Jump and Volatility Risk Premiums Implied by VIX

Jin-Chuan Duan* and Chung-Ying Yeh†

(January 22, 2007)

Abstract

An estimation method is developed for extracting the latent stochastic volatility from VIX, a volatility index for the S&P 500 index return produced by the Chicago Board Options Exchange (CBOE) using the so-called model-free volatility construction. Our model specification encompasses all mean-reverting stochastic volatility option pricing models with a constant-elasticity of variance and those allowing for price jumps under stochastic volatility. Our approach is made possible by linking the latent volatility to the VIX index via a new theoretical relationship under the risk-neutral measure. Because option prices are not directly used in estimation, we can avoid the computational burden associated with option valuation for stochastic volatility/jump option pricing models. Our empirical findings are: (1) incorporating a jump risk factor is critically important; (2) the jump and volatility risks are priced; and (3) the popular square-root stochastic volatility process is a poor model specification irrespective of allowing for price jumps or not.

Keywords: Model-free volatility, stochastic volatility, jump, options, VIX, Constant elasticity of variance.

JEL classification code: G12, G13.

*Duan is with Joseph L. Rotman School of Management, University of Toronto. E-mail: jcduan@rotman.utoronto.ca; Tel: 416-946 5653; Fax: 416-971 3048. The author acknowledges support received as the Manulife Chair in Financial Services and research funding from the Social Sciences and Humanities Research Council of Canada.

†Yeh is a doctoral student at Department of Finance, National Taiwan University.

1 Introduction

The development of stochastic volatility models with jumps has come a long way in many dimensions in recent recent years. The importance of incorporating jumps has long been advocated in the empirical option pricing literature, such as Bakshi, Cao and Chen (1997), Bates (2000), Chernov and Ghysel (2000), Duffie, Pan, and Singleton (2000), Pan (2002), Eraker (2004), and Broadie, Chernov and Johannes (2006). Stochastic volatility being a latent variable, however, poses a significant methodological challenge to model testing and applications. In this paper, we devise an estimation method that conveniently extracts the latent stochastic volatility from VIX, a volatility index provided by Chicago Board Options Exchange (CBOE) for the S&P 500 index. The VIX index is based on forming a portfolio of European options with a target maturity of 30 days. The value of such an option portfolio has been shown in the model-free volatility literature to represent the risk-neutral expected realized volatility over the horizon defined by the maturity of the option contracts in the portfolio. Interestingly, we are able to derive a closed-form expression that further links the risk-neutral expected realized volatility to the latent stochastic volatility for the class of stochastic volatility models (with or without jumps) whose elasticity of variance is constant. This new theoretical link between the VIX index and the latent stochastic volatility allows us to in effect view the VIX index, after a proper transformation, as the latent stochastic volatility. Consequently, the estimation task for a large class of stochastic volatility models with or without jumps can be dramatically simplified.

The new theoretical link between the VIX index and the latent stochastic volatility has an added benefit. Since the linking function depends on the volatility and jump risk premiums, the values for these critical risk premiums can be inferred without directly using option prices in estimation. This stands in sharp contrast to the existing estimation methods in the literature that in one way or another requires of repeated option valuations under some stochastic volatility/jump model. Our proposed estimation method thus avoids costly numerical option valuations and significantly reduces the computational burden associated with model testing and applications.

A joint consideration of the volatility and jump risk factors is expected to better capture the dynamics of equity returns, which in turn better reconciles the theoretical model with the observed volatility smile/smirk, and as a result, improves option valuation. However, the empirical results in the literature appear to be mixed. Anderson, Benzoni and Lund (2002) and Eraker, Johannes and Polson (2003) conclude that allowing jumps in prices can improve the fitting for the time-series of equity returns. However, Bakshi, Cao and Chen (1997), Bates (2000), Pan (2002) and Eraker (2004) offer different and inconsistent results in terms of improvement on option pricing. There is no joint significance in the volatility and jump risk premium estimates in most cases. Broadie, Chernov, and Johannes (2006) provided one plausible explanation for these diverse findings, which they attributed to the short sample period and/or limited option contracts used in those papers. Practically speaking, using options over a wide range of strike prices over a long time span in estimation will quickly create an unmanageable computational burden. Our approach of using the VIX index is a joint estimation method suitable for the data sample over a long time span that can avoid costly option valuations. In effect, the VIX index has summarized all critical information in options over the entire spectrum of strike prices, and it is also informative about the time series behavior of the latent stochastic volatility.

Our analysis is based on the maximum likelihood principle applied to the transformed data setting as proposed by Duan (1994). The derived link between the VIX index and the latent stochastic volatility serves as the critical transformation from the unobserved risk factor (the latent stochastic volatility) to the observed VIX (the value of an option portfolio). We conduct the maximum likelihood estimation and inference on the observed S&P 500 and VIX index values from the first trading day of 1990 to the last trading day of 2006. Our empirical findings are: (1) incorporating a jump risk factor is critically important; (2) the jump and volatility risks are priced; and (3) the popular square-root stochastic volatility process is a poor model specification irrespective of allowing for price jumps or not.

The balance of this paper is organized as follows. In section 2, we propose under the physical probability measure a constant-elasticity-of-variance stochastic volatility model that allows jumps in the price. Then we proceed to derive the corresponding system under a risk-neutral pricing measure. The critical link between the VIX index and the latent volatility is also established there. In section 3, the likelihood function for the model is then derived and presented. The empirical results are reported and discussed in Section 4. Some concluding remarks follow in Section 5.

2 A stochastic volatility model with jumps in asset prices

The asset price is assumed to follow a jump-diffusion model and the asset volatility is allowed to be stochastic. Specifically, the dynamics under the physical probability measure P are

$$d\ln S_t = \left[r - q + \delta_S V_t - \frac{V_t}{2}\right] dt + \sqrt{V_t} dW_t + J_t dN_t - \lambda\mu_J dt \quad (1)$$

$$dV_t = \kappa(\theta - V_t)dt + vV_t^{\gamma} dB_t \quad (2)$$

where W_t and B_t are two correlated Wiener processes with the correlation coefficient equal to ρ; N_t is a Poisson process with intensity λ and independent of W_t and B_t; J_t is an independent normal random variable with mean μ_J and standard deviation σ_J. Note that dW_t and $J_t dN_t$ have respective variances equal to dt and $\lambda(\mu_J^2 + \sigma_J^2)dt$. Thus, $V_t + \lambda(\mu_J^2 + \sigma_J^2)$ is the variance rate of the asset price process. The price and volatility processes are dependent through two correlated diffusive terms – W_t and B_t. In the above equation, the risk-free rate, the dividend yield and the asset risk premium are r, q and δ_S, respectively. The term $\lambda\mu_J dt$ is used to center the Poisson innovation so that $J_t dN_t - \lambda\mu_J dt$ has its mean equal to 0.

