



American Finance Association

Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices

Author(s): Bjorn Eraker

Source: *The Journal of Finance*, Vol. 59, No. 3 (Jun., 2004), pp. 1367-1403

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/3694740>

Accessed: 23/11/2009 12:57

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices

BJØRN ERAKER*

ABSTRACT

This paper examines the empirical performance of jump diffusion models of stock price dynamics from joint options and stock markets data. The paper introduces a model with discontinuous correlated jumps in stock prices and stock price volatility, and with state-dependent arrival intensity. We discuss how to perform likelihood-based inference based upon joint options/returns data and present estimates of risk premiums for jump and volatility risks. The paper finds that while complex jump specifications add little explanatory power in fitting options data, these models fare better in fitting options and returns data simultaneously.

THE STATISTICAL PROPERTIES of stock returns have long been of interest to financial decision makers and academics alike. In particular, the great stock market crashes of the 20th century pose particular challenges to economic and statistical models. In the past decades, there have been elaborate efforts by researchers to build models that explicitly allow for large market movements, or “fat tails” in return distributions. The literature has mainly focused on two approaches: (1) time-varying volatility models that allow for market extremes to be the outcome of normally distributed shocks that have a randomly changing variance, and (2) models that incorporate discontinuous jumps in the asset price.

Neither stochastic volatility models nor jump models have alone proven entirely empirically successful. For example, in the time-series literature, the models run into problems explaining large price movements such as the October 1987 crash. For stochastic volatility models, one problem is that a daily move of –22% requires an implausibly high-volatility level both prior to, and after the crash. Jump models on the other hand, can easily explain the crash of 1987 by a parameterization that allows for a sufficiently negative jump. However, jump models typically specify jumps to arrive with constant intensity. This assumption poses problems in explaining the tendency of large movements to cluster over time. In the case of the 1987 crash, for example, there were large movements both prior to, and following the crash. With respect to option prices,

*Eraker is from Duke University. I am indebted to Rick Green (the editor) and an anonymous referee for their helpful comments. I also thank Fredrico Bandi, Lars Hansen, Nick Polson, Pietro Veronesi, Mike Johannes, David Bates, Ken Singleton, and Luca Benzoni as well as seminar participants at the University of Chicago, Duke University, McGill University, University de Montreal, University of Toronto, London Business School, Norwegian School of Economics, Northwestern University, University of Iowa, and University of Minnesota for helpful comments and discussions. Remaining errors are mine alone.

this modeling assumption implies that a jump has no impact on relative option prices. Hence, a price jump cannot explain the enormous increase in implied volatility following the crash of 1987. In response to these issues, researchers have proposed models that incorporate both stochastic volatility and jumps components.

In particular, recent work by Bates (2000) and Pan (2002) examines combined jump diffusion models. Their estimates are obtained from options data and joint returns/options data, respectively.¹ They conclude that the jump diffusions in question do not adequately describe the systematic variations in option prices. Results in both papers point toward models that include jumps to the volatility process.

In response to these findings, Eraker, Johannes, and Polson (EJP) (2000) use returns data to investigate the performance of models with jumps in volatility as well as prices using the class of jump-in-volatility models proposed by Duffie, Pan, and Singleton (2000) (henceforth DPS). The DPS class of models generalizes the models in Merton (1976), Heston (1993), and Bates (1996). The results in EJP show that the jump-in-volatility models provide a significantly better fit to the returns data. EJP also provide decompositions of various large market movements which suggest that large returns, including the crash of 1987, are largely explained *ex post* by a jump to volatility. In a nutshell, the volatility-based explanation in EJP is driven by large subsequent moves in the stock market following the crash itself. This clustering of large returns is inconsistent with the assumption of independently arriving price jumps, but consistent with a temporarily high level of spot volatility caused by a jump to volatility. Overall, the results in EJP are encouraging with respect to the models that include jumps to volatility. In the current paper, therefore, we put the volatility jumping model to a more stringent test and ask whether it can explain price changes in both stock markets and option markets simultaneously.

The econometric technique in EJP is based on returns data only. By contrast, the empirical results presented in this paper are based on estimates obtained from joint returns and options data—an idea pursued in Chernov and Ghysels (2000) and Pan (2002). This is an interesting approach because even if option prices are not one's primary concern, their use in estimation, particularly in conjunction with returns, offers several advantages. A primary advantage is that risk premiums relating to volatility and jumps can be estimated. This stands in contrast to studies that focus exclusively on either source of data. Secondly, the one-to-one correspondence of options to the conditional returns distribution allows parameters governing the shape of this distribution to potentially be very accurately estimated from option prices.² For example, EJP report fairly wide posterior standard deviations for parameters that determine

¹ Stochastic volatility option pricing has been considered in Hull and White (1987), Wiggins (1987), Scott (1987), Chesney and Scott (1989), Melino and Thornbull (1990), Stein and Stein (1991), and Amin and Ng (1993) among others.

² A number of papers obtain non-parametric estimates of return distributions using options. See Jackwerth and Rubenstein (1996), Dumas, Fleming, and Whaley (1998) and among others.

the jump sizes and jump arrival intensity. Their analysis suggests that estimation from returns data alone requires fairly long samples to properly identify all parameters. Hopefully, the use of option prices can lead to very accurate estimates, even in short samples. Moreover, the use of option prices allows, and in fact requires, the estimation of the latent stochastic volatility process. Since volatility determines the time variation in relative option prices, there is also a strong potential for increased accuracy in the estimated volatility process. Finally, joint estimation also raises an interesting and important question: Are estimates of model parameters and volatility consistent across both markets? This is the essential question to be addressed in this paper.

Previously, papers by Chernov and Ghysels (2000), and Pan (2002) have proposed GMM-based estimators for joint options/returns data using models similar to the ones examined here, but without the jump to volatility component. In this paper we develop an approach based on Markov Chain Monte Carlo (MCMC) simulation. MCMC allows the investigator to estimate the posterior distributions of the parameters as well as the unobserved volatility and jump processes. Recent work by Jacquier and Jarrow (2000) points to the importance of accounting for estimation risk in model evaluation. This is potentially even more important in our setting because the parameter space is so highly dimensional. For example, one practical implication of MCMC is that the filtered volatility paths obtained by MCMC methods tend to be more erratic than estimates obtained by other methods (see Jacquier, Polson, and Rossi (1994)). This is important because previous studies find that estimates of the “volatility-of-volatility” parameter governing the diffusion term in the volatility process, is too high to be consistent with time-series estimates of the volatility process.

The empirical findings reported in this paper can be summarized as follows. Parameter estimates obtained for the (Heston) stochastic volatility model as well as the (Bates (1996)) jump diffusion with jumps in prices, are similar to those in Bakshi, Cao, and Chen (1997). In particular, our estimates imply a jump every other year on average, which compared to estimates in EJP from returns data alone, is very low. The volatility-of-volatility estimates are higher than those found from returns data alone in EJP. However, we do not conclude that they are inconsistent with the latent volatility series. Our posterior simulations of the latent volatility series have “sample volatility-of-volatility” that almost exactly matches those estimated from the joint returns and options data. This evidence contrasts with the findings of model violations reported elsewhere.

Evidence from an in-sample test reported in this paper shows surprisingly little support for the jump components in option prices. The overall improvement in in-sample option price fit for the general model with jumps in both prices and volatility, is less than 2 cents relative to the stochastic volatility model. There is some evidence to suggest that this rather surprising finding can be linked to the particular sample period used for estimation. In particular, we show that the models tend to overprice long dated options out-of-sample—a finding that can be linked to the high volatility embedded in options prices during the estimation period. If the mean volatility parameter is adjusted to match its historical average, out-of-sample pricing errors drop dramatically, and the

jump models perform much better than the SV model. The option pricing models also seem to perform reasonably well whenever the calculations are based on parameter estimates obtained using only time-series data of returns on the underlying index.

The jump models, and particularly the jump-in-volatility model, are doing a far better job of describing the time-series dimension of the problem. In particular, the general model produces return residuals with a sample kurtosis of a little less than four (one unit in excess of the hypothesized value under the assumed normal distribution). However, all models in question produce too heavily tailed residuals in the volatility process for it to be consistent with its assumed square root, diffusive behavior. This obtains even when the volatility process is allowed to jump. This model violation is caused primarily by a sequence of large negative outliers following the crash of 1987, and corresponds to the fall in the implied spot volatility process following the huge increase (jump) on the day of the crash. Since the jump-in-volatility model allows for positive jumps to volatility, it does explain the run-up in relative option prices on the day of the crash, but it has problems explaining the subsequent drop. In conclusion, the jump-in-volatility model does improve markedly on the simpler models, but its dynamics do not seem sufficiently general to capture variations in both returns and options markets simultaneously.

The rest of the paper is organized as follows: In the next section, we present the general model for the stock price dynamics as formulated in Duffie et al. (2000) special cases of this model, and discuss implications for option pricing. Section II discusses the econometric design and outlines a strategy for obtaining posterior samples by MCMC. Section III presents the data, while Section IV contains the empirical results. Section V summarizes the findings and suggests directions for further investigation.

I. Models and Pricing

We now outline the diffusion and jump dynamics that form the basis for the option pricing analysis in Duffie et al. (2000). Under the objective probability measure, the dynamics of stock prices, S , are assumed to be given by

$$\frac{dS_t}{S_t} = a dt + \sqrt{V_t} dW_t^S + dJ_t^S, \quad (1)$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma_V \sqrt{V_t} dW_t^V + dJ_t^V, \quad (2)$$

where V is the volatility process. The Brownian increments, dW^S and dW^V are correlated and $E(dW_t^S dW_t^V) = \rho dt$. The parameter a measures the expected return.³

The jump term, $dJ_t^i = Z_t^i dN_t^i$, $i = \{S, V\}$, has a jump-size component Z_t^i and a component given by a Poisson counting process N_t^i . In the case of a common

³ Attempts to fit a more elaborate drift under the objective measure were undertaken in EJP and shown to be unimportant.

Poisson process, $N_t^V = N_t^S$, jump sizes Z_t^S and Z_t^V can be correlated. The specification of jump components determines the models to be considered. We discuss the various possible specifications and their implications for option prices next.

A. SV Model

The basic stochastic volatility (SV) model with a square root diffusion driving volatility, was initially proposed for option pricing purposes in Heston (1993). It obtains as a special case of the general model in this paper with jumps restricted to zero ($dJ_t^S = dJ_t^V = 0$). The volatility process, V_t , captures serial correlation in volatility. The κ and θ parameters measure the speed of mean reversion and the (mean) level of volatility, respectively. The parameter σ is commonly referred to as the “volatility-of-volatility.” Higher values of σ_V imply that the stock price distribution will have fatter tails. The correlation parameter, ρ , is typically found to be negative which implies that a fall in prices usually will be accompanied by an increase in volatility which is sometimes referred to as the “leverage effect” (Black (1976)). A negative ρ implies that the conditional (on initial stock price, S_t , and volatility, V_t) returns distribution is skewed to the left.

