The Nonlinearity and Jumps in Stochastic Volatility: Evidence from Returns and Options Data

Karim Mimouni *

American University in Dubai, United Arab Emirates

We estimate the Constant Elasticity of Variance (CEV) model in order to study the level of nonlinearity in the volatility dynamic. We also estimate a CEV process combined with a jump process (CEVJ), and analyze the effects of the jump component on the nonlinearity coefficient. We investigate whether there is complementarity or competition between the jumps and the CEV specification since both are intended to address the misspecification of existing linear models. Estimation is performed using the particle-filtering technique on a long series of S&P500 returns and on options data. Our results show that both returns and returns and options favor nonlinear specifications for the volatility dynamic, suggesting that the extensive use of linear models is not supported empirically. We also find that the inclusion of jumps does not lower the level of nonlinearity and does not improve the CEV model fit.

JEL classification: G12

Keywords: CEV model; CEVJ model; maximum likelihood; particle filter; Monte-Carlo; approximation; options; nonlinearity; jumps

1. Introduction

The performance of any option pricing model is evaluated by its ability to fit at-the-money call and put prices which represent liquid instruments actively traded in financial markets. Furthermore, in response to the increasing demand for deep in-the-money and deep out-of-the-money options by hedgers and speculators, more complex option pricing models were introduced. A convenient way to evaluate their performance is to compute the model-implied Black and Scholes volatilities and compare them to their market counterparts. We observe that the market-implied volatility is typically higher for in-and out-of-the-money call options compared to at-the-money calls. Plotting the strike versus implied volatility produces therefore a U-shaped curve, known as the smirk. There has been limited success in the literature to find a model which perfectly fits the smirk. One surmises that this is due to the fact that all popular options pricing models share the common feature of having linear specifications for the volatility dynamic. In fact, the existing literature on stochastic volatility offers scant evidence on nonlinear models. In particular, the degree of nonlinearity implied from returns and options data and the role of jump processes in a nonlinear context are not investigated in a consistent manner which would allow for comparison. We attempt to understand these issues using a set of S&P500 returns and European call options.

We should note that the existing literature uses linear specifications in the volatility dynamic because they allow closed-form solutions for option prices facilitating their empirical implementation. The most popular empirical implementations include the original version of Heston (1993) to which jumps in returns and volatility can be added. See, for example, Bakshi, Cao and Chen (1997), Chernov and Ghysels (2000), Pan (2002), Eraker (2004), and Eraker, Johannes and Polson (2003).

---

* Corresponding author: School of Business Administration, American University in Dubai, Dubai, P.O. Box 28282, United Arab Emirates; Email: kmimouni@aud.edu

I thank FQRSC and IFM² for financial support. I am grateful for helpful comments from participants at the Northern Finance Association Meetings 2009 and an anonymous referee. Any errors are my responsibility.

Nonlinear models, in particular the CEV and CEVJ models, have not been the preferred choice for empirical investigations. Therefore, few researchers implement them compared to the extensive literature on linear models. New evidence shows that nonlinear specifications may lead to a better fit for option prices. Recent findings using a set of options and daily returns conclude that the Heston (1993) model, while more convenient and computationally easy, is dominated by a continuous time stochastic volatility model where the diffusion term is quadratic. See Christoffersen et al. (2006), Ait-Sahalia and Kimmel (2007), and Jones (2003). We adopt these findings and use nonlinear models as building blocks to explore better specifications.

We investigate two nonlinear models: the CEV model which has been previously estimated in a small number of empirical papers, and the CEV model with jumps i.e., CEVJ. The (CEV) specification has been investigated in interest rate literature in a number of papers, and has been applied to stochastic volatility processes by Chacko and Viceira (2003), Lewis (2000), Jones (2003) and Ait-Sahalia and Kimmel (2007).

One drawback to CEV models is that they do not generally satisfy some sufficient conditions for a number of important results including the global growth and Lipschitz conditions. However, Jones (2003) shows that violations of growth and Lipschitz conditions outside the range sample data are less critical to certain estimation techniques, whose calculations are all conditional on the observed sample. The estimation methodology we employ in this paper belongs to these techniques.

Our paper makes two main contributions. First, it estimates the models on S&P500 returns and on multiple cross sections of European call options. This options data set is richer than the one used by Jones (2003) and Ait-Sahalia and Kimmel (2007), and the nonlinearity coefficient is estimated using options with different maturities. Our empirical implementation uses the particle-filtering technique in order to conduct a fair comparison between the estimates obtained using returns and those obtained using options. We therefore assess whether nonlinearity is an option phenomenon which is not present in returns, or if it is a characteristic of both data. Second, we investigate the effects of jumps on the degree of nonlinearity and on the model fit. To our knowledge, Chacko and Viceira (2003) is the only study to include jumps in a nonlinear volatility model, but they use returns only. Hence, we believe that the estimation of the CEVJ model using options could prove very informative. The principal reason why the existing literature does not study the CEVJ model using options is the computational challenge to compute option prices by Monte Carlo simulation. This burden is greatly reduced thanks to our estimation methodology which extracts easily the unobserved variance needed to obtain option prices.

The Heston (1993) model is also estimated and will serve as a reference. This model has a “quasi” closed-form solution for option prices. However, to ensure consistency, we estimate all models by Monte-Carlo simulation. Therefore, the three models we study will be directly comparable and differences across models cannot be related to the estimation methodology.

Our empirical results show clearly that nonlinearity is confirmed by returns and options alike, and that the level of nonlinearity obtained from returns and options is of the same order of magnitude. We also find evidence that the inclusion of jumps does not affect the degree of nonlinearity. It is therefore more likely that the two features are complementary rather than competitive, as concluded by Chacko and Viceira (2003).

We employ two different sources of data to estimate the models. First, we estimate the parameters on S&P500 returns. Although stochastic volatility models are motivated by the need to fit option prices, estimation on returns only is very common. Indeed, we typically aim to avoid overfitting by using the return estimates to price options. Second, we use a combination of daily returns and at-the-money European call options to estimate the models.

When we estimate the models on returns, we use the Maximum Likelihood Importance Sampling (MLIS) technique introduced by Pitt (2002) who proposes a likelihood approximation and shows its efficiency in the presence of unobserved states. His likelihood estimator is a by-product of the particle filter which uses the true dynamic of returns to compute the approximate likelihood. This method belongs to the Simulated Maximum Likelihood techniques.

To estimate the CEV and CEVJ models on options, we use Christoffersen et al. (2006) methodology based on an iterative Nonlinear Least Squares (NLS) procedure. Since neither the CEV nor the CEVJ models admit a closed-form solution, option prices are computed by Monte Carlo simulations. This adds considerably to the computational burden in estimation.

In both data sets, the variance path is filtered from daily returns using the Sampling Importance Resampling (SIR) particle filter of Gordon et al. (1993) which is suitable for nonlinear state space applications. This method allows the models to be estimated using returns only, and using returns and options. In addition, it is easy to implement empirically.

The paper is organized as follows. Section 2 presents the CEV and CEVJ models. Section 3 describes the particle-filtering technique used to obtain the variance conditional densities. It then presents the estimation methodology based on those conditional densities using returns and using returns and options jointly. In section 4 we present the empirical results. Finally, section 5 concludes.

2. The CEV and the CEVJ models

The most general model that we investigate, the CEVJ model, is defined by the following two equations under the physical measure

$$d\log(S_t) = \left(r + \lambda_S\left(1-\rho^2\right)V_t - \frac{1}{2}V_t - \lambda_J\overline{\mu}_J\right)dt + \sqrt{V_t}\,dB_{1t} + dJ_tN_t \tag{1}$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma V_t^{\beta}\left(\rho dB_{1t} + \sqrt{1-\rho^2}\,dB_{2t}\right),$$

where S_t is the price of the underlying asset, V_t is the variance, J_t is the jump intensity, N_t is the jump size, and $corr(dB_{1t}, dB_{2t}) = 0.$ The parameter κ represents the speed of mean reversion of the volatility to its long-run mean, θ is the stationary value for the volatility process known also as the long-run mean, and σ determines the level of the volatility of volatility. The parameter λ_S determines the risk premium required to compensate investors for holding the underlying asset, and ρ represents the correlation between returns and variance leading to a skewed returns distribution. As in most of the existing literature, we assume that B_{1t} and B_{2t} are two standard Brownian motions, $dJ_t \to Poisson(\lambda_J)$, $N_t \to N\left(\mu_J, \sigma_J^2\right)$. The jump compensator is $\overline{\mu}_J = \exp\left(\mu_J + \frac{1}{2}\sigma_J^2\right) - 1$.