The specification in equations (1) and (2) contains many well-known stochastic volatility models with or without jumps. If there are no jumps (i.e., $\lambda = 0$), then the Hull and White (1987) or Heston (1993) stochastic models follow by further setting γ to 1 or 1/2.[1] If jumps are allowed, the price innovation becomes that of Bates (2000) and Pan (2002). Note that the asset return's jump size is not related to volatility and thus does not exhibit any dynamic behavior. One can introduce a dynamic feature to the jump component by, for example, making λ time-varying. The joint price-volatility model is more general than that that of Bates (2000) and Pan (2002) because their specification corresponding to the special case of $\gamma = 1/2$, i.e, a square-root volatility process.

For option pricing, we follow the standard approach of using the risk-neutral pricing idea, which implies that the discounted asset price process is a martingale with respect to an equivalent martingale measure, Q. Note that the asset price is subject to jumps and the volatility is not a traded asset. Either feature makes the market incomplete. Although no arbitrage implies the existence of an equivalent martingale measure, but it is not unique. The choice made below is consistent with that of Hull and White (1987), Heston (1993), Bates (2000) and Pan (2002) for the volatility risk premium in terms of dealing with the incompleteness due to stochastic volatility. To deal with jumps in both price and volatility, we follow that of Bates (2000) and Pan (2002) to restrict to our attention to the equivalent martingale measures under which the jump dynamic remains in the same form but the jump intensity and the mean of the jump size are allowed to differ from those under the physical measure, i.e., from λ to λ^* and from μ_J to μ_J^*. This can be accomplished by adopting a particular form for the pricing kernel. We refer readers to Appendix A of Pan (2002) for details. The corresponding system under measure Q becomes

$$d\ln S_t = \left[r - q - \frac{V_t}{2} + \lambda^*\left(\mu_J^* - e^{\mu_J^* + \frac{\sigma_J^2}{2}}\right)\right]dt + \sqrt{V_t}dW_t^* + J_t^* dN_t^* - \lambda^*\mu_J^* dt \tag{3}$$

$$dV_t = (\kappa\theta - \kappa^* V_t)\,dt + vV_t^\gamma dB_t^* \tag{4}$$

where $\kappa^* = \kappa + \delta_V$ and $B_t^* = B_t + \frac{\delta_V}{v}\int_0^t V_s^{1-\gamma}ds$ with δ_V being interpreted as the volatility risk premium. W_t^* and B_t^* are two correlated Wiener processes under measure Q and their correlation coefficient remains at ρ; N_t^* is a Poisson process with intensity λ^* and independent of W_t^* and B_t^*; J_t^* is an independent normal random variable under measure Q with a new mean μ_J^* but its standard deviation remains unchanged at σ_J. It can be easily verified by Ito's lemma that equation (3) leads to $E_t^Q\left(\frac{dS_t}{S_t}\right) = (r-q)dt$ so that the expected return under measure Q is indeed the risk-free rate minus the dividend yield.

[1] The Hull and White (1987) model was originally formulated without the volatility reversion feature (i.e., $\theta = 0$). In this paper, we interpret that model as one with volatility reversion. Note that a specification without volatility reversion will have quite poor empirical performance.

Note that $V_t + \lambda^*(\mu_J^{*2} + \sigma_J^2)$ becomes the variance rate of the asset price process under measure Q, which may be different from $V_t + \lambda(\mu_J^2 + \sigma_J^2)$ when jumps are allowed. An interesting consequence of introducing jumps is that the local volatility of the asset return is no longer invariant to the change of measures.

This complex stochastic volatility-jump model retains a useful feature; that is, we can derive a closed-form expression for the risk-neutral expected cumulative variance over any horizon. First, for any $\kappa^* \neq 0$ and $\tau > 0$,

$$E_t^Q(V_{t+\tau}) = \frac{\kappa\theta}{\kappa^*} + \left(V_t - \frac{\kappa\theta}{\kappa^*}\right) e^{-\kappa^*\tau}. \tag{5}$$

Thus, the risk-neutral expected cumulative variance becomes

$$\int_t^{t+\tau} E_t^Q(V_s)\, ds = \frac{\kappa\theta}{\kappa^*}\left(\tau - \frac{1 - e^{-\kappa^*\tau}}{\kappa^*}\right) + \frac{1 - e^{-\kappa^*\tau}}{\kappa^*} V_t. \tag{6}$$

If $\kappa^* = 0$, the corresponding formulas for equations (5) and (6) should be replaced with their limiting results, which are $\kappa\theta\tau$ and $\frac{\kappa\theta\tau^2}{2} + \tau V_t$, respectively. Moreover, equation (3) can be used to obtain

$$\begin{aligned} & E_t^Q\left(\ln\frac{S_{t+\tau}}{S_t}\right) \\ &= (r-q)\tau - \frac{1}{2}\int_t^{t+\tau} E_t^Q(V_s)\, ds + \int_t^{t+\tau} \lambda^* E_t^Q\left(\mu_J^* + 1 - e^{\mu_J^* + \frac{\sigma_J^2}{2}}\right) ds \\ &= \left[r - q - \lambda^*\left(e^{\mu_J^* + \frac{\sigma_J^2}{2}} - (\mu_J^* + 1)\right)\right]\tau - \frac{1}{2}\int_t^{t+\tau} E_t^Q(V_s)\, ds. \end{aligned} \tag{7}$$

A well-known fact in the model-free volatility literature is that $\ln\frac{S_{t+\tau}}{S_t}$ can be replicated by a portfolio of European options. First, define an option portfolio value at time t with its component options expiring at time $t+\tau$:

$$\Pi_{t+\tau}(K_0, t+\tau) \equiv \int_0^{K_0} \frac{P_{t+\tau}(K; t+\tau)}{K^2} dK + \int_{K_0}^{\infty} \frac{C_{t+\tau}(K; t+\tau)}{K^2} dK. \tag{8}$$

By the generic payoff expansion result of Carr and Madan (2000), we have

$$\Pi_{t+\tau}(K_0, t+\tau) = \frac{S_{t+\tau} - K_0}{K_0} - \ln\frac{S_t}{K_0} - \ln\frac{S_{t+\tau}}{S_t}, \tag{9}$$

which can in turn be translated to a relationship at time t as

$$e^{r\tau}\Pi_t(K_0, t+\tau) = \frac{F_t(t+\tau) - K_0}{K_0} - \ln\frac{S_t}{K_0} - E_t^Q\left(\ln\frac{S_{t+\tau}}{S_t}\right). \tag{10}$$

where $F_t(t+\tau)$ denotes the forward price at time t with a maturity at time $t+\tau$.

