Estimation of the CEV and the CEVJ Models on Returns and Options

Karim Mimouni

Desautels Faculty of Management, McGill University*

November 13, 2006

Abstract

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; nonlinear; jumps.

*I am grateful to IFM^2 and FQRSC for financial support.

Correspondence to: Karim Mimouni, Desautels Faculty of Management, McGill University, 1001 Sherbrooke Street West, Montreal, Quebec, Canada, H3A 1G5; Tel: (514) 398-4000; E-mail: karim.mimouni@mcgill.ca.

1 Introduction

The performance of any option pricing model is measured by its ability to fit far from at-the-money prices. In that respect, a convenient way to evaluate the performance of different processes is to compute the model-implied Black and Scholes volatilities and compare them to their market counterparts retrieved from options data. In fact, 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 also as the smile. However, despite many attempts in the literature to find a model which replicates the smile, there has been mitigated success in fitting deep in-and out-of-the-money option prices. 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 available 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 endeavor 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 diffusion term because they allow for closed-form solutions for option prices and facilitate empirical implementations.[1] Nonlinear models, in particular the CEV and CEVJ models, were not the preferred choice for empirical investigations. Therefore, few empirical studies implement them relative to the extensive literature on linear models. However, 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 as a function of the spot volatility.[2] These new findings suggest the use of nonlinear models as building blocks to explore better specifications which is the primary concern of this paper.

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.[3] This paper makes two main contributions. First, we perform the estimation of the CEV model on S&P500 returns and on a multiple cross-section of European call options. This options data set is richer than the one used by Jones (2003) and Ait-Sahalia and Kimmel (2006) as it allows the nonlinearity coefficient to be estimated using options with different maturities. Our empirical implementation

[1] The most popular empirical implementations include the original version of Heston (1993) or by adding jumps in returns and volatility. See, for example, Bakshi, Cao and Chen (1997), Chernov and Ghysels (2000), Pan (2002), Eraker (2004), and Eraker, Johannes and Polson (2003).

[2] See Christoffersen, Jacobs and Mimouni (2006), Ait-Sahalia and Kimmel (2006), and Jones (2003).

[3] See Jones (2003) and Ait-Sahalia and Kimmel (2006) for empirical implementations of the CEV model. See Chacko and Viceira (2003) for empirical implementation of the CEVJ model using returns only.

builds on the particle-filtering technique in order to conduct a fair comparison between the estimates obtained using returns and those obtained using options. Specifically, we aim to determine 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 allowing for jumps on the degree of nonlinearity and on the model fit. To our knowledge, Chacko and Viceira (2003) are unique in including jumps in the nonlinear context using 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 challenge posed estimating the model by Monte Carlo simulations. One advantage of our estimation methodology is that it allows the latent variable to be easily extracted and, hence, reduces the estimation burden of a CEVJ model using options.

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 was the case in the findings of Chacko and Viceira (2003).

We use two different sources of data to estimate the models. First, we estimate the parameters on S&P500 returns only. Although stochastic volatility models were motivated by fitting options, estimation of these models on returns only is very common because we typically want to be able to use the estimates obtained from returns to price options. Second, we use a combination of daily returns and at-the-money European call options in order to estimate the model parameters.

To estimate the model on returns, we use the technique presented 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 used to estimate the volatility path, and hence it uses the true dynamic of returns to compute the approximate likelihood. This method is attributed to the Simulated Maximum Likelihood techniques. The Simulated Maximum Likelihood and Quasi-Maximum Likelihood approaches have been used extensively in the context of stochastic volatility mainly because the inclusion of spot volatility as a state variable considerably increases the dimension of the true likelihood and, therefore, does not yield a closed-form solution. However, most of these techniques are difficult to implement and computationally intensive. We propose therefore to use the particle-filtering technique to estimate the CEV and CEVJ models. The advantage of this method is that it allows the models to be estimated using returns only, and using returns and options. This ensures consistency in methodology and allows for comparison between the parameter estimates. In addition, this method is easy to implement empirically.

To estimate the CEV and CEVJ models on options, we use the methodology presented by Christoffersen et al. (2006) based on an iterative Nonlinear Least Squares (NLS) procedure. This

NLS optimization technique has not been used in the context of the CEV and CEVJ models because, since neither model admits a closed-form solution, option prices have to be computed by Monte Carlo simulations. This adds considerably to the computational complexity of the estimation.

In both estimations, the volatility path is filtered from daily returns using the particle filter technique of Gordon et al. (1993) which is suitable for nonlinear state space applications.

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 volatility conditional densities. It then exposes 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

$$\begin{aligned} d\log(S_t) &= \left(r + \lambda_1\left(1-\rho^2\right)V_t - \frac{1}{2}V_t - \lambda_J\mu_J\right)dt + \sqrt{V_t}dB_{1t} + J_t dN_t \\ dV_t &= \kappa\left(\theta - V_t\right)dt + \sigma V_t^{\beta}\left(\rho dB_{1t} + \sqrt{1-\rho^2}dB_{2t}\right), \end{aligned} \tag{1}$$

where S_t is the price of the underlying asset, V_t is the volatility, 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. σ determines the level of the volatility of volatility, λ_1 determines the risk premium required to compensate investors for holding the underlying asset, and ρ represents the correlation between returns and volatility 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, $J_t \sim Poisson(\lambda_J)$, and $dN_t \sim N\left(\mu_J, \sigma_J^2\right)$.

In this paper, we consider jumps only in returns. The effects of jumps in the volatility dynamic are left for future research.[4]

When $J_t dN_t = 0$ in the CEVJ model, the jump component vanishes and we obtain the CEV model. Hence, the CEV model is defined by the following two equations

[4] Eraker (2004) estimates a model with correlated jumps in returns and volatility. Because this model is not parsimonious and because there is no empirical evidence on the role of jumps in volatility, we do not include the model for estimation.

$$
\begin{aligned}
d\log(S_t) &= \left(r+\lambda_1\left(1-\rho^2\right)V_t-\frac{1}{2}V_t\right)dt+\sqrt{V_t}dB_{1t} \qquad (2)\\
dV_t &= \kappa\left(\theta-V_t\right)dt+\sigma V_t^{\beta}\left(\rho dB_{1t}+\sqrt{1-\rho^2}dB_{2t}\right).
\end{aligned}
$$

Models (1) and (2) share the same expression for the volatility 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.[5] 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.[6] 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.

The most general representation of the Euler disretization of any stochastic volatility model is given by

$$
\begin{aligned}
S(t) &= F\left(S_{1:(t-1)},V_{1:(t-1)},\xi,\varepsilon_t^S\right) \qquad (3)\\
V(t) &= G(V_{1:(t-1)},\xi,\varepsilon_t^V), \qquad (4)
\end{aligned}
$$

where F and G are some given nonlinear functions of the previous states, ξ is the vector of model parameters, and ε_t^S and ε_t^V are two standard normal distributions.