In this paper, we consider jumps only in returns. The effects of jumps in the variance dynamic are left for future research. Eraker (2004) estimates a model with correlated jumps in returns and variance. Because this model is not parsimonious and because there is no empirical evidence on the role of jumps in variance, we do not include the model for estimation.

When we set $dJ_tN_t = 0$ in the CEVJ model, the jump component vanishes and we obtain the Lewis (2000) CEV model. Hence, the CEV model is defined by the following two equations

$$d\log(S_t) = \left(r + \lambda_S\left(1-\rho^2\right)V_t - \frac{1}{2}V_t\right)dt + \sqrt{V_t}\,dB_{1t} \tag{2}$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma V_t^{\beta}\left(\rho dB_{1t} + \sqrt{1-\rho^2}\,dB_{2t}\right).$$

When $dJ_tN_t = 0$ in the CEVJ model and β is deliberately set equal to 1/2, then the CEVJ model reduces to the Heston (1993) defined by the following two equations

$$d\log(S_t) = \left(r + \lambda_S\left(1-\rho^2\right)V_t - \frac{1}{2}V_t\right)dt + \sqrt{V_t}\,dB_{1t} \tag{3}$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}\left(\rho dB_{1t} + \sqrt{1-\rho^2}\,dB_{2t}\right).$$

Models (1) and (2) share the same expression for the variance process. The only difference is the inclusion of jumps in the price dynamic of the CEVJ model. In what follows, we study the impact of including jumps on the estimate of the nonlinearity coefficient β, and its implications on the model fit when we use returns and when we use options. In fact, while most of the literature on linear models shows the importance of jumps in prices, estimation of those models based on options offers mixed results. Some infer that there are no economic benefits to including jumps, whereas others find tremendous improvements in fit. Eraker (2004) reports an improvement in fit of 1%. Bates (2000) finds an improvement of circa 2%. Broadie, Chernov and Johannes (2007) find a 50% improvement in fit by adding jumps in prices to the SV model, but their empirical setup is different from ours and from most of the literature. It is therefore interesting to investigate these effects in a nonlinear context.

3. The Estimation Methodology

In order to estimate the models on returns and on options we ascertain to know the conditional distribution of the volatility at each time step. To this end, we apply the particle-filtering technique. For more details on the particle-filtering technique and its applications in the context of the estimation on returns and options, see Christoffersen et al. (2006).

In what follows, we describe how we can derive the conditional densities and outline, in the appendix, all the steps required to obtain them in the context of the CEV and CEVJ models.

3.1. Model Estimation Using Returns

We first examine the MLIS approach of estimating the CEV and CEVJ models. Pioneered by Pitt (2002), the method computes an approximate likelihood when the state is unobserved. Not only does the technique apply to general models, but it does not require efforts, and it is not model specific. It is fully consistent with the returns dynamic since, in this setup the variance is treated as endogenous, and is estimated at the same time as the parameters.

Pitt (2002) shows that, in the context of particle filters, the likelihood is given by the following equation

$$p(S(t) \mid \theta, S_{1:t-1}, V_{1:t-1}) = \int p(S(t) \mid \theta, V(t)) p(V(t) \mid \theta, S_{1:t-1}, V_{1:t-1}) dV(t). \quad (4)$$

This likelihood can be approximated by

$$\hat{p}(S(t) \mid \theta, S_{1:t-1}, V_{1:t-1}) = \frac{1}{N} \sum_{k=1}^{N} \varpi_t^k, \quad (5)$$

where $\{\varpi_k\}_{k=1}^N$ represent the unnormalized weights obtained from the particle filter. The particle filter approximates the conditional density of the variance by a set of N discrete particles (see the appendix for further details on the particle filter). It can be shown by applying Kolmogorov's strong law of large numbers that $\hat{p}(S(t) \mid \theta, S_{1:t-1}, V_{1:t-1}) \overset{a.s.}{\rightarrow} p(S(t) \mid \theta, S_{1:t-1}, V_{1:t-1})$ as N tends to infinity. See Gallant (1997) and Geweke (1989) for further details on Kolmogorov's strong law of large numbers. The computation of the Log likelihood is therefore a by-product of the particle filter, and extra computation is not incurred.

The objective function to be maximized is therefore given by

$$L_{PF} = \sum_{t=2}^{T} \ln\left(\frac{1}{N} \sum_{k=1}^{N} \varpi_t^k \right). \quad (6)$$

The estimation of the models using returns requires the following three steps. First, for a given set of candidate parameters, we compute the weights $\{\varpi_t^k\}_{t=1}^N$ using the particle filter approach described in the appendix. Second, we evaluate the objective function given by (6). Third, the optimizer proposes a new set of parameters and the procedure restarts until the objective function (6) is maximized.

3.2. Model Estimation Using Options

The risk-neutral dynamic of the CEVJ model implied by equation (1) is given by

$$d\log(S_t) = \left(r - \frac{1}{2} V_t - \lambda_J^* \overline{\mu}_J^* \right) dt + \sqrt{V_t}\, dB_{1t}^* + dJ_t^* N_t^* \quad (7)$$

$$dV_t = \kappa^*\left(\theta^* - V_t\right)dt + \sigma V_t^{\beta}\left(\rho dB_{1t}^* + \sqrt{1-\rho^2}\, dB_{2t}^*\right),$$

where dB_{1t}^* and dB_{2t}^* are two uncorrelated standard Brownian motions under the risk-neutral measure Q, $dJ_t^* \to Poisson(\lambda_J^*)$, and $N_t^* \to N\left(\mu_J^*, \sigma_J^2\right)$ under Q. $\kappa^* = \kappa + \lambda_V$ and $\theta^* = \dfrac{\kappa\theta}{\kappa + \lambda_V}$.

λ_J^* and $\overline{\mu}_J^*$ are allowed to be different from λ_J and $\overline{\mu}_J$ under the risk-neutral measure. Note that we have assumed that the volatility risk premium is linear in V_t.

Discretizing equation (7) using the Euler discretization yields

$$\log(S_{t+\Delta}) = \log(S_t) + \left(r - \frac{1}{2}V_t - \lambda_J^* \overline{\mu}_J^*\right)\Delta + \sqrt{V_t\Delta}\varepsilon_{1,t+\Delta}^* + J_{t+\Delta}^* N_{t+\Delta}^* \tag{8}$$

$$V_{t+\Delta} = \kappa^*\left(\theta^* - V_t\right)\Delta + \sigma\sqrt{\Delta}V_t^{\beta}\left(\rho\varepsilon_{1,t+\Delta}^* + \sqrt{1-\rho^2}\,\varepsilon_{2,t+\Delta}^*\right),$$

where $corr\left(\varepsilon_{1,t+\Delta}^*, \varepsilon_{2,t+\Delta}^*\right) = 0$.

The CEVJ model with jumps does not admit a closed form solution. Therefore, option prices must be computed by Monte Carlo simulations. Estimating the CEVJ model by NLS requires the following steps. First, we choose a set of starting points for the parameters of the model and filter the volatility using the Gordon et al. (1993) particle filter described in the appendix. Next, option prices are computed by Monte Carlo simulations. Finally, the following objective function is evaluated.

$$SSE = \sum_{t=1}^{T}\sum_{i=1}^{n_t}\left(C_{t,i}^{Model} - C_{t,i}^{Market}\right)^2, \tag{9}$$

The parameter T is the total number of days where option prices are observed, n_t is the number of contracts in day t, C^{Model} is the model price obtained by Monte Carlo simulation, and C^{Market} is the market price. This procedure is repeated until an optimum is reached.