The CBOE introduced the new VIX index in 2003, intending to capture the risk-neutral expected cumulative volatility; that is,

$$\mathrm{VIX}_t^2(\tau) \equiv \frac{2}{\tau} e^{r\tau} \Pi_t(F_t(t+\tau), t+\tau). \tag{11}$$

Applying equations (7) and (10), we establish a critical theoretical link:

$$\mathrm{VIX}_t^2(\tau) = 2\phi^* + \frac{1}{\tau} \int_t^{t+\tau} E_t^Q (V_s)\, ds, \tag{12}$$

where $\phi^* = \lambda^* \left(e^{\mu_J^* + \sigma_J^2/2} - 1 - \mu_J^* \right)$. If there are no jumps, then $\mathrm{VIX}_t^2(\tau)$ obviously equals the standardized risk-neutral expected cumulative variance or the risk-neutral expected realized variance over the horizon τ, which is a well-known result and serves as the theoretical basis underlying the VIX index. The model-free realized volatility literature, such as Britten-Jones and Neuberger (2000), Demeterfi, Derman, Kamal and Zhou (1999) and Jiang and Tian (2004), in essence, deals with this relationship for generic models without jumps.

When there are jumps, $\mathrm{VIX}_t^2(\tau)$ becomes a jump-adjusted risk-neutral expected cumulative variance over the horizon τ and it could be reduced to the standardized risk-neutral expected cumulative variance or the risk-neutral expected realized variance only when both μ_J^* and σ_J are small enough to justify that ϕ^* is negligible. When the jump size is small, the statement that the VIX index approximately equals the risk-neutral expected realized variance was first made in Jiang and Tian (2004). Although the result pertaining to the relationship between the VIX index and the risk-neutral expected realized variance is generic to the stochastic volatility models, it is not true for models with jumps. The result in equation (2) serves a concrete counterexample. As to what the specific form of the relationship applies, it will inevitably depend on how the jump model is specified.

The CBOE sets $K_0 = F_t(t+\tau)$. Such a choice is not a theoretical necessity, however. If one sets $K_0 = 0$ ($K_0 = \infty$), the option portfolio consists of only call (put) options. Arguably, the CBOE's choice is more natural because out-of-the-money options tend to be more liquid contracts.

The operational reality is that one can never have a complete set of options. Therefore, using a proxy becomes a must. The CBOE VIX index is based on approximating the right-hand side of equation (8) using the available out-of-the-money S&P 500 index options. The CBOE VIX index specifically targets the 30-day maturity. On any given day, the target maturity is expected to be sandwiched by two adjacent maturities, $\tau_t^l \le \tau_t^u$, and a linear interpolation of two option portfolios is then used to represent the index. Let n_t^l and n_t^u be the numbers of out-of-the-money options used in the CBOE approximation that correspond to τ_t^l and τ_t^u, respectively. To differentiate the CBOE VIX index from the theoretical VIX index, we denote the CBOE approximation by $\mathrm{VIX}_t(\tau_t^l, \tau_t^u, n_t^l, n_t^u)$. Combining equations (6) and (12) gives rise to

$$\mathrm{VIX}_t^2(\tau_t^l, \tau_t^u, n_t^l, n_t^u) \simeq \mathrm{VIX}_t^2(\tau) = \frac{\kappa\theta}{\kappa^*}\left(1 - \frac{1 - e^{-\kappa^*\tau}}{\kappa^*\tau}\right) + 2\phi^* + \frac{\left(1 - e^{-\kappa^*\tau}\right)}{\kappa^*\tau}V_t. \tag{13}$$

The above result shows that the CBOE VIX index can still be linked in closed-form to the latent volatility, V_t, for this complex model with stochastic volatility and jumps, and thus provides a simple way to deal with the estimation challenge posed by our inability to observe the latent volatility. In fact, equation (13) gives rise to our econometric specification with which the volatility and jump risk premiums can be estimated without actually performing option valuation based on a specific option pricing model such as Pan (2002). Our approach thus substantially simplifies the estimation task and avoids using option data directly.

Similar to an observation made in Pan (2002), λ^* and μ_J^* can be be separately identified. Pan (2002) simply assumed $\lambda^* = \lambda$. Equally acceptable is to assume $\mu_J^* = \mu_J$. Instead of forcing an equality on a specific pair of parameters, we find it convenient to use a composite parameter ϕ^* to define the jump risk premium. Specifically, the jump risk premium is regarded as $\delta_J = \phi^* - \phi$, where $\phi = \lambda\left(e^{\mu_J + \sigma_J^2/2} - 1 - \mu_J\right)$. The jump risk premium is meant to reflect the compensation term in the expected return for the jump risk. If the jump risk is priced, the compensation term will change by the amount equal to δ_J, which is induced by changing from the physical probability measure P to the risk-neutral pricing measure Q.

The volatility and jump risk premiums along with other parameters can be estimated using the summarized information about the option prices as reflected in the VIX index. These parameter values can naturally be used to assess the performance of an option pricing model, the general one or its special cases, on pricing individual options with different strike prices and maturities.

3 The econometric specification

The log-likelihood function for the observed time series $(\ln S_t, \mathrm{VIX}_t)$ can be constructed using the transformed data idea as in Duan (1994). In short, we can view the observed variable VIX_t as the transformed data of the latent volatility V_t. The log-likelihood function for the observed data pair $(\ln S_t, \mathrm{VIX}_t)$ will then be composed of two components. The first component is the standard log-likelihood function associated with a time series of $(\ln S_t, V_t)$ by acting as if V_t could be observed. The second component deals with the transformation, which turns out to be the logarithm of the Jacobian for the transformation. Of course, the eventual expression contains the unobserved V_t, which needs to be replaced with its implied value obtained via inverting VIX_t at some parameter value. Since VIX_t is observed and fixed, V_t also becomes a function of the unknown parameters. Needless to say, the inversion must be unambiguous. Our model is clearly the case, because the relationship linking V_t to VIX_t in equation (13) is always invertible at all parameter values.

