
- [29] Diebold, F.X. and J. Lopez (1995), Modeling Volatility Dynamics, in *Macroeconometrics: Developments, Tensions and Prospects*, Kevin Hoover (ed.). Boston: Kluwer Academic Press.
- [30] Diebold, F.X., and R.S. Mariano (1995), Comparing Predictive Accuracy, *Journal of Business and Economic Statistics*, 13, 253-265.
- [31] Duan, J.-C. (1995), The GARCH Option Pricing Model, *Mathematical Finance*, 5, 13-32.
- [32] Duan, J.-C., (1997), Augmented GARCH(p,q) Process and its Diffusion limit, *Journal of Econometrics*, 79, 97-127.
- [33] Duan, J.-C., P. Ritchken, and Z. Sun. (2004), Jump Starting GARCH: Pricing and Hedging Options with Jumps in Returns and Volatilities, Manuscript, University of Toronto.
- [34] Duan, J.-C. and J.-G. Simonato (1998), Empirical Martingale Simulation for Asset Prices, *Management Science*, 44, 1218-1233.
- [35] Duffie, D., and R. Kan (1996), A Yield-Factor Model of Interest Rates, *Mathematical Finance*, 6, 379-406.
- [36] Dumas, B., J. Fleming, and R. Whaley (1998), Implied Volatility Functions: Empirical Tests, *Journal of Finance*, 53, 2059-2106.
- [37] Engle, R. (1982), Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation, *Econometrica*, 50, 987-1008.
- [38] Engle, R. and C. Mustafa (1992), Implied ARCH Models from Options Prices, *Journal of Econometrics*, 52, 289-311.
- [39] Engle, R. and V. Ng (1993), Measuring and Testing the Impact of News on Volatility, *Journal of Finance*, 48, 1749-1778.
- [40] Eraker, B. (2001), MCMC Analysis of Diffusion Models with Application to Finance, *Journal of Business and Economic Statistics*, 19, 177-191.
- [41] Eraker, B. (2004), Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, *Journal of Finance*, 59, 1367-1403.
- [42] Eraker, B., M. Johannes, and N. Polson (2003) The Role of Jumps in Returns and Volatility, *Journal of Finance*, 58, 1269-1300.
- [43] Granger, C. (1969), Prediction with a generalized cost of error function, *Operations Research Quarterly* 20, 199-207.


- [44] 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.
- [45] Heston, S. and S. Nandi (2000), A Closed-Form GARCH Option Pricing Model, *Review of Financial Studies*, 13, 585-626.
- [46] Hsieh, K. and P. Ritchken (2000), An Empirical Comparison of GARCH Option Pricing Models, manuscript, Case Western Reserve University
- [47] 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.
- [48] Johannes, M., Polson, N. and J. Stroud (2002), Nonlinear Filtering of Stochastic Differential Equations with Jumps, Working Paper, University of Chicago.
- [49] Jones, C. (2003), The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets, *Journal of Econometrics* 116, 181-224.
- [50] Lewis, A. (2000), *Option Valuation Under Stochastic Volatility*, Finance Press, Newport Beach, CA.
- [51] 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.
- [52] Nelson, D. (1990), Conditional Heteroskedasticity in Asset Returns: A New Approach, *Econometrica*. 59, 347-370.
- [53] Pan, J. (2002), The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, *Journal of Financial Economics*, 63, 3-50.
- [54] Pitt, M., and N. Shephard (1999), Filtering via Simulation: Auxiliary Particle Filters, *Journal of the American Statistical Association*, 94, 590-599.
- [55] Ritchken, P. and R. Trevor (1999), Pricing Options under Generalized GARCH and Stochastic Volatility Processes, *Journal of Finance*, 54, 377-402.
- [56] Rubinstein, M. (1976), The Valuation of Uncertain Income Streams and the Pricing of Options, *Bell Journal of Economics*, 7, 407-425.
- [57] Schroder, M. (2004), Risk Neutral Parameter Shifts and Derivatives Pricing in Discrete Time, *Journal of Finance*, 59, 2375-2402.
- [58] Weiss, A., and A. Andersen (1984), Estimating time series models using the relevant forecast evaluation criterion, *Journal of the Royal Statistical Society, Series A* 147, 484-487

Figure 1: QQ Plot of Daily Volatility Changes and Scatter Plots of Daily Volatility Changes and Absolute Changes Versus Volatility Level



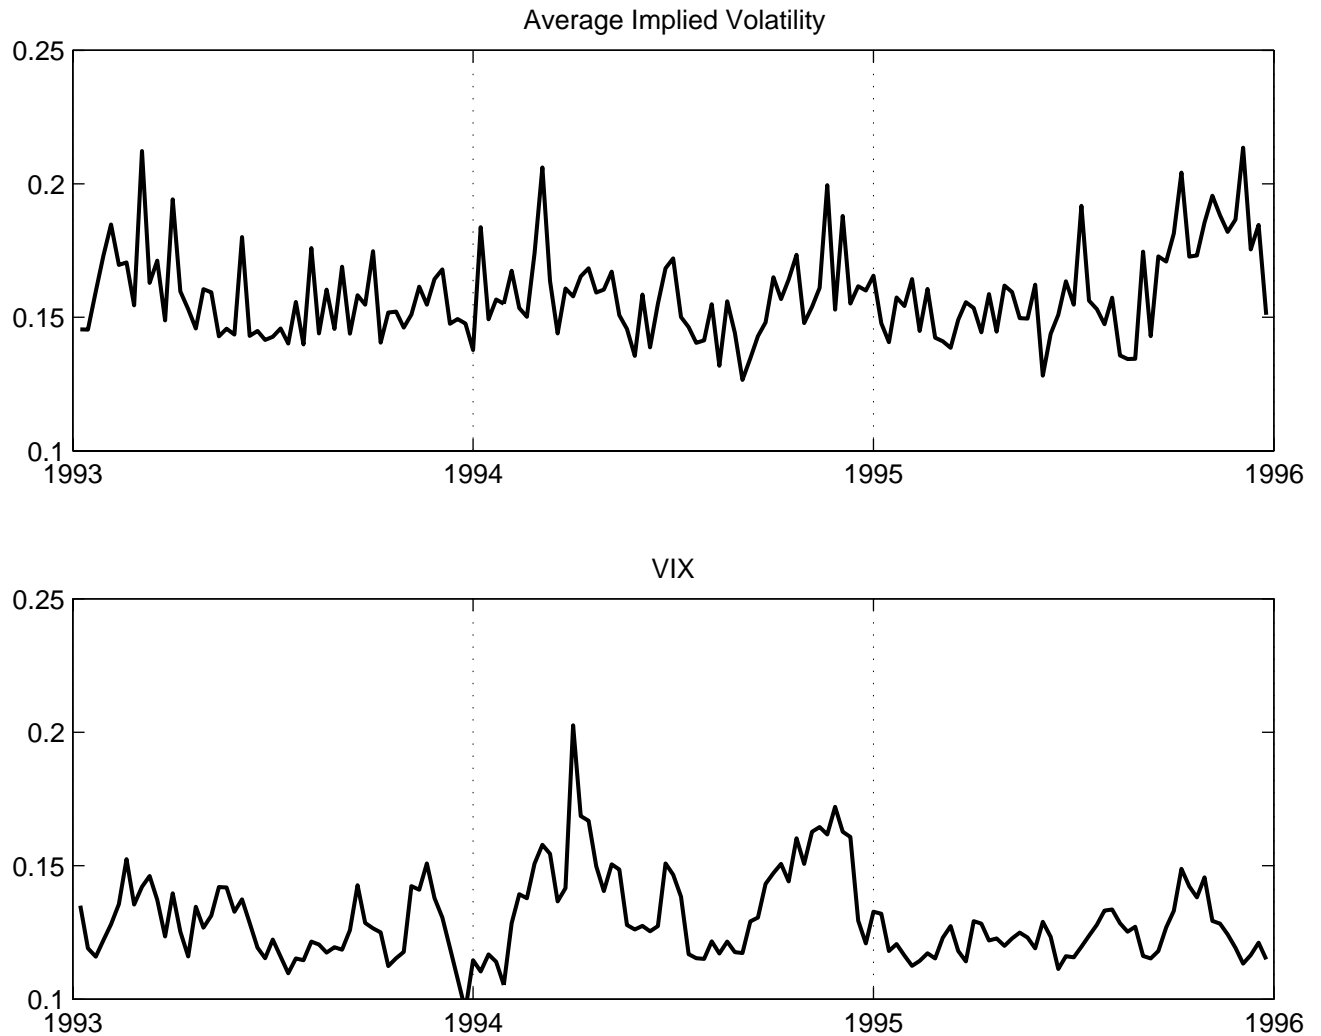
Notes to figure: Using realized volatilities from 1990 to 2002 we compute the daily changes and compare their empirical quantiles with quantiles from the normal distribution (top panel). We then scatter plot the daily volatility changes against the daily volatility level (middle panel). In the bottom panel we scatter plot the absolute daily volatility changes against the daily volatility level and show an OLS regression line for reference.