Similar results can be obtained for the CEV model by assuming that $dJ_t^* N_t^* = 0$. Hence, for the CEV model we have

$$d\log(S_t) = \left(r - \frac{1}{2}V_t\right)dt + \sqrt{V_t}\,dB_{1t}^* \tag{10}$$

$$dV_t = \kappa^*\left(\theta^* - V_t\right)dt + \sigma V_t^{\beta}\left(\rho dB_{1t}^* + \sqrt{1-\rho^2}\, dB_{2t}^*\right),$$

and its discretized version gives

$$\log(S_{t+\Delta}) = \log(S_t) + \left(r - \frac{1}{2}V_t\right)\Delta + \sqrt{V_t\Delta}\varepsilon_{1,t+\Delta}^* \tag{11}$$

$$V_{t+\Delta} = \kappa^*\left(\theta^* - V_t\right)\Delta + \sigma\sqrt{\Delta}V_t^{\beta}\left(\rho\varepsilon_{1,t+\Delta}^* + \sqrt{1-\rho^2}\,\varepsilon_{2,t+\Delta}^*\right).$$

Finally, the discretized version of the Heston (1993) model is

$$\log(S_{t+\Delta}) = \log(S_t) + \left(r - \frac{1}{2}V_t\right)\Delta + \sqrt{V_t\Delta}\varepsilon_{1,t+\Delta}^* \tag{12}$$

$$V_{t+\Delta} = \kappa^*\left(\theta^* - V_t\right)\Delta + \sigma\sqrt{\Delta}\sqrt{V_t}\left(\rho\varepsilon_{1,t+\Delta}^* + \sqrt{1-\rho^2}\,\varepsilon_{2,t+\Delta}^*\right).$$

Again, the objective function (9) is minimized using Monte Carlo simulations until the set of optimal parameters is reached.

4. Empirical Results

In this section, we evaluate the models described by equations (1), (2) and (3), initially using returns and then using European call options and returns. We propose to investigate the level of nonlinearity implied by options and by returns, and then include jumps to study the effects on the nonlinearity coefficient.

4.1. Data

We evaluate the models on returns using two S&P500 return sample periods. The first sample relates to the period from January 2, 1987 to December 31, 2004, and includes the 1987 crash whereas in the second sample, relating to the period from January 2, 1990 to December 31, 2004, we exclude the year 1987 as well as the three subsequent years which might have been indirectly affected by the extreme volatility recorded around the crash. We utilize closing prices from the CRSP database. Table 1 contains some statistics about the sample periods. The chosen samples are representative of previous empirical studies using returns. In fact, the standard deviation, skewness, and kurtosis of returns are of the same order of magnitude as any typical sample used in the literature.

Table 1
Summary Statistics for Daily S&P500 Returns

Statistics	1987-2004	1990-2004	1990-1995
Mean	8.5299	7.7710	8.9412
Volatility	17.3837	16.0826	11.4270
Skewness	-2.0894	-0.1020	-0.0997
Kurtosis	44.4628	3.7922	2.4402
Min	-22.8997	-7.1139	-3.7272
Max	8.7089	5.5732	3.6642

Notes: We provide summary statistics for daily S&P500 index for the two samples used in the MLIS estimation form January 2, 1987 to December 31, 2004, and from January 2, 1990 to December 31, 2004. We provide the same summary statistics for the sample used to estimate the models on options.

To evaluate the model on returns and options we use at-the-money (ATM) European call options on the S&P500 index for the period 1990-1995. We apply the same filters to the data as in Bakshi, Cao and Chen (1997). We use Wednesday data since it is the day of the week least likely to be a holiday. A call option is considered ATM if the forward stock price $F(t,T)$ divided by the strike price K, is equal to 1. Since this equality is not typically fulfilled for the available set of options for each Wednesday, we choose the options closest to ATM.

We use a volatility updating rule on the 252 days predating the first Wednesday used in the evaluation sample. We initialize this volatility updating rule as the model's unconditional variance.

Table 2 presents descriptive statistics of the options data by maturity for the period 1990-1995.

Table 2
S&P500 Index Call Option Data, 1990-1995.

	Number of Call Option Contracts	Average Call Price	Average Implied Volatility from Call Options
DTM<20	282	4.350	0.146
20<DTM<80	1,170	8.510	0.142
80<DTM<180	722	14.480	0.150
DTM>180	1,101	23.010	0.155
All	3,275	14.340	0.149

Notes: The sample contains At the Money (ATM) European call options on the S&P500 index. We use quotes within 30 minutes from closing on every Wednesday during the January 1, 1990 to December 31, 1995 period. The moneyness is determined as defined in the data section.

There are 3,275 contracts, the largest group among them of maturities ranging from 20 days to 80 days. The average call price is 14.34 dollars and the average volatility is around 15%, somewhat higher than the sample volatility of the S&P500 index for the period 1990-1995 as reported in Table 1. These features indicate that the sample at hand is standard. The top panel of Figure I gives some indication about the pattern of implied volatility over time. We present the average implied volatility of the options on each Wednesday. It is evident from Figure I that substantial clustering occurs in implied volatilities. It is also evident that volatility is higher in the early part of the sample.

Figure I
Average Weekly Implied Volatility in the S&P500 Option Data and the CBOE VIX

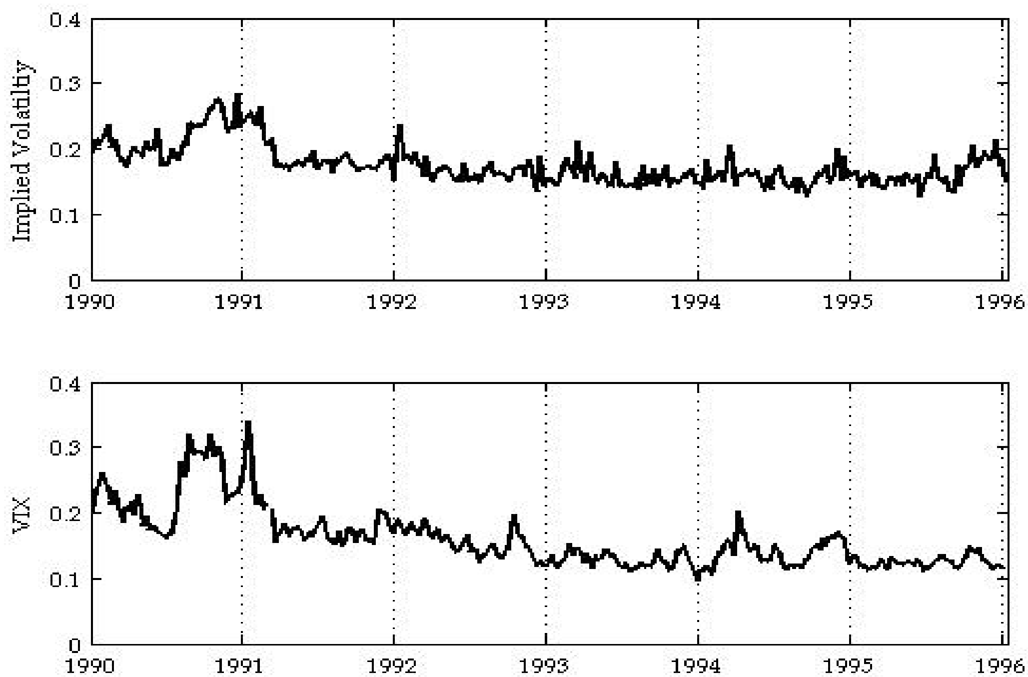

Notes: The top panel plots the average implied Black-Scholes volatility each Wednesday during 1990-1995. The average is taken across maturities and strike prices using the call options in our data set. For comparison, the bottom panel shows the one-month, at-the-money VIX volatility index retrieved from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

4.2. Discussion of the Results

4.2.1. Estimation Using Historical Returns

Table 3 contains the parameter estimates and their standard deviations for the sample period January 2, 1987 to December 31, 2004.[2] Column 3 of Table 3 presents the results for the CEV model from which it is clear that the nonlinearity coefficient β differs significantly from 0.5. This suggests that use of returns data rejects the Heston (1993) model in favor of a more general CEV specification.