Denote the parameters by $\Theta = (\kappa, \theta, \lambda, \mu_J, \sigma_J, \upsilon, \rho, \gamma, \delta_S, \kappa^*, \phi^*)$. The observed data sample consists of N observations with each data point being denoted by $X_{t_i} = (\ln S_{t_i}, \mathrm{VIX}_{t_i})$. Let $\widehat{Y}_{t_i}(\Theta) = (\ln S_{t_i}, \widehat{V}_{t_i}(\Theta))$ where $\widehat{V}_{t_i}(\Theta)$ is the inverted value evaluated at parameter value Θ according to equation (13).

Since our model has jumps, the conditional density function for $\widehat{Y}_{t_i}(\Theta)$ will be a Poisson mixture of the bivariate normal densities in the following form:

$$f\left(\widehat{Y}_{t_i}(\Theta)\big|\widehat{Y}_{t_{i-1}}(\Theta);\Theta\right) = \sum_{j=0}^{\infty} \frac{e^{-\lambda h_i}(\lambda h_i)^j}{j!} g\left(\mathbf{w}_{t_i}(j,\Theta); \mathbf{0}, \mathbf{\Omega}_{t_i}(j,\Theta)\right), \tag{14}$$

where

$$\mathbf{w}_{t_i}(j,\Theta) = \left[\begin{array}{c} \ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right) - \left[r - q + (\delta_s - \frac{1}{2})\widehat{V}_{t_{i-1}}(\Theta)\right] h_i - (j - \lambda h_i)\mu_J \\ \widehat{V}_{t_i}(\Theta) - \widehat{V}_{t_{i-1}}(\Theta) - \kappa\left(\theta - \widehat{V}_{t_{i-1}}(\Theta)\right) h_i \end{array}\right], \tag{15}$$

$h_i = t_i - t_{i-1}$, and $g(\cdot; \mathbf{0}, \mathbf{\Omega}_{t_i}(j,\Theta))$ is a bivariate normal density function with mean $\mathbf{0}$ and variance-covariance matrix:

$$\mathbf{\Omega}_{t_i}(j,\Theta) = \left[\begin{array}{cc} \widehat{V}_{t_{i-1}}(\Theta)h_i + j\sigma_J^2 & \rho\upsilon\widehat{V}_{t_{i-1}}^{0.5+\gamma}(\Theta)h_i \\ \rho\upsilon\widehat{V}_{t_{i-1}}^{0.5+\gamma}(\Theta)h_i & \upsilon^2\widehat{V}_{t_{i-1}}^{2\gamma}(\Theta)h_i \end{array}\right]. \tag{16}$$

Thus, the log-likelihood function corresponding to the asset prices and the VIX indices can be written as

$$\mathcal{L}\left(\Theta; X_{t_1}, \cdots, X_{t_N}\right) = \sum_{i=1}^{N} \ln f\left(\widehat{Y}_{t_i}(\Theta)\big|\widehat{Y}_{t_{i-1}}(\Theta);\Theta\right) - N\ln\left(\frac{1-e^{-\kappa^*\tau}}{\kappa^*\tau}\right), \tag{17}$$

In the above, the first component of the right-hand side is the log-likelihood function associated with $(\ln S_{t_i}, V_{t_i})$ whereas the second component corresponds to the Jacobian for the transformation from VIX_t to V_t. Note that the above log-likelihood function has been derived using the Euler discretization to equations (1) and (2). If h_i is small such as a sample of daily data, the discretization bias is expected to be negligible.

4 Empirical analysis

4.1 Data description

The data set consists of the S&P 500 index values, the CBOE's VIX index values and the risk-free rates on the daily frequency over the period from January 2, 1990 to December 29,

2006. The VIX index measures the market's expectation of the 30-day (or 22 trading days) forward S&P 500 index volatility implicit in the index option prices. The CBOE launched the VIX index in 1993 and switched to the new VIX index in September 2003. The VIX index values used in this study are the new VIX index series provided by the CBOE.[2] Our proxy for the risk-free rate is the continuously compounded one-month LIBOR rate.

Table 1: Summary Statistics (January 2, 1990 – December 29, 2006)

	S&P500 return	VIX
Mean	0.00032	19.0566
Standard deviation	0.0099	6.4286
Skewness	-0.1021	0.9840
Excess Kurtosis	3.9154	0.7975
Maximum	0.0557	45.7400
Minimum	-0.0711	9.3100

Table 1 provides some basic statistics on the S&P 500 index returns and the VIX index values. The index return is clearly negatively skewed and with heavy tails. Note that the VIX index is stated as percentage points per annum. The summary statistics indicate that based on the VIX index, the S&P 500 index return had about 19% annualized volatility over the sample period. Volatility should be naturally skewed in the positive direction, which is indeed the feature of the VIX index. The result shows that the VIX index also reveals a minor degree of heavy tails. The stochastic volatility phenomenon is also fairly clear with the volatility ranging from 9.3% to 45.7% over the sample period.

Figure 1 plots the time series of the S&P 500 and VIX indices over the 17-year period. Added as a comparison is the S&P 500 index return's realized volatility calculated from the subsequent 22 trading days for which the VIX index is intended to measure. There are several noticeable features. The market experienced a steady run-up in the 90's, and became jittery towards the end of 90's. Then the Dot-Com bubble burst which brought down the market until its recovery in 2003. Since then the market volatility has been in a steady decline until reaching the middle of 2006 with a noticeable spike. Comparing the VIX index to the realized volatility reveals an interesting and important fact; that is, the VIX has consistently been higher than the realized volatility throughout the sample period. Since the VIX index is meant to be the risk-neutral expected realized volatility, it suggests

[2]The old VIX index (the current ticker symbol is VXO) uses 8 implied volatilities of S&P 100 (OEX) options to approximate a hypothetical at-the-money OEX option with 30 days to maturity. The new VIX index is constructed by all valid out-the-money S&P 500 index (SPX) calls and puts. Both data series (VIX and VXO) are available on the CBOE's website. The new VIX tracks VXO reasonably closely, but the new VIX index tends to be slightly lower based on a chart in the CBOE's white paper.