The option pricing implications of the SV specification have been carefully examined. We summarize its properties in the following:

1. For finite and strictly positive return horizons, the excess kurtosis of the conditional returns distribution is everywhere positive (Das and Sundaram (1999)) and increasing in both σ_V and $|\rho|$. Positive kurtosis implies concave, U shaped implied Black–Scholes volatility (IV) curves. Hence, the concavity of the IV curves is increasing in σ_V and $|\rho|$.
2. Skewness of the conditional returns distribution is positive/zero/negative for positive/zero/negative values of ρ . Correspondingly, for negative ρ , long maturity contracts have a downward sloping IV curve across strikes (far in-the-money contracts tend to be relatively more expensive) and vice versa for positive ρ (again, see Das and Sundaram (1999)).
3. Depending on the initial volatility, V_t , the IV curves will be upward (low V_t) or downward (high V_t) sloping in the maturity of the contracts.
4. The conditional returns distribution converges to a normal as the holding period approaches either zero or infinity (assuming proper regularity). As a consequence, the model predicts increasingly flatter IV curves for both very long and very short maturity options. Contracts with a moderately long time to maturity have both steeper and more concave Black–Scholes implied volatility curves across different strikes.

B. SVJ Model

The stochastic volatility with jumps (SVJ) model is an extension to the SV model that allows random jumps to occur in the prices. Specifically, J_t^S is a pure Poisson process, and $J_t^V = 0 \forall t$. It is assumed that jump sizes are distributed

$$Z_t^S \sim N(\mu_y, \sigma_y^2). \quad (3)$$

It is natural to think of this model as one that adds a mixture component to the returns distribution. Essentially, this component adds mass to the tails of the returns distribution. Increasing σ_y adds tail mass to both tails while a negative (say) μ_y implies relatively less mass in the right tail, and vice versa.

The SVJ model inherits the option pricing implications one through three for the SV model. Noticeably however, at short time horizons, the transition density is non-Gaussian. This implies, relative to the SV model, that the IV curves for short maturity options will be steeper whenever the jump size mean is negative. Moreover the IV curves are typically more U shaped for larger values of the jump size variance, σ_y^2 .

C. SVCJ Model

In this model volatility is allowed to jump. Jumps to volatility and prices are driven by the same Poisson process, $J^S = J^V$. This allows jump sizes to be correlated

$$Z_t^V \sim \exp(\mu_V), \quad (4)$$

$$Z_t^S | Z_t^V \sim N(\mu_y + \rho_J Z_t^V, \sigma_S^2). \quad (5)$$

Thus, this model is labeled stochastic volatility with correlated jumps (SVCJ).

In essence, this model assumes that price jumps will simultaneously impact both prices and volatility. Notice that the leverage effect built into the basic SV model, is cluttered in the SVJ model because only small price movements resulting from Brownian shocks will have an impact on volatility. Large price moves stemming from jumps on the other hand, have no impact on volatility in the SVJ model.

The SVCJ specification corrects this shortcoming in the SVJ model. Whenever the ρ_J parameter is negative, the larger a market crash, the more its volatility will increase. Moreover, it is possible in this model for small price changes to have no discernible impact on volatility while large price changes (jumps) do (i.e., $\rho = 0$ and $|\rho_J| > 0$). It is also possible that this model attributes large market movements entirely to increases in volatility by setting the parameters in the price jump distribution, μ_y , ρ_J , and σ_y , to zero.

The option pricing implications for the SVCJ model are quite similar to those of the SVJ model. The added volatility jump component, however, will typically add right skewness in the distribution of volatility, and hence, overall fatten the tails of the returns distribution. Pan (2002) argues that the addition of a jumping volatility component might explain her findings of a severely pronounced increase in the volatility smile at short maturity, far in the money put options. Although adding volatility jumps into the model primarily increases kurtosis in the returns distribution, it is also possible for the SVCJ model to add skewness into the conditional returns distribution through the parameter ρ_J .

D. SVSCJ Model

The stochastic volatility with state-dependent, correlated jumps (SVSCJ) is the most general model considered in this paper. It generalizes the correlated jump model described above to allow the jump frequency to depend on volatility,

$$\lambda_0 + \lambda_1 V_t. \quad (6)$$

Allowing for volatility jumps, this model generalizes models with state-dependent price jumps considered by Bates (1996) and Pan (2002). A property of the previous models is that in high-volatility regimes, the diffusive, Gaussian part of the system will tend to dominate the conditional distribution. As a result, as the spot volatility V_t increases, short-term option smiles will tend to flatten. The SVSCJ model relaxes this restriction by allowing jumps to arrive more frequently in high-volatility regimes. A previous draft of this paper presented empirical evidence suggesting that this generalization is potentially empirically important.

E. Pricing

In discussing the model implications for implied Black–Scholes volatility above, we implicitly assumed that option prices were obtained as the expected payoff under the observed, “objective” probability measure, P . In the well established theory of arbitrage pricing, there is an equivalence between the absence of arbitrage and the existence of a “risk neutral” probability measure, Q . Option prices are computed as

$$C_t = E^Q \{ e^{-r_t(T-t)} (S_T - X)^+ \mid \mathcal{F}_t \}.$$

See Duffie et al. (2000) for details. The difference between the two measures follows from model assumptions. Here, we impose standard risk premium assumptions in the literature and parameterize the models under the equivalent measure Q , as

$$\frac{dS_t}{S_t} = (r - \mu^*) dt + \sqrt{V_t} dW(Q)_t^S + dJ(Q)_t^S, \quad (7)$$

$$dV_t = (\kappa(\theta - V_t) + \eta_V V_t) dt + \sigma_V \sqrt{V_t} dW(Q)_t^V + dJ(Q)_t^V, \quad (8)$$

where μ^* is the jump compensator term, and where $dW(Q)_t$ is a standard Brownian motion under Q defined by $dW(Q)_t^i = \eta_i dt + dW_t^i$ for $i = V, J$. Notice that the volatility is reverting at the rate $\kappa^Q := \kappa - \eta^V$ where η_V is a risk premium parameter associated with shocks to the volatility process. A similar risk premium, η_J , is assumed to be associated with jumps. This gives the price jump distribution under Q as

$$Z_t^S \mid Z_t^V \sim N(\mu_y^Q + \rho_J Z_t^V, \sigma_S^2). \quad (9)$$

Notice that it is generally very difficult to identify these risk premium parameters. Consider for example the parameter η_J : while we can use the relationship $\eta_J = \mu_y^Q - \mu_y$ to estimate η_J , μ_y is only identified through time series observations on the underlying security. The problem, in a nutshell, is that we only observe a fraction λ of a T long sample of jump dates for which we have data available to estimate μ_y . Hence, whenever λ approaches zero, the likelihood function is uninformative about μ_y .

In order to compute theoretical options prices using the proposed stock price dynamics above, we must proceed to invert characteristic functions of the transition probability for the stock price. The characteristic function takes on an exponential affine form $\phi(u, \tau) = \exp(\alpha(u, \tau) + \beta(u, \tau)V_t + u \ln S_t)$ where the coefficients β and α solve ordinary differential equations (ODE). For the SVSCJ model, general form for these ODEs can be found in Duffie et al. (2001). For the SVCJ model, DPS give analytical expressions for these coefficients.

The advantage of the SVCJ model and its simplifications, is that option prices are available in semiclosed form through a numerical Fourier inversion. This involves a numerical integration of the imaginary part of the complex Fourier transform. This numerical integration needs to be carried out repeatedly in our MCMC sampler. Hence, it is crucial to the performance of the algorithm that this integration can be rapidly performed.

II. Econometric Methodology

While a number of methods have been proposed for the estimation of diffusion processes,⁴ the latent nature of volatility as well as the jump components complicates estimation. In recent work Singleton (2001), Pan (2002), and Viceira and Chacko (1999) develop GMM-based procedures exploiting the known characteristic functions of affine models. Typically, however, methods based on simulation have been employed for analysis. These methods include the method of simulated moments (Duffie and Singleton (1993)), indirect inference methods (Gourieroux, Monfort, and Renault (1993)) and the efficient method of moments (EMM) (Gallant and Tauchen (1996)) among others. The generality of simulation-based methods offers obvious advantages. For instance, Andersen, Benzoni, and Lund (2002) use EMM to estimate jump diffusion models from equity returns. Building on work for general state space models in Carlin, Polson, and Stoffer (1992), Jacquier et al. (1994) pioneered a method for estimating discrete time stochastic volatility models from returns data. Their work showed that MCMC is particularly well suited to deal with stochastic volatility models. Extensions into multivariate models can be found in Jacquier, Polson, and Rossi (2003). Eraker (2001) proposes a general

⁴ See Aït-Sahalia (2002) (analytic approximation to the likelihood function), Duffie and Glynn (1996), Hansen and Scheinkman (1995) and Conley, Hansen, Luttmer, and Scheinkman (1997) (moment based methods), and Aït-Sahalia (1996), Jiang and Knight (1997), Stanton (1997), and Bandi and Phillips (1999, 2000), (non-parametric methods) among others.

approach to estimating diffusion models with arbitrary drift and diffusion functions and possibly latent, unobserved state variables (e.g., stochastic volatility). MCMC-based estimation of jump diffusion models based on returns data, has been considered in Johannes, Kumar, and Polson (1999) and Eraker et al. (2003).

In the following, we outline the principles on which we construct an MCMC sampler to estimate the option price parameters and the pricing errors. Although the discussion is primarily concerned with option prices, it can be applied to virtually any estimation problem involving derivative assets for which theoretical model prices are explicitly available. For example, the analysis would carry over straightforwardly to bond pricing/term structure applications.

At the current level of generality, assume that option prices are determined by a set of state variables, $X_t = \{S_t, V_t\}$, a parameter vector Θ , in addition to other arguments such as time to maturity etc., χ , through a known function F :

$$P_t = F(X_t, \chi, \Theta). \quad (10)$$

The theoretical prices P_t are computed through the Fourier inversion mentioned above. The theoretical prices do not depend on the jump times or jump sizes but only the current price S_t and volatility V_t .

We assume that there are observations available on n different assets recorded over a period $[0, T]$. There are m_i recorded prices for contract i so that $M = \sum_{i=1}^n m_i$ is the total number of observed option prices. Let s_i and T_i denote the first and the last observation times for contract i .

There are two essential possibilities for the design of the database to be used:

1. Observations are recorded at discrete, integer times so that the total number of time periods under considerations is T . For example, we could record daily or weekly closing prices on, say, n assets to obtain a total of $n \times T$ data points.
2. Observations are available at transaction times. Typically, the use of transactions data will result in unequally spaced observations at times $t_i, i = 0, \dots, N$. Let $S(t)$ denote the collection of assets for which we have observations at time t . Although we assume that this is not the case, we will typically only observe one asset trading at the time, in which case N equals the total sample size, M .