We now review the values of the other parameters of the model. We see that the speed of mean reversion is around 2.18; this is expected, since many empirical studies have shown the volatility to be very persistent. Our estimate of the mean reversion implies a daily persistence of around 99.13%. The annualized long-run mean volatility $\sqrt{\theta}$ is around 20.42%. This value is also not surprising because our sample period is characterized overall by several volatile periods including the 1987 crash. The

[2]The standard errors are computed using conventional first-order techniques

volatility of the volatility parameter is 2.21. The correlation between returns and volatility is around -0.67 confirming most empirical findings in the literature according to which negative skewness is present in the distribution of the S&P500 index; see, for example, Benzoni (2002), and Pan (2002) for empirical evidence on the skewness of returns distribution. Finally, the value of the parameter associated with the risk premium is high. This parameter has been poorly estimated in the literature which is also the case in this paper, and is confirmed by the relatively high standard deviation. Hence, we will not arrive at conclusions from this estimate except to observe that, using a conventional confidence interval, this parameter is positive which suggests that investors expect a risk premium to hold the index.

Table 3
Parameter Estimates Using S&P500 Returns Data, 1987-2004.

Parameters	Heston		CEV		CEVJ	
	Estimate	Standard Error	Estimate	Standard Error	Estimate	Standard Error
κ	4.5381	0.5295	2.1752	0.6102	1.4200	0.5914
θ	0.0302	0.0021	0.0417	0.0080	0.0350	0.0199
σ	0.4184	0.0103	2.2131	0.2723	8.8611	1.0256
ρ	-0.5546	0.0290	-0.6676	0.0269	-0.6016	0.0911
λ_S	3.1879	1.4289	4.1721	1.7687	8.4424	2.5400
β			0.9300	0.0425	1.3378	0.0371
μ_J					-6.0672	4.3123
σ_J					0.2653	9.1210
λ_J					2.2024	0.5338
Log Likelihood	15,780.37		15,818.15		15,834.55	
Annualized volatility (%)	17.39		20.41		18.71	
Daily persistence (%)	98.20		99.14		99.44	

Notes: We estimate the Heston, the CEV and CEVJ models using daily S&P500 returns from January 2, 1987 to December 31, 2004. Columns 1, 3 and 5 contain the parameter estimates for the Heston, the CEV and the CEVJ models respectively. Columns 2, 4 and 6 contain their corresponding standard deviation.

Column 5 of Table 3 contains the estimates for the CEVJ model. It is apparent that, when we add jumps to the model, the coefficient of nonlinearity rises even higher and changes from 0.93 in the CEV model to 1.34 in the CEVJ model. Our result posits that the inclusion of jumps does not rule out nonlinearity and directly contradicts the findings of Chacko and Viceira (2003). Therefore, evaluation of the CEVJ model using options data is aimed at investigating the robustness of the results obtained using daily S&P500 returns.

We also find a slightly higher persistence of approximately 99.44% in line with reports in most of the stochastic volatility literature. The unconditional volatility drops to around 18.71%; this implies that the data become less demanding on this parameter in the presence of jumps. The parameter determining the volatility of volatility is higher and the correlation is of the same order of magnitude as in the CEV model, although somewhat lower.

Turning now to the jump process parameters, we find that the jump size has a negative mean of around -2.41% daily and that the jump intensity is very small, at around 2.2 jumps per year. This low intensity confirms the infrequent occurrence of jumps in the financial data. Figure II-A presents the estimated jump sizes and jump probabilities using the estimates of the CEVJ model from Table 3. See

Johannes, Polson and Stroud (2006) for details on how to estimate the jump sizes and jump probabilities using the SIR particle filter. The top panel clearly displays the large negative drop in returns that occurred in October 1987. The middle- and bottom panels show that the particle filter is able to detect this jump. Overall, we conclude from Figure II-A that almost all jumps are of negative size and that jumps are very infrequent. In fact, adopting the yardstick of Johannes et al. (2006) that a jump is present if its estimated probability is greater than 0.5, we count a mere 12 jumps in 18 years.

Figure II-A
Estimated Jump Sizes and Probabilities Using the SIR Particle Filter for Historical Returns Data 1987-2004

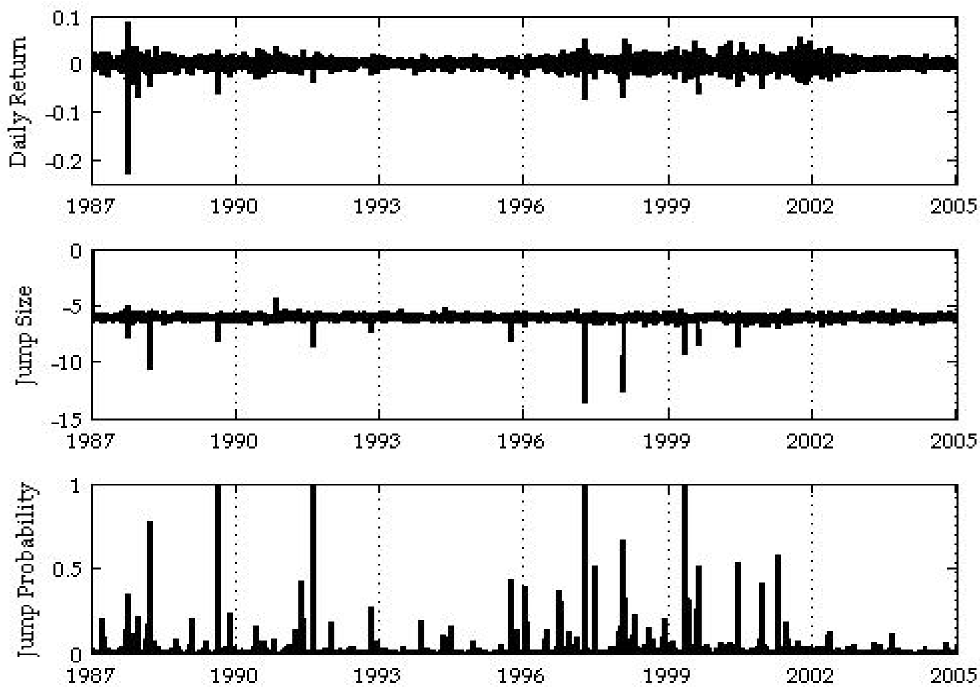

Notes to Figure: The top panel plots the daily S&P500 returns for the period January 2, 1987 to December 31, 2004. The middle panel plots the estimated jump sizes obtained using the particle filter. Finally, the bottom panel represents the jump probabilities obtained by applying the same particle filter. The middle and bottom panels are obtained using the returns estimates in Table 3.

Table 4 contains the estimates of parameters for the three models when the 1987 crash and the three subsequent years are not included. Comparing the CEV and CEVJ models, it is noticed that, even with this set of returns data, the Heston (1993) model is rejected in favor of a nonlinear specification. All the other parameter estimates move in the expected direction. In fact, we obtain lower persistence, lower long run volatility, and lower level for the nonlinearity coefficient. We obtain almost the same correlation as in the 1987-2004 sample. The jump size and jump intensity are remarkably lower than when the 1987 crash is included. Figure II-B highlights the extent by which the estimated jump sizes and probabilities are smaller for the sample 1990-2004, indicating the nontrivial impacts of excluding the 1987 crash on the estimates of the jump process parameters.

For reference, we may compare our estimates to the existing results in the literature. Indeed, the value of the mean reversion parameter is similar not only to the value obtained by Ait-Sahalia and Kimmel (2007) using the VIX index as a proxy for the daily spot volatility, but also to the value obtained by Jones (2003) which is around 4. We should stress, however, that the results in Table 3 are not directly comparable to their findings.