Figure 1: The S&P 500 index, the VIX index and the Corresponding Realized Volatility

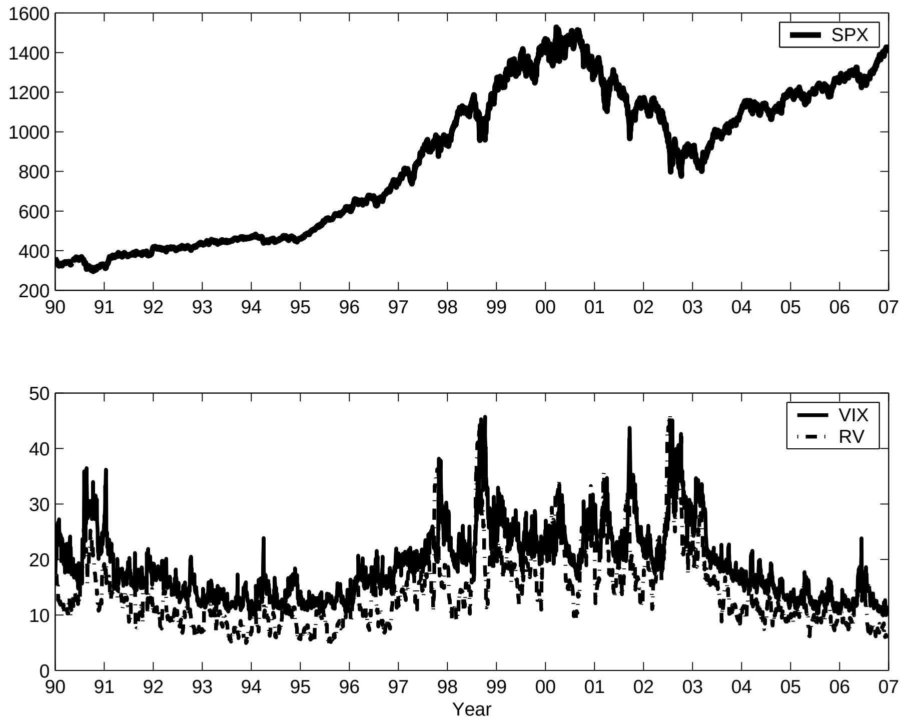

that the volatility dynamic under the risk-neutral pricing measure must be different from that under the physical probability measure. In other words, the volatility risk has mostly likely been priced by the market.

4.2 Empirical results

Table 2 summarizes our maximum likelihood estimation and inference results on three versions of the stochastic volatility model, where SV0 denotes the stochastic volatility model with an unconstrained CEV parameter γ, SV1 is the Hull and White (1987) stochastic volatility model with $\gamma = 1$, and SV2 corresponds to the Heston (1993) stochastic volatility model with $\gamma = 1/2$. The parameter estimates along with their corresponding standard errors inside the parentheses are reported in this table. LR denotes the likelihood ratio test statistic with its corresponding p value given inside the parentheses.

When the CEV parameter γ is unconstrained (SV0), its estimate is 0.9141 for the entire data sample. Using sub-samples, the estimates range from 0.9534 to 0.9825. These estimates indicate that the popular square-root specification for the volatility dynamic is strongly at odd with the data, whereas $\gamma = 1$, a specification adopted in Hull and White (1987), appears to be a better constraint to use. In comparison to the results reported in the literature, we note that Jones (2003) has estimates from 0.84 to 1.5, Ait-Sahalia and Kimmel (1995) have an estimate around 0.65, and Bakshi, Ju and Ou-yang (2006) have estimates from 1.2 to 1.5. The difference can of course be attributed to different methodologies and data samples. With the exception of Ait-Sahalia and Kimmel (1995) perhaps, all results strongly point to the inappropriateness of the square-root volatility specification. Using the formal likelihood ratio test, we have found the square-root volatility specification ($\gamma = 1/2$) is resoundingly rejected in all cases. In contrast, the stochastic volatility model with $\gamma = 1$ is rejected in the whole data sample, but passed the likelihood ratio test at the 10% level in all sub-samples.

Surprisingly perhaps, the estimates for the volatility risk premium turn out to be fairly stable and highly significant in all cases. The estimated volatility risk premium is negative, a result reflective of the fact that the VIX index has been higher than the corresponding realized volatility as shown in Figure 1. In fact, the magnitude of the estimated volatility risk premium is so large that it makes the volatility process not to be mean-reverting under the risk-neutral probability measure (i.e., $\kappa^* < 0$) even though the volatility process under the physical probability measure is mean-reverting.

The correlation between the price and volatility innovations is found to be significantly negative, a well-known empirical fact. The conclusion is robust over different models and data periods, and the estimates are fairly stable as well. An interesting issue to note is the estimated mean jump is positive, meaning that jumps are on average in the positive direction. As a controlled comparison, we forced the correlation between the price and volatility innovations to zero and re-estimated the jump model. The result indicates a negative mean jump. By allowing a correlation between the price and volatility innovations, we have in effect removed the negative return asymmetry in returns. We can therefore conclude that the appearance of negative jumps can be induced when one fails to properly remove the effect of stochastic volatility.

Table 3 summarizes the maximum likelihood estimation results for the stochastic volatility models with jumps on the whole data sample. We denote the model with an unconstrained γ by SVJ0. The model corresponding to $\gamma = 1$ is SVJ1 and the one corresponding to $\gamma = 1/2$ is SVJ2. The reported results clearly indicate the presence of jumps. The estimates associated with jumps – λ, μ_J and σ_J – are significant in all cases. The log-likelihood value increases substantially moving from SV0 to SVJ0. Although the table does not provide the likelihood ratio test statistic on the presence of jumps, the difference in the log-likelihood values clearly reveals that the test based on 4 degrees of freedom (four more parameters) would be highly significant. This finding is consistent with Bates (2000), Anderson, Benzoni and Lund (2002), Pan (2002) and Eraker, Johannes and Polson (2003). Our jump intensity estimate indicates 82 price jumps per annum (based on SVJ0). As compared to the results in Bate (2002), Pan (2002) and Eraker (2004), our result implies more frequent small jumps because their estimates are from 3 to 27 jumps per year.