Equation (3) represents the dynamic of the observed measurement, whereas equation (4) represents the dynamic of the state that is usually unobserved. The filtering problem arises when estimating sequentially the new state (the volatility) using the history of returns. Therefore, solving the sequential filtering problem is equivalent to finding a way to sample from the true posterior

[5] Eraker (2004) reports an improvement in fit of 1%. Bates (2000) finds an improvement of circa 2%. Broadie, Chernov and Johannes (2006) find a 50% improvement in fit by adding jumps in prices to the SV model.

[6] For more details on the particle-filtering technique and its applications in the context of the estimation on returns and options, see Christoffersen, Jacobs, and Mimouni (2006).

density of the state, $p\left(V(t)|S\right)$. Unfortunately, in most applications, it is impossible to sample directly from the posterior density $p\left(V(t)|S\right)$. We need therefore to define an approximate density $q\left(V(t)|S\right)$ that allows for easy sampling, and that plays the role of proposal density.

Typically, we need to evaluate the mean of the true distribution $p\left(V(t)|S\right)$ or some confidence interval around the mean. This is easily achieved in the context of particle filtering by noting that

$$
\begin{aligned}
E\left(V(t)\right) &= \int V(t) p\left(V(t)|S\right) dV(t) \\
&= \int V(t) \frac{p\left(V(t)|S\right)}{q\left(V(t)|S\right)} q\left(V(t)|S\right) dV(t) \\
&= \int V(t) \frac{\frac{p(S|V(t))p(V(t))}{p(S)}}{q\left(S|V(t)\right)} q\left(V(t)|S\right) dV(t) \\
&= \frac{1}{p(S)} \int V(t) w_t\left(V(t)\right) q\left(V(t)|S\right) dV(t),
\end{aligned} \tag{5}
$$

where $w_t\left(V(t)\right) = \frac{p(S|V(t))p(V(t))}{q(S|V(t))}$.

$$
\begin{aligned}
p(S) &= \int p(S|V(t)) p(V(t)) dV(t) \\
&= \int \frac{p(S|V(t))p(V(t))}{q\left(S|V(t)\right)} q\left(S|V(t)\right) dV(t) \\
&= \int w_t\left(V(t)\right) q\left(S|V(t)\right) dV(t).
\end{aligned} \tag{6}
$$

Combining (5) and (6) we conclude that

$$
E\left(V(t)\right) = \frac{E_q\left(V(t) w_t\left(V(t)\right)\right)}{E_q\left(w_t\left(V(t)\right)\right)}. \tag{7}
$$

Equation (7) describes how it is possible to compute indirectly the mean of the distribution of $V(t)$ by drawing a sample of size N from the proposal and computing the expectation defined in (7) as follows

$$
E_d\left(V(t)\right) = \frac{\frac{1}{N}\sum_{k=1}^{N} V^k(t) w_t\left(V^k(t)\right)}{\frac{1}{N}\sum_{k=1}^{N} w_t\left(V^k(t)\right)}, \tag{8}
$$

where $V^k(t)$ is a draw from $q\left(V(t)|S\right)$.

It can be shown that $E_d(V(t)) \to E(V(t))$ as $N \to \infty$. Note that

$$E(V(t)) \simeq E_d(V(t)) = \sum_{k=1}^{N} V^k(t)\Psi(V^k(t)), \tag{9}$$

where

$$\Psi(V^k(t)) = \Psi^k = \frac{w_t\left(V^k(t)\right)}{\sum_{k=1}^{N} w_t\left(V^k(t)\right)}. \tag{10}$$

We can see from (10) that $\sum_{k=1}^{N}\Psi^k = 1$. Hence, the sample $\left\{V^k\right\}_{k=1}^{N}$ with the corresponding weights $\left\{\Psi^k\right\}_{k=1}^{N}$ can be considered as a draw from $p(V(t)|S)$. Therefore, $\left\{V^k, \Psi^k\right\}_{k=1}^{N}$ converges to the true posterior density $p(V(t)|S)$ as N tends to infinity.

Equation (9) shows how we can utilize the history of returns and the filtered states to obtain a sample from $p(V(t)|S)$. First, we need to define an importance density q. The choice of $q(V(t)|S)$ is crucial for the performance of the particle filter. Next, we draw a sample of size N from $q(V(t)|S)$ and assign weights to each particle given by (10). Hence, $\left\{V^k, \Psi^k\right\}_{k=1}^{N}$ is an approximate draw from $p(V(t)|S)$. We can then compute the mean of the distribution using equation (9). These steps are repeated for $t = 1, .., T$. A description of the particle filter applied to the cases of the CEV and CEVJ models is provided in the appendix.

3.1 Model Estimation Using Returns

We now examine the Importance Sampling Maximum Likelihood (ISML) approach of estimating the CEV and CEVJ models. Pioneered by Pitt (2002), the method allows us to compute an approximate likelihood when the state is unobserved. Not only does the technique apply to general models, but it does not require efforts or extra computations to switch from one model to another. It is fully consistent with the returns dynamic since, in this setup the volatility 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

$$p(S(t)|\theta, S_{1:t-1}, V_{1:t-1}) = \int p(S(t)|\theta, V(t))\, p(V(t)|\theta, S_{1:t-1}, V_{1:t-1})\, dV(t). \tag{11}$$

This likelihood can be approximated by

$$\widehat{p}(S(t)|\theta, S_{1:t-1}, V_{1:t-1}) = \frac{1}{N}\sum_{k=1}^{N}\varpi_t^k, \tag{12}$$

where $\{\varpi_k\}_{k=1}^N$ represent the unnormalized weights obtained from the particle filter. It can be shown by applying the Kolmogorov's strong law of large numbers that $\widehat{p}(S(t)|\theta, S_{1:t-1}, V_{1:t-1}) \overset{a.s.}{\rightarrow} p(S(t)|\theta, S_{1:t-1}, V_{1:t-1})$ as N tends to infinity.[7] Hence, the computation of the Log likelihood is a by-product of the particle filter, and extra computation is not incurred.

The objective function to be maximized is therefore given by

$$L_{PF} = \sum_{t=2}^{T} \ln\left(\frac{1}{N}\sum_{k=1}^{N} \varpi_t^k\right). \tag{13}$$

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 (13). Third, the optimizer proposes a new set of parameters and the procedure restarts until the objective function (13) is maximized.