Table 4
Parameter Estimates Using S&P500 Returns Data, 1990-2004.

Parameters	Heston		CEV		CEVJ	
	Estimate	Standard Error	Estimate	Standard Error	Estimate	Standard Error
κ	3.3857	0.5442	2.5818	0.9208	2.6070	0.9116
θ	0.0254	0.0026	0.0287	0.0056	0.0268	0.0135
σ	0.3097	0.0156	1.8667	0.5322	2.0274	0.4992
ρ	-0.6561	0.0337	-0.6739	0.0338	-0.6885	0.0828
λ_S	3.4885	2.1915	3.2887	2.1914	3.1763	3.2103
β			0.9283	0.0843	0.9488	0.0798
μ_J					-2.0574	3.2133
σ_J					1.5252	0.7954
λ_J					1.0215	0.5623
Log Likelihood	13,417.08		13,433.60		13,434.92	
Annualized volatility (%)	15.94		16.95		16.36	
Daily persistence (%)	98.66		98.98		98.97	

Notes: We estimate the Heston, the CEV and CEVJ models using daily S&P500 returns from January 2, 1990 to December 31, 2004. Columns 1, 3 and 5 contain the parameter estimates for the Heston, the CEV and the CEVJ models respectively. Columns 2, 4 and 6 contain their corresponding standard deviation.

Figure II-B
Estimated Jump Sizes and Probabilities Using the SIR Particle Filter for Historical Returns Data 1990-2004

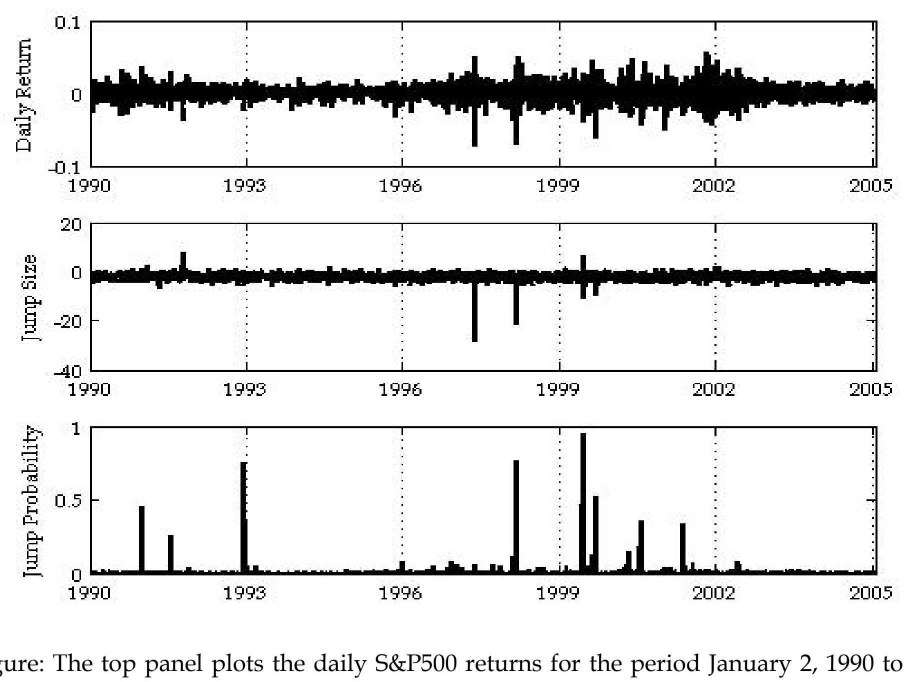

Notes to Figure: The top panel plots the daily S&P500 returns for the period January 2, 1990 to December 31, 2004. The middle panel plots the estimated jump sizes obtained using the particle filter. Finally, the bottom panel represents the jump probabilities obtained by applying the same particle filter. The middle and bottom panels are obtained using the return estimates in Table 4.

Figure III
Daily Log likelihood (L) Difference: $\mathbf{L}_t\,\mathbf{(CEVJ)} - \mathbf{L}_t\,\mathbf{(CEV)}$

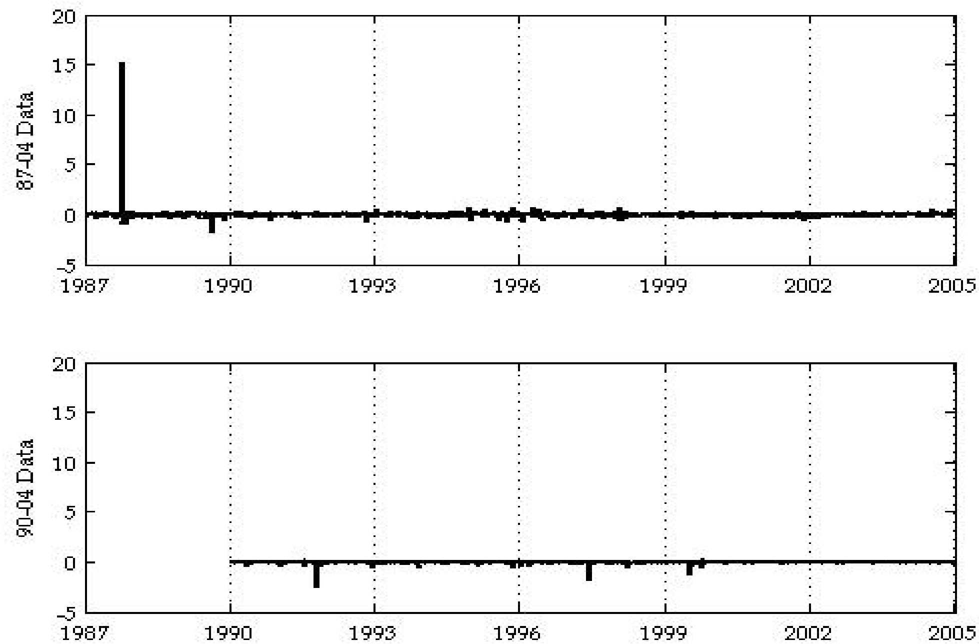

Notes to Figure: We plot the difference in Log likelihood observation by observation for the periods January 2, 1987 to December 31, 2004, and January 2, 1990 to December 31, 2004. The difference represents the Log likelihood of the CEVJ model less the Log likelihood of the CEV model.

-itude. However, a closer look at the top panel of Figure III, which plots the difference between daily Log likelihood over the period January 2, 1987 to December 31, 2004, reveals that the difference stems from one observation corresponding to the October 1987 crash. This result is similar to the findings of Christoffersen et al. (2006) where they compare different models for S&P500 dynamics. In fact, they find that some of the differences in Log likelihood across models vanish when one observation is removed from their sample. We conclude therefore that the difference of 16 points in Log likelihood between the CEV and the CEVJ models is fully explained by the 1987 crash. Table 4 and the bottom panel of Figure III confirm this finding since when we estimate the models excluding the 1987 crash we obtain almost the same Log likelihood.

Tables 3 and 4 show clearly that the CEV and CEVJ models have considerably better fit compared to the Heston (1993) model for the two estimation periods.

Overall, we find that returns seem to favor a nonlinear specification regardless of inclusion of the 1987 crash. We also ascertain that jumps and nonlinearity are complementary in the sense that the presence of jumps does not rule out the importance of nonlinearity.

4.2.2. Estimation Using Options Data 1990-1995

Table 5 exhibits the results of the estimation of the CEV and the CEVJ models using options. The third column in Table 5 contains the estimates for the CEV model. The coefficient of nonlinearity is slightly lower than the estimates obtained from returns, suggesting that options may require less nonlinearity. However, our estimate of the nonlinearity coefficient indicates clearly that options data rejects the linear specification.

Table 5
Parameter Estimates Using European Call Options on the S&P500 Index, 1990-1995.