The volatility risk premium continues to be significantly negative and reasonably stable across models. For the jump risk premium, the results are quite different. In the case of the model with an unconstrained γ (SVJ0) and that with $\gamma = 1$ (SVJ1), the jump risk premium is significantly negative. But in the case of square-root volatility specification (SVJ2), the jump risk premium becomes positive and insignificant, suggesting that the jump risk premium is sensitive to how the stochastic volatility process is specified. The log-likelihood value for SVJ2 is so much smaller than that of SVJ0, suggesting that the square-root volatility specification with the presence of price jumps continues to be resoundingly rejected based on the likelihood ratio criterion. The model specification used by Pan (2002) and Eraker (2004) thus appears to be at odd with the data. Jones (2003) argued that the volatilities generated from the square-root volatility process are too smooth to reconcile with the reality. By introducing jumps, the burden on the stochastic volatility to generate returns on extreme tails is significantly lessened, but it appears to be not enough to just allow for smooth stochastic volatilities. This is not at all surprising, however, knowing that the VIX index has been volatile throughout the sample period. Without the VIX data, jumps can perhaps alleviate the deficiency associated with the square-root volatility specification. In a way, the presence of the VIX series refutes the square-root volatility specification simply because such a dynamic is at odd with the time series feature of the VIX index.

Tables 4, 5 and 6 respectively summarize the maximum likelihood estimation results for the stochastic volatility models with jumps on three sub-samples. The main findings for the whole sample continue to be valid in three sub-samples, suggesting that our earlier conclusions are quite robust. The parameter estimates for the jump component are significant in all cases but their magnitudes depend on the sample used. The jump risk premium is found not to be significant in sub-samples although their signs remain negative which is consistent with the result on the whole sample.

5 Conclusion

We have devised a new method to estimate the stochastic volatility model with/without jumps via the use of the VIX index. Applying the method to the data sample of the S&P 500 index and the VIX index over a period of 17 years, we have obtained the following findings: 1) incorporating a jump risk factor is critically important; (2) the jump and volatility risks are priced; and (3) the popular square-root stochastic volatility process is a poor model specification irrespective of allowing for price jumps or not.

Our estimation method is based on the transformed data technique, which in its present form can only accommodate one VIX-like data series. Although the CBOE has only produced the VIX series for the 30-day maturity, it is conceivable that one can generate another VIX for, say, the 90-day maturity and add it to the data set for analysis. Conceptually, using several VIX's corresponding to different maturities can help better pin down the risk-neutral price and volatility dynamics. This is certainly an area deserving further exploration.

References

1. Ait-Sahalia, Y. and R. Kimmel, 2005, "Maximum Likelihood Estimation of Stochastic Volatility Models," *Journal of Financial Economics*, forthcoming.

2. Anderson, T., L. Benzoni and J. Lund, 2002, "Towards an Empirical Foundation for Continuous-Time Equity Returns Models," *Journal of Finance*, 57, 1239-1284.

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

4. Bakshi, G., N. Ju and H. Ou-Yang, 2006, "Estimation of Continuous-Time Models with An Application to Equity Volatility Dynamics," *Journal of Financial Economics*, 82, 227-249.

5. Bates, D., 2000, "Post-'87 Crash Fears in S&P500 Future Options," *Journal of Econometrics*, 94, 181-238.

6. Britten-Jones, M. and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, 55, 839-866.

7. Broadie, M., M. Chernov and M. Johannes, 2006, "Model Specification and Risk Premia: Evidence From Future Options," *Journal of Finance*, forthcoming.

8. CBOE white paper on VIX, CBOE website.

9. Chernov, M. and E. Ghysel, 2000, "A Study Towards A Unified Approach to the Joint Estimation of Objective and Risk-Neutral Measures for the Purpose of Options Valuation," *Journal of Financial Economics*, 56, 407-458.

10. Demeterfi, K., E. Derman, M. Kamal and J. Zhou, 1999,"More Than You Ever Wanted to Know about Volatility Swaps," Goldman Sachs Quantitative Strategies Research Notes.

11. Duan, J., 1994, "Maximum Likelihood Estimation Using Price Data of the Derivative Contract," *Mathematical Finance*, 4, 155-167.

12. Duffie, D., J. Pan and K. Singleton, 2000, "Transform Analysis and Asset Pricing for Affine Jump-Diffusions," *Econometrica*, 68, 1343-1376.

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

14. Eraker, B., M. Johannes and N. Polson, 2003, "The Impact of Jumps in Equity Index Volatility and Returns," *Journal of Finance*, 58, 1269-1300.

15. Jiang, G. and Y. Tian, 2005, "The Model-Free Implied Volatility and its Information Content," *Review of Financial Studies*, 18, 1305-1342.

16. Jones, C., 2003, "The Dynamics of the Stochastic Volatility: Evidence from Underlying and Options Markets," *Journal of Econometrics*, 116, 181-224.

17. Pan, J., 2002, "The Jump-Risk Premia Implicit in Options: Evidence From an Integrated Time-Series Study," *Journal of Financial Economics*, 63, 3-50.