Figure 2: QQ Plot of Daily Log Variance Changes
and Scatter Plots of Daily Log Variance Changes and Absolute Changes Versus Log Variance



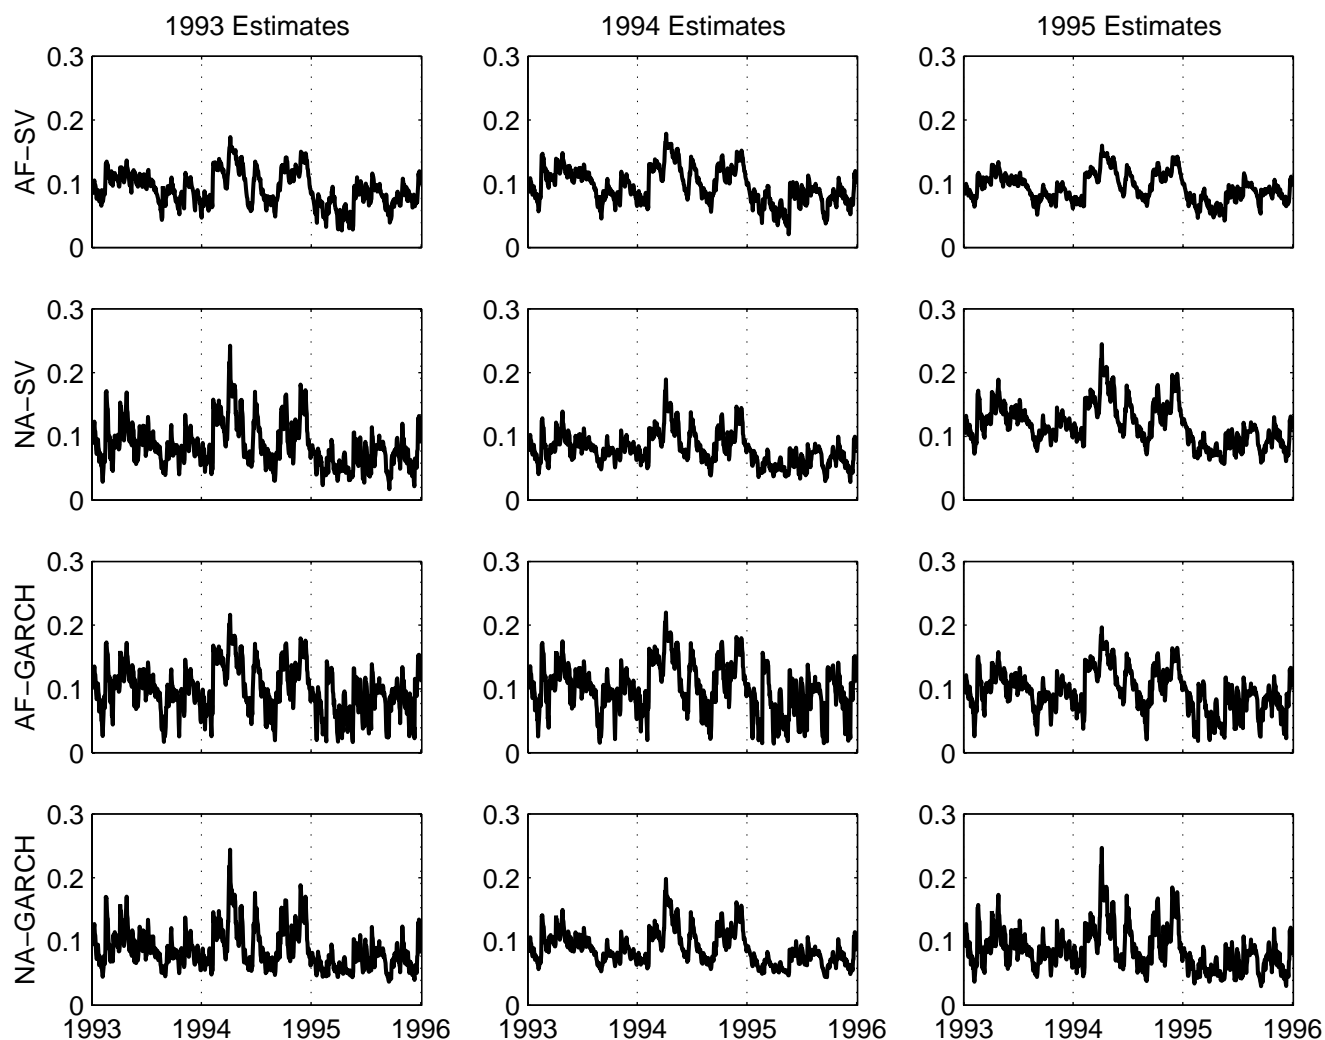
Notes to figure: Using realized variances from 1990 to 2002 we compute the daily log changes and compare their empirical quantiles with quantiles from the normal distribution (top panel). We then scatter plot the daily log variance changes against the daily log variance level (middle panel). In the bottom panel we scatter plot the absolute daily log variance changes against the daily log variance level and show an OLS regression line for reference.

Figure 3: Average Implied Volatility in S&P500 Option Data and the CBOE VIX




Notes to figure: The top panel plots the average implied Black-Scholes volatility each Wednesday during 1993-1995. 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 www.cboe.com.

Figure 4: Spot Volatility Paths from Each Model. 1993-1995.




Notes to figure: For each model we plot the daily spot volatility path (annualized) during 1993-1995. The three columns use the estimates from 1993, 1994 and 1995 respectively as reported in Table 3. The top row shows the volatility paths from the affine SV model (AF-SV), the second row shows the non-affine SV model (NA-SV), the third row shows the affine GARCH model (AF-GARCH), and the bottom row shows the non-affine GARCH model (NA-GARCH).

Figure 5: Weekly In Sample Root Mean Squared Error (RMSE) for each Model. 1993-1995.



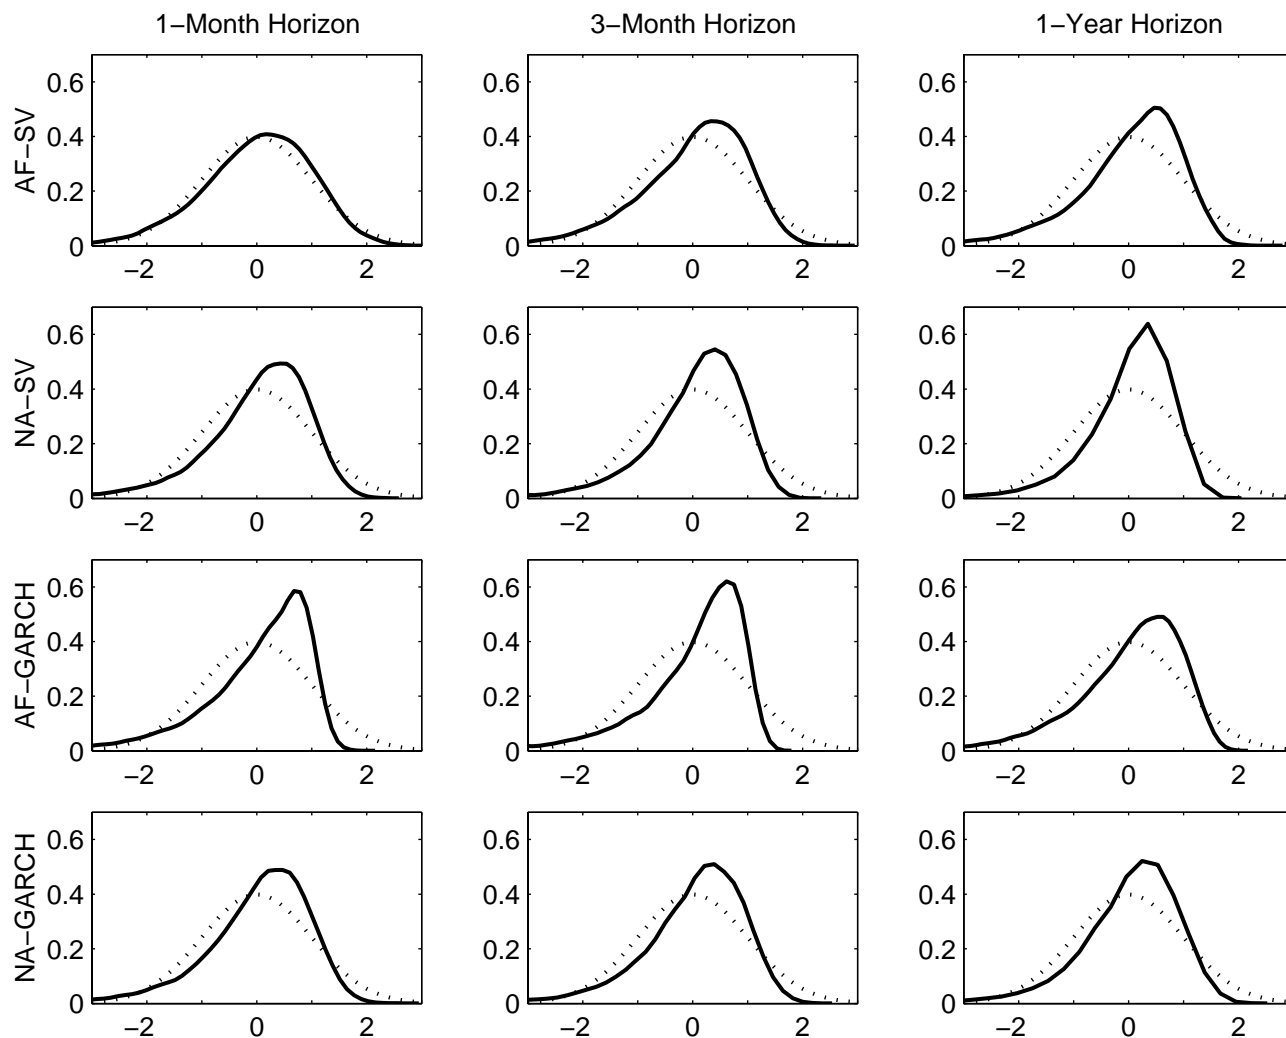
Notes to figure: The four panels show the weekly root mean squared option valuation error (RMSE) for the AF-SV, NA-SV, AF-GARCH, and NA-GARCH models respectively.

