THE JOURNAL OF FINANCE • VOL. LVIII, NO. 3 • JUNE 2003

The Impact of Jumps in Volatility and Returns

BJØRN ERAKER, MICHAEL JOHANNES, and NICHOLAS POLSON*

Abstract

This paper examines continuous-time stochastic volatility models incorporating jumps in returns and volatility. We develop a likelihood-based estimation strategy and provide estimates of parameters, spot volatility, jump times, and jump sizes using S&P 500 and Nasdaq 100 index returns. Estimates of jump times, jump sizes, and volatility are particularly useful for identifying the effects of these factors during periods of market stress, such as those in 1987, 1997, and 1998. Using formal and informal diagnostics, we find strong evidence for jumps in volatility and jumps in returns. Finally, we study how these factors and estimation risk impact option pricing.

A SURPRISING RECENT FINDING INDICATES that models with both diffusive stochastic volatility and jumps in returns are incapable of fully capturing the empirical features of equity index returns or option prices (see, e.g., Bakshi, Cao, and Chen (1997), Bates (2000), and Pan (2002)). Empirical evidence indicates that the conditional volatility of returns rapidly increases, a movement difficult to generate using a diffusion specification for volatility and jumps in returns.[1] In this paper, we propose to remedy this problem by incorporating jumps in stochastic volatility, and we provide empirical evidence supporting the presence and importance of this additional factor.

Intuitively, what is it that jumps in volatility provide that jumps in returns and diffusive stochastic volatility cannot? Jumps in returns can generate large movements such as the crash of 1987, but the impact of a jump is transient: A jump in returns today has no impact on the future distribution of returns. On the other hand, diffusive volatility is highly persistent, but its dynamics are driven by a Brownian motion. For this reason, diffusive stochastic volatility can only in-

*Eraker is at the Economics Department, Duke University; Johannes is at the Graduate School of Business, Columbia University; and Polson is at the Graduate School of Business, University of Chicago. For helpful comments, we thank David Bates, Mike Chernov, Paul Conway, Anne Gron, Lars Hansen, an anonymous referee, and seminar participants at the 2001 Western Finance Assocation conference, Columbia, Duke, NYU, and Wharton. Rob Engle deserves special thanks for his discussion of the paper. All remaining errors are our own.

[1] Bates (2000), Duffie, Pan, and Singleton (2000) and Pan (2002) provide evidence for the presence of positive jumps in volatility. For example, Bates (p. 219), referring to large increases observed in volatility, argues that "the high improbability of those outliers given the diffusion assumptions indicates that the true conditional transition distribution is far more leptokurtotic than hypothesized, and suggests that the factors underlying option prices follow jump processes."

crease gradually via a sequence of small normally distributed increments. Jumps in volatility fill the gap between jumps in returns and diffusive volatility by providing a rapidly moving but persistent factor that drives the conditional volatility of returns.

We focus on the role of jumps in volatility and returns in S&P 500 and Nasdaq 100 index returns, two prominent indices with actively traded futures and European option contracts. Our empirical approach departs from the usual routine of estimating parameters and performing specification tests in that we also estimate the unobserved jump times, jump sizes, and spot volatilities. These estimates provide a dynamic picture of the roles these factors play and are useful for analyzing periods of market stress. It is especially important to determine the contribution of jumps to periods of market stress because jump risk, either in returns or in volatility, cannot typically be hedged away, and investors may demand a large premia to carry these risks.[2]

We consider two models with jumps in volatility and returns, one with contemporaneous arrivals and correlated jump sizes and another with independent arrivals and sizes, both introduced by Duffie, Pan, and Singleton (2000). In these models, we find strong evidence for the presence of both jumps in volatility and returns. First, adding jumps in returns to the square-root stochastic volatility model dramatically changes the behavior of stochastic volatility. At certain points in time, the addition of jumps reduces annualized spot volatility by as much as 20 percent, from 50 percent to 30 percent. While jumps in returns are infrequent events (one to two per year), they are typically large and explain 8 to 15 percent of the total variance of returns.