3.2 Model Estimation Using Options

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

$$\begin{aligned} d\log(S_t) &= \left(r - \frac{1}{2}V_t - \lambda_J^*\mu_J^*\right)dt + \sqrt{V_t}dB_{1t}^* + J_t^* dN_t^* \\ dV_t &= \kappa^*\left(\theta^* - V_t\right)dt + \sigma V_t^\beta \left(\rho dB_{1t}^* + \sqrt{1-\rho^2}dB_{2t}^*\right), \end{aligned} \tag{14}$$

Where dB_{1t}^* and dB_{2t}^* are two uncorrelated standard Brownian motions under the risk-neutral measure Q, $J_t^* \sim Poisson(\lambda_J^*)$, and $dN_t^* \sim N\left(\mu_J^*, \sigma_J^2\right)$ under Q. $\kappa^* = \kappa + \lambda$ and $\theta^* = \frac{\kappa\theta}{\kappa+\lambda}$. λ_J^* and μ_J^* are allowed to be different from λ_J and μ_J under the risk-neutral measure. Note that we have assumed that the volatility risk premium is linear in V_t.

Discretizing equation (14) using the Euler discretization yields

$$\begin{aligned} \log(S_{t+\Delta}) &= \log(S_t) + \left(r - \frac{1}{2}V_t - \lambda_J^*\mu_J^*\right)\Delta + \sqrt{V_t\Delta}\varepsilon_{1,t+\Delta}^* + J_{t+\Delta}^* N_{t+\Delta}^* \\ V_{t+\Delta} &= \kappa^*\left(\theta^* - V_t\right)\Delta + \sigma\sqrt{\Delta}V_t^\beta\left(\rho\varepsilon_{1,t+\Delta}^* + \sqrt{1-\rho^2}\varepsilon_{2,t+\Delta}^*\right), \end{aligned} \tag{15}$$

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

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

[7] See Gallant (1997) and Geweke (1989) for further details on Kolmogorov's strong law of large numbers.

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.

$$SSE = \sum_{t=1}^{T} \sum_{i=1}^{N_t} \left(C_{t,i}^{Model} - C_{t,i}^{Market} \right)^2, \tag{16}$$

where 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 $J_t^* dN_t^* = 0$. Hence, for the CEV model we have

$$\begin{aligned} d\log(S_t) &= \left(r - \frac{1}{2}V_t\right) dt + \sqrt{V_t} dB_{1t}^* \\ dV_t &= \kappa^* \left(\theta^* - V_t\right) dt + \sigma V_t^{\beta} \left(\rho dB_{1t}^* + \sqrt{1-\rho^2} dB_{2t}^*\right), \end{aligned} \tag{17}$$

and its discretized version gives

$$\begin{aligned} \log(S_{t+\Delta}) &= \log(S_t) + \left(r - \frac{1}{2}V_t\right) \Delta + \sqrt{V_t \Delta} \varepsilon_{1,t+\Delta}^* \\ V_{t+\Delta} &= \kappa^* \left(\theta^* - V_t\right) \Delta + \sigma \sqrt{\Delta} V_t^{\beta} \left(\rho \varepsilon_{1,t+\Delta}^* + \sqrt{1-\rho^2} \varepsilon_{2,t+\Delta}^*\right). \end{aligned} \tag{18}$$

Again, the objective function (16) is minimized using Monte Carlo simulations until the set of optimal parameters is reached.

4 Empirical Results

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

4.1 Data

To estimate the CEV and CEVJ models on returns we use two S&P500 return samples. The first sample, from January 2, 1987 to December 31, 2004, includes the 1987 crash whereas in the second sample, from January 2, 1990 to December 31, 2004, we exclude the year 1987 as well as the three subsequent years which might be indirectly affected by the extremely high volatility levels 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.

To estimate 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. Hence, we identify ATM options by checking the following condition

$$\frac{F(t,T)}{K} = \frac{S_t e^{r(T-t)}}{K} = 1. \tag{19}$$

Since the equality in (19) is not typically fulfilled for the available set of options at each Wednesday, we construct the ATM sample as follows. For each Wednesday, we choose two options such that their forward price divided by their strike price bracket one.

If, for a certain date, all forward prices normalized by the strike prices are less than one, we choose the option corresponding to $F_j = Max\left(\{F_i(t,T)\}_{i=1}^{N_t}\right)$, where N_t is the total number of options contracts at date t, as the ATM option for that date. If, however, all forward prices normalized by the strike prices are greater than one, we choose the option corresponding to $F_k = Min\left(\{F_i(t,T)\}_{i=1}^{N_t}\right)$ as the ATM option for date t.

We use a volatility updating rule on the 252 days predating the first Wednesday used in the estimation sample. This volatility updating rule is initialized at the model's unconditional variance.

Table 2 presents descriptive statistics of the options data for the period 1990-1995 by maturity. There are 3,275 contracts, the largest group among them having a maturity ranging from 20 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 above features indicate that the sample at hand is standard. The top panel of Figure 1 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 1 that there is substantial clustering in implied volatilities. It can also be seen that volatility is higher in the early part of the sample.

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 January 2, 1987 to December 31, 2004. Column 1 of Table 3 presents the results for the CEV model. It is clear that the nonlinearity coefficient β is significantly different from 0.5. This suggests that the Heston (1993) model is rejected by using returns data in favor of a more general CEV specification.

We now review the values of the other parameters of the model. We can 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 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 volatility of volatility parameter is 2.21. The correlation is negative around -0.67 confirming most findings in the literature and the observed empirical skewed distribution of the S&P500 index.[8] Finally, the value of the parameter associated with the risk premium is high. This parameter has been poorly estimated in the literature and this is also the case in this paper, as confirmed by the relatively high standard deviation. Hence, we will not want to offer conclusions from this estimate except noticing that, in the range of any conventional confidence interval, this parameter is positive. The latter suggests that investors ask for a risk premium to hold the index.

Column 3 of Table 3 contains the estimates for the CEVJ model. It is striking that, when we add jumps to the model, the coefficient of nonlinearity becomes 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 results found by Chacko and Viceira (2003). Therefore, estimation of the CEVJ model using options data is in order to test the results obtained using daily S&P500 returns.

We also find a slightly higher persistence of approximately 99.44% in line with 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 in 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 2-A presents the estimated jump sizes and jump probabilities using the estimates of the CEVJ model from Table 3.[9] 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 can conclude from Figure 2-A that almost all jumps are of negative size and that jumps are very infrequent. In fact, following Johannes et al. (2006) who consider that there

[8]See, for example, Benzoni (2002), and Pan (2002) for empirical evidence of the skewness of the distribution of returns.

[9]See Johannes, Polson and Stroud (2006) for details on how to estimate the jump sizes and jump probabilities using the SIR particle filter.

is a jump if its estimated probability is greater than 0.5, we count a mere 12 jumps in 18 years.