Figure 6: Weekly In Sample Bias for each Model. 1993-1995.



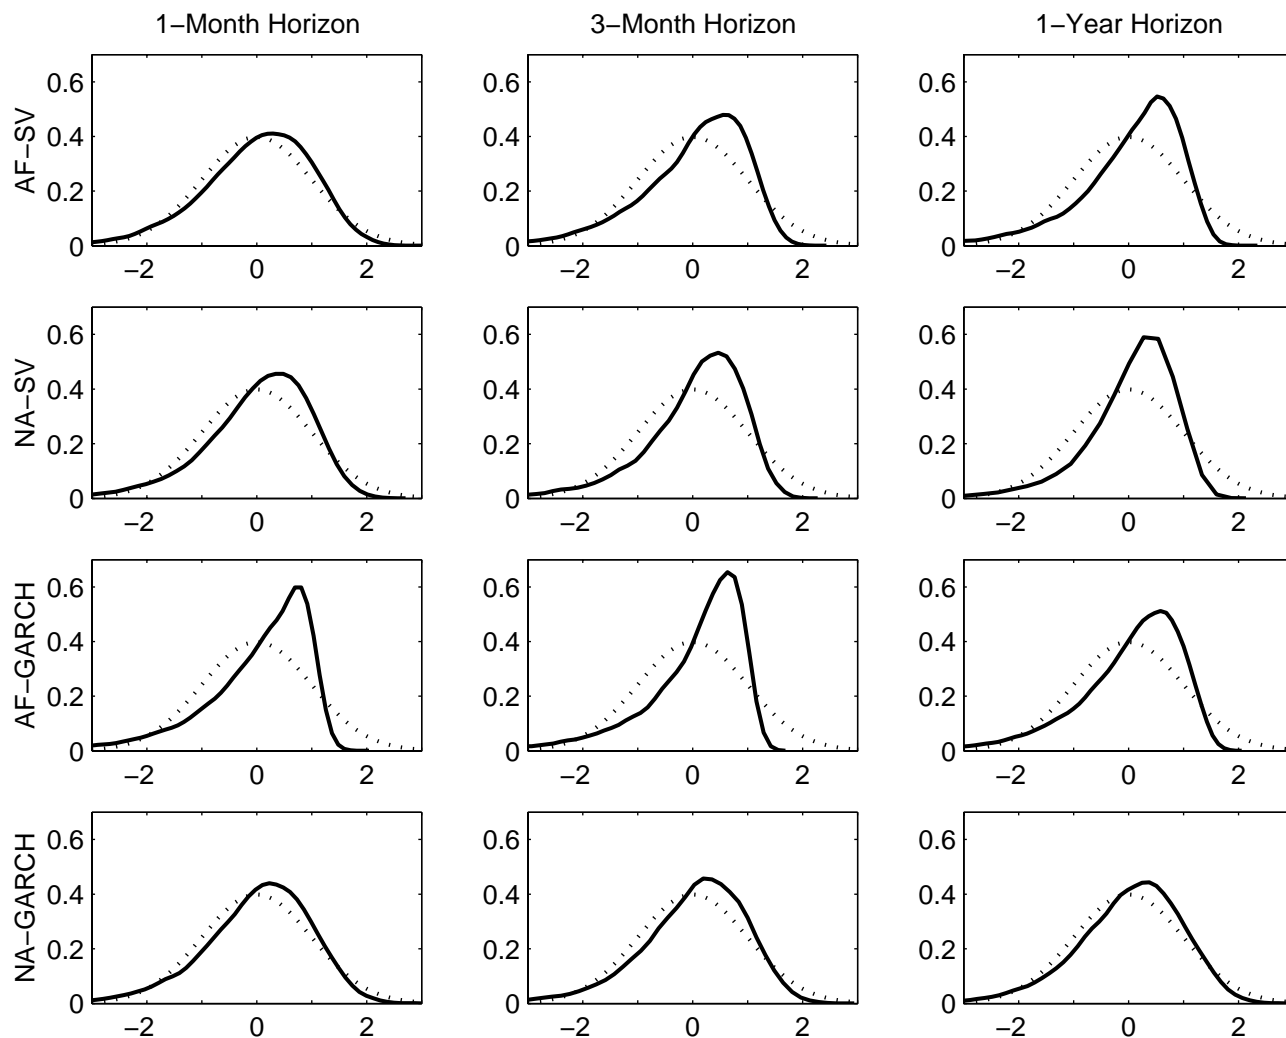
Notes to figure: The four panels show the weekly option valuation bias for the AF-SV, NA-SV, AF-GARCH, and NA-GARCH models respectively. The bias is calculated as average data price less average model price.

Figure 7.A: Model Implied State Price Densities. 1993 Estimates
 1-Month, 3-Month and 1-Year Horizon