Jumps in volatility are also important as they allow volatility to rapidly increase. For example, in the market stress of Fall 1987, volatility jumped up from roughly 20 percent to over 50 percent. Once at this high level, volatility mean reverts back to its long-run level, which shows the persistent effect of jumps in volatility on the distribution of returns. We find little, if any, misspecification in the models with jumps in volatility. This provides formal evidence supporting Bates (2000), Duffie, Pan, and Singleton (2000), and Pan (2002), who suggest that jumps in volatility may remove the misspecification documented in models with diffusive stochastic volatility.

It is important to note that the presence of jumps in volatility does not eliminate the need for jumps in returns. With both types of jumps, jumps in returns occur less often, but they still play an important role, as they generate the large, though infrequently observed, crashlike movements. For example, in both of the models with jumps in volatility and returns, jumps in returns generate more than half of the crash in 1987 while high volatility explains the rest. An analysis of the three periods of market stress in our sample, 1987, 1997, and 1998, indicates that jumps in volatility and returns play a greater role than diffusive stochastic volatility in generating these episodes. This suggests that jump components should command relatively larger risk premia than diffusive ones, as their contribution to periods of market stress is greater.

[2] Pan (2002) finds evidence for large jump risk premia.

Specification diagnostics provide insight into exactly why models without jumps in volatility are misspecified. Heston's (1993) square-root stochastic volatility model requires implausibly large shocks to generate the largest observed movements in returns. For example, the square-root model requires almost an eight standard deviation shock in returns to generate the crash. This is not due to the square-root specification: Jacquier, Polson, and Rossi (2001) find the same misspecification using a more flexible log-variance model.

Diagnostics indicate that a model with diffusive stochastic volatility and only jumps in returns is also misspecified. Estimates indicate that jump times are clustered, evidence in contrast to the constant arrival intensity assumption. For example, in the model with jumps in returns and diffusive stochastic volatility, we estimate three jump arrivals during the week of the crash in 1987. Although the evidence is slightly different, these results reinforce the conclusions of Bakshi et al. (1997), Bates (2000), and Pan (2002), who, using additional information in option prices, find strong evidence for misspecification in models with diffusive volatility and jumps in returns. The fact that we arrive at the same conclusion using only returns data is not a coincidence. Due to the absolute continuity of changes in probability measures, returns and options data should contain the same information about the sources of risks, although their impact may be altered due to risk premia.

All of the models we consider generate near closed-form option prices, and we next examine how the different factors impact option prices.[3] Using implied volatility as a metric, we find that the two types of jumps induce important differences in both the term structure and cross section of implied volatilities. Compared to the stochastic volatility model, adding jumps in returns steepens the slope of the implied volatility curves. Jumps in volatility further steepen implied volatility curves and increase implied volatility for in-the-money options. This latter effect found by Duffie, Pan, and Singleton (2000) and Pan (2002) was labeled as the "hook" or "tipping at the end" effect. Although the motivation for jumps in volatility was to improve on the dynamics of volatility, the results indicate that jumps in volatility also have an important cross-sectional impact on option prices.

Finally, we evaluate the effect of estimation risk on option prices. We decompose estimation risk into two components: parameter and spot volatility estimation risk. These sources of estimation risk have very different effects. For short maturity options, volatility uncertainty is primarily an at-the-money effect and has little impact on out-of-the-money options. Parameter uncertainty, on the other hand, has little at-the-money effect for short maturity options, but generates nearly all of the out-of-the-money uncertainty. For long maturity options, parameter uncertainty dominates as the uncertainty regarding the average level of volatility, speed of mean reversion in volatility, and volatility of volatility play a more important role. Spot volatility uncertainty has little impact on long-dated options.

[3] Liu, Longstaff, and Pan (2003) consider related models with jumps in returns and/or jumps in volatility and find that these factors have important implications for optimal portfolio allocation.

The rest of the paper proceeds as follows. Section I introduces the models, and Section II describes our estimation approach. Section III summarizes the empirical results for the S&P 500 and Nasdaq 100 indices. Section IV analyzes the option pricing implications, and Section V concludes.

I. Jump-Diffusion Models of Returns and Volatility

A number of recent papers examine equity price models with jumps in returns and stochastic volatility (see Bakshi, Cao, and Chen (1997), Bates (2000), Andersen, Benzoni, and Lund (2002), Chernov et al. (2002), and Pan (2002)). While it is clear that both stochastic volatility and jumps in returns are important components, Bakshi , Cao, and Chen (1997), Bates (2000), and Pan (2002) find strong evidence for misspecification in the volatility process.