Market prices, $Y_{t_i,j}$, are assumed to be observed with pricing errors $\varepsilon_{t_i,j}$

$$Y_{t_i,j} = F(X_{t_i}, \chi_i, \Theta) + \varepsilon_{t_i,j}.$$

We wish to allow for a serial dependence in the pricing errors, $\varepsilon_{t_i,j}$, of each asset j . This reflects the prior belief that if an asset is misspriced at time t , it

is also likely to be misspriced at $t + 1$. For consecutive observations on option j at times $t - 1$ and t , as in case (1) above, we assume that the pricing errors are distributed

$$\varepsilon_{t,j} \sim N(\rho_j \varepsilon_{t-1,j}, s_j^2)$$

and that $\varepsilon_{t,i}$ is independent of $\varepsilon_{t,j}$, for $j \neq i$.

Whenever transactions data are recorded at random transaction times t_i (case 2), the corresponding distributional assumption with respect to the pricing errors is

$$\varepsilon_{t_i,j} \sim N\left(\varepsilon_{t_{i-1},j} e^{\rho_j \nabla t_i}, s_j^2 \frac{1}{2\rho_j} (1 - e^{-2\rho_j \nabla t_i})\right). \quad (11)$$

This corresponds to the simple AR(1) model for the pricing errors as in case (1), but with observations recorded at fractional, random times. Notice that this specification is observationally equivalent to assuming that the pricing errors follow independent Ornstein–Uhlenbeck processes. Accordingly, we must also impose a prior restriction on the autoregressive parameters ρ_j to be strictly negative.

It is, in principle, possible to relax the assumption of cross-sectionally independent pricing errors. Multiple complications arise, however. First, as the number of contracts pairs which have overlapping observations is limited, correlations, say $r_{i,j}$, between pricing errors of contracts i and j would only be identified through prior assumptions whenever the sampling intervals of the two do not overlap. Bayesian analysis would of course allow this, arguably less restrictively so by assuming a hyper-distribution on $r_{i,j}$'s. Second, an equally serious problem would involve the joint identification of all model quantities if we were to allow for such cross-sectional correlation in the pricing errors. A latent factor structure $\varepsilon_{t,i} = C_i e_t + e_{t,i}$ with e_t being a common factor illustrates the problem: Here, e_t would “compete” with the volatility factor, V_t , in explaining cross-sectional shifts in prices. As such, the introduction of a common factor in the residual terms is likely to increase estimation error in V_t , as well as parameters that govern its dynamics.

A. Joint Posterior Density

The specification of the error distribution, together with a specification of the full prior of all model parameters Θ , completes the specification of joint density for observed data, model parameters, and the state variables. This joint density is the key quantity needed to derive an MCMC sampler for the problem. In the following, we therefore outline the specification of this density.

First, the specification of the pricing errors allows us to write the conditional density for the joint observations, $Y = \{Y_{t_i,j} : i = 1, \dots, N, j = 1, \dots, n\}$ as

$$\begin{aligned} p(Y | X, \Theta) &\propto \prod_{i=1}^N \prod_{j \in S(t)} \phi(Y_{t_i,j}; F(X_{t_i}, \chi_j, \theta) + \rho_j \varepsilon_{t_i,j}, s_j^2) \\ &=: \prod_{i=1}^N p(Y_{t_i} | X_{t_i}, \varepsilon_{t_i}, \Theta), \end{aligned}$$

where $\phi(x; m, s^2)$ denotes a normal density with mean m and variance s^2 evaluated at x and $\Theta := \{\theta, \rho_j, s_j, j = 1, \dots, n\}$.

Even though the option prices do not depend on the jump times and the jump sizes, the dynamics of the prices and volatility, jointly labeled X_t , of course do depend on these quantities.

The joint posterior density of options data Y , state variables, X , jump times, J , and jump sizes Z , and parameters Θ can now be written

$$\begin{aligned} p(Y, X, Z, J, \Theta) &\propto p(Y | X, \Theta) p(X | \theta) p(Z | J, \Theta) p(J) p(\Theta) \\ &= \prod_{i=1}^N p(Y_{t_i} | X_{t_i}, \varepsilon_{t_i}, \Theta) p(X_{t_i} | X_{t_{i-1}}, Z_{t_i}, J_{t_i}, \Theta) \\ &\quad \times p(Z_{t_i} | J_{t_i}, \Theta) p(J_{t_i}, \Theta) p(\Theta), \end{aligned} \quad (12)$$

where $p(X_{t_i} | X_{t_{i-1}}, Z_{t_i}, J_{t_i}, \Theta)$ is the transition density of the jump diffusion process and $p(\Theta)$ is the prior for Θ .

The density, $p(X_{t_i} | X_{t_{i-1}}, Z_{t_i}, J_{t_i}, \Theta)$, is generally unknown. There are two ways of estimating this density: (1) the Fourier transform of the transition density $p(X_{t_i} | X_{t_{i-1}}, \theta)$ can be found in closed form so the Levy inversion formula can be used to evaluate. This is potentially time consuming because, unlike evaluation of the option prices, this Fourier inversion requires a two-dimensional numerical integration to obtain the joint density for $X_t = (S_t, V_t)$. Also, working with the transition density directly eliminates the need to simulate J and Z . A second possibility is to use a Gaussian approximation corresponding to the Euler discretization of the process:

$$X_{t_i} - X_{t_{i-1}} = \mu(X_{t_{i-1}}; \theta) \nabla t_i + \sigma(X_{t_{i-1}}; \theta) \nabla W_{t_i} + J_{t_i} Z_{t_i}. \quad (13)$$

This is reasonable as long as the observation intervals, t_i , are not too far apart. In a simulation study, Eraker et al. (2003) show that the discretization bias arising from the use of daily intervals is negligible. In the setup here, these errors are potentially even smaller whenever intradaily data are available. Finally, it is also possible to base estimation on the Gaussian, discretized transition density whenever observations are not finely spaced by filling in missing data in between large observation intervals (see Eraker (2001)).

B. MCMC Sampling

The purpose of MCMC sampling is to obtain a sample of parameters, Θ , latent volatilities, V_{t_i} , jump times, J_{t_i} , and jump sizes, Z_{t_i} , from the their joint posterior density $p(V, J, Z, \Theta | Y)$. This accomplishes the essential problem of estimating the marginal posterior density of the parameters, $p(\Theta | Y)$. Notice that we can accomplish this task by designing a sampling scheme that produces random draws, $\Theta^{(1)}, \Theta^{(2)}, \dots, \Theta^{(G)}$, whose density is $p(V, J, Z, \Theta | Y)$ since, by a simple application of Bayes theorem, $p(\Theta | Y) \propto p(V, J, Z, \Theta | Y)$. In other words, a random sample from the joint posterior density is also implicitly a sample from the marginal density of each of its components.

Therefore, all methods of Bayesian inference through MCMC have the common goal of designing a scheme to sample from a high-dimensional joint density. Notice that, unfortunately, there are no known methods that allow us to sample from this density directly. There are two problems. First, the density's dimension is a multiple of the sample size, making it far too highly dimensional to sample from directly. Second, the density is nonstandard, making even the drawing of one particular element a difficult task, since direct sampling methods only apply to problems where the density is known. We will deal with these problems in a fashion that has become rather mainstream in the numerical Bayes literature. The dimensionality problem is handled by employing a Gibbs sampler that essentially allows all random variables of the joint posterior to be simulated sequentially, one at a time. Gibbs sampling calls for the derivation of the conditional posterior densities, to be discussed below. Second, the problem of nonstandard conditional densities, will be solved by implementing a series of Metropolis Hastings draws. The specific details of how these sampling algorithms are implemented, are given next.

C. Gibbs Sampling

The Gibbs sampler employed here calls for the following simulation scheme:

For $g = 1, \dots, G, i = 1, \dots, N$, simulate

$$V_{t_i}^{(g)} \text{ from } p(V_{t_i} | V_{t_i}^{(g-1)}, J_{t_i}^{(g-1)}, Z_{t_i}^{(g-1)}, \Theta^{(g-1)}, Y)$$

$$J_{t_i}^{(g)} \text{ from } p(J_{t_i} | V_{t_i}^{(g)}, J_{t_i}^{(g-1)}, Z_{t_i}^{(g-1)}, \Theta^{(g-1)}, Y)$$

$$Z_{t_i}^{(g)} \text{ from } p(Z_{t_i} | V_{t_i}^{(g)}, J_{t_i}^{(g)}, Z_{t_i}^{(g-1)}, \Theta^{(g-1)}, Y)$$

and

$$\Theta^{(g)} \text{ from } p(\Theta | V^{(g)}, J^{(g)}, Z^{(g)}, Y^{(g)}, \rho^{(g-1)}, s^{(g-1)}).$$

This recursive scheme is started from an appropriate set of starting values for $g = 0$.

Seemingly, the Gibbs sampling scheme above calls for the derivation of the four different conditional posterior densities. While this might seem difficult, notice that these conditionals are all proportional to the joint posterior $p(V, J, Z, \Theta | Y)$ which follows from an easy application of Bayes theorem. Since the

normalizing constant is unimportant, it suffices to know this joint density to run this sampler. In practice, it is useful to study each and one of the four densities in question to see if simplifications can speed up the algorithm.

The densities in the above Gibbs sampling scheme do indeed simplify. Hence, it is not necessary to use the entire expression in (12). We now go through these simplifications in turn. The conditional posterior for spot volatility at time t_i can be written,

$$p(V_{t_i} | V_{t_i}^{(g-1)}, J^{(g-1)}, Z^{(g-1)}, \Theta^{(g-1)}, Y) \\ \propto p(Y_{t_i} | X_{t_i}, \varepsilon_{t_i}, \Theta) p(X_{t_i} | X_{t_{i-1}}, Z_{t_i}, J_{t_i}, \Theta) p(X_{t_{i+1}} | X_{t_i}, Z_{t_{i+1}}, J_{t_{i+1}}, \Theta). \quad (14)$$

The simplification here stems from the Markov property of the process. Only the terms in (12) where V_{t_i} enters directly are relevant and the remaining terms in the product sum are absorbed into the normalizing constant. Similar simplifications are commonly encountered in state-space models where the state variable (here: volatility) is Markov. Notice that this is a nonstandard density. As such, simulation from this density requires a metropolis step.

Next, we turn to the densities for the jump times and sizes. The jump indicator, J_{t_i} , is a binary random variable (taking on 0 or 1). Thus, it is sufficient to find the probability $\Pr(J_{t_i} = 1 | X, Z, Y, \Theta)$, which again simplifies to

$$\Pr(J_{t_i} = 1 | X, Z, Y, \Theta) \\ = \frac{p(X_{t_i} | X_{t_{i-1}}, Z_{t_i}, J_{t_i} = 1, \theta) p(Z_{t_i} | J_{t_i} = 1) p(J_{t_i} = 1)}{\sum_{s=0,1} p(X_{t_i} | X_{t_{i-1}}, Z_{t_i}, J_{t_i} = s, \theta) p(Z_{t_i} | J_{t_i} = s) p(J_{t_i} = s)}.$$

Notice that it does not depend on the option prices directly. This follows from the fact that the option prices do themselves not depend on the jump indicator and, hence, vice versa. The conditional draw for the jump indicator, J_{t_i} , is identical to the one in Eraker et al. (2003).