Table 4 contains the parameter estimates for the CEV and CEVJ models when we do not include the 1987 crash and the three subsequent years. It is noticeable 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 toward the expected directions. In fact, we obtain a lower persistence, a lower long run volatility, and a 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 2-B highlights how much smaller the estimated jump sizes and probabilities are for the sample 1990-2004, indicating 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 (2006) 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 described in Table 3 are not directly comparable to their findings. First, Jones (2003) and Ait-Sahalia and Kimmel (2006) perform a joint estimation using options and returns data whereas, here, we use returns only. Besides, Ait-Sahalia and Kimmel (2006) restrict the value of the power on the volatility in the diffusion term to less than 1 and we do not impose any restrictions.[10] The level of correlation is in the same range as that generated by other empirical studies using returns[11] but less than the correlation obtained using options or a combination of returns and options.[12]

In terms of Log likelihood, adding jumps to the CEV model improves the Log likelihood criterion by more than 16 points for the sample 1987-2004. This difference seems important in terms of magnitude. However, a closer look at the top panel of Figure 3, 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) when they compare different models for the S&P500 dynamics. The authors find that some of the differences in Log likelihood across models vanish when they remove one observation from their sample. Therefore, we conclude 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 3 confirm this finding since, when we estimate the model excluding the 1987 crash, we obtain almost the same Log likelihood.

[10] When $\beta \geq 1$ the Euler approximation given by equations (15) and (18) may diverge. We follow Jones (2003) who shows that these Euler discretizations are accurate for typical estimates of stochastic volatility parameters. Ait-Sahalia and Kimmel (2006) impose the restriction $0 \leq \beta \leq 1$ to obtain the unicity of options prices.

[11] Eraker, Johannes and Polson (2003) find a correlation of circa -0.7. Jacquier, Polson and Rossi (1994) find a correlation of around -0.5.

[12] See, for example, Christoffersen et al. (2006) and Eraker (2004).

Overall, we find that returns seem to favor a nonlinear specification for the model regardless of whether we include 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 in the model.

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. It is clear that when we compare Table 5 and Table 3, there is some consistency between the estimates obtained using returns and those obtained using returns and options. The first 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 the linear specification is rejected by the options data.

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 negative at around -0.86. The negative correlation is a standard result in the literature which is observed when we estimate stochastic volatility models on any set of data. The long-run volatility is around 27%, close to the results obtained with returns but somewhat higher than the volatility in Table 1. We should stress that, even though the estimate of this parameter varies considerably in the stochastic volatility literature, it always falls within a certain interval around the sample variance of returns.[13] The risk premium associated with the volatility dynamic λ is small and statistically not significant. Therefore, setting this parameter to zero in many empirical papers seems to be a realistic assumption.[14] Next, the risk premium coefficient λ_1 associated with returns is quite variable but it always remains positive, suggesting again that investors ask for a premium to hold risky assets. Its value is in the same range as 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 the correlation implied from returns. This result is confirmed by Eraker (2004) and by Christoffersen et al. (2006). Our estimate is lower than the correlation obtained by Christoffersen et al. (2006) using the same estimation technique. We believe that this paper's use of long samples of options data permits more accurate identification of the level of correlation. However, the estimate reported in Eraker (2004) using options but a

[13] Ait-Sahalia and Kimmel, for example, find an unconditional variance of around 21%. Eraker, Johannes and Polson (2003) find it equal to 15%.

[14] See, for example, Ait-Sahalia and Kimmel (2006).

different estimation methodology yields a correlation ranging from -0.57 to -0.59, even lower than our results. But, as noticed by Eraker (2004) there is no consensus in the literature as to the level of this parameter. In the case of the level of nonlinearity, a lack of empirical studies does not allow to make valid comparison with the existing literature. The only known exceptions who estimated the CEV model on returns and options are Jones (2003) and Ait-Sahalia and Kimmel (2006). Our estimate is higher than the one obtained by Ait-Sahalia and Kimmel (2006) and lower than the estimate of Jones (2003).

Column 3 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 can 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 below, jumps do not add much to the model in terms of fit. We suspect therefore that this poor performance in fit for the CEVJ model is due to the fact that jumps do not improve the volatility process itself. 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 around 0.84 jumps per year.

As we pointed out, the inclusion of jumps does not improve the model fit. In fact, the RMSE decreases from 1.38 to 1.36, which cannot be considered as a large benefit. This result confirms the findings of Bates (2000) and Eraker (2004) but contradicts those of Broadie et al. (2006). However, the results of Broadie et al. (2006) are not directly comparable to 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, to be different under the objective and risk-neutral measure, whereas we use a much more parsimonious specification. Finally, the options data used and the periods covered differ from those used in our sample.

Figure 4 elaborates on the potential reasons why we obtain similar performances for the CEV and CEVJ models. The top panel of Figure 4 shows that the weekly RMSE from the CEVJ model and the CEV model are almost indistinguishable. The bottom panel further investigates the difference between the RMSE obtained from the two models. We see that we cannot infer an obvious pattern. These findings, along with the results obtained from returns, stress the fact that the difference in fit, when we add jumps to the CEV model, is very small. This is not surprising when we examine the top and bottom panels of Figure 5. In fact, the first and second columns of Figure 5 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 seem to depend on whether we estimate the model on returns or on options. Moreover, when we compare the residuals of the CEV model to those of the CEVJ model we observe from the right column of Figure 5 that their empirical CDFs are almost the same. The latter may explain why these models do equally well in fitting options data.

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 option are consistent. We also find that adding jumps to nonlinear models does 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 has been previously documented in the literature. This is important since Christoffersen et al. (2006) used the same technique and found that the correlation sometimes approaches the prespecified boundary of -0.999 when they estimated the model using options. In this paper, we obtain a lower correlation by using a longer sample, and by estimating the CEV model instead of using quadratic volatility in the diffusion term as in Christoffersen et al. (2006).

Although we find in this paper that adding jumps to nonlinear models does 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 may not be rich enough in terms of volatility dynamics to show 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. Indeed, including a full cross section of options data along with the time series dimension may 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 effect including extra parameters on the out-of-sample performance for these models is not obvious and should be studied in future work.

References

[1] Ait-Sahalia, Y. and R. Kimmel (2006), Maximum Likelihood Estimation of Stochastic Volatility Models, Forthcoming Journal of Financial Economics.

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

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

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

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

[6] Broadie, M., M. Chernov and M. S. Johannes (2006), Model Specification and Risk Premiums: The Evidence from the Futures Options, Journal of Finance, forthcoming.

[7] Chacko, G. and L. Viceira (2003), Spectral GMM Estimation of Continuous-Time Processes, Journal of Econometrics 116, 259-292.

[8] 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.

[9] Christoffersen, P. F., K. Jacobs and K. Mimouni (2006), Models for S&P500 Dynamics: Evidence from Realized Volatility, Daily Returns, and Option Prices, Working Paper, McGill University.

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

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

[12] Gallant, A., R. (1997), An Introduction to Econometric Theory. Princeton: Princeton University Press.