Parameters	Heston		CEV		CEVJ	
	Estimate	Standard Error	Estimate	Standard Error	Estimate	Standard Error
κ	0.1970	0.0145	0.1824	0.0190	0.1784	0.0175
θ	0.0907	0.0072	0.0740	0.0064	0.0756	0.0107
σ	0.1043	0.0041	0.2107	0.0314	0.2267	0.0267
ρ	-0.8690	0.0088	-0.8604	0.0140	-0.8847	0.0184
λ_S	1.7879	3.4485	1.7735	3.8425	1.9111	11.3695
λ_V	0.0174	0.0369	0.0157	0.0504	0.0153	0.1040
β			0.8200	0.0414	0.7958	0.0376
μ_J					-4.9965	9.9098
σ_J					0.4843	16.7820
λ_J					0.8368	0.2996
μ_J^*					-5.0030	7.8746
λ_J^*					0.9864	0.2020
RMSE	1.4555		1.3833		1.3584	
Annualized volatility (%)	30.12		27.20		27.49	
Daily persistence (%)	99.92		99.93		99.93	

Notes: We estimate the models using Wednesday Options on the S&P500 Index for the period 1990 to 1995. Columns 1, 3 and 5 contain the parameter estimates for the Heston, the CEV and the CEVJ models respectively. Columns 2, 4 and 6 contain their corresponding standard deviation.

For the other parameters of the model, we notice that the speed of mean reversion is lower when we estimate the model on options compared to the estimate in Table 3. Hence, we may conclude that option data imply strong persistence in the volatility; at around 99.9% slightly higher than the persistence obtained using returns only. The correlation coefficient is approximately -0.86. The negative correlation is a standard result in the literature that we observe when we estimate stochastic volatility models on any set of data. The long-run volatility is around 27%, and is close to the results obtained with returns but somewhat higher than the volatility in Table 1. We should stress that, even though this parameter varies considerably in the stochastic volatility literature, it always falls within a reasonable interval around the sample volatility. Ait-Sahalia and Kimmel (2007), for example, find an unconditional volatility of around 21%. Eraker, Johannes and Polson (2003) find it 15%. The risk premium associated with the volatility dynamic λ_V is small and statistically not significant. The fact that many empirical papers set this parameter to zero seems, therefore, to be a realistic assumption; see, for example, Ait-Sahalia and Kimmel (2007). Next, the risk premium coefficient λ_S related to returns is quite variable but always remains positive, suggesting again that investors expect a premium for holding risky assets. Its value is in line with the estimate obtained using returns.

At this stage, a few remarks on some of the parameters values are in order. In fact, the correlation implied from options is higher than that implied from returns. This finding is confirmed by Eraker (2004) and by Christoffersen et al. (2006). Our estimate is lower than the correlation obtained in Christoffersen et al. (2006) using the same estimation technique. We believe that this paper's use of longer samples of options data permits more accurate identification of the level of correlation. However, the estimate reported in Eraker (2004) using options but a different estimation methodology yields a correlation ranging from -0.57 to -0.59, which is even lower than that of our results. But, as observed by Eraker (2004) there is no consensus in the literature on the level of this parameter. Few papers estimate the level of nonlinearity β from returns and options. Our estimate is higher than the one obtained by Ait-Sahalia and Kimmel (2007) and lower than that by Jones (2003). Column 5 of Table 5 contains the estimates of the CEVJ model. The options data confirm our findings using returns since the coefficient of nonlinearity remains almost unchanged when we add jumps to the CEV model. This result further supports the complementary nature of nonlinearity and jumps.

At this stage, a few remarks on some of the parameters values are in order. In fact, the correlation implied from options is higher than that implied from returns. This finding is confirmed by Eraker (2004) and by Christoffersen et al. (2006). Our estimate is lower than the correlation obtained in Christoffersen et al. (2006) using the same estimation technique. We believe that this paper's use of longer samples of options data permits more accurate identification of the level of correlation.

However, the estimate reported in Eraker (2004) using options but a different estimation methodology yields a correlation ranging from -0.57 to -0.59, which is even lower than that of our results. But, as observed by Eraker (2004) there is no consensus in the literature on the level of this parameter. Few papers estimate the level of nonlinearity β from returns and options. Our estimate is higher than the one obtained by Ait-Sahalia and Kimmel (2007) and lower than that by Jones (2003). Column 5 of Table 5 contains the estimates of the CEVJ model. The options data confirm our findings using returns since the coefficient of nonlinearity remains almost unchanged when we add jumps to the CEV model. This result further supports the complementary nature of nonlinearity and jumps.

Turning now to the other model parameters we see that the inclusion of jumps increases the persistence. This result is in line with the findings of Eraker (2004) in the context of linear models. What is surprising is that the unconditional estimate of the variance is higher than in the CEV model. In fact, we expect that the inclusion of jumps will lower the unconditional variance since the data becomes less demanding on this parameter in the presence of jumps. However, as we are going to see later in this paper, jumps do not add much to the model in terms of improving the fit. The parameter σ and the coefficient of nonlinearity are in the same order of magnitude as in the CEV model. Finally, jumps have a large negative mean around -1.98% daily and are very infrequent at around 0.84 jumps per year. As we pointed out, including jumps does not improve the model fit. In fact, the RMSE declines from 1.38 to 1.36, which cannot be considered a large benefit. This result confirms the findings of Bates (2000) and Eraker (2004) but contradicts those of Broadie et al. (2007). However, the results of Broadie et al. (2007) are not directly comparable with our results. First, they use a linear specification for the volatility process, whereas we use a nonlinear specification. Second, their model parameterization allows all the parameters to have a risk premium and, therefore, differ under objective and risk-neutral measures, whereas we use a much more parsimonious specification. Finally, the options data they use and the periods they cover differ from those we use in our sample.

Figure IV elaborates on the potential reasons for the similarity in performances between the CEV and CEVJ models. The top panel of Figure IV shows that the weekly RMSE from the CEVJ- and the CEV models are almost indistinguishable. The bottom panel investigates further the difference between the RMSEs obtained from the two models. An obvious pattern cannot be inferred. These findings, coupled with the results obtained from returns, stress that the difference in fit, when we add jumps to the CEV model, is very small. This is to be expected when we examine the top and bottom panels of Figure V. In fact, the first- and second columns of Figure V show that the residuals obtained from the CEV- and CEVJ models generate empirical distributions having both tails similar to the standard normal. This finding does not appear to depend on whether we estimate the model on returns or on options. Moreover, when comparing the residuals of the CEV model to those of the CEVJ we observe from the right column of Figure V that their empirical CDFs are almost the same. This similarity may explain why these models do equally well in fitting options data. Finally, Table 5 clearly shows that the fit of the CEV and CEVJ models is much better than that of the Heston (1993) model.

Figure IV
Weekly RMSE Difference: RMSE(CEVJ) – RMSE(CEV)

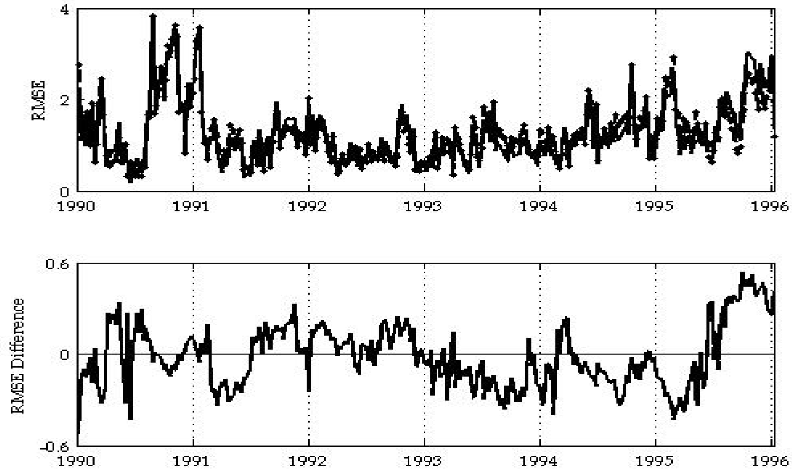

Notes: We plot the difference in weekly RMSE for the period January 2, 1990 to December 31, 1995. The difference represents the weekly RMSE of the CEVJ model less the weekly RMSE of the CEV model.