The conditional distributions for the jump sizes, $Z_{t_i} = (Z_{t_i}^S, Z_{t_i}^V)$, also do not depend on the option prices. Again this follows from the fact that option prices depend only on the state, $X = (S, V)$. As a result, the conditional distribution for Z_{t_i} is the same truncated normal distribution as given in Eraker et al. (2003), and the reader is referred to this paper for details.

Finally, we note that the conditional distribution for the parameter vector Θ is also nonstandard. This makes it necessary to do a metropolis step. We use a normal proposal density centered at the current draw and with covariance matrix estimated by a preliminary run of the algorithm. Notice that this accomplishes drawing of simultaneous, highly correlated variables which reduces serial correlation in output of the sampler relative to the optional method of drawing an element at a time. The latter is also more computationally expensive as it requires a recomputation of all option prices (through recomputing the likelihood) for each of the parameters.

The distribution of the pricing error parameters, ρ_j and s_j are Gaussian/gamma respectively, whenever the data are collected at integer, equally sized

intervals, but not whenever sampling times are random. Still, whenever sampling times are random, we can factor the conditional density into a Gaussian, and a non-Gaussian part which enables an easy metropolis sampler to be set up by proposing from the Gaussian part and accepting with the non-Gaussian part.

III. Data

The empirical analysis in this paper is based on a sample of S&P 500 options contracts. The sample consists of daily CBOE closing prices, and has previously been used by Ait-Sahalia, Wang, and Yared (2001) and David and Veronesi (1999), among others. Ait-Sahalia et al. (2001) point out that the lack of timeliness between the closing prices of potentially illiquid options and the underlying index introduces noise in the observations. In response to this, they suggest backing out the value of the underlying asset through a put-call parity. Essentially, this approach is equivalent to computing the put/call prices without relying on their theoretical relation to the underlying asset, but rather their theoretical relation to the corresponding put or call. We apply the same adjustment here, for the same reasons. We also delete from the database observations with prices less than one dollar because the discreteness of quotes has a large impact on these observations.

The full sample of options available in the estimation period contains some 22,500 option prices. Using the entire data set call for excessive computing times, likely in the order of months. In order to facilitate estimation in a timely fashion, we need to select a subsample from the larger database. The subsampling scheme is constructed by randomly choosing a contract trading at the first day of the sample. All recorded observations of this particular contract are then included. As of the first trading day for which this contract where not trading, we randomly pick another contract trading on that day. The procedure is repeated so as to construct a subsample for which there is at least one contract available each trading day. In total, this subsampling procedure produces 3,270 call options contracts recorded over 1,006 trading days, covering the period of January 1, 1987 to December 31, 1990. The remaining sampling period, January 1, 1991 to March 1, 1996, is kept for out-of-sample testing purposes. This leaves some 36,890 transactions to be used in the out-of-sample analysis.

The subsampling procedure described above mimics the characteristics of the larger database as closely as possible. Notice also that our sample includes a random number of contracts at any one particular point in time, making it a panel data set with a randomly sized cross section. Sample statistics describing the data set are reported in Tables I and II.

IV. Empirical Results

A. Parameter Estimates

Table III reports the posterior means, posterior standard deviations, and the 1 and 99 posterior percentiles of the parameters in the various models. The

Table I
Descriptive Statistics for Options Data

The table reports mean and standard deviation for option prices, implied volatilities (IV), time to maturity (τ), and moneyness (strike/spot).

	Strike/Spot	τ	Price	IV
1987–1991				
Mean	0.98	50.33	10.4833	20.2%
SD	0.05	37.60	7.425	7.8%
1991–1996				
Mean	0.99	45.91	17.88	14%
SD	0.06	44.27	23.13	6%

Table II
Descriptive Statistics for Returns Data

The table reports means, standard deviations, skewness, and kurtosis for S&P 500 returns data collected from January 1, 1987 to December 31, 1990.

	All Observations				Deleting 1987 Crash			
	Mean	SD	Skew.	Kurt.	Mean	SD	Skew.	Kurt.
1987–1991								
Daily	0.03	1.35	−5.06	87.50	0.05	1.15	−0.50	12.67
Weekly	0.14	3.28	−4.14	41.32	0.26	2.45	−0.70	3.99
Monthly	0.58	6.26	−1.42	7.92	1.08	4.52	0.17	4.23
1970–1991								
Daily	0.03	0.98	−2.32	60.26	0.03	0.93	−0.05	8.44
Weekly	0.13	2.29	−0.48	7.46	0.15	2.22	−0.05	4.14
Monthly	0.53	4.52	−0.67	4.93	0.62	4.36	−0.35	3.92

parameters are quoted using a daily time interval following the convention in the time-series literature. Notice that the parameters need to be annualized to be comparable to typical results in the option pricing literature (e.g., Bates (2000), Pan (2002)).

There are several interesting features of the parameter estimates in Table III. We start with an examination of the estimates in the SV model. The long-term mean of the volatility process, θ is 1.93, which is relatively high. It corresponds to an annualized long-run volatility of 22%. The estimate is slightly higher than the unconditional sample variance of 1.832 (see Table II). Typically, estimates reported elsewhere of the unconditional variance of S&P 500 returns are somewhat below 1%, corresponding to an annualized value somewhat less than 15%. Our estimates of 1.933/1.832 are indicative of the relatively volatile sample period used. For the SVJ and SVCJ models, the estimates of θ are much smaller, suggesting that the jump components are explaining a significant portion of the unconditional return variance.

Table III
Parameter Estimates

The table reports posterior means and standard deviations (in parenthesis), and 99 credibility intervals (in square brackets) for parameters in the jump diffusion models based on joint options and returns data. The parameters, η_V and η_J , denote the market premiums for volatility and jump risk, respectively. Parameter estimates correspond to a unit of time defined to be one day, and returns data scaled by 100.

	SV	SVJ	SVCJ	SVSCJ
θ	1.933 (0.048) [1.843, 2.064]	1.652 (0.053) [1.523, 1.777]	1.353 (0.067) [1.243, 1.560]	0.943 (0.065) [0.767, 1.080]
κ	0.019 (0.007) [0.004, 0.035]	0.019 (0.006) [0.005, 0.034]	0.023 (0.007) [0.010, 0.038]	0.023 (0.007) [0.010, 0.040]
κ^Q	0.009 (0.000) [0.009, 0.010]	0.011 (0.000) [0.010, 0.011]	0.011 (0.000) [0.010, 0.012]	0.006 (0.000) [0.005, 0.008]
η_V	0.010 (0.007) [-0.005, 0.026]	0.009 (0.006) [-0.006, 0.024]	0.013 (0.007) [-0.001, 0.028]	0.017 (0.007) [0.003, 0.033]
ρ	-0.569 (0.014) [-0.601, -0.535]	-0.586 (0.027) [-0.652, -0.526]	-0.582 (0.024) [-0.646, -0.528]	-0.542 (0.034) [-0.620, -0.462]
σ_V	0.220 (0.007) [0.203, 0.240]	0.203 (0.007) [0.187, 0.218]	0.163 (0.007) [0.148, 0.181]	0.137 (0.007) [0.122, 0.154]
μ_y		-0.388 (3.456) [-8.560, 7.631]	-6.062 (2.274) [-11.566, -0.881]	-1.535 (0.184) [-1.997, -1.120]
μ_y^Q		-2.002 (1.867) [-6.079, 2.022]	-7.508 (0.932) [-9.725, -5.527]	-7.902 (0.843) [-9.703, -5.984]
η_J		1.613 (3.789) [-7.320, 10.478]	1.446 (2.474) [-4.516, 6.939]	6.367 (0.842) [4.472, 8.117]
ρ_J			-0.693 (0.096) [-0.856, -0.449]	-2.214 (0.099) [-2.435, -2.013]
σ_y		6.634 (1.081) [4.972, 9.508]	3.630 (1.106) [1.403, 6.261]	2.072 (0.302) [1.439, 2.717]
μ_V			1.638 (0.790) [0.833, 3.231]	1.503 (0.279) [1.078, 2.065]
λ_0		0.002 (0.001) [0.001, 0.003]	0.002 (0.001) [0.001, 0.004]	0.002 (0.001) [0.001, 0.004]
λ_1				1.298 (0.141) [1.086, 1.606]

The κ parameter is low under both the objective and risk neutral measures. As a rule of thumb, estimates of κ can be interpreted as approximately one minus autocorrelation of volatility. Hence, estimates ranging from 0.019 to 0.025 imply volatility autocorrelations in the range 0.975–0.981 which is in line with the voluminous time-series literature on volatility models.

The speed of mean reversion parameter under the risk neutral measure, κ^Q , is a parameter of particular interest. The difference between these parameters across the two measures, $\eta_V = \kappa - \kappa^Q$, is the risk premium associated with volatility risk. These volatility risk premiums are estimated to be positive across all models. For the SVSCJ model, the estimate of this parameter is “significant”, in the sense that its 1 percentile is greater than zero. The volatility risk premium is similarly found to be “significant” at the 5% level for the SVCJ model, and insignificant for the SVJ and SV models. A positive value for this parameter implies that investors are averse to changes in volatility. For option prices, this implies that whenever volatility is high, options are more expensive than what is implied by the objective measure as the investor commands a higher premium. Conversely, in low-volatility periods, the options are less expensive, *ceteris paribus*. Figure 1 quantifies this effect in terms of annualized Black–Scholes volatility for an ATM option across maturities for different values of the initial volatility, V_t . As can be seen by the figure, whenever volatility is high, the risk premium is positive at all maturities and peaks at about 2 percentage points (annualized) at 130 days to maturity. Conversely, the low initial




Figure 1. Volatility premium. The figure shows the increase in Black–Scholes volatility due to volatility risk premium, η_V , for the SVCJ model.

spot volatility gives a decreasing and then increasing premium. The limiting premium (as maturity increases) is zero since the limiting (unconditional) stock price distribution does not depend on the speed of mean reversion under either measure.

Next, we examine the volatility of volatility, σ_V , and correlation, ρ , between Brownian increments. Our estimates of σ_V are a little more than two times those obtained for the same models in Eraker et al. (2003) from time-series analysis on a longer sample of S&P 500 returns (see below). The correlation coefficient is also larger in magnitude. There is a certain disagreement in the literature as to the magnitude of both these parameters, as well as whether estimates obtained previously, are reasonable. We examine these issues in more detail below.