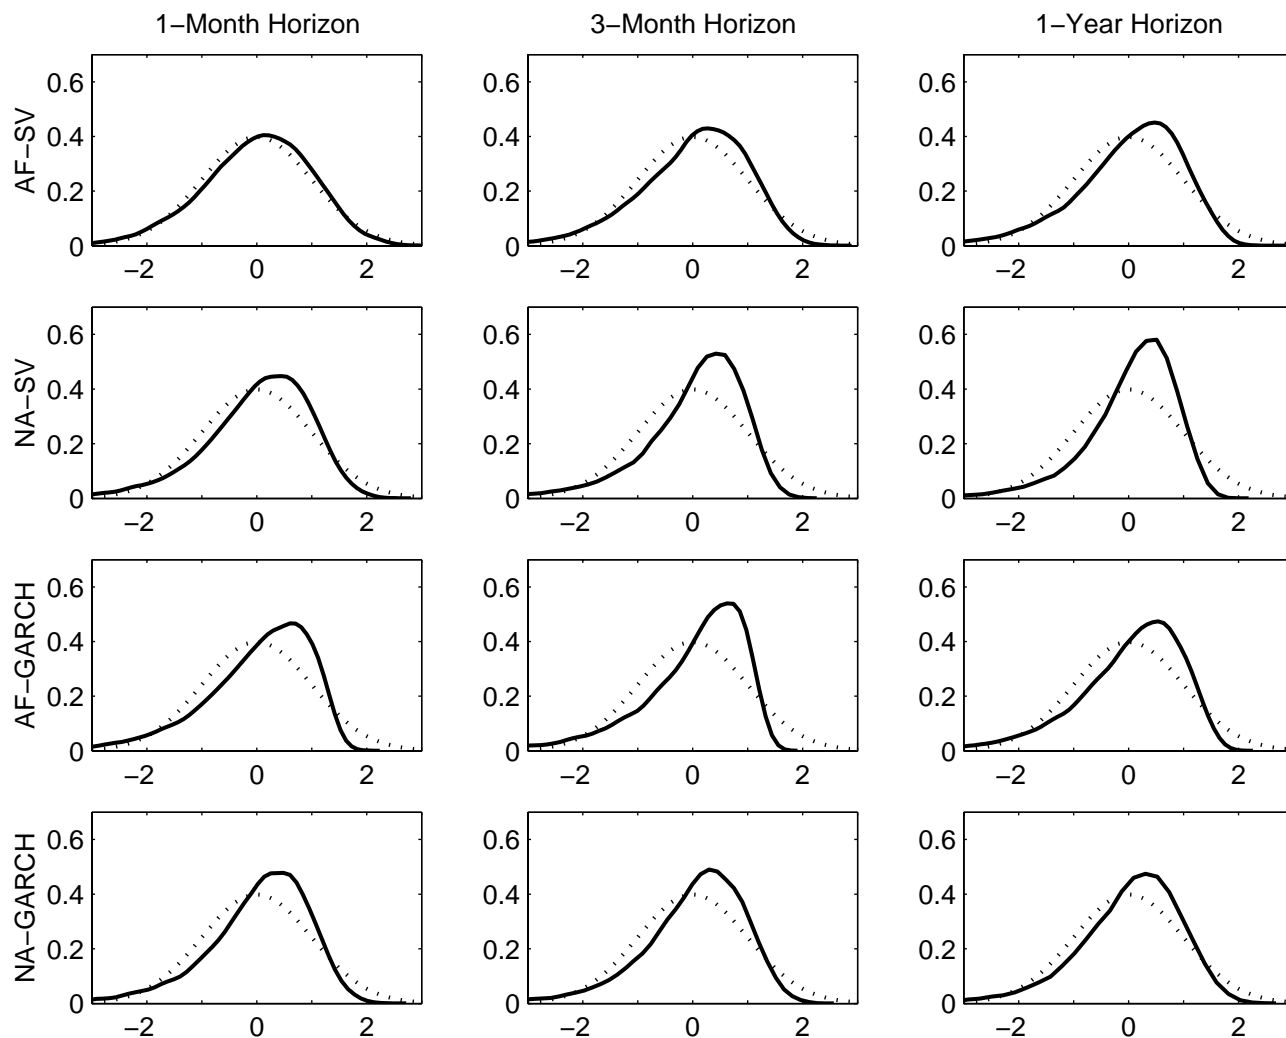
Notes to figure: Each row of panels reports the risk neutral distribution of the index return according to the AF-SV, NA-SV, AF-GARCH and NA-GARCH model respectively. The normal distribution corresponding to the Black-Scholes model is shown in dots. The left column reports the 1-month horizon, the center column the 3-month horizon, and the right column shows the 1-year horizon distribution. The distributions are constructed by simulating 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. The 1993 parameter estimates from Table 3 converted to the risk neutral measure are used to simulate the risk neutral returns.

Figure 7.B: Model Implied State Price Densities. 1994 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-SV, NA-SV, AF-GARCH and NA-GARCH model respectively. The normal distribution corresponding to the Black-Scholes model is shown in dots. The left column reports the 1-month horizon, the center column the 3-month horizon, and the right column shows the 1-year horizon distribution. The distributions are constructed by simulating 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. The 1994 parameter estimates from Table 3 converted to the risk neutral measure are used to simulate the risk neutral returns.

Figure 7.C: Model Implied State Price Densities. 1995 Estimates
 1-Month, 3-Month and 1-Year Horizon



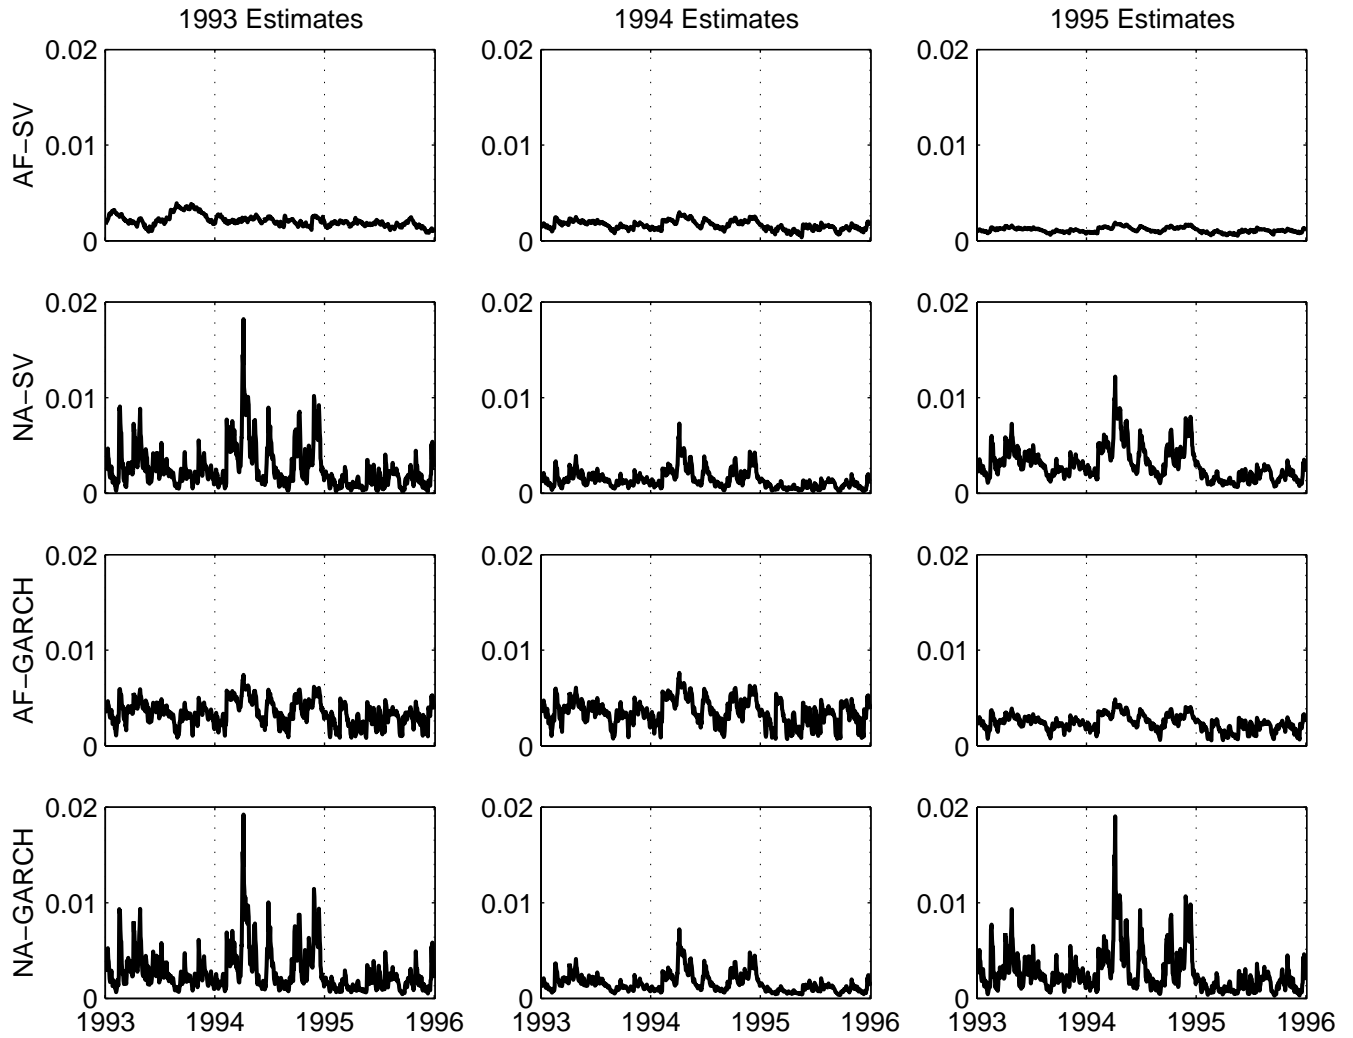
Notes to figure: Each row of panels reports the risk neutral distribution of the index return according to the AF-SV, NA-SV, AF-GARCH and NA-GARCH model respectively. The normal distribution corresponding to the Black-Scholes model is shown in dots. The left column reports the 1-month horizon, the center column the 3-month horizon, and the right column shows the 1-year horizon distribution. The distributions are constructed by simulating 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. The 1995 parameter estimates from Table 3 converted to the risk neutral measure are used to simulate the risk neutral returns.

Figure 8: Conditional Leverage Paths from Each Model. 1993-1995.