Figure V
Empirical Cumulative Distribution Function of the CEV and CEVJ Models Implied from S&P500 Returns and from ATM Options.

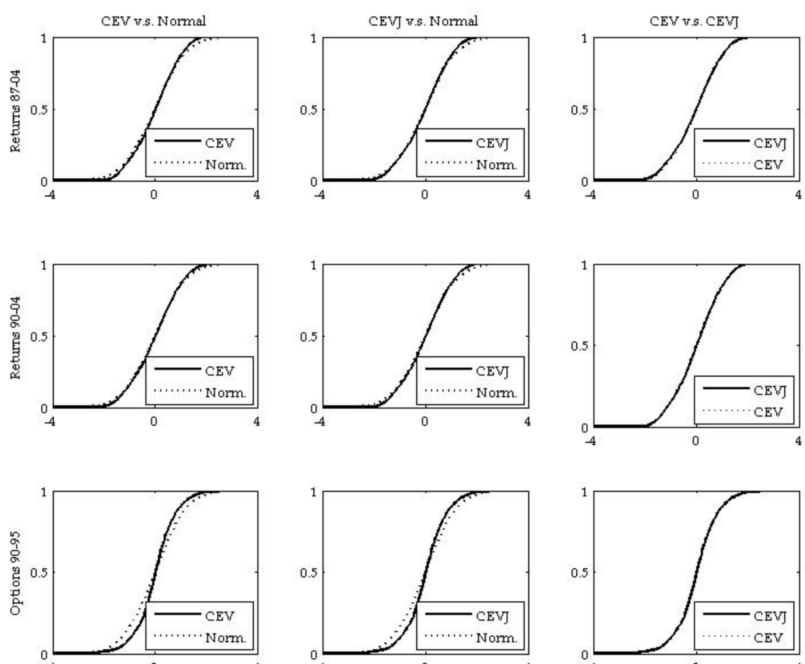

Notes: We plot the empirical CDFs for the CEV model and CEVJ model using the residuals evaluated at the optimal parameters. To plot the CDFs in the top panel, we use the estimates in Table 3 obtained from returns data for the period 87-04. To plot the CDFs in the middle panel, we use the estimates in Table 4 obtained from returns data for the period 90-04. The CDFs in the bottom panel were generated using the estimates in Table 5 obtained from ATM options.

5. Conclusion

We investigate the degree of nonlinearity implied by returns and options, and the impact of including jump processes on this parameter. We find that both returns and options data favor nonlinear specifications and that the coefficient of nonlinearity is between 0.93 and 1.34 when we use returns and between 0.80 and 0.82 when we use returns and options. Our findings are significant since they show that estimations based on returns and on returns and options are consistent. We also find that adding jumps to nonlinear models did not minimize the importance of nonlinearity in the models' specifications. Hence, nonlinearity and jumps seem to be complementary rather than competitive.

Nonlinear models are therefore good building blocks for models that include jumps. We also obtained reasonable correlation that fell within the range of what was previously documented in the literature.

Although we find in this study that adding jumps to nonlinear models did not improve the model fit, this does not imply that we should exclude them from stochastic volatility models. First, jumps are infrequent in the sense that our sample might not be rich enough in terms of volatility dynamics to reveal their importance in improving the model fit. Second, because of the computational burden, we only use ATM call options which define a moneyness interval where almost all stochastic volatility models perform the best. Including a full cross-section of options data together with the time series dimension might lead to more favorable results for jump processes.

Finally, while the CEV- and the CEVJ models are certainly better models in-sample compared to the typical linear model, the implications of including additional parameters for the out-of-sample performance of these models are not obvious and should be studied in future work.

6. Appendix

6.1. Appendix 1: The SIR Particle Filter (PF) of the CEV model

We illustrate the implementation of the particle filter technique in the context of the CEV model in which the Euler discretization is given by

$$\log(S_{t+\Delta}) = \log(S_t) + \left(r + \lambda_S \left(1-\rho^2\right) V_t - \frac{1}{2} V_t \right)\Delta + \sqrt{V_t \Delta}\,\varepsilon_{1,t+\Delta} \tag{13}$$

$$V_{t+\Delta} = \kappa(\theta - V_t)\Delta + \sigma\sqrt{\Delta} V_t^{\beta}\left(\rho\varepsilon_{1,t+\Delta} + \sqrt{1-\rho^2}\,\varepsilon_{2,t+\Delta}\right).$$

Filtering the state variable consists of the following 3 steps.

6.1.1. Step 1: Simulating the state forward: Sampling

This is done by computing $V_{t+\Delta}^j$ from the original set of particles $\{V_t^j\}_{j=1}^N$ assumed to be known at time t using equation (13) and taking the correlation into account.[3] We have

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

where $\mu_t = r + \lambda_S\left(1-\rho^2\right) V_t^j$.

The above equation gives

[3]We initialize the variance in the first period to equal the model-implied unconditional variance, that is, $V_0^j = \theta$, for all j. In the MLIS estimation, $t = 0$ is simply the first day of observed returns that is January 2, 1987 for the first sample, and January 2, 1990 for the second sample. In the NLS estimation, $t = 0$ is January 2, 1989 corresponding to one year prior to the first available option quote.

$$\varepsilon^{j}_{1,t+\Delta} = \frac{\ln\left(\frac{S_{t+\Delta}}{S_t}\right) - \left(\mu_t - \frac{1}{2}V_t^j\right)\Delta}{\sqrt{V_t^j \Delta}}.$$

Since

$$w^{j}_{t+\Delta} = \rho\varepsilon^{j}_{1,t+\Delta} + \sqrt{1-\rho^2}\,\varepsilon^{j}_{2,t+\Delta},$$

where $corr(\varepsilon^{j}_{1,t+\Delta}, \varepsilon^{j}_{2,t+\Delta}) = 0$, we get

$$V^{j}_{t+\Delta} = V_t^j + \kappa\left(\theta - V_t^j\right)\Delta + \sigma V_t^{j\,\beta}\sqrt{\Delta}\left(\rho\frac{\ln\left(\frac{S_{t+\Delta}}{S_t}\right) - \left(\mu_t - \frac{1}{2}V_t^j\right)\Delta}{\sqrt{V_t^j\Delta}} + \sqrt{1-\rho^2}\,\varepsilon^{j}_{2,t+\Delta}\right).$$

We simulate N particles which describe the set of possible values of $V_{t+\Delta}$.

6.1.2. Step 2: Computing and normalizing the weights: Importance Sampling

At this point, we have a vector of N possible values of $V_{t+\Delta}$ and we know, according to equation (13), that, given the other available information, $V_{t+\Delta}$ is sufficient to generate $\ln(S_{t+2\Delta})$. Therefore, equation (13) offers a simple way to evaluate the likelihood that the observation $S_{t+2\Delta}$ was generated by $V_{t+\Delta}$. Hence, we compute the weight assigned to each particle (or the likelihood or probability that the particle has generated $S_{t+2\Delta}$). The likelihood is computed as follows:

$$W^{j}_{t+\Delta} = \frac{1}{\sqrt{V^{j}_{t+\Delta}\Delta}}\exp\left(-\frac{1}{2}\frac{\left(\ln\left(\frac{S_{t+2\Delta}}{S_{t+\Delta}}\right) - \left(\mu_t - \frac{1}{2}V^{j}_{t+\Delta}\right)\Delta\right)^2}{V^{j}_{t+\Delta}}\right)$$

for $j = 1,..,N$. Finally, because nothing guarantees that $\sum_{j=1}^{N} W^{j}_{t+\Delta} = 1$, we have to normalize and set

$$W^{j}_{t+\Delta} = \frac{W^{j}_{t+\Delta}}{\sum_{j=1}^{N} W^{j}_{t+\Delta}}.$$

6.1.3. Step 3: Resampling

This step is necessary for propagating high probability particles often and vice versa. We use a simple technique to resample the particles, which eliminates the particles of low probability and replicates those of high probability. Therefore, we construct a set of integer variables $\{\iota^{j}_{t+\Delta}\}_{j=1}^{N}$ which can be done in different ways. Our method uses the resampling scheme proposed by Pitt (2002) that allows us to obtain a smooth objective function in the parameters' space.