We now turn to discuss the jump size parameters and the related jump frequencies. First, we note that the jumps occur extremely rarely: The λ estimates in Table III indicate that one can expect about two to three jumps in a stretch of 1,000 trading days. The unconditional jump frequency is only marginally higher for the state dependent SVSCJ model. Whenever spot volatility is high, say a daily standard deviation of 3%, the estimate of λ_1 is indicative of an instantaneous jump probability of about 0.003—a 50% increase over the constant arrival intensity specifications. According to SVJ estimates, observed jumps will be in the range $[-13.4, 12.6]$ with 95% probability when they occur.

A notable implication of the estimates of the mean jump sizes, μ_y and μ_y^Q in Table III, is that under both measures these parameters are difficult to identify. In particular, this parameter is difficult to estimate accurately under the objective risk probability measure. The reason is simple: Option prices do not depend on μ_y , so this parameter is only identified through the returns data. To illustrate how the difficulties arise, assume a hypothetical estimator based on knowledge of when the jumps occurred and by how much the process jumped. Then μ_y would be identified as the mean of these observations' jump sizes. But, since there are very few observations in the sample for which jumps are estimated to have occurred, this hypothetical estimator becomes very noisy. In reality, it is even more difficult to identify μ_y because jump times and sizes are not known. Not only does this introduce more noise than the hypothetical estimator, but it gives an improper posterior distribution if the investigator does not impose an informative (proper) prior.⁵

There are equivalent problems in estimating the other parameters in the jump distribution for low jump-frequencies. This manifests itself by the fairly high posterior standard deviations in jump-size parameters, particularly for the SVJ model. This contrasts with the parameters which govern the dynamics in the volatility process, which, without exception, are extremely sharply estimated. The posterior standard deviations are somewhat smaller in the

⁵ The same problem should be expected for maximum likelihood and other less informationally efficient non-Bayesian methods and such methods should give standard errors associated with μ_y that are proportional to $1/\sqrt{\lambda}$.

SVCJ model. The reason is that this model incorporates simultaneous jumps in volatility and returns, and since the latent volatility is very accurately estimated (see below), so are the jump times. Since jump times are accurately estimated, it is easier for the method to find the jump sizes, and hence jump size parameters under the objective measure. As a consequence, for the SVCJ model, we estimate μ_y more accurately than for the SVJ model.

The difference between the mean jump sizes under the respective probability measures, η_J , is the premium associated with jump risk, and is therefore of particular interest. The premium is estimated to be positive for all models, but with a very sizable credibility interval for both the SVJ and the SVCJ models. Hence, we cannot conclude that there is a significant market premium to jump risk for these models. For the SVSCJ model, the posterior 1 percentile for η_J is 4.47, indicating that for this model the jump risk premium is significant. The estimates reported here are much smaller in magnitude than corresponding estimates in Pan (2002).

The parameter estimates in Table III are interesting in light of estimates obtained elsewhere. The results in BCC are a particularly interesting reference because they are obtained purely by fitting the option prices, and are not conditioned on time-series properties of returns and volatility. The BCC estimates are indeed very similar to the ones obtained here. Notably, BCC estimate the jump frequency for the SVJ model is $0.59/252 = 0.0023$ daily jump probability, and the jump-size parameters μ_y^Q and σ_y are -5 and -7% , respectively. Hence, using only options data, BCC obtain quite similar estimates from those reported here for the jump component of the model. Their volatility process parameters are also close to the ones reported here: Their estimate of κ and σ_V are 0.008 and 0.15, respectively, when converted into daily frequencies.⁶

Pan (2002) and Bates' (2000) model specifications differ from the simple price-jump model (SVJ) considered here, in that the jump frequency, λ , depends on the spot volatility, and hence the jump size and jump frequency parameters are not really comparable to those reported in BCC and here. Interpreting the differences with this in mind, the average jump intensity point estimates in Pan are in the range $[0.0007, 0.003]$ across different model specifications. Hence, her estimates are in the same ballpark as those reported in BCC and here. Interestingly, Bates obtains quite different results, with an average jump intensity of 0.005.⁷ Bates also reports a jump size mean ranging from -5.4 to -9.5% and standard deviations of about 10–11%. Hence, Bates' estimates imply more frequent and more severe crashes than the parameter estimates reported in BCC and in this paper. The practical implication of the difference

⁶ The numbers refer to BCC's Table III, "all options" column. To convert their parameters into daily counterparts, divide λ and κ by 252. The jump size parameters are unaffected by the difference in frequency.

⁷ The papers by Pan and Bates model the jump frequencies as $\lambda = \lambda_1 V_t$ where V_t is the spot volatility. The average jump intensities reported here are obtained from Pan and Bates' papers as the multiple of the long-run volatility mean, θ , and the proportionality parameter, λ_1 . The calculations are at this author's own account.

is that his estimates will generate more skewness and kurtosis in the conditional returns distributions, and consequently also steeper Black–Scholes implied volatility smiles across all maturity option contracts. In particular, since the jump intensity in Bates depends on volatility, his model will generate particularly steep implied volatility curves whenever the spot volatility, V_t , is high.

Finally, to put the estimates in perspective, Table IV shows parameter estimates obtained by using returns data only. Estimates are shown for the SV, SVJ, and SVCJ models only. The estimates were obtained using returns collected

Table IV
Parameter Estimates from Returns Data

The table reports posterior means and standard deviations (in parenthesis) and 99% credibility intervals (in square brackets) for parameters in the jump diffusion models based on returns data only. Parameter estimates were obtained using the estimation procedure in Eraker et al. (2003) using 5,307 time-series observations of the S&P 500 index from January 1970 to December 1990.

	SV	SVYJ	SVCJ
α	0.026 (0.011) [0.000, 0.051]	0.026 (0.011) [0.000, 0.050]	0.030 (0.011) [0.004, 0.056]
θ	0.881 (0.098) [0.692, 1.163]	0.834 (0.122) [0.590, 1.221]	0.573 (0.078) [0.397, 0.750]
κ	0.017 (0.005) [0.008, 0.030]	0.012 (0.006) [0.004, 0.022]	0.016 (0.003) [0.009, 0.023]
ρ	-0.373 (0.056) [-0.500, -0.242]	-0.468 (0.065) [-0.601, -0.295]	-0.461 (0.073) [-0.616, -0.237]
σ_V	0.108 (0.011) [0.082, 0.137]	0.079 (0.011) [0.061, 0.104]	0.058 (0.012) [0.030, 0.078]
μ_y		-3.661 (2.486) [-10.752, 1.281]	-3.225 (2.523) [-10.086, 2.436]
ρ_J			0.312 (1.459) [-3.580, 3.833]
σ_y		6.628 (1.697) [3.714, 11.742]	4.918 (1.272) [2.880, 9.295]
μ_V			1.250 (0.381) [0.681, 2.523]
λ		0.003 (0.001) [0.001, 0.006]	0.004 (0.001) [0.001, 0.007]

over the period January 1970 to December 1990, so comparisons with the estimates in Table III should be done with this in mind. Not too surprisingly, the estimates in Table IV are close to those reported in EJP although some sensitivity to the sampling period in use is inevitable: the two papers use about 20 years of data, with the 1980s being common. So while the sample used in EJP includes the -7% returns on October 27, 1997 and August 31, 1998, comparably large moves did not occur in the 1970s. This leads to a different estimate of jump frequency in this paper ($\lambda = 0.004$ versus 0.0055 , respectively). The other parameters are in the same ballpark as those reported in EJP. Notice again how the volatility-of-volatility parameter, σ_V , is noticeably smaller than in Table III.

B. Option Price Fit

In this section, we discuss the empirical performance of the various models in fitting the historical option prices. Table V reports the posterior means of absolute option pricing errors for the different models, conditional upon moneyness and maturity.

The results in Table V may seem surprising at first. The pricing errors for all four models are about 47 cents on average. Information about bid/ask spreads is missing from this database. However, similar data were used by BCC who reported that the spread ranged from 6 to 50 cents. Consequently, the pricing errors reported here are likely to be somewhat larger than the average bid/ask spread.

At first glance, the fact that the simple SV model fits the options data as well as the SVSCJ model seems implausible. After all, the SVSCJ model incorporates six more parameters, so how is it possible that it does not improve on the simple SV model? The answer is a combination of explanations: First, our MCMC approach does not minimize pricing errors. Simply, there are no objective functions that are optimized, but rather the results in Table V are mean errors over a large range of plausible (in the sense that they have positive posterior probability mass), parameter values. This way of conducting Bayesian model comparisons differs fundamentally from classical methods. An analogy is the computation of Bayes factors which are the ratio of likelihood functions that are averaged over the posterior distributions. This will sometimes give (marginalized) likelihoods which favor the most parsimonious model. Moreover, the jump models do not really increase the degrees of freedom by a notable amount. The number of parameters increases by six in the SVSCJ model, but there are 3,270 sample option prices, so the increase in degrees of freedom is marginal and practically negligible. This is an important difference between this study and the one by Bakshi et al. (1997). In their paper, since the model parameters are recalibrated every day, adding one parameter increases the degrees of freedom by the number of time-series observations. Notice that the improvement in fit for the SVJ model over the SV model reported in BCC, is in the order of a few cents only. By contrast, BCC report that the SV model improves more than 90 cents on the Black–Scholes model. BCC concludes that

Table V
Absolute Pricing Errors

The table reports mean absolute pricing errors for different option models conditional on time to maturity and moneyness. The errors are not corrected for serial correlation. All pricing errors in dollars.

Maturity		Moneyness (Strike/Spot)						All
		<0.93	0.93–0.97	0.97–1.0	1.0–1.03	1.03–1.07	>1.07	
<1 m	#	13	86	289	272	71	9	740
	SV	0.67	0.35	0.35	0.48	0.53	0.47	0.42
	SVJ	0.61	0.35	0.36	0.46	0.51	0.55	0.42
	SVCJ	0.69	0.35	0.32	0.43	0.48	0.51	0.39
	SVSCJ	0.76	0.35	0.29	0.40	0.43	0.38	0.36
1–2 m	#	13	105	248	312	257	76	1011
	SV	0.33	0.42	0.32	0.43	0.54	1.14	0.48
	SVJ	0.32	0.41	0.34	0.44	0.54	1.21	0.49
	SVCJ	0.35	0.45	0.33	0.44	0.52	1.19	0.49
	SVSCJ	0.30	0.49	0.38	0.46	0.54	1.20	0.51
2–3 m	#	10	74	125	188	127	79	603
	SV	0.37	0.39	0.47	0.43	0.48	0.58	0.46
	SVJ	0.37	0.40	0.50	0.45	0.51	0.63	0.49
	SVCJ	0.36	0.40	0.48	0.39	0.48	0.55	0.45
	SVSCJ	0.36	0.36	0.41	0.39	0.51	0.49	0.43
3–6 m	#	22	68	110	194	140	176	710
	SV	1.09	0.38	0.38	0.57	0.50	0.39	0.48
	SVJ	0.98	0.37	0.39	0.61	0.51	0.39	0.49
	SVCJ	0.95	0.35	0.36	0.57	0.50	0.42	0.48
	SVSCJ	1.14	0.46	0.38	0.49	0.50	0.45	0.48
> 6m	#	5	14	23	42	40	82	206
	SV	0.74	0.64	0.39	0.63	0.74	0.45	0.56
	SVJ	0.69	0.68	0.43	0.67	0.78	0.46	0.58
	SVCJ	0.73	0.70	0.42	0.65	0.72	0.41	0.55
	SVSCJ	0.84	0.71	0.42	0.60	0.64	0.39	0.52
All	#	63	347	795	1008	635	422	3270
	SV	0.71	0.40	0.37	0.48	0.53	0.57	0.47
	SVJ	0.65	0.40	0.38	0.49	0.54	0.60	0.47
	SVCJ	0.66	0.41	0.36	0.46	0.52	0.58	0.46
	SVSCJ	0.74	0.43	0.35	0.44	0.52	0.58	0.46

“once stochastic volatility is modeled, adding other features will usually lead to second-order pricing improvements.”