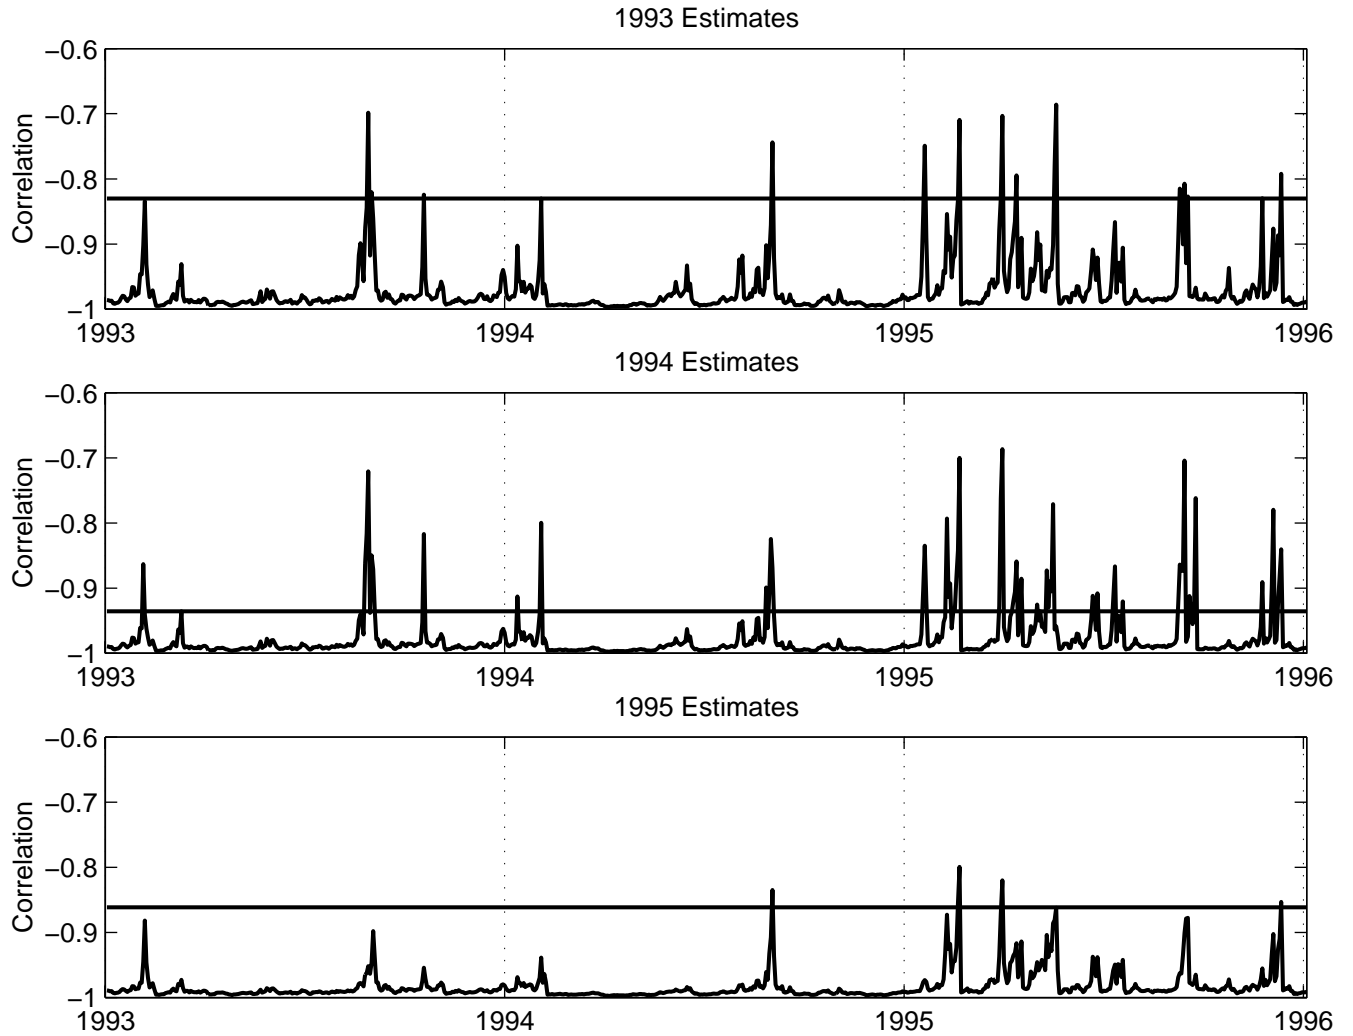
Notes to figure: For each model we plot the daily conditional leverage path defined as the conditional covariance between shocks to returns and shocks to variance. The paths are annualized by multiplying by 252 and plotted during 1993-1995. The three columns use estimates from 1993, 1994, and 1995 respectively all of which are reported in Table 3. The top row shows the leverage paths from the affine SV model (AF-SV), the second row shows the non-affine SV model (NA-SV), the third row shows the affine GARCH model (AF-GARCH), and the bottom row shows the non-affine GARCH model (NA-GARCH).

Figure 9: Conditional Volatility of Variance Paths from Each Model. 1993-1995.



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 variance of returns. The paths are annualized by multiplying by 252 and plotted during 1993-1995. The three columns use estimates from 1993, 1994, and 1995 respectively all of which are reported in Table 3. The top row shows the volatility of variance paths from the affine SV model (AF-SV), the second row shows the non-affine SV model (NA-SV), the third row shows the affine GARCH model (AF-GARCH), and the bottom row shows the non-affine GARCH model (NA-GARCH).

Figure 10: Time-Varying Shock Correlation in AF-GARCH Model. 1993-1995
and Constant AF-SV Shock Correlation



Notes to figure: Each panel plots the time-varying shock correlation implied by the affine GARCH (AF-GARCH) model during 1993-1995. The top panel uses the 1993 estimates, the middle panel uses the 1994 estimates, and the bottom panel uses the 1995 estimates from Table 3. The horizontal line in each plot corresponds to the correlation estimation for the affine SV (AF-SV) model for the relevant year. These estimates are reported in Table 3 as well.

Table 1: S&P500 Index Call Option Data. 1993-1995. In Sample**Panel A. Number of Call Option Contracts**

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 16 | 902 | 874 | 907 | 2,699 |
| 0.975<S/X<1 | 151 | 762 | 374 | 276 | 1,563 |
| 1<S/X<1.025 | 168 | 722 | 399 | 310 | 1,599 |
| 1.025<S/X<1.05 | 147 | 699 | 359 | 274 | 1,479 |
| 1.05<S/X<1.075 | 136 | 637 | 330 | 249 | 1,352 |
| <u>S/X>1.075</u> | <u>345</u> | <u>1620</u> | <u>1076</u> | <u>830</u> | <u>3,871</u> |
| All | <u>963</u> | <u>5,342</u> | <u>3,412</u> | <u>2,846</u> | <u>12,563</u> |

Panel B. Average Call Price

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|--------------|
| S/X<0.975 | 0.56 | 1.91 | 5.07 | 10.87 | 5.94 |
| 0.975<S/X<1 | 1.92 | 6.48 | 14.57 | 28.69 | 11.90 |
| 1<S/X<1.025 | 8.62 | 14.03 | 22.81 | 36.14 | 19.94 |
| 1.025<S/X<1.05 | 19.22 | 23.40 | 31.51 | 44.69 | 28.90 |
| 1.05<S/X<1.075 | 30.28 | 33.17 | 40.62 | 51.46 | 38.07 |
| <u>S/X>1.075</u> | <u>56.29</u> | <u>57.05</u> | <u>63.55</u> | <u>71.87</u> | <u>61.97</u> |
| All | <u>29.19</u> | <u>27.46</u> | <u>32.85</u> | <u>39.95</u> | <u>31.89</u> |

Panel C. Average Implied Volatility from Call Options

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.1134 | 0.1005 | 0.1075 | 0.1171 | 0.1086 |
| 0.975<S/X<1 | 0.1088 | 0.1112 | 0.1235 | 0.1377 | 0.1190 |
| 1<S/X<1.025 | 0.1342 | 0.1282 | 0.1354 | 0.1440 | 0.1338 |
| 1.025<S/X<1.05 | 0.1817 | 0.1477 | 0.1465 | 0.1521 | 0.1519 |
| 1.05<S/X<1.075 | 0.2329 | 0.1671 | 0.1580 | 0.1595 | 0.1714 |
| <u>S/X>1.075</u> | <u>0.3879</u> | <u>0.2276</u> | <u>0.1818</u> | <u>0.1740</u> | <u>0.2253</u> |
| All | <u>0.2568</u> | <u>0.1575</u> | <u>0.1408</u> | <u>0.1410</u> | <u>0.1596</u> |