Table 2: Maximum Likelihood Estimation Results for Stochastic Volatility Models

	q	κ	θ	υ	ρ	γ	δ_S	κ^*	δ_V	LR
Sample period: 1990/1/2 – 2006/12/29										
SV0	-0.0499	1.5730	0.0311	1.4826	-0.6787	0.9141	-0.4166	-10.7051	-12.2781	
	(0.0385)	(0.6405)	(0.0113)	(0.0565)	(0.0060)	(0.0117)	(2.1312)	(0.5041)	(0.6339)	
SV1	-0.1167	0.0164	1.3751	1.9697	-0.6772	1	-3.6972	-11.9005	-11.9169	23.5638
	(0.0407)	(0.5755)	(48.1698)	(0.0204)	(0.0061)		(2.1168)	(0.4256)	(0.6086)	($p < 0.01$)
SV2	0.0887	5.5222	0.0262	0.3831	-0.6574	1/2	5.2136	-5.5810	-11.1032	975.76454
	(0.0259)	(0.5039)	(0.0024)	(0.0069)	(0.0069)		(2.1465)	(0.5630)	(0.6850)	($p < 0.01$)
Sample period: 1990/1/2 – 1995/12/29										
SV0	-0.1311	3.0107	0.0169	2.1952	-0.5499	0.9827	-7.0219	-13.1478	-16.1586	
	(0.0677)	(1.4468)	(0.0067)	(0.2446)	(0.0133)	(0.0304)	(6.1019)	(0.9775)	(1.2648)	
SV1	-0.1380	2.6131	0.0179	2.3439	-0.5499	1	-7.6747	-13.4143	-16.0274	0.1806
	(0.0677)	(1.2288)	(0.0080)	(0.0373)	(0.0133)		(6.0921)	(0.7857)	(1.2408)	($p = 0.67$)
SV2	-0.0123	7.4908	0.0183	0.3495	-0.5268	1/2	0.0281	-6.8919	-14.3827	265.2769
	(0.0515)	(1.1096)	(0.0027)	(0.0115)	(0.0141)		(5.7508)	(1.0534)	(1.3469)	($p < 0.01$)
Sample period: 1996/1/2 – 2000/12/29										
SV0	-0.2837	0.5674	0.1218	1.7260	-0.7512	0.9534	-6.4322	-9.3486	-9.9160	
	(0.1291)	(1.4202)	(0.2672)	(0.1681)	(0.0093)	(0.0355)	(4.3871)	(1.0086)	(1.1206)	
SV2	-0.3319	0.0215	2.2824	1.9825	-0.7512	1	-7.6197	-9.8729	-9.8945	1.4046
	(0.1166)	(1.1480)	(21.2755)	(0.0573)	(0.0092)		(3.9205)	(0.8033)	(1.0802)	($p = 0.24$)
SV2	-0.0520	6.1297	0.0360	0.4467	-0.7317	1/2	1.7280	-4.9276	-11.0574	224.6914
	(0.0698)	(1.0077)	(0.0059)	(0.0140)	(0.0104)		(3.6582)	(0.9386)	(1.1891)	($p < 0.01$)
Sample period: 2001/1/2 – 2006/12/29										
SV0	0.0153	1.9877	0.0266	1.6738	-0.7701	0.9662	-0.0044	-8.7431	-10.7309	
	(0.0625)	(1.0929)	(0.0117)	(0.0765)	(0.0101)	(0.0165)	(2.8592)	(1.0387)	(1.1467)	
SV1	0.0027	1.3327	0.0313	1.8632	-0.7695	1	-0.2373	-9.2721	-10.6049	2.3682
	(0.0570)	(0.9107)	(0.0199)	(0.0379)	(0.0101)		(2.6204)	(0.8917)	(1.1199)	($p = 0.12$)
SV2	0.2175	5.7710	0.0305	0.4031	-0.7207	1/2	6.8256	-1.7986	-7.5697	636.1014
	(0.0370)	(0.8136)	(0.0046)	(0.0160)	(0.0135)		(3.1963)	(1.2489)	(1.4362)	($p < 0.01$)

Note: SV0 denotes the stochastic volatility model with unconstrained γ; SV1 denotes the stochastic volatility model with $\gamma = 1$; SV2 denotes the stochastic volatility model with fixed $\gamma = 1/2$. The standard errors are inside the parentheses. The volatility risk premium δ_V is computed as $\kappa^* - \kappa$ and its standard error follows from the standard calculation. LR denotes the likelihood ratio test statistic with its corresponding p value.

Table 3: Maximum Likelihood Estimation Results for the Stochastic Volatility Models with Jumps (the Whole Sample)

	SV0	SVJ0	SVJ1	SVJ2
q	-0.0499	-0.0462	-0.0384	-0.0389
	(0.0385)	(0.0508)	(0.0568)	(0.0379)
κ	1.5730	2.5092	2.2654	3.0025
	(0.6405)	(0.9208)	(0.9282)	(0.6463)
θ	0.0311	0.0227	0.0240	0.0218
	(0.0113)	(0.0058)	(0.0063)	(0.0043)
λ		82.1814	46.2522	100.5712
		(13.4456)	(8.7741)	(12.8075)
$\mu_J(\%)$		0.2925	0.4114	0.1653
		(0.0465)	(0.0719)	(0.0307)
$\sigma_J(\%)$		0.5838	0.7207	0.5364
		(0.0347)	(0.0459)	(0.0319)
v	1.4826	1.3587	1.9272	0.3625
	(0.0565)	(0.0548)	(0.0208)	(0.0062)
ρ	-0.6787	-0.7915	-0.7779	-0.7524
	(0.0060)	(0.0089)	(0.0084)	(0.0084)
γ	0.9141	0.8816	1	1/2
	(0.0117)	(0.0116)		
δ_S	-0.4166	-0.4256	-0.0042	0.3522
	(2.1312)	(2.6864)	(2.8262)	(2.4268)
κ^*	-10.7052	-12.8872	-14.6401	-9.7069
	(0.5041)	(0.5551)	(0.5122)	(0.5945)
$\phi^*(\%)$		0.0350	-0.1354	0.2523
		(0.0326)	(0.0331)	(0.0384)
δ_V	-12.2781	-15.3964	-16.9056	-12.7094
	(0.6339)	(1.0160)	(0.9807)	(0.8746)
$\delta_J(\%)$		-0.1407	-0.2953	0.0936
		(0.0604)	(0.0610)	(0.0546)
Log-Lik	37313.0192	37463.93725	37452.8324	37067.0472

Note: The reported estimates for μ_J, σ_J, ϕ^* and δ_J have been multiplied by 100. SVJ0 denotes the stochastic volatility model with jumps and an unconstrained γ; SVJ1 denotes the stochastic volatility model with jumps and $\gamma = 1$; SVJ2 denotes the stochastic volatility model with jumps and $\gamma = 1/2$. δ_V and δ_J are computed by $\kappa^* - \kappa$ and $\phi^* - \lambda(e^{\mu_J + \sigma_J^2/2} - 1 - \mu_J)$, respectively, and their standard errors follow from the standard calculation.