The lack of improvements in pricing results in Table V is somewhat at odds with the results in Bates (2000) and Pan (2002) who both conclude that jumps (in prices) are important in capturing systematic variations in Black–Scholes volatilities. Figure 2 presents further evidence on the fit of the various models. The figure plots the model and market implied Black–Scholes volatilities (IV) for high, medium, and low initial spot-volatility, and for different maturity classes. Notice that the IVs are very similar for all three models, and on




Figure 2. Model and market smiles conditional upon initial volatility and maturity. The figure shows implied Black and Scholes volatility smiles for model generated prices and market prices for different degrees of moneyness (strike/spot) on the x -axis. The plots are constructed conditional upon the initial (spot) market volatility, V_t , (top to bottom), and for different maturity contracts (left to right).

average, they tend to fit the observed data reasonably well. There are a few exceptions:

- For the short and medium maturity contracts during days of high or average spot volatility, the models do not generate sufficiently steep IVs.
- For long maturity contracts trading on days with low spot volatility, all models except the SVSCJ model overprice.

The first observation motivated the generalization into state-dependent jump intensity incorporated by the SVSCJ model. Models with such state dependency will potentially generate steeper volatility smiles in high-volatility environments because the chance of observing large, negative jumps increases. However, the estimate of λ_1 is not large enough to yield a significantly steeper volatility smile for this model relative to the simpler ones.

It is important to recognize that the results in Table V could be sensitive to the dollar/cent denomination specification for the pricing errors. For example, if the errors were measured in percent of the option price, the results would place heavier weight on short-term OTM or ITM contracts and thus, put relatively more emphasis on the tails of the return distributions at short horizons. The same effect is likely to occur if the errors are measured in terms of Black–Scholes implied volatilities. This is not only true for the pricing errors in Table V, but also for the parameter estimates in Table III. In particular, it is true that the values of the joint posterior distribution of the SVSCJ model are not too different for different parameter constellations for which the importance of jumps increases relative to that of the volatility component. In particular, parameter constellations involving more frequent jumps tend to improve the option price fit at short horizons, but deteriorate the fit for long-term contracts. Hence, judging by measures that would lend more weight to shorter contracts, such parameter constellations would potentially be found important. It may be that a different benchmark measure of performance could alter our conclusions. To examine the reasonableness of this conjecture, it is important to realize that the distribution of the pricing errors should be the guideline for the specification. To see this, assume that there were no pricing errors. We should then be able to identify all model parameters as well as latent volatility with arbitrarily good precision *regardless of whether the likelihood function is defined over dollar values, implied volatilities, or relative pricing errors*. Since this assumption is obviously violated, the all-important issue becomes the distribution of pricing errors. For example, assume that we modeled the errors in terms of implied Black–Scholes volatilities. We know that implied volatilities become increasingly variable for contracts in or out of the money (see Jackwerth and Rubenstein (1996)). Hence, if we were to specify the likelihood function over implied Black–Scholes volatility, we would have to incorporate this heteroskedastic feature of the data into the specification of the likelihood function. In doing so, we would in fact specify a likelihood function which would lend equal weight to all observations across different moneyness categories. Hence, it is not clear that short ITM/OTM contracts are given greater weight in a correctly

specified likelihood function defined over relative errors or implied Black–Scholes volatilities.

Our choice of error distribution based on dollars and cents can be motivated from two observations. First, errors caused by discreteness of quotes should be uniformly distributed in dollars across moneyness and maturity. Second, and perhaps more important, the mean absolute errors in Table V are reasonably similar across moneyness and maturities. Thus, there do not seem to be systematic patterns of heteroscasticity across different option classes.

A final note of caution on the performance of the models: The lack of improvement for the jump models does not indicate that these models are incorrect. Indeed, jumps may very well be warranted to model the time-series behavior of the returns. I elaborate on this below.

C. Out-of-Sample Performance

The results discussed so far reflect in-sample fit obtained over the period January 1987 to December 1990. This section presents results of out-of-sample fit for the period January 1991 to March 1996. This leaves a total of 35,890 observations to be used in the out-of-sample performance study.

The procedure used to construct the out-of-sample errors is as follows. Given parameter estimates in Table III, a so-called particle filtering approach is used to estimate volatility V_t for each day t using information available at $t - 1$. This approach has similarities with the popular method of inserting yesterday's Black–Scholes implied volatility into the Black–Scholes model to obtain today's prices. Pricing errors resulting from this approach will generally be "close" to the ones obtained by actual minimization (with respect to V_t only). The reader should bear in mind that since the parameters are fixed, there is no way in which we can calibrate the shape of the term structure/volatility smile as market conditions change. This is therefore a very restrictive exercise, and results should be interpreted with this in mind.

Table VI presents the pricing errors broken down in maturity and moneyness categories. The average pricing errors are not dauntingly large and vary between 0.51 cents (SVSCJ) and 0.67 cents (SVCJ) on average. Hence, pricing errors are larger out-of than in-sample. No uniform relationship between the model complexity and the out-of-sample performance is evident. The SVSCJ model produces prices which on average differ from market prices by 4 cents less than the SV model. The SVJ and SVCJ models perform the worst.

Figure 3 plots the pricing errors for different maturity contracts for the different models. As can be seen from the figure, the pricing errors are moderately small in the first year or two following the estimation period. They then become progressively larger. This is particularly obvious for the first three models.

The mechanical explanation for the above findings is as follows. The volatility of the out-of-sample period is significantly below that of the estimation period. During the out-of-sample period, markets experienced unusually low volatility and spot volatility (as well as implied Black–Scholes volatility) can be shown

Table VI
Out-of-Sample Absolute Pricing Errors

The table reports mean absolute pricing errors for different option models conditional on time to maturity and moneyness. Results are based on parameters estimated in Table II for the subsequent period January 1991 to March 1996. For each day, the model re-estimates spot volatility estimates using a particle filtering method. All pricing errors in dollars.

Maturity		Moneyness (Strike/Spot)						All
		<0.93	0.93–0.97	0.97–1.0	1.0–1.03	1.03–1.07	>1.07	
<1 m	#	1213	2623	3306	3159	1398	161	11860
	SV	0.19	0.29	0.40	0.40	0.17	0.07	0.32
	SVJ	0.26	0.26	0.44	0.45	0.16	0.07	0.34
	SVCJ	0.10	0.16	0.43	0.54	0.16	0.08	0.33
	SVSCJ	0.76	0.38	0.46	0.36	0.43	0.50	0.45
1–2 m	#	964	1621	2830	3159	1675	244	10493
	SV	0.37	0.50	0.52	0.52	0.33	0.17	0.46
	SVJ	0.37	0.48	0.51	0.51	0.32	0.18	0.45
	SVCJ	0.21	0.32	0.46	0.58	0.35	0.25	0.43
	SVSCJ	0.35	0.36	0.55	0.41	1.14	0.45	0.55
2–3 m	#	552	668	1446	2027	998	303	5994
	SV	0.36	0.47	0.42	0.41	0.38	0.29	0.40
	SVJ	0.39	0.45	0.39	0.41	0.39	0.30	0.40
	SVCJ	0.31	0.39	0.45	0.48	0.42	0.41	0.43
	SVSCJ	0.29	0.30	1.20	0.39	0.46	0.48	0.58
3–6 m	#	635	482	861	1430	905	474	4787
	SV	0.35	0.55	0.64	0.74	0.82	0.74	0.67
	SVJ	0.37	0.62	0.81	0.97	1.09	0.88	0.84
	SVCJ	0.57	0.89	1.09	1.20	1.25	1.04	1.06
	SVSCJ	0.40	0.49	0.52	0.51	0.38	0.84	0.50
>6 m	#	354	289	390	598	516	609	2756
	SV	0.93	1.59	1.66	2.37	2.53	2.66	2.10
	SVJ	0.91	1.66	1.82	2.70	2.93	3.05	2.36
	SVCJ	1.35	2.20	2.33	3.23	3.40	3.46	2.84
	SVSCJ	0.44	0.38	0.36	0.49	0.49	0.71	0.50
All	#	3718	5683	8833	10373	5492	1791	35890
	SV	0.36	0.46	0.52	0.60	0.59	1.18	0.56
	SVJ	0.39	0.44	0.55	0.66	0.66	1.35	0.61
	SVCJ	0.36	0.40	0.59	0.79	0.75	1.56	0.67
	SVSCJ	0.49	0.37	0.61	0.41	0.65	0.65	0.51

to move as low as 5 to 7%. This change of market conditions is hard to predict from the volatile estimation sample which implicitly incorporates a long-term, unconditional volatility which of course matches that of the sample average over the estimation period.

There is a simple mechanical explanation for the above findings: The first period used for parameter estimation included the crash of 1987, and the subsequent high-volatility. This is reflected, for instance, in the relatively high parameter estimates for θ . For the SV model, the estimate of θ implies an average




Figure 3. Out-of-sample pricing errors. The plot depicts pricing errors in dollars (model less market) for short, medium, and long maturity contracts. Model prices are computed using parameter estimates in Table III, along with rolling estimates of spot volatility, conditional upon information available up to, but not including, the sample date.

volatility of 22%. Comparably, the annualized volatility during the out-of-sample period is 15.3% which is very close to the long-term sample standard deviation of S&P 500 returns. Since the average volatility will determine the relative expensiveness of the longer term contracts, a too high value of θ will lead to a systematic overpricing of long-term contracts. Conversely, short-term contracts depend less heavily on the value of θ .

The premiums on long-dated contracts become increasingly sensitive to the initial volatility as the value of κ^Q decreases. Smaller values of κ^Q imply that the term structure of option premiums will shift parallel in response to volatility changes. In the out-of-sample period, both short- and long-dated contracts become much less expensive than in the estimation period, so the out-of-sample data indicate a parallel shift consistent with a small value of κ^Q . Since κ^Q was estimated to be much lower for the SVSCJ model, this again explains why this model outperforms the other models out-of-sample.