[13] Geweke, J. (1989), Bayesian Inference in Econometric Models Using Monte-Carlo Integration, Econometrica, 57, 1317-1339.

[14] 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.

[15] Jacquier, E., N. G. Polson, P. E. Rossi (1994), Bayesian Analysis of Stochastic Volatility Models, Journal of Business & Economic Statistics, American Statistical Association, 12, 371-389.

[16] Johannes, M. S., N. G. Polson (2003), MCMC methods for Financial Econometrics, forthcoming in the Handbook of Financial Econometrics, edited by Yacine Aït-Sahalia and Lars Hansen.

[17] Johannes, M. S., N. G. Polson and J. R. Stroud (2006), Optimal Filtering of Jump Diffusions: Extracting Latent States from Asset Prices, Working Paper.

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

[19] Pitt M. K. (2002), Smooth Particle Filters for Likelihood Evaluation and Maximization, Working Paper, University of Warwick.

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 which has the Euler discretization given by

$$\begin{aligned} \log(S_{t+\Delta}) &= \log(S_t) + \left(r + \lambda_1 \left(1-\rho^2\right) V_t - \frac{1}{2} V_t\right)\Delta + \sqrt{V_t \Delta}\varepsilon_{1,t+\Delta} \\ V_{t+\Delta} &= \kappa\left(\theta - V_t\right)\Delta + \sigma\sqrt{\Delta} V_t^{\beta}\left(\rho\varepsilon_{1,t+\Delta} + \sqrt{1-\rho^2}\varepsilon_{2,t+\Delta}\right). \end{aligned} \tag{20}$$

Filtering the state variable consists in the following 3 steps.

6.2 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 (20) and taking the correlation into account.[15] 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 $\mu_t = r + \lambda_1 \left(1-\rho^2\right) V_t^j$.
which gives

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

Since

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

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

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

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

[15] 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.

6.3 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 (20) , that, given the other available information, $V_{t+\Delta}$ is sufficient to generate $\ln(S_{t+2\Delta})$. Therefore, equation (20) offers a simple way to evaluate the likelihood that the observation $S_{t+2\Delta}$ has been generated by $V_{t+\Delta}$. Hence, we are able to compute the weight given to each particle (or the likelihood or probability that the particle has generated $S_{t+2\Delta}$). The likelihood is computed as follows:

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

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

6.4 Step 3: Resampling

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

First, the adjusted weights obtained in Step 2, $W_{t+\Delta}^{j}$, are mapped into a set of integer variables $\{\iota_{t+\Delta}^{j}\}_{j=1}^{N}$, using an algorithm that takes into account that the weights are not multiples of $1/N$. This algorithm is based on the empirical CDF of V 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 $\{V(\iota)_{t+\Delta}^{j}\}_{j=1}^{N}$ by replicating each particle in the original set $\{V_{t+\Delta}^{j}\}_{j=1}^{N}$ $\iota_{t+\Delta}^{j}$ times. Therefore, the particles in the original set are either eliminated, or included one or multiple times according to their adjusted weights $\{W_{t+\Delta}^{j}\}_{j=1}^{N}$. The higher the weight, $W_{t+\Delta}^{j}$, the higher the integer variable $\iota_{t+\Delta}^{j}$, and the more often the original particle $V_{t+\Delta}^{j}$ is included in the resampled set $\{V(\iota)_{t+\Delta}^{j}\}_{j=1}^{N}$.

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

6.5 Appendix 2: Adaptation of the PF to the CEVJ model

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

$$f(R_t|V_{t-1}) = \sum_{x=0}^{\infty} N\left(R_t|x\mu_J, \int_t^{t+\Delta} V_{t-1}\Delta + x\sigma_J^2\right)\frac{(\Delta\lambda_J)^x e^{-\Delta\lambda_J}}{x!}. \tag{21}$$

Proof of this approximation result:

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

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

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

Then we have

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

This converges quickly (we typically can ignore terms beyond three or four terms i.e. x$>$ 4). So now we have the form of $f(R_t|V_{r-1})$ which is more heavy tailed than Gaussian as it is a mixture.

This form of the density given by equation (21) allows us to do the smooth resampling as it was previously carried in the filtering algorithm (see step 3 in appendix1). 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 can be shown to be

$$
\begin{aligned}
\log(S_{t+\Delta}) &= \log(S_t) + \left(\mu_t - \frac{1}{2}V_t - \lambda_J\mu_J + \sum_{x=0}^{3} x \Pr\left(J_x = x\right)\mu_J\right)\Delta \\
&\quad + \sqrt{\left(V_t + \sum_{x=0}^{3} x \Pr\left(J_x = x\right)\sigma_J^2\right)\Delta}\, z_{1,t+\Delta} \qquad (22) \\
V_{t+\Delta} &= \kappa\left(\theta - V_t\right)\Delta + \sigma\sqrt{\Delta}V_t^{\beta}\left(\rho z_{1,t+\Delta} + \sqrt{1-\rho^2}\varepsilon_{2,t+\Delta}\right),
\end{aligned}
$$

where $z_{1,t+\Delta} \sim N(0,1)$, $corr(z_{1t}, \varepsilon_{2t}) = 0$ and $\Pr\left(J_x = x\right) = \frac{(\Delta\lambda_J)^x e^{-\Delta\lambda_J}}{x!}$.

We then proceed with the filtering exercise exactly as with the CEV model.