First, the adjusted weights obtained in Step 2, $W^{j}_{t+\Delta}$, are mapped into a set of integer variables $\{\iota^{j}_{t+\Delta}\}_{j=1}^{N}$, using an algorithm that considers the weights that are not multiples of $1/N$. This algorithm is based on the empirical CDF of V and smoothed using linear interpolation as suggested by Pitt (2002). The smoothing enables gradient-based optimization and the computation of standard errors using conventional first-order techniques.

Next, we construct the new set of particles $\{V(\iota)^{j}_{t+\Delta}\}_{j=1}^{N}$ by replicating each particle in the original set $\{V^{j}_{t+\Delta}\}_{j=1}^{N}$ $\iota^{j}_{t+\Delta}$ times. Therefore, the particles in the original set are either eliminated, or included once or several times according to their adjusted weights $\{W^{j}_{t+\Delta}\}_{j=1}^{N}$. The greater the weight, $W^{j}_{t+\Delta}$, the higher the integer variable $\iota^j_{t+\Delta}$, and the more often the original particle $V^j_{t+\Delta}$ is included in the resampled set $\{V(t)^j_{t+\Delta}\}^N_{j=1}$.

We now have a new set of N particles and weights $\{V(t)^j_{t+\Delta}, V(t)^j_{t+\Delta}\}^N_{j=1}$ which are implicitly functions of the variable $\iota_{t+\Delta}$ and all of which have weights $1/N$. We are thus ready to return to Step 1 to move the filter forward.

6.2. Appendix 2: Adaptation of the PF to the CEVJ model

Note that the jumps in equation (1) create further discontinuities in the objective function besides those generated by the particle filter. One possible solution to this problem is to approximate the density of returns by the following expression

$$f(R_t \mid V_{t-1}) = \sum_{x=0}^{\infty} N\left(R_t \mid x\mu_J, \int_t^{t+\Delta} V_{t-1}\Delta + x\sigma_J^2\right)\frac{(\Delta\lambda_J)^x e^{-\Delta\lambda_J}}{x!}. \quad (14)$$

Proof of this approximation result is:

$$R_t = \log(S(t+\Delta)) - \log(S(t))$$
$$= \left(\mu_t - \frac{1}{2}V_t - \lambda_J \overline{\mu}_J\right)\Delta + \sqrt{V_t}\,dB(t) + \sum_{x=1}^{N(t+\Delta)-N(t)} J_x,$$

where $N(t+\Delta) - N(t) \to Poisson\,(\int_t^{t+\Delta}\lambda_J(u)du)$. Assuming that the jump intensity λ_J is constant, $N(t+\Delta) - N(t) \to Poisson\,(\Delta\lambda_J)$ and that each jump $J_x \to N(\mu_J, \sigma_J^2)$, then we may write

$$R_t = \left(\mu_t - \frac{1}{2}V_{t-1} - \lambda_J \overline{\mu}_J\right)\Delta + \sqrt{V_{t-1}\Delta}\,\varepsilon_t + \sum_{x=1}^{Poisson(\Delta\lambda_J)} J_x$$

Then we have

$$f(R_t \mid V_{t-1}) = \sum_{x=0}^{\infty} f(R_t \mid x, V_{t-1})\Pr(x)$$
$$= \sum_{x=0}^{\infty} N(R_t \mid x\mu_J, V_{t-1}\Delta + x\sigma_J^2)\Pr(x)$$
$$= \sum_{x=0}^{\infty} N(R_t \mid x\mu_J, \int_t^{t+\Delta} V_{t-1}\Delta + x\sigma_J^2))\,\frac{(\Delta\lambda_J)^x e^{-\Delta\lambda_J}}{x!}.$$

This converges quickly (we can normally ignore terms in excess of three or four terms i.e. x> 4). Consequently, we have the form $f(R_t \mid V_{r-1})$, which is more heavy-tailed than the Gaussian as it is a combination.

This form of the density given by equation (14) enables us to do smooth resampling as was previously carried out in the filtering algorithm (see step 3 in Appendix 1). Note that if the density is not written in the above form, then the optimization using the particle-filtering technique will be infeasible.

The Euler discretization of the model after applying the density approximation is shown to be

$$\log(S_{t+\Delta}) = \log(S_t) + \left(\mu_t - \frac{1}{2}V_t - \lambda_J \overline{\mu}_J + \sum_{x=0}^{3} x\Pr(J_x = x)\mu_J\right)\Delta + \sqrt{\left(V_t + \sum_{x=0}^{3} x\Pr(J_x = x)\sigma_J^2\right)\Delta}\,z_{1,t+\Delta} \quad (15)$$

$$V_{t+\Delta} = \kappa(\theta - V_t)\Delta + \sigma\sqrt{\Delta}V_t^{\beta}\left(\rho z_{1,t+\Delta} + \sqrt{1-\rho^2}\,\varepsilon_{2,t+\Delta}\right),$$

where $z_{1,t+\Delta} \to N(0,1)$, $corr\,(z_{1t}, \varepsilon_{2t}) = 0$ and $\Pr(J_x = x) = \frac{(\Delta\lambda_J)^x e^{-\Delta\lambda_J}}{x!}$. We then proceed with the filtering exercise exactly as was done with the CEV model.

References

Ait-Sahalia, Y., Kimmel, R., 2007, Maximum Likelihood Estimation of Stochastic Volatility Models. Journal of Financial Economics, 83, 413-452.

Bakshi, G., Cao, C., Chen, Z., 1997, Empirical Performance of Alternative Option Pricing Models. Journal of Finance, 52, 2003-2049.

Bates, D.S., 2000, Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics, 94, 181–238.

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

Benzoni, L., 2002, Pricing Options Under Stochastic Volatility: An Empirical Investigation. Working Paper, University of Minnesota.

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

Broadie, M., Chernov, M., Johannes, M., 2007, Model Specification and Risk Premiums: Evidence from Futures Options. Journal of Finance, 62, 1453-1490.

Chacko, G., Viceira, L.M., 2003, Spectral GMM estimation of continuous-time processes. Journal of Econometrics, 116, 259–292.

Chernov, M., Ghysels, E., 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.

Christoffersen, P., Jacobs, K., Mimouni, K., 2006, An Empirical Comparison of Affine and Non-Affine models For Equity Index Options. Working Paper, McGill University.

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

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

Gallant, A.R., 1997, An Introduction to Econometric Theory. Princeton University Press, Princeton.

Geweke, J., 1989, Bayesian Inference in Econometric Models Using Monte-Carlo Integration. Econometrica, 57, 1317-1339.

Gordon, N.J., Salmond, D.J., Smith, A.F.M., 1993, A Novel Approach to nonlinear/non-Gaussian Bayesian state estimation. IEE-Proceedings F, 140, 107-133.

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.

Johannes, M., Polson, N., 2006, MCMC methods for Financial Econometrics. The Handbook of Financial Econometrics, edited by Yacine Ait-Sahalia and Lars Hansen.

Jones, C.S., 2003, The dynamics of stochastic volatility: evidence from underlying and options markets. Journal of Econometrics, 116, 181–224.

Koopman, S. J., Creal, D., Shephard, N., 2009, Testing the Assumptions Behind Importance Sampling. Journal of Econometrics, 149, 2-11.

Lewis, A.L., 2000, Option Valuation under Stochastic Volatility. Finance Press, Newport Beach, CA.

Li, M., Pearson, N.D., Poteshman, A.M., 2004, Conditional Estimation of Diffusion Processes. Journal of Financial Economics, 74, 31-66.

Pitt, M.K., 2002, Smooth Particle Filters for Likelihood Evaluation and Maximization. Working Paper, University of Warwick.