Table 4: Maximum Likelihood Estimation Results for the Stochastic Volatility Models with Jumps (Sample Period: 1990/1/2 – 1995/12/29)

	SVJ0	SVJ1	SVJ2
q	-0.0351	-0.0306	-0.0964
	(0.0794)	(0.0809)	(0.0412)
κ	2.4165	2.4548	5.0656
	(2.4145)	(2.4431)	(1.4423)
θ	0.0189	0.0187	0.0128
	(0.0129)	(0.0125)	(0.0036)
λ	25.0972	24.3863	53.8559
	(8.3344)	(8.2510)	(12.7624)
$\mu_J(\%)$	0.9321	0.9431	0.2645
	(0.1462)	(0.1505)	(0.0791)
$\sigma_J(\%)$	0.0603	0.0562	0.6261
	(0.4862)	(0.5433)	(0.0695)
v	2.0315	2.3341	0.3607
	(0.2608)	(0.0380)	(0.0114)
ρ	-0.6465	-0.6468	-0.6374
	(0.0168)	(0.0168)	(0.0180)
γ	0.9635	1	1/2
	(0.0328)		
δ_S	-0.0211	-0.1051	-0.0907
	(7.7812)	(7.8709)	(5.2670)
κ^*	-17.1315	-17.6932	-8.7684
	(1.0321)	(0.8904)	(1.1469)
$\phi^*(\%)$	0.0092	-0.0234	0.2605
	(0.0638)	(0.0584)	(0.0477)
δ_V	-19.5480	-20.1480	-13.8341
	(2.5257)	(2.4588)	(1.7453)
$\delta_J(\%)$	-0.1005	-0.1326	0.1357
	(0.0878)	(0.0825)	(0.0769)
Log-Lik	13785.3385	13784.4074	13686.8489

Note: The reported estimates for μ_J, σ_J, ϕ^* and δ_J have been multiplied by 100. SVJ0 denotes the stochastic volatility model with jumps and an unconstrained γ; SVJ1 denotes the stochastic volatility model with jumps and $\gamma = 1$; SVJ2 denotes the stochastic volatility model with jumps and $\gamma = 1/2$. δ_V and δ_J are computed by $\kappa^* - \kappa$ and $\phi^* - \lambda(e^{\mu_J + \sigma_J^2/2} - 1 - \mu_J)$, respectively, and their standard errors follow from the standard calculation.

Table 5: Maximum Likelihood Estimation Results for the Stochastic Volatility Models with Jumps (Sample Period: 1996/1/2 – 2000/12/29)

	SVJ0	SVJ1	SVJ2
q	-0.0029	0.1220	0.0002
	(0.1749)	(0.1659)	(0.1195)
κ	5.0226	6.4439	6.7005
	(1.9356)	(1.9163)	(1.6097)
θ	0.0316	0.0289	0.0278
	(0.0064)	(0.0045)	(0.0049)
λ	83.3850	89.3408	193.2140
	(31.5592)	(32.2494)	(51.6809)
$\mu_J(\%)$	0.2388	0.2286	0.1651
	(0.0894)	(0.0841)	(0.0485)
$\sigma_J(\%)$	0.7213	0.7036	0.5150
	(0.0878)	(0.0821)	(0.0518)
v	2.0922	1.9266	0.4283
	(0.3042)	(0.0595)	(0.0144)
ρ	-0.8377	-0.8395	-0.8314
	(0.0159)	(0.0159)	(0.0154)
γ	1.0310	1	1/2
	(0.0547)		
δ_S	2.7338	8.4133	3.8556
	(6.4255)	(6.1649)	(5.2990)
κ^*	-14.6516	-14.1660	-9.0575
	(1.3586)	(1.0441)	(1.1278)
$\phi^*(\%)$	-0.4935	-0.5040	0.1872
	(0.1689)	(0.1142)	(0.1064)
δ_V	-19.6743	-20.6100	-15.7580
	(2.0562)	(1.9535)	(1.7552)
$\delta_J(\%)$	-0.7348	-0.7490	-0.0958
	(0.2482)	(0.1922)	(0.1679)
Log-Lik	10408.3277	10407.9546	10304.2353

Note: The reported estimates for μ_J, σ_J, ϕ^* and δ_J have been multiplied by 100. SVJ0 denotes the stochastic volatility model with jumps and an unconstrained γ; SVJ1 denotes the stochastic volatility model with jumps and $\gamma = 1$; SVJ2 denotes the stochastic volatility model with jumps and $\gamma = 1/2$. δ_V and δ_J are computed by $\kappa^* - \kappa$ and $\phi^* - \lambda(e^{\mu_J + \sigma_J^2/2} - 1 - \mu_J)$, respectively, and their standard errors follow from the standard calculation.

Table 6: Maximum Likelihood Estimation Results for the Stochastic Volatility Models with Jumps (Sample Period: 2001/1/2 – 2006/12/29)

	SVJ0	SVJ1	SVJ2
q	0.0004	-0.0180	0.0319
	(0.0812)	(0.0841)	(0.0534)
κ	1.7367	2.6626	2.9840
	(1.3290)	(1.3265)	(0.9519)
θ	0.0291	0.0212	0.0213
	(0.0162)	(0.0069)	(0.0069)
λ	24.3230	18.5726	288.2319
	(13.7381)	(12.4050)	(60.9738)
$\mu_J(\%)$	0.6452	0.7915	0.1003
	(0.3833)	(0.5472)	(0.0279)
$\sigma_J(\%)$	0.6303	0.6892	0.3216
	(0.1718)	(0.2575)	(0.0377)
v	1.6638	1.8396	0.3761
	(0.1025)	(0.0372)	(0.0137)
ρ	-0.8200	-0.8178	-0.7987
	(0.0118)	(0.0117)	(0.0127)
γ	0.9627	1	1/2
	(0.0228)		
δ_S	-0.5198	-0.1991	1.1725
	(3.5663)	(3.6307)	(3.5289)
κ^*	-9.9338	-10.7613	-7.3038
	(1.1248)	(0.9512)	(1.1569)
$\phi^*(\%)$	0.0117	-0.0620	0.3274
	(0.0549)	(0.0429)	(0.0388)
δ_V	11.6706	-13.4240	-10.2878
	(1.6704)	(1.5430)	(1.4932)
$\delta_J(\%)$	-0.0876	-0.1648	0.1637
	(0.1196)	(0.1289)	(0.0767)
Log-Lik	13435.4140	13435.3897	13186.0219

Note: The reported estimates for μ_J, σ_J, ϕ^* and δ_J have been multiplied by 100. SVJ0 denotes the stochastic volatility model with jumps and an unconstrained γ; SVJ1 denotes the stochastic volatility model with jumps and $\gamma = 1$; SVJ2 denotes the stochastic volatility model with jumps and $\gamma = 1/2$. δ_V and δ_J are computed by $\kappa^* - \kappa$ and $\phi^* - \lambda(e^{\mu_J + \sigma_J^2/2} - 1 - \mu_J)$, respectively, and their standard errors follow from the standard calculation.