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

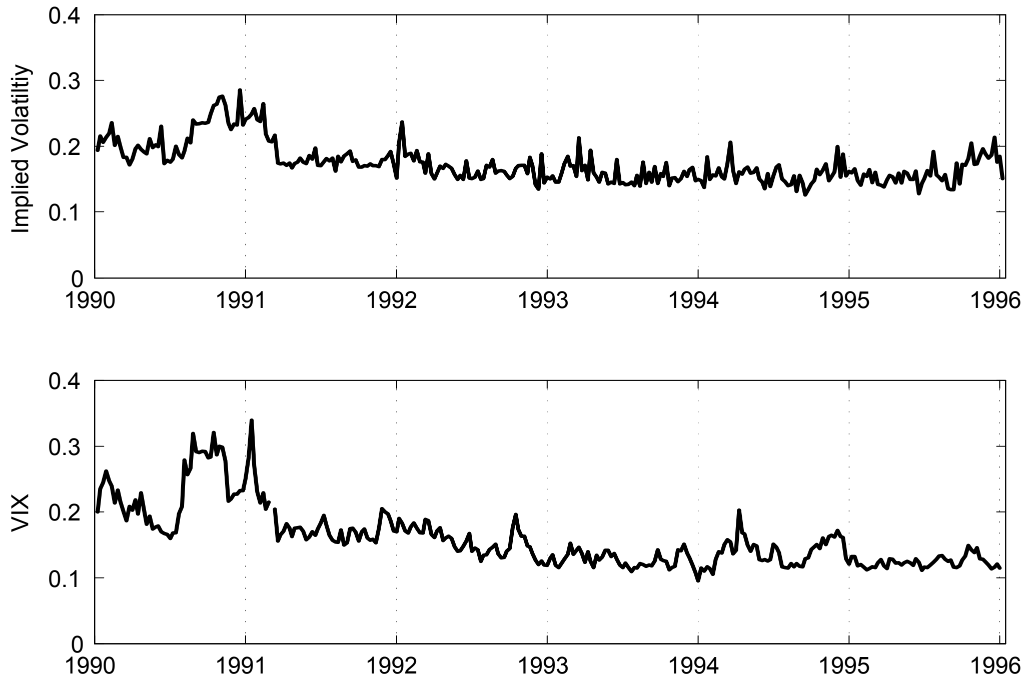

Notes to Figure: The top panel plots the average implied Black-Scholes volatility each Wednesday during 1990-1995. The average is taken across maturities and strike prices using the call options in our data set. For comparison, the bottom panel shows the one-month, at-the-money VIX volatility index retrieved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

Figure 2-A: Estimated Jump Sizes and Probabilities Using the SIR Particle Filter for Historical Returns Data 1987-2004

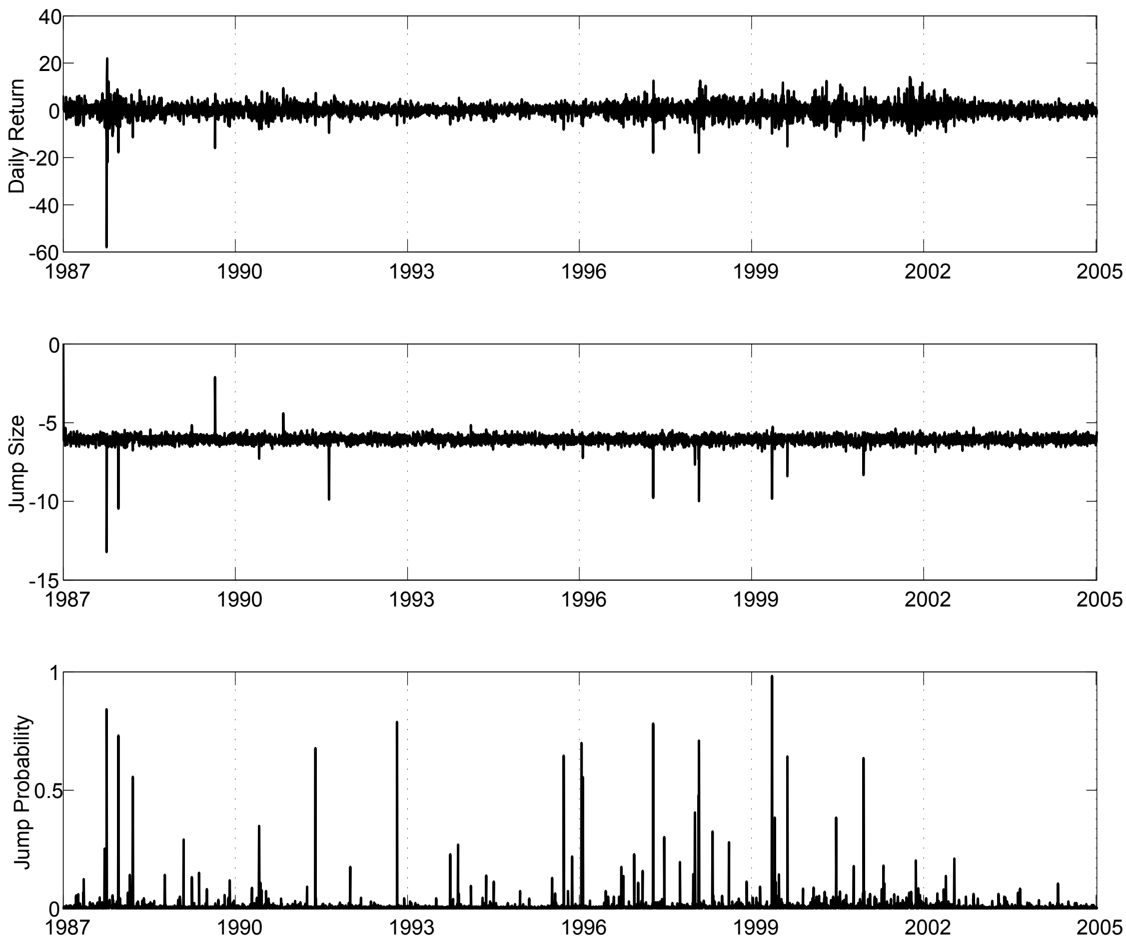

Notes to Figure: The top panel plots the S&P500 returns for the period January 2, 1987 to December 31, 2004. The middle panel plots the estimated jump sizes obtained using the particle filter. Finally, the bottom panel represents the jump probabilities obtained by applying the same particle filter. The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

Figure 2-B: Estimated Jump Sizes and Probabilities Using the SIR Particle Filter for Historical Returns Data 1990-2004