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, 1993 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. 1993-1995. Out of Sample**Panel A. Number of Call Option Contracts**

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|--------------|
| S/X<0.975 | 19 | 858 | 826 | 855 | 2,558 |
| 0.975<S/X<1 | 142 | 731 | 390 | 256 | 1,519 |
| 1<S/X<1.025 | 154 | 702 | 374 | 274 | 1,504 |
| 1.025<S/X<1.05 | 137 | 662 | 332 | 219 | 1,350 |
| 1.05<S/X<1.075 | 102 | 533 | 268 | 124 | 1,027 |
| <u>S/X>1.075</u> | <u>132</u> | <u>602</u> | <u>304</u> | <u>172</u> | <u>1,210</u> |
| All | <u>686</u> | <u>4,088</u> | <u>2,494</u> | <u>1,900</u> | <u>9,168</u> |

Panel B. Average Call Price

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|--------------|
| S/X<0.975 | 0.59 | 1.86 | 4.98 | 10.91 | 5.88 |
| 0.975<S/X<1 | 1.95 | 6.29 | 14.28 | 28.63 | 11.70 |
| 1<S/X<1.025 | 8.49 | 13.72 | 22.72 | 35.74 | 19.44 |
| 1.025<S/X<1.05 | 18.55 | 23.15 | 31.08 | 44.68 | 28.13 |
| 1.05<S/X<1.075 | 29.65 | 32.92 | 40.14 | 51.41 | 36.71 |
| <u>S/X>1.075</u> | <u>52.09</u> | <u>51.73</u> | <u>57.07</u> | <u>73.10</u> | <u>56.15</u> |
| All | <u>20.46</u> | <u>19.53</u> | <u>22.70</u> | <u>29.04</u> | <u>22.43</u> |

Panel C. Average Implied Volatility from Call Options

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.1137 | 0.0991 | 0.1066 | 0.1169 | 0.1078 |
| 0.975<S/X<1 | 0.1058 | 0.1086 | 0.1221 | 0.1366 | 0.1170 |
| 1<S/X<1.025 | 0.1306 | 0.1247 | 0.1344 | 0.1427 | 0.1312 |
| 1.025<S/X<1.05 | 0.1728 | 0.1433 | 0.1435 | 0.1485 | 0.1474 |
| 1.05<S/X<1.075 | 0.2203 | 0.1615 | 0.1536 | 0.1523 | 0.1653 |
| <u>S/X>1.075</u> | <u>0.3348</u> | <u>0.1975</u> | <u>0.1686</u> | <u>0.1645</u> | <u>0.2065</u> |
| All | <u>0.2031</u> | <u>0.1390</u> | <u>0.1325</u> | <u>0.1346</u> | <u>0.1422</u> |

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, 1993 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-SV**

| | μ | κ | θ | σ | ρ | λ | Risk Neutral Properties | |
|-------------|----------|----------|----------|----------|---------|-----------|--------------------------------|------------|
| | | | | | | | κ^* | θ^* |
| 1993 | 6.54E-04 | 0.0122 | 0.1066 | 1.01E-03 | -0.8301 | -0.0042 | 0.0081 | 0.1313 |
| 1994 | 6.57E-04 | 0.0150 | 0.0958 | 1.05E-03 | -0.9359 | -0.0068 | 0.0081 | 0.1301 |
| 1995 | 6.16E-04 | 0.0132 | 0.1044 | 7.46E-04 | -0.8616 | -0.0039 | 0.0093 | 0.1244 |

NA-SV

| | μ | κ | θ | σ | ρ | λ | κ^* | θ^* |
|-------------|----------|----------|----------|----------|---------|-----------|------------|------------|
| 1993 | 2.49E-06 | 0.0223 | 0.1378 | 3.10E-01 | -0.9990 | -0.0003 | 0.0219 | 0.1388 |
| 1994 | 5.96E-04 | 0.0412 | 0.0771 | 2.03E-01 | -0.9990 | -0.0315 | 0.0097 | 0.1592 |
| 1995 | 1.27E-03 | 0.0142 | 0.1140 | 2.04E-01 | -0.9730 | -0.0055 | 0.0087 | 0.1457 |

AF-GARCH

| | p | κ | θ | a | c | κ^* | θ^* |
|-------------|----------|----------|----------|----------|-------|------------|------------|
| 1993 | 1.16E+00 | 0.0399 | 0.1128 | 1.65E-06 | 652.1 | 0.0363 | 0.1183 |
| 1994 | 9.64E-01 | 0.0359 | 0.1155 | 1.43E-06 | 761.7 | 0.0327 | 0.1211 |
| 1995 | 4.66E+00 | 0.0363 | 0.1027 | 9.43E-07 | 822.0 | 0.0283 | 0.1164 |

NA-GARCH

| | p | κ | θ | σ | ρ | κ^* | θ^* |
|-------------|----------|----------|----------|----------|---------|------------|------------|
| 1993 | 4.67E-02 | 0.0356 | 0.1093 | 3.23E-01 | -0.9433 | 0.0212 | 0.1416 |
| 1994 | 1.40E-01 | 0.0341 | 0.0801 | 1.84E-01 | -0.9176 | 0.0095 | 0.1516 |
| 1995 | 8.96E-02 | 0.0505 | 0.0912 | 3.12E-01 | -0.9788 | 0.0227 | 0.1359 |

Notes: For each model, we estimate the parameters using Nonlinear Least Squares (NLS) on the option valuation errors. We use Wednesday option prices in each of the years 1993, 1994 and 1995 to estimate the models. The spot volatility is computed within the NLS routine using the auxiliary particle filter on daily returns. The unconditional variance, θ , is reported in annual standard deviations.

Table 4: RMSE In and Out of Sample

| Panel A. RMSE In Sample | | | | |
|--------------------------------|---------------|---------------|-----------------|-----------------|
| | AF-SV | NA-SV | AF-GARCH | NA-GARCH |
| 1993 | 0.9395 | 0.7087 | 0.8118 | 0.7015 |
| 1994 | 1.0339 | 0.7993 | 0.9591 | 0.8761 |
| 1995 | <u>0.9371</u> | <u>0.6933</u> | <u>0.8728</u> | <u>0.7654</u> |
| Overall | <u>0.9875</u> | <u>0.7354</u> | <u>0.8865</u> | <u>0.7881</u> |
| Normalized | 1.0000 | 0.7447 | 0.8977 | 0.7981 |

| Panel B. RMSE Out of Sample | | | | |
|------------------------------------|---------------|---------------|-----------------|-----------------|
| | AF-SV | NA-SV | AF-GARCH | NA-GARCH |
| 1993 | 0.8275 | 0.5818 | 0.6925 | 0.5728 |
| 1994 | 1.0451 | 0.6805 | 0.9577 | 0.7686 |
| 1995 | <u>0.8971</u> | <u>0.7735</u> | <u>0.9449</u> | <u>0.8214</u> |
| Overall | <u>0.9540</u> | <u>0.6920</u> | <u>0.8811</u> | <u>0.7383</u> |
| Normalized | 1.0000 | 0.7254 | 0.9235 | 0.7739 |

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

Table 5: Diebold-Mariano Test on Average Weekly RMSE

| Panel A. RMSE In Sample | | | | |
|--------------------------------|--------------|--------------|-----------------|-----------------|
| | AF-SV | NA-SV | AF-GARCH | NA-GARCH |
| Average Weekly RMSE | 0.9442 | 0.7057 | 0.8476 | 0.7429 |
| DM Test versus AF-SV | -- | 5.7555 | 3.2676 | 5.7719 |
| DM Test versus NA-SV | -5.7555 | -- | -5.0359 | -1.8319 |

| Panel B. RMSE Out of Sample | | | | |
|------------------------------------|--------------|--------------|-----------------|-----------------|
| | AF-SV | NA-SV | AF-GARCH | NA-GARCH |
| Average Weekly RMSE | 0.8987 | 0.6447 | 0.8189 | 0.6783 |
| DM Test versus AF-SV | -- | 4.2112 | 1.6507 | 4.8494 |
| DM Test versus NA-SV | -4.2112 | -- | -4.5825 | -1.3306 |

Notes: We compute the average weekly RMSE for each model during the 1993 to 1995 period. Using the Diebold and Mariano (1995) test we compare the average weekly RMSE of each of the two SV models to all the other models. We allow for autocorrelation of up to four weeks in the test.