To further shed light on how the high values of θ affect the pricing errors, I conducted another out-of-sample exercise where the values of θ were adjusted to match the historical average volatility of 15%. Not surprisingly, the models do a lot better after this adjustment, and the average pricing error for the SVCJ model is only 36 cents, a mere half of the previous error. The SVCJ model is particularly accurate in predicting the prices of the short-term contracts and the average error in the less than one month category is only 22 cents, almost half that of the SV model. Hence, these numbers support the notion that jump diffusion models give better descriptions of the relative pricing of short-term option contracts.⁸

The out-of-sample results described in this section nevertheless illustrate one important shortcoming of stochastic volatility/jump models: The term structure of implied volatility cannot be matched with constant parameter values throughout the nine year history considered here. In essence, the expensiveness of options in the first period implies a value of θ that is too high to be consistent with the relatively cheaper prices of long-term options found in the latter period.

The above evidence is consistent with various types of model misspecification including structural shifts or other parameterization errors. It seems most likely that a model which allows for more complex volatility dynamics could capture the pricing errors. For instance, in a previous section it was argued that the option pricing errors were consistent with a higher speed of mean reversion in high-volatility states, and vice versa. This could potentially also explain the failure of the parameter estimates to capture the term structure of implied volatility in the out-of-sample period because if mean reversion were systematically lower in low volatility states, the term structure of implied volatility would be flatter in those states. Hence, the impact of the too high values of θ would diminish.

The evidence is also consistent with two-factor volatility models where the second factor determines the long-run volatility. Such models have been

⁸ These additional results are available from the author upon request.

introduced to deal with the overly simple interest rate models which run into similar problems as the volatility models considered here (see Dai and Singleton (2000) and Andersen and Lund (1997) for examples). Two factor volatility specifications have been considered by Bates (2000) and Chernov et al. (1999) in the context of pure time-series analysis. Both papers consider models where the (stock) diffusion term take on the form $\sqrt{V_{t,1}}dW_{t,1} + \sqrt{V_{t,2}}dW_{t,2}$. Given a sufficiently generous correlation structure, factor rotations can make such specifications observationally equivalent to the stochastic mean volatility specification suggested above.

D. Do Option Prices Matter for Estimation?

There may be circumstances in which prior options data are not available for estimation, or one does not want to use options data for estimation for other reasons. In this case, parameter estimates obtained from historical returns data alone can be used to price the option contracts, under simplifying assumptions about risk premia. For instance, investors typically buy OTM put options for (“crash”) insurance purposes. A seller of such an option may be indifferent to the occurrence of large negative moves, and want to collect the positive jump risk premia. Such an issuer might consider the fair value of the option to equal that obtained under a zero risk premium assumption.

Table VII considers the performance of the option pricing models when the parameters governing the price dynamics are set equal to those obtained from pure historical returns data in Table IV. As can be seen from the table, the returns-based parameter estimates produce overall pricing errors that are smaller than those in Table VI, but larger than those obtained when the long-term volatility parameter is adjusted to match the S&P 500 historical average. Hence, by the latter comparison, the use of parameter estimates based on joint options and returns data provide economically significant performance enhancements.

E. Time-Series Fit

In what follows, we examine an essential question: Do option prices imply stock price dynamics consistent with time-series data? The first topic of interest is a re-examination of the evidence in BCC, Bates (2000) and Pan (2002) suggesting that the volatility of volatility, σ_V implied by option prices, cannot be reconciled with time-series estimates. In particular, BCC and Bates show that their estimated volatility paths are too smooth to be consistent with the relatively high σ_V estimated from option prices.

Table VIII reports estimates of σ_V from the simulated values of the historical volatilities, V_t^g , $t = 1, \dots, T$, $g = 1, \dots, G$ analogous to those in BCC and Bates. The numbers match almost exactly those reported in Table III. The point estimates in Table VIII are slightly smaller than the posterior means reported in Table III, however the posterior credibility intervals in Table III do indeed overlap with the point estimates in Table VIII. Hence, the mismatch between

Table VII
Out-of-Sample Absolute Pricing Errors Using Parameters Based on
Time-series Estimates

The table reports mean absolute pricing errors for different option models conditional on time to maturity and moneyness. Results are based on parameters estimated from returns data only (Table IV) and are obtained out-of-sample, January 1991 to March 1996. For each day, the model re-estimates spot volatility estimates using a particle filtering method.

Maturity		Moneyness (Strike/Spot)						All
		<0.93	0.93–0.97	0.97–1.0	1.0–1.03	1.03–1.07	>1.07	
<1 m	#	1213	2623	3306	3159	1398	161	11860
	SV	0.20	0.33	0.36	0.43	0.25	0.08	0.34
	SVJ	0.10	0.17	0.31	0.34	0.17	0.06	0.25
	SVCJ	0.13	0.16	0.37	0.37	0.22	0.19	0.28
1–2 m	#	964	1621	2830	3159	1675	244	10493
	SV	0.50	0.72	0.47	0.46	0.58	0.36	0.52
	SVJ	0.29	0.46	0.43	0.39	0.42	0.29	0.41
	SVCJ	0.31	0.36	0.48	0.44	0.48	0.58	0.44
2–3 m	#	552	668	1446	2027	998	303	5994
	SV	0.75	1.04	0.54	0.41	0.71	0.55	0.60
	SVJ	0.49	0.72	0.45	0.36	0.53	0.44	0.47
	SVCJ	0.46	0.51	0.46	0.42	0.66	0.84	0.51
3–6 m	#	635	482	861	1430	905	474	4787
	SV	1.12	1.31	0.73	0.62	0.92	0.97	0.87
	SVJ	0.74	0.89	0.56	0.57	0.83	0.88	0.70
	SVCJ	0.60	0.71	0.65	0.76	1.15	1.42	0.85
>6 m	#	354	289	390	598	516	609	2756
	SV	1.61	1.55	1.32	0.94	1.23	1.53	1.33
	SVJ	0.92	0.95	0.95	1.00	1.26	1.67	1.17
	SVCJ	0.78	0.99	1.10	1.46	1.79	2.47	1.56
All	#	3718	5683	8833	10373	5492	1791	35890
	SV	0.65	0.67	0.51	0.49	0.63	0.93	0.58
	SVJ	0.39	0.42	0.43	0.43	0.52	0.92	0.46
	SVCJ	0.37	0.35	0.48	0.52	0.68	1.46	0.54

the option implied volatility of volatility and the variability in the estimated volatility series reported elsewhere, cannot be replicated here.

Some reflections on this result are in order. First, the MCMC estimator explicitly imposes the time-series constraint of the volatility dynamics through the likelihood function of the volatility path. This is also true for the full likelihood estimates in Bates (2000), however, Bates uses about 10 times as many option prices as here and about 1/3 as many time-series observations. Hence, in his analysis, the likelihood function is much more heavily influenced by option prices than the time-series dynamics. The converse is true here so the restrictions are much more likely to hold true. Second, Jacquier et al. (1994) show that their MCMC method provides much more erratically behaving volatility paths than other methods based on Kalman filtering, and quasi-maximum likelihood

Table VIII
Estimates of σ_V from Filtered Volatility Series

The table reports posterior means and standard deviations (in parenthesis) of estimates of σ_V from posterior simulations of the latent spot volatility, V_t . Estimates are obtained as the mean square errors of

$$e_t^g = \frac{V_t^g - V_{t-1}^g - \kappa^g (\theta^g - V_{t-1}^g) - z_t^{V,g}}{\sqrt{V_{t-1}^g}}$$

across posterior simulations, $g = 1, \dots, G$. Parameter estimates correspond to a unit of time defined to be one day, and returns data scaled by 100.

	SV	SVJ	SVCJ	SVSCJ
Mean	0.202	0.198	0.151	0.134
SD	0.005	0.004	0.005	0.007

methods. Still, this is likely to play less of a role in explaining the difference between the results reported here and those in Bates (2000) because the volatility paths are estimated much more precisely than what is typically the case from returns data only.

The MCMC estimator employed here enables the investigator to obtain historical estimates of the Brownian increments $\Delta W_t^S = W_t^i - W_{t-1}^i$ for $i = \{S, V\}$, the continuously arriving shocks to prices and volatility. These can be interpreted as the model standardized “residuals.” Figure 4 plots these residual for returns. It is evident from these plots that the jump models fare far better in explaining large stock price movements than do the simple SV model. For example, the -22% crash of October 19, 1987 and the 6% drop on October 13, 1989 produce too large return residuals relative to the prevailing market volatility at the times, to be consistent with the SV model. The SVJ model also has problems explaining the large market movements, and produces surprisingly large residuals for the same dates. In a nutshell, the reason for the large residuals is that the SVJ model partly fails to identify the large negative returns on these dates as jumps. This is again related to the fact that the unconditional jump probability, λ , is estimated to be so low. For the SVCJ and SVSCJ models, this changes because the large simultaneous moves in prices and volatility on these dates makes the algorithm identify these dates as “jump dates” in spite of the low unconditional jump probability. Hence, the return residuals shown for the SVCJ model in Figure 4 are not too different from what can be expected under the assumed $N(0, 1)$ distribution.

Figure 5 plots the estimated (Brownian) shocks to the volatility process. As is well known, the infamous crash of 1987 had a huge impact on option prices across all maturities, especially in the days following the crash. The effect on the estimated, latent spot volatility is a huge jump in volatility. Regardless of whether such jumps are specified as part of the model, or not, the estimated spot volatility increases dramatically on the day of the crash. The SVSCJ and SVCJ models attribute this large increase to jumps in volatility and consequently




Figure 4. Return residuals. The figures show standardized innovations in log returns from estimates based on the joint data set of options and returns.




Figure 5. Volatility residuals. The figures show standardized innovations in volatility from estimates based on the joint data set of options and returns.

Table IX
Residual Skewness and Kurtosis

The table reports skewness and kurtosis for estimated residuals (i.e., daily Brownian increments) for returns and volatility, respectively. The table contains posterior means, standard deviations (in parenthesis), and 99% credibility intervals (in square brackets).

	SV	SVJ	SVCJ	SVSCJ
Return Residuals				
Skewness	-1.891 (0.088) [-2.078, -1.709]	-0.692 (0.294) [-1.633, -0.364]	-0.129 (0.060) [-0.240, -0.016]	0.138 (0.707) [-0.068, 0.532]
Kurtosis	20.352 (1.121) [18.083, 22.615]	7.223 (3.282) [4.525, 18.630]	3.959 (0.267) [3.563, 4.620]	3.757 (1.321) [3.244, 5.798]
Volatility Residuals				
Skewness	2.487 (0.290) [1.940, 3.194]	2.045 (0.249) [1.475, 2.638]	-0.300 (0.134) [-0.564, -0.068]	-0.446 (0.220) [-0.782, -0.058]
Kurtosis	27.758 (4.515) [19.876, 38.507]	19.256 (3.076) [13.687, 28.204]	6.017 (0.693) [4.337, 7.700]	5.616 (1.596) [4.190, 8.036]

produce plausible residual values. The SV and SVJ models produce residuals which are about 10 standard deviations from zero. Hence, these models require implausibly large Brownian volatility increments to deal with the market data from October 1987.