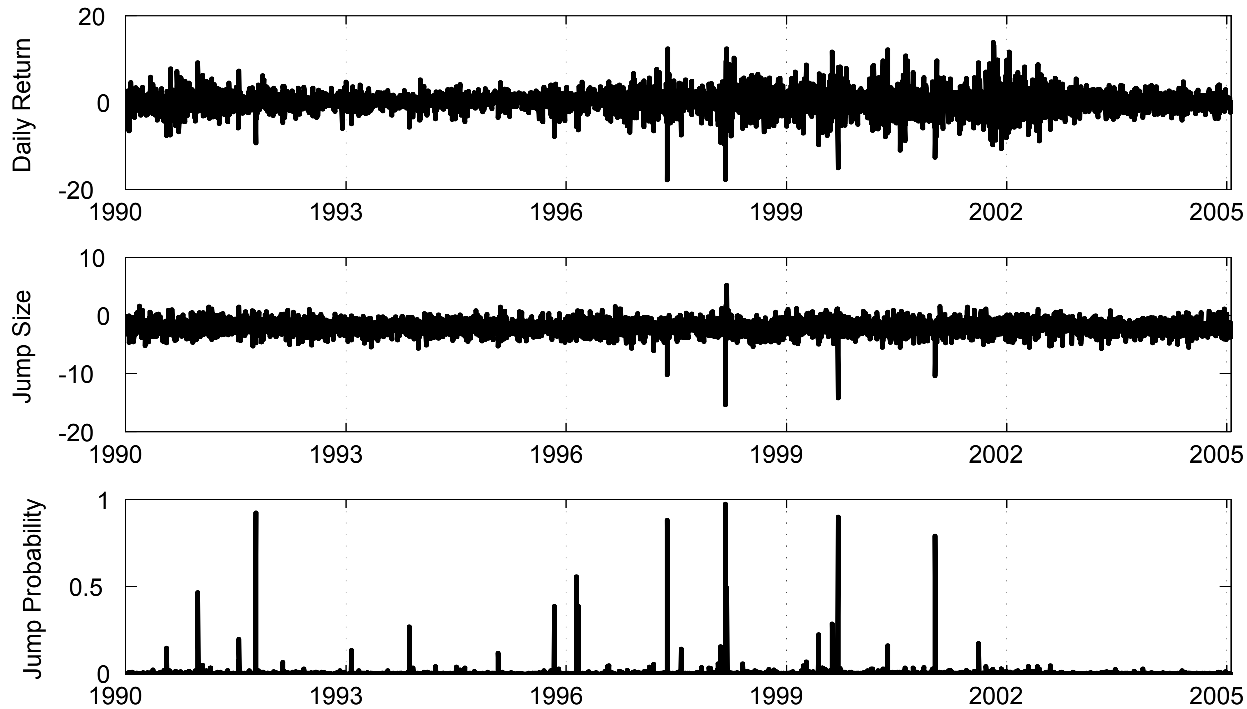

Notes to Figure: The top panel plots the S&P500 returns for the period January 2, 1990 to December 31, 2004. The middle panel plots the estimated jump sizes obtained using the particle filter. Finally, the bottom panel represents the jump probabilities obtained by applying the same particle filter. The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

Figure 3: Daily Log likelihood (L) Difference: $\mathbf{L}_t(\mathbf{CEVJ}) - \mathbf{L}_t(\mathbf{CEV})$

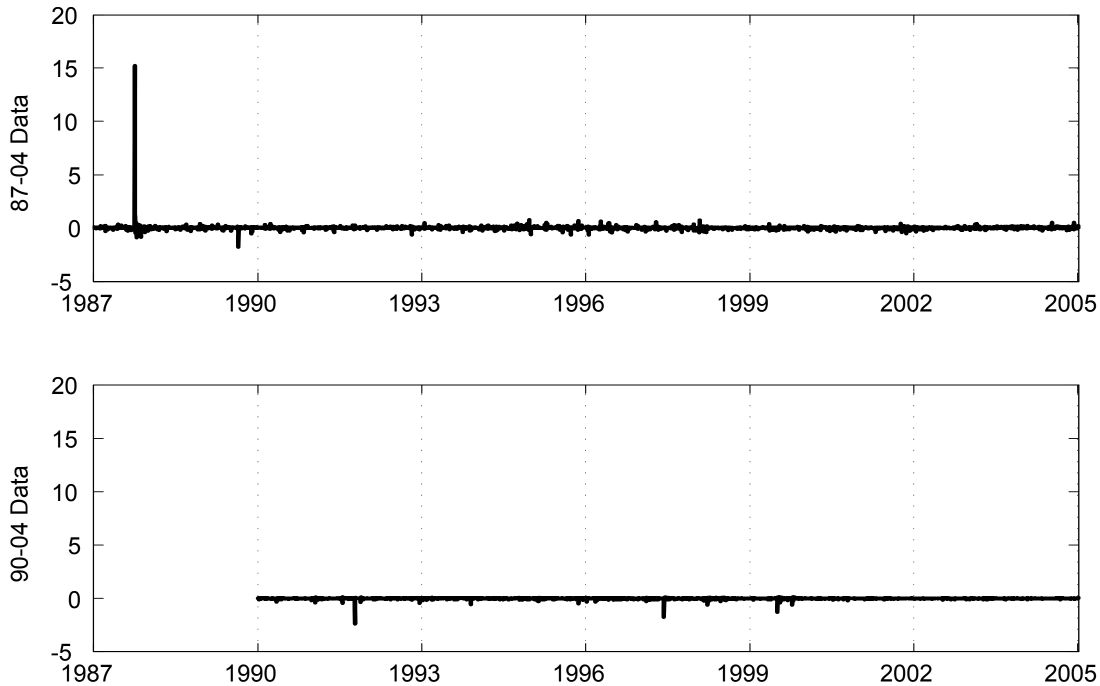

Notes to Figure: We plot the difference in Log likelihood observation by observation for the periods January 2, 1987 to December 31, 2004, and January 2, 1990 to December 31, 2004. The difference represents the Log likelihood of the CEVJ model less the Log likelihood of the CEV model.

Figure 4: Weekly RMSE Difference: RMSE(CEVJ) – RMSE(CEV)

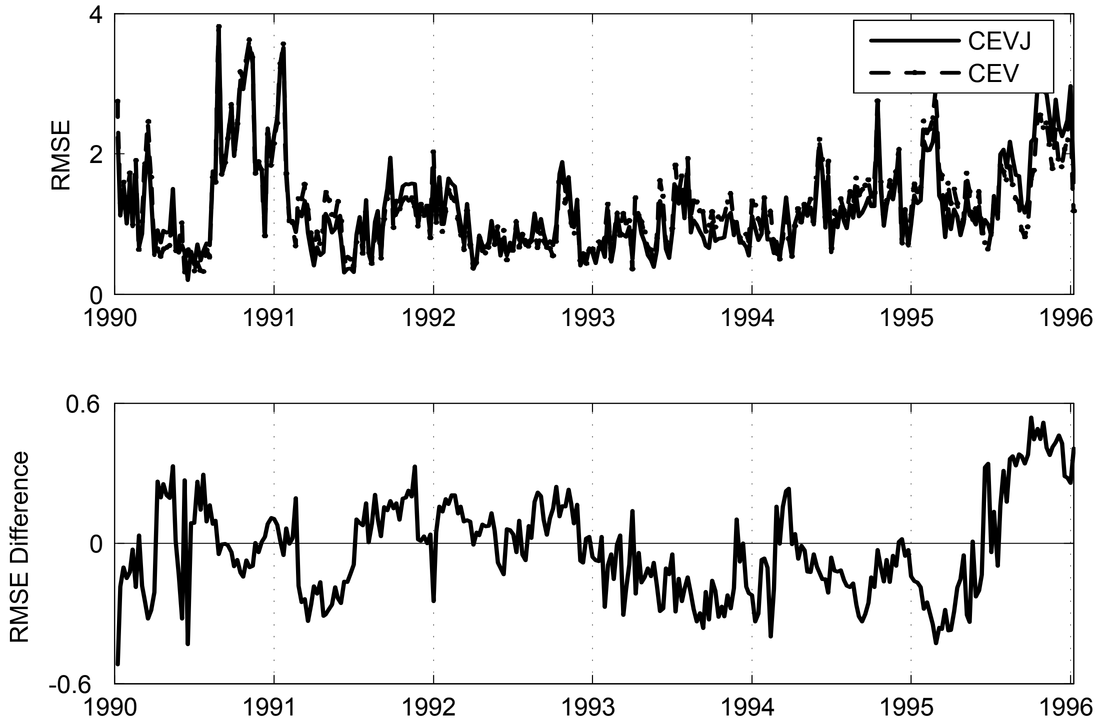

Notes to Figure: We plot the difference in weekly RMSE for the period January 2, 1990 to December 31, 1995. The difference represents the weekly RMSE of the CEVJ model less the weekly RMSE of the CEV model.