Table 6: RMSE by Moneyness and Maturity. 1993-1995. In Sample**Panel A. AF-SV**

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.4419 | 0.8693 | 0.8587 | 1.0598 | 0.9327 |
| 0.975<S/X<1 | 0.7366 | 1.1045 | 0.9845 | 0.9633 | 1.0212 |
| 1<S/X<1.025 | 0.6869 | 0.9981 | 0.9626 | 0.9489 | 0.9514 |
| 1.025<S/X<1.05 | 0.6263 | 0.9313 | 0.9059 | 1.0048 | 0.9139 |
| 1.05<S/X<1.075 | 0.8439 | 0.8730 | 0.9560 | 1.1290 | 0.9427 |
| <u>S/X>1.075</u> | <u>0.8417</u> | <u>0.9728</u> | <u>1.0010</u> | <u>1.3594</u> | <u>1.0649</u> |
| All | <u>0.7648</u> | <u>0.9631</u> | <u>0.9456</u> | <u>1.1374</u> | <u>0.9875</u> |

Panel B. NA-SV

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.3604 | 0.5089 | 0.5778 | 0.8110 | 0.6453 |
| 0.975<S/X<1 | 0.5515 | 0.7089 | 0.6607 | 0.9098 | 0.7246 |
| 1<S/X<1.025 | 0.6375 | 0.7386 | 0.6490 | 0.8099 | 0.7220 |
| 1.025<S/X<1.05 | 0.6050 | 0.7015 | 0.5937 | 0.8350 | 0.6955 |
| 1.05<S/X<1.075 | 0.5749 | 0.6733 | 0.5997 | 0.9123 | 0.6985 |
| <u>S/X>1.075</u> | <u>0.8061</u> | <u>0.7290</u> | <u>0.6634</u> | <u>1.1419</u> | <u>0.8260</u> |
| All | <u>0.6753</u> | <u>0.6847</u> | <u>0.6271</u> | <u>0.9389</u> | <u>0.7354</u> |

Panel C. AF-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.6214 | 0.7399 | 0.7101 | 0.9949 | 0.8253 |
| 0.975<S/X<1 | 1.0279 | 1.0521 | 0.8322 | 1.0501 | 1.0011 |
| 1<S/X<1.025 | 0.9180 | 0.9869 | 0.7580 | 0.9116 | 0.9126 |
| 1.025<S/X<1.05 | 0.6336 | 0.8422 | 0.6306 | 0.9633 | 0.8022 |
| 1.05<S/X<1.075 | 0.8824 | 0.7171 | 0.5752 | 1.1122 | 0.7932 |
| <u>S/X>1.075</u> | <u>0.8022</u> | <u>0.8393</u> | <u>0.8197</u> | <u>1.2265</u> | <u>0.9282</u> |
| All | <u>0.8491</u> | <u>0.8658</u> | <u>0.7470</u> | <u>1.0720</u> | <u>0.8865</u> |

Panel D. NA-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.4916 | 0.6447 | 0.6865 | 0.9670 | 0.7792 |
| 0.975<S/X<1 | 0.6729 | 0.8080 | 0.7347 | 1.0142 | 0.8202 |
| 1<S/X<1.025 | 0.7118 | 0.7797 | 0.6818 | 0.8825 | 0.7712 |
| 1.025<S/X<1.05 | 0.6295 | 0.7108 | 0.6067 | 0.9010 | 0.7197 |
| 1.05<S/X<1.075 | 0.5731 | 0.6631 | 0.6123 | 1.0015 | 0.7186 |
| <u>S/X>1.075</u> | <u>0.8052</u> | <u>0.7203</u> | <u>0.6550</u> | <u>1.1884</u> | <u>0.8349</u> |
| All | <u>0.7089</u> | <u>0.7221</u> | <u>0.6667</u> | <u>1.0297</u> | <u>0.7881</u> |

Notes: We use the NLS estimates from Table 3 to compute the 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 1993-1995 in-sample period, which consists of Wednesday option prices.

Table 7: RMSE by Moneyness and Maturity. 1993-1995. Out of Sample

Panel A. AF-SV

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.5150 | 0.8711 | 0.9006 | 1.1122 | 0.9653 |
| 0.975<S/X<1 | 0.7278 | 1.1917 | 1.0774 | 1.0482 | 1.1028 |
| 1<S/X<1.025 | 0.6754 | 1.0168 | 1.0055 | 0.9105 | 0.9653 |
| 1.025<S/X<1.05 | 0.5227 | 0.8882 | 0.8474 | 0.8678 | 0.8446 |
| 1.05<S/X<1.075 | 0.7840 | 0.7974 | 0.8520 | 0.9915 | 0.8362 |
| <u>S/X>1.075</u> | <u>0.5492</u> | <u>0.8749</u> | <u>1.0257</u> | <u>1.1069</u> | <u>0.9232</u> |
| All | <u>0.6507</u> | <u>0.9561</u> | <u>0.9506</u> | <u>1.0422</u> | <u>0.9540</u> |

Panel B. NA-SV

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.3311 | 0.5046 | 0.6140 | 0.8399 | 0.6661 |
| 0.975<S/X<1 | 0.5483 | 0.7736 | 0.7550 | 0.9336 | 0.7806 |
| 1<S/X<1.025 | 0.6177 | 0.7483 | 0.7536 | 0.8314 | 0.7534 |
| 1.025<S/X<1.05 | 0.5210 | 0.6882 | 0.6476 | 0.6979 | 0.6648 |
| 1.05<S/X<1.075 | 0.4553 | 0.6465 | 0.6189 | 0.7772 | 0.6406 |
| <u>S/X>1.075</u> | <u>0.5084</u> | <u>0.6197</u> | <u>0.6367</u> | <u>0.6367</u> | <u>0.6154</u> |
| All | <u>0.5345</u> | <u>0.6666</u> | <u>0.6675</u> | <u>0.8166</u> | <u>0.6920</u> |

Panel C. AF-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.6796 | 0.7347 | 0.7636 | 1.0784 | 0.8726 |
| 0.975<S/X<1 | 1.0162 | 1.1190 | 0.9471 | 1.1366 | 1.0710 |
| 1<S/X<1.025 | 0.9447 | 1.0352 | 0.8791 | 0.9650 | 0.9764 |
| 1.025<S/X<1.05 | 0.5747 | 0.8611 | 0.7037 | 0.8861 | 0.8039 |
| 1.05<S/X<1.075 | 0.8066 | 0.6691 | 0.6081 | 0.8397 | 0.6918 |
| <u>S/X>1.075</u> | <u>0.5757</u> | <u>0.7001</u> | <u>0.6747</u> | <u>0.9396</u> | <u>0.7210</u> |
| All | <u>0.8083</u> | <u>0.8786</u> | <u>0.7813</u> | <u>1.0238</u> | <u>0.8811</u> |

Panel D. NA-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.4552 | 0.6134 | 0.7019 | 0.9907 | 0.7842 |
| 0.975<S/X<1 | 0.6343 | 0.8474 | 0.8056 | 1.0137 | 0.8503 |
| 1<S/X<1.025 | 0.6806 | 0.7576 | 0.7592 | 0.8658 | 0.7715 |
| 1.025<S/X<1.05 | 0.5378 | 0.6825 | 0.6278 | 0.7587 | 0.6694 |
| 1.05<S/X<1.075 | 0.4620 | 0.6262 | 0.5677 | 0.7937 | 0.6204 |
| <u>S/X>1.075</u> | <u>0.5103</u> | <u>0.6056</u> | <u>0.5885</u> | <u>0.6384</u> | <u>0.5965</u> |
| All | <u>0.5767</u> | <u>0.6975</u> | <u>0.6931</u> | <u>0.9125</u> | <u>0.7383</u> |

Notes: We use the NLS estimates from Table 3 to compute the 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 1993-1995 out-of-sample period, which consists of Thursday option prices.