Specifically, Bakshi, Cao, and Chen (1997), using a test developed in Bates (1996), find that the implied structural volatility parameters are inconsistent with time-series estimates using implied volatilities. Additionally, Bates (2000) and Pan (2002) find that the higher moments of volatility changes are inconsistent with the diffusion specification. Together, these results point to the presence of an additional, rapidly moving factor driving conditional volatility, which, unlike jumps in returns, has a persistent component.[4] Jumps in volatility provide such a factor.

We assume that the logarithm of asset's price, $Y_t = \log(S_t)$, solves

$$\begin{pmatrix} dY_t \\ dV_t \end{pmatrix} = \begin{pmatrix} \mu \\ \kappa(\theta - V_{t-}) \end{pmatrix} dt + \sqrt{V_{t-}} \begin{pmatrix} 1 & 0 \\ \rho\sigma_v & \sqrt{(1-\rho^2)} \cdot \sigma_v \end{pmatrix} dW_t + \begin{pmatrix} \xi^y dN_t^y \\ \xi^v dN_t^v \end{pmatrix} \quad (1)$$

where $V_{t-} = \lim_{s \uparrow t} V_s$, W_t is a standard Brownian motion in $\mathbb{R}^2$, N_t^y and N_t^v are Poisson processes with constant intensities λ_y and λ_v, and ξ^y and ξ^v are the jump sizes in returns and volatility, respectively.[5] We assume that the parameters and initial conditions have sufficient regularity for the solution of (1) to be well defined.

This specification nests many of the popular models used for option pricing and portfolio allocation applications. Without jumps, $\lambda_y = \lambda_v = 0$, (1) reduces to Heston's (1993) square-root stochastic volatility model, the SV model. Bates' (1996) SVJ model has normally distributed jumps in returns, $\xi^y \sim N(\mu_y, \sigma_y^2)$, but no jumps in volatility, $\lambda_v = 0$. Duffie, Pan, and Singleton (2000) introduced the models with jumps in volatility. The SVIJ model has independently arriving jumps in volatility, $\xi^v \sim \exp(\mu_v)$, and jumps in returns, $\xi^y \sim N(\mu_y, \sigma_y^2)$. The SVCJ model has contemporaneous arrivals, $N_t^y = N_t^v = N_t$, and correlated jump sizes, $\xi^v \sim \exp(\mu_v)$ and $\xi^y | \xi^v \sim N(\mu_y + \rho_J \xi^v, \sigma_y^2)$.

[4] Andersen et al. (2001) and Alizadeh, Brandt and Diebold (2002) argue, for U.S. equities and a number of foreign exchange rates, that there are two factors generating volatility, one highly persistent and slowly moving, the other rapidly moving. This behavior is nicely captured in the models we consider with volatility driven by diffusive and jump components.

[5] We originally included a variance risk premia term in the return drift, $\mu + \beta V_t$. It was insignificant and was therefore dropped from the analysis, consistent with Andersen, Benzoni, and Lund (2002) and Pan (2002), who also found this parameter insignificant in models with jumps in returns.

Table I

Conditional Moments

This table summarizes the instananeous conditional moments for the four models under consideration. In the case of the SVCJ model, the second moment of the jump sizes is $E[(\xi^y)^2] = \mu_y^2 + 2\mu_y\mu_v\rho_J + \rho_J^2\mu_v^2 + \sigma_y^2$.

	SV	SVJ	SVCJ	SVIJ
$\frac{1}{dt}var_t(dY_t)$	V_t	$V_t + \lambda_y(\mu_y^2 + \sigma_y^2)$	$V_t + \lambda_y E\left((\xi^y)^2\right)$	$V_t + \lambda_y\left(\mu_y^2 + \sigma_y^2\right)$
$\frac{1}{dt}var_t(dV_t)$	$\sigma_v^2 V_t$	$\sigma_v^2 V_t$	$\sigma_v^2 V_t + \lambda_y\mu_v$	$\sigma_v^2 V_t + \lambda_v\mu_v$
$\frac{1}{dt}cov_t(dY_t, dV_t)$	$\rho\sigma_v V_t$	$\rho\sigma_v V_t$	$\rho\sigma_v V_t + \lambda_y\rho_J\mu_v^2$	$\rho\sigma_v V_t$

In the SV and SVJ models, θ controls the long-run mean of V_t as $\bar{V} = E(V_t) = \theta$. However, in the presence of jumps in volatility, θ is only the diffusive, and not total, long-run mean of V_t. With jumps in volatility, $E(V_t) = \theta + (\mu_v\lambda_v)/\kappa$ is the long-run mean.[6] To see how jumps in returns and volatility affect the second moments of returns and volatility, Table I provides the instantaneous variance and covariance of Y_t and V_t for each of the models.[7] The exponentially distributed jumps in V_t capture the large positive outliers in volatility documented in Bates (2000) and guarantee that V_t is positive. The SVCJ and SVIJ models have three factors (diffusive stochastic volatility, jumps in returns, and jumps in volatility), and we now discuss how these factors impact the distribution of returns.

Although jumps in returns and diffusive stochastic volatility can both generate realistic patterns of *unconditional* nonnormalities in returns, they generate very different patterns of *conditional* nonnormalities. Jumps in returns result in a discrete mixture of normal distributions for returns, which easily generates unconditional and conditional nonnormalities over short frequencies such as daily or weekly. Over longer intervals, a central limit effect results in decreasing amounts of excess skewness and kurtosis (see Duffie and Pan (1997) or Das and Sundaram (1999)).

Diffusive stochastic volatility models induce a different distributional structure. Conditional on the current level of volatility, returns are approximately

[6] To calculate this, integrate the SDE to get

$$E(V_t) = E(V_0) + E\left(\int_0^t \kappa(\theta - V_s)ds\right) + E\left(\int_0^t \sigma_v\sqrt{V_s}\,dW_s\right) + E\left(\sum\nolimits_{j=1}^{N_t^v} \xi_j^v\right).$$

Since the stochastic integral against the Brownian motion is mean zero, Fubini's theorem implies that

$$\bar{V} = \bar{V} + \kappa(\theta - \bar{V})t + E\left[\sum\nolimits_{j=1}^{\infty} 1_{[N_t^v \le j]}\xi_j^v\right].$$

Since $E(\xi_j^v) = \mu_v$ and the jump arrivals are Poisson, we have that $E\left[\sum_{j=1}^{\infty} 1_{[N_t^v = j]}\xi_j^v\right] = \mu_v\lambda_v t$, which implies that $\bar{V} = \theta + \frac{\mu_v\lambda_v}{\kappa}$.

[7] The notation $\frac{1}{dt}var_t(dY_t)$ is a heuristic for $\lim_{\Delta\to 0} E\left[(\log(S_{t+\Delta}/S_t))^2 | V_t, S_t\right]$. Since our initial conditions, drift, diffusion, jump intensity, and jump size distributions have sufficient regularity, the limiting operation is valid.

normally distributed over short time horizons such as daily or even weekly. Das and Sundaram (1999) find that diffusive stochastic volatility models, with reasonable parameters, can generate realistic amounts of *conditional* nonnormalities in returns only over longer horizons, such as months or even years. This explains why diffusive stochastic volatility models (square-root or otherwise) generate very flat implied volatility curves for short-dated options with plausible parameters (see Das and Sundaram (1999) or Jones (2002)).

Models with only jumps in returns and diffusive stochastic volatility can generate realistic patterns of both unconditional and conditional nonnormalities, but they have difficulty capturing the dynamics of the conditional volatility of returns. In the SVJ model, the conditional variance of returns is $V_t+({\mu_y}^2+{\sigma_y}^2)\ \lambda_y$ (see Table I). When V_t is a diffusion, the conditional volatility of returns is time varying and persistent, but moves slowly as it is driven by normally distributed shocks. As pointed out by Bates (2000) and Pan (2002), this creates misspecification, as they find that volatility needs to increase rapidly.

Jumps in volatility provide a factor that combines features from both jumps in returns and diffusive stochastic volatility. Like jumps in returns and unlike stochastic volatility, jumps in volatility are a rapidly moving factor driving returns. Like diffusive stochastic and unlike jumps in returns whose impact on returns is transient, a jump in volatility persists. Thus, jumps in volatility provide a rapidly moving but persistent factor driving volatility. The fact that each factor generates very different behavior is helpful for econometric identification.

There are alternative explanations for the failure of models with only jumps in returns and square-root stochastic volatility. These include additional square-root volatility factors, more flexible, parametric, single-factor stochastic volatility specifications, or combinations of these two. More flexible single-factor specifications such as the log and CEV models allow for the volatility of volatility to be state dependent, a property absent in square-root models. These models are limited, because they do not provide closed-form option prices. Andersen, Benzoni, and Lund (2002) and Chernov et al. (2002) find that the log-volatility and square-root models provide a near identical fit to the data and that neither model can capture the fat tails in the return distribution. Chacko and Viceira (2001) and Jones (2002) consider a CEV stochastic volatility specification. While more flexible than the log model, returns in the CEV model are still conditionally normal over short time intervals. Chacko and Viceira find nonlinearities in the variance of V_t, but the effect disappears when jumps in returns are added. Consistent with the arguments above, Jones shows that the CEV model generates realistic unconditional nonnormalities, but over short time intervals, the model offers only a modest improvement over the square-root model in generating conditional nonnormalities.

Bates (2000) and Chernov et al. (2002) consider and reject two-factor square-root stochastic volatility models. In fact, Bates (p. 218) argues that "the postulated square-root process for implicit factor evolution is fundamentally misspecified," that the "two-factor models have even greater difficulties than the one-factor models in generating sample paths consistent with the postulated process and the implicit parameters" (p. 214), and additionally argues that these models are overfit.

To investigate any improvements generated by additional volatility factors, we estimated a two-factor independent square-root stochastic volatility model where $dV_{t,i} = \kappa_i(\theta_i - V_{t,i})dt + \sigma_i\sqrt{V_{t,i}}dW_t^i$.[8] Consistent with the findings of Bates (2000) and Chernov et al. (2002), we found little evidence that this model provided any substantive improvement over the single-factor square-root model. In fact, it suffered from the same misspecification as the square-root and log-volatility models: It cannot capture the tails of the return distribution. This is not surprising, as the multifactor model is instantaneously Gaussian, conditional on total volatility.[9] For these reasons, we focus on a single-factor square-root model of stochastic volatility and extensions incorporating jumps in returns and volatility.

II. Estimating Stochastic Volatility Jump Diffusions

This section develops a likelihood-based estimation approach for estimating multivariate jump-diffusion models using Markov Chain Monte Carlo (MCMC) methods. Robert and Casella (1999) provide a general discussion of these methods, and Johannes and Polson (2002) provide an overview of MCMC estimation of continuous-time models. This approach has four advantages over other estimation methods:[10] (1) MCMC provides estimates of the latent volatility, jump times, and jumps sizes; (2) MCMC accounts for estimation risk; (3) MCMC methods have been shown in related settings to have superior sampling properties to competing methods;[11] and (4) MCMC methods are computationally efficient so that we can check the accuracy of the method using simulations.

Our approach uses only returns data to estimate and test the models, although it can be extended in a straightforward manner to include option price data (see Eraker (2002)). Due to the absolute continuity of the change in measure from objective to risk neutral, the presence of jumps in returns or volatility under one

[8] We thank the referee for suggesting this.

[9] For example, in a model with two independent square root volatility factors, the instantaneous distribution of price changes,

$$\frac{dS_t}{S_t} = \mu dt + \sqrt{V_{t,1}}dW_{t,1}^s + \sqrt{V_{t,2}}dW_{t,2}^s$$

is still Gaussian:

$$\frac{dS_t}{S_t} \sim N\big(\mu dt, \big(V_{t,1} + V_{t,2}\big)dt\big).$$

[10] Other methods that have been used to estimate models with stochastic volatility and jumps include EMM, simulated maximum likelihood (Brandt and Santa-Clara (2002), Durham and Gallant (2001), Piazessi (2001)), calibration (Bates (1996, 2000), Bakshi, Cao, and Chen (1997)), and the implied-state GMM method of Pan (2002).

[11] Jacquier, Polson, and Rossi (1994) find in simulations that MCMC outperforms GMM and QMLE in estimation of stochastic volatility models, and Andersen, Chung, and Sorensen (1999) find that MCMC outperforms EMM.

measure implies their presence under the other, although due to risk