Interestingly, all four models generate too large *negative* movements in the days following the 1987 crash. Indeed, allowing for jumps to volatility in the SVCJ model does not explain the fall in spot volatility following the huge increase on the day of the crash. The reason is that the volatility jumps are restricted to be positive.

Finally, Table IX quantifies the magnitudes of the model violations discussed above through the posterior distributions of sample skewness and sample kurtosis in the residual series. The numbers confirm our conclusions from studying the residual plots: the residuals do not conform to the assumed normal distribution, although the magnitude of the violations is less for the jump models and least for the SVCJ model. Notice that the SVCJ model does markedly better in capturing the tails of both returns and volatility, and produces a residual kurtosis with a lower first percentile of 3.56 and 4.33, respectively.

V. Concluding Remarks

The models considered in this paper do a reasonable job of fitting option prices in-sample. Somewhat surprisingly, models with complicated jump components

do not seem to improve markedly upon the simpler stochastic volatility (Heston) model. This is reflected in the jump arrival intensities that are estimated to be relatively small.

Out-of-sample, pricing errors are found to increase in magnitude as time increases. This is related to a high estimate of the unconditional, long-term volatility reflected in parameter estimates obtained over the highly volatile estimation period. An exception here is the SVSCJ model for which the option implied mean reversion speed of volatility is found to be less than for the other models. This negates the effect of the high long-term volatility, and consequently improves out-of-sample performance. Ad hoc parameter adjustments to the long-term volatility coefficients, or the mean reversion speed, will improve upon the out-of-sample performance. Such adjustments, of course, are not compatible with the underlying structural models, and should consequently be viewed as model diagnostics only. To this end, the reported model violations are suggestive of models, which allow for multiple factors to explain the variations in the moneyness and term structure of option premiums. These findings may not seem too surprising: In the fixed income literature, it is widely upheld that one needs at least three factors to accurately capture the term structure of bond yields over time. If this is viewed as an exercise in “shooting at a moving two-dimensional object,” option pricing is an even more difficult exercise if the object of interest is a three-dimensional one with the relative prices over different strikes (moneyness) being the additional dimension. With this perspective, it is not unreasonable to imagine that multifactor volatility representations may provide better fit to the data.

Another model extension, which may prove empirically warranted, is one in which the volatility is allowed to mean revert at a different speed depending on its level. Such models have shown some success in modeling interest rates (i.e., Ait-Sahalia (1996)). The sharp decline of option implied spot volatility following the extreme peak caused by the 1987 crash would be indicative of such a model. This is also consistent with near unit root behavior of volatility in periods of relatively calm markets implicit in the out-of-sample results presented here.

The models considered in this paper are fairly successful in fitting the time-series dimension of the data. In fact, the results presented here suggest that the price and volatility dynamics implicit in the joint data are fairly consistent with the more general models under consideration. The models that allow for volatility jumps do a very decent job of explaining the time variations in volatility, as well as the dynamics of volatility itself. With respect to the latter, these results are at odds with previous evidence suggesting that options imply a volatility-of-volatility is too high to be consistent with time-series dynamics of estimated volatility series. This finding cannot be attributed to the addition sophistication of the jump in volatility models: The finding holds equally for the Heston SV model. While different data may explain some of these differences, this paper uses posterior simulations for the latent volatility. This is a very important difference from previous work.

REFERENCES

- Aït-Sahalia, Yacine, 1996, Testing continuous-time models of the spot interest rate, *Review of Financial Studies* 9, 385–426.
- Aït-Sahalia, Yacine, 2002, Maximum likelihood estimation of discretely sampled diffusions: A closed-form approach, *Econometrica* 70, 223–262.
- Aït-Sahalia, Yacine, and Andrew Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.
- Aït-Sahalia, Yacine, Yubo Wang, and Francis Yared, 2001, Do options markets correctly assess the probabilities of movement in the underlying asset? *Journal of Econometrics* 102, 67–110.
- Amin, Kaushik I., and Victor Ng, 1993, Option valuation with systematic stochastic volatility, *Journal of Finance* 48, 881–910.
- Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2002, An empirical investigation of continuous-time models for equity returns, *Journal of Finance* 57, 1239–1284.
- Andersen, Torben G., and Jesper Lund, 1997, Estimating continuous-time stochastic volatility models of the short-term interest rate, *Journal of Econometrics* 7, 343–377.
- Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003–2049.
- Bakshi, Gurdip, and Delip Madan, 1999, Crash discovery in stock and options markets, Working paper, University of Maryland.
- Bandi, Federico M., and Peter C. B. Phillips, 2000, Nonstationary continuous time processes, Working paper, University of Chicago.
- Bandi, Federico M., and Peter C. B. Phillips, 2003, Fully nonparametric estimation of scalar diffusion models, *Econometrica* 71, 241–283.
- Bates, David, 1996, Jump and stochastic volatility: Exchange rate processes implicit in Deutsche mark options, *Review of Financial Studies* 9, 69–107.
- Bates, David, 2000, Post-'87 Crash fears in S&P 500 Futures Options, *Journal of Econometrics* 94, 181–238.
- Benzoni, Luca, 2000, Pricing options under stochastic volatility: an empirical investigation, Working paper, University of Minnesota.
- Black, Fisher, 1976, Studies in stock price volatility changes, in *Proceedings of the 1976 Meeting of the Business and Economic Statistics Section, American Statistical Association*, pp. 177–181.
- Carlin, Bradley P., Nicholas G. Polson, and David S. Stoffer, 1992, A Monte Carlo approach to nonnormal and nonlinear state-space modeling, *Journal of the American Statistical Association* 87, 493–500.
- Chernov, Mikhail, A. Ron Gallant, Eric Ghysels, and George Tauchen, 1999, A new class of stochastic volatility models with jumps: Theory and estimation, Working Paper, Duke University.
- Chernov, Mikhail, and Eric Ghysels, 2000, 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.
- Chesney, Marc, and Louis Scott, 1989, Pricing European currency options: A comparison of the modified Black–Scholes model and a random variance model, *Journal of Financial and Quantitative Analysis* 24, 267–284.
- Conley, Tim, Lars P. Hansen, Erzo G. J. Luttmer, and Jose A. Scheinkman, 1997, Estimating subordinated diffusions from discrete time data, *Review of Financial Studies* 10, 525–577.
- Dai, Qiang, and Kenneth S. Singleton, 2000, Specification analysis of affine term structure models, *Journal of Finance* LV, 1943–1978.
- Das, Sanjiv R., and Rangarajan K. Sundaram, 1999, Of smiles and smirks: A term-structure perspective, *Journal of Financial and Quantitative Analysis* 34, 211–239.
- David, Alexander, and Pietro Veronesi, 1999, Option prices with uncertain fundamentals: Theory and evidence on the dynamics of implied volatility and over/underreaction in the options market, Working paper, University of Chicago.
- Duffie, Darrel, and Peter Glynn, 1996, Estimation of continuous-time Markov processes sampled at random time intervals, Manuscript, Stanford University.
- Duffie, Darrel, Jun Pan, and Kenneth J. Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

- Duffie, Darrel, and Kenneth Singleton, 1993, Simulated moments estimation of Markov models of asset prices, *Econometrica* 61, 929–952.
- Dumas, Bernard, Jeff Fleming, and Robert W. Whaley, 1998, Implied volatility functions: Empirical tests, *Journal of Finance* 53, 2059–2106.
- Eraker, Bjørn, 2001, MCMC analysis of diffusion models with application to finance, *Journal of Business and Economic Statistics* 19, 177–191.
- Eraker, Bjørn, Michael S. Johannes, and Nicholas G. Polson, 2003, The impact of jumps in returns and volatility, *Journal of Finance* 53, 1269–1300.
- Gallant, A. Ronald, and George Tauchen, 1996, Which moments to match? *Econometric Theory* 12, 657–681.
- Gourieroux, Christian, Alan Monfort, and Eric Renault, 1993, Indirect inference, *Journal of Applied Econometrics* 8, S85–S118.
- Hansen, Lars P., and Jose Scheinkman, 1995, Back to the future: Generating moment implications for continuous-time Markov processes, *Econometrica* 63, 767–804.
- Heston, Steven, 1993, Closed-form solution of options with stochastic volatility with application to bond and currency options, *Review of Financial Studies* 6, 327–343.
- Hull, John, and Alan White, 1987, The pricing of options with stochastic volatilities, *Journal of Finance* 42, 281–300.
- Jackwerth, Jens C., and Mark Rubenstein, 1996, Recovering probability distributions from option Prices, *Journal of Finance* 51, 1611–1631.
- Jacquier, Eric, and Robert Jarrow, 2000, Bayesian analysis of contingent claim model error, *Journal of Econometrics* 94, 145–180.
- Jacquier, Eric, Nicholas G. Polson, and Peter Rossi, 1994, Bayesian analysis of stochastic volatility models, *Journal of Business and Economic Statistics* 12, 371–389.
- Jacquier, Eric, Nicholas G. Polson, and Peter Rossi, 2004, Bayesian analysis of stochastic volatility models with leverage effect and fat tails, *Journal of Econometrics* (forthcoming).
- Jiang, George J., and John L. Knight, 1997, A nonparametric approach to the estimation of diffusion processes, with an application to a short-term interest rate model, *Econometric Theory* 13, 615–645.
- Johannes, Michael, 2004, The statistical and economic role of jumps in interest rates, *Journal of Finance* 59, 227–260.
- Johannes, Michael S., Rohit Kumar, and Nicholas G. Polson, 1999, State dependent jump models: How do U.S. equity indices jump? Working paper, University of Chicago.
- Melino, Angelo, and Stuart M. Thornbull, 1990, Pricing foreign currency options with systematic stochastic volatility, *Journal of Econometrics* 45, 239–265.
- Merton, Robert, 1976, Option pricing when the underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.
- Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.
- Rubenstein, Mark, 1985, Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August, 1978, *Journal of Finance* 40, 455–480.
- Rubenstein, Mark, 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.
- Scott, Louis O., 1987, Option pricing when the variance changes randomly: Theory, estimation, and an application, *Journal of Financial and Quantitative Analysis* 22, 419–483.
- Singleton, Kenneth J., 2001, Estimation of affine asset pricing models using the empirical characteristic function, *Journal of Econometrics* 102, 111–141.
- Stanton, Richard, 1997, A nonparametric model of term structure dynamics and the market price of interest rate risk, *Journal of Finance* 52, 1973–2002.
- Stein, E. M., and J. C. Stein, 1991, Stock price distributions with stochastic volatility: an analytic approach, *Review of Financial Studies* 4, 727–752.
- Viceira, Luis M., and George Chacko, 1999, Spectral GMM estimation of continuous-time processes, *Journal of Econometrics* 116, 259–292.
- Wiggins, James B., 1987, Option values under stochastic volatility: Theory and empirical estimates, *Journal of Financial Economics* 19, 351–372.