Table 8: Bias by Moneyness and Maturity. 1993-1995. In Sample**Panel A. AF-SV**

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | -0.1929 | -0.1947 | -0.1819 | -0.0545 | -0.1434 |
| 0.975<S/X<1 | 0.1159 | -0.2876 | -0.2338 | -0.0374 | -0.1916 |
| 1<S/X<1.025 | 0.1648 | -0.0349 | 0.0452 | 0.0377 | 0.0201 |
| 1.025<S/X<1.05 | 0.3746 | 0.3643 | 0.2477 | 0.2227 | 0.3108 |
| 1.05<S/X<1.075 | 0.7274 | 0.5065 | 0.4701 | 0.4212 | 0.5041 |
| <u>S/X>1.075</u> | <u>0.5178</u> | <u>0.7417</u> | <u>0.6268</u> | <u>0.5921</u> | <u>0.6577</u> |
| All | <u>0.3891</u> | <u>0.2544</u> | <u>0.2023</u> | <u>0.2141</u> | <u>0.2414</u> |

Panel B. NA-SV

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.0528 | 0.1193 | 0.0858 | -0.1507 | 0.0173 |
| 0.975<S/X<1 | 0.0862 | -0.0376 | 0.0300 | 0.0580 | 0.0074 |
| 1<S/X<1.025 | 0.1916 | -0.0349 | 0.0449 | -0.0824 | -0.0004 |
| 1.025<S/X<1.05 | 0.3851 | 0.0535 | 0.0200 | -0.0173 | 0.0652 |
| 1.05<S/X<1.075 | 0.4198 | 0.1149 | 0.0643 | 0.0145 | 0.1147 |
| <u>S/X>1.075</u> | <u>0.5589</u> | <u>0.3452</u> | <u>0.0670</u> | <u>-0.0029</u> | <u>0.2123</u> |
| All | <u>0.3661</u> | <u>0.1354</u> | <u>0.0600</u> | <u>-0.0526</u> | <u>0.0900</u> |

Panel C. AF-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | -0.0503 | 0.2278 | -0.0341 | -0.2576 | -0.0218 |
| 0.975<S/X<1 | 0.2631 | -0.2754 | -0.4201 | 0.0796 | -0.1953 |
| 1<S/X<1.025 | -0.0962 | -0.4349 | -0.2984 | 0.1483 | -0.2522 |
| 1.025<S/X<1.05 | 0.1812 | -0.1304 | -0.1686 | 0.3712 | -0.0158 |
| 1.05<S/X<1.075 | 0.6991 | -0.0315 | 0.0569 | 0.5275 | 0.1665 |
| <u>S/X>1.075</u> | <u>0.4835</u> | <u>0.4848</u> | <u>0.2601</u> | <u>0.6664</u> | <u>0.4612</u> |
| All | <u>0.3232</u> | <u>0.0666</u> | <u>-0.0199</u> | <u>0.2180</u> | <u>0.0971</u> |

Panel D. NA-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | 0.0123 | 0.0262 | -0.0763 | -0.3148 | -0.1217 |
| 0.975<S/X<1 | 0.2425 | 0.0044 | -0.0122 | 0.0712 | 0.0352 |
| 1<S/X<1.025 | 0.3219 | 0.0473 | 0.0480 | -0.0034 | 0.0665 |
| 1.025<S/X<1.05 | 0.4273 | 0.0869 | 0.0185 | 0.0902 | 0.1047 |
| 1.05<S/X<1.075 | 0.4246 | 0.1012 | 0.0741 | 0.1672 | 0.1393 |
| <u>S/X>1.075</u> | <u>0.5582</u> | <u>0.3130</u> | <u>0.0547</u> | <u>0.1846</u> | <u>0.2355</u> |
| All | <u>0.4196</u> | <u>0.1298</u> | <u>0.0111</u> | <u>-0.0166</u> | <u>0.0866</u> |

Notes: We use the NLS estimates from Table 3 to compute the bias for various moneyness and maturity bins for the four models. The option prices used in the table are for the 1993-1995 in-sample period, which consists of Wednesday option prices.

Table 9: Bias by Moneyness and Maturity. 1993-1995. Out of Sample**Panel A. AF-SV**

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|---------------|
| S/X<0.975 | -0.2725 | -0.2497 | -0.3017 | -0.1617 | -0.2372 |
| 0.975<S/X<1 | -0.0318 | -0.4691 | -0.3933 | -0.1822 | -0.3604 |
| 1<S/X<1.025 | 0.0166 | -0.2366 | -0.1447 | 0.0323 | -0.1388 |
| 1.025<S/X<1.05 | 0.2639 | 0.2030 | 0.1795 | 0.2722 | 0.2146 |
| 1.05<S/X<1.075 | 0.6815 | 0.4605 | 0.4535 | 0.3657 | 0.4692 |
| <u>S/X>1.075</u> | <u>0.3388</u> | <u>0.6779</u> | <u>0.7273</u> | <u>0.8129</u> | <u>0.6725</u> |
| All | <u>0.2088</u> | <u>0.0158</u> | <u>-0.0218</u> | <u>0.0362</u> | <u>0.0242</u> |

Panel B. NA-SV

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|----------------|
| S/X<0.975 | 0.0543 | 0.0543 | -0.0072 | -0.1906 | -0.0474 |
| 0.975<S/X<1 | -0.0628 | -0.1916 | -0.1148 | -0.0657 | -0.1386 |
| 1<S/X<1.025 | -0.0113 | -0.2012 | -0.0903 | -0.1249 | -0.1403 |
| 1.025<S/X<1.05 | 0.2644 | -0.0828 | -0.0863 | -0.1145 | -0.0536 |
| 1.05<S/X<1.075 | 0.3233 | 0.0315 | -0.0695 | -0.3235 | -0.0087 |
| <u>S/X>1.075</u> | <u>0.4278</u> | <u>0.2151</u> | <u>-0.0006</u> | <u>-0.1938</u> | <u>0.1260</u> |
| All | <u>0.1692</u> | <u>-0.0350</u> | <u>-0.0529</u> | <u>-0.1645</u> | <u>-0.0514</u> |

Panel C. AF-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|----------------|
| S/X<0.975 | -0.2124 | 0.1924 | -0.0460 | -0.2080 | -0.0214 |
| 0.975<S/X<1 | 0.0983 | -0.3799 | -0.4476 | 0.0533 | -0.2796 |
| 1<S/X<1.025 | -0.2952 | -0.5686 | -0.3273 | 0.1685 | -0.3463 |
| 1.025<S/X<1.05 | 0.0452 | -0.2410 | -0.1990 | 0.2876 | -0.1159 |
| 1.05<S/X<1.075 | 0.6456 | -0.0498 | 0.0016 | 0.1908 | 0.0617 |
| <u>S/X>1.075</u> | <u>0.2694</u> | <u>0.4003</u> | <u>0.2517</u> | <u>0.6181</u> | <u>0.3796</u> |
| All | <u>0.1051</u> | <u>-0.1118</u> | <u>-0.1299</u> | <u>0.0394</u> | <u>-0.0692</u> |

Panel D. NA-GARCH

| | <u>DTM<20</u> | <u>20<DTM<80</u> | <u>80<DTM<180</u> | <u>DTM>180</u> | <u>All</u> |
|---------------------|------------------|------------------------|-------------------------|-------------------|----------------|
| S/X<0.975 | -0.0557 | -0.0463 | -0.1753 | -0.3635 | -0.1940 |
| 0.975<S/X<1 | 0.0974 | -0.1672 | -0.1746 | -0.0760 | -0.1290 |
| 1<S/X<1.025 | 0.1396 | -0.1527 | -0.1142 | -0.0761 | -0.0992 |
| 1.025<S/X<1.05 | 0.2935 | -0.0669 | -0.0863 | -0.0103 | -0.0259 |
| 1.05<S/X<1.075 | 0.3361 | 0.0255 | -0.0677 | -0.2161 | 0.0029 |
| <u>S/X>1.075</u> | <u>0.4307</u> | <u>0.1634</u> | <u>-0.0354</u> | <u>0.0337</u> | <u>0.1242</u> |
| All | <u>0.2414</u> | <u>-0.0493</u> | <u>-0.1256</u> | <u>-0.1970</u> | <u>-0.0789</u> |

Notes: We use the NLS estimates from Table 3 to compute the bias for various moneyness and maturity bins for the four models. The option prices used in the table are for the 1993-1995 out-of-sample period, which consists of Thursday option prices.

Table 10: Parameter Estimates on Daily S&P500 Returns. 1980 - 1999

| AF-SV | | | | | |
|-----------------|----------|----------|----------|---------|------------------------|
| μ | κ | θ | σ | ρ | ln(L) Objective |
| 5.55E-04 | 0.0198 | 0.1342 | 1.06E-03 | -0.5393 | 17,226.0 |
| NA-SV | | | | | |
| μ | κ | θ | σ | ρ | ln(L) Objective |
| 6.62E-03 | 0.0066 | 0.1322 | 2.73E-01 | -0.4751 | 17,613.3 |
| AF-GARCH | | | | | |
| p | κ | θ | a | c | ln(L) Objective |
| 4.29E+00 | 0.0497 | 0.1496 | 4.42E-06 | 95.0 | 16,743.3 |
| NA-GARCH | | | | | |
| p | κ | θ | σ | ρ | ln(L) Objective |
| 7.32E-02 | 0.0181 | 0.1566 | 1.26E-01 | -0.6308 | 16,803.6 |

Notes: Using a normal likelihood objective function we estimate the four models on daily S&P500 returns from January 1, 1980 through December 31, 1999. For the SV models, the spot volatility is computed using the auxiliary particle filter.