Figure 5: Empirical Cumulative Distribution Function of the CEV and CEVJ Models Implied from S&P500 Returns and from ATM Options.

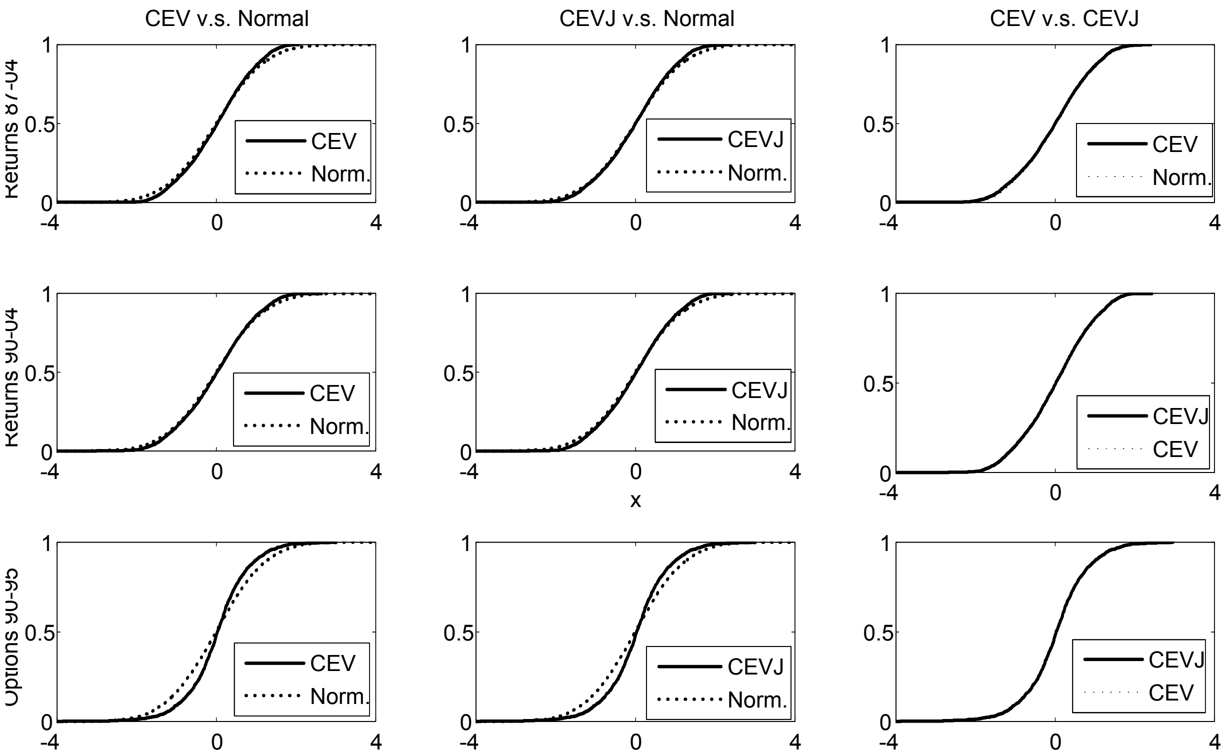

Notes to Figure: We plot the empirical CDFs for the CEV model and CEVJ model using the residuals evaluated at the optimal parameters. To plot the CDFs in the top panel, we use the estimates in Table 3 obtained from Returns for the period 87-04. To plot the CDFs in the middle panel, we use the estimates 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.

Table 1: Summary Statistics for Daily S&P500 Returns.

	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

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 in the NLS estimation from January 2, 1990 to December 31, 1995.These statistics include the annualized mean, volatility, skewness and kurtosis. We also provide the minimum and the maximum daily returns.

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

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
Number of Call Option Contracts	282	1170	722	1101	3275
Average Call Price	4.35	8.51	14.48	23.01	14.34
Average Implied Volatility from Call Options	0.146	0.142	0.150	0.155	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. When choosing the options, we use the same filters as in Bakshi, Cao and Chen (1997).

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

	CEV		CEVJ	
	Estimate	**Standard Error**	**Estimate**	**Standard Error**
κ	2.1752	0.6102	1.4200	0.5914
θ	0.0417	0.0080	0.0350	0.0199
σ	2.2131	0.2723	8.8611	1.0256
ρ	-0.6676	0.0269	-0.6016	0.0911
β	0.9300	0.0425	1.3378	0.0371
λ_1	4.1721	1.7687	8.4424	2.5400
μ_J			-6.0672	4.3123
σ_J			0.2653	9.1210
λ_J			2.2024	0.5338
Log Likelihood	15818.15		15834.55	
Annualized volatility (%)	20.41		18.71	
Daily persistence (%)	99.14		99.44	

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

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

	CEV		CEVJ	
	Estimate	**Standard Error**	**Estimate**	**Standard Error**
κ	2.5818	0.9208	2.6070	0.9116
θ	0.0287	0.0056	0.0268	0.0135
σ	1.8667	0.5322	2.0274	0.4992
ρ	-0.6739	0.0338	-0.6885	0.0828
β	0.9283	0.0843	0.9488	0.0798
λ_1	3.2887	2.1914	3.1763	3.2103
μ_J			-2.0574	3.2133
σ_J			1.5252	0.7954
λ_J			1.0215	0.5623
Log Likelihood	13433.60		13434.92	
Annualized volatility (%)	16.95		16.36	
Daily persistence (%)	98.98		98.97	

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

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

	CEV		CEVJ	
	Estimate	**Standard Error**	**Estimate**	**Standard Error**
κ	0.1824	0.0190	0.1784	0.0175
θ	0.0740	0.0064	0.0756	0.0107
σ	0.2107	0.0314	0.2267	0.0267
ρ	-0.8604	0.0140	-0.8847	0.0184
β	0.8200	0.0414	0.7958	0.0376
λ_1	1.7735	3.8425	1.9111	11.3695
λ	0.0157	0.0504	0.0153	0.1040
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	1.3833		1.3584	
Annualized volatility (%)	27.20		27.49	
Daily persistence (%)	99.93		99.93	

Note: We estimate the CEV model using Wednesday Options on the S&P500 Index for the period 1990 to 1995. Columns 2 and 4 contain the parameter estimates for the CEV model and the CEVJ model, respectively. Columns 3 and 5 contain their corresponding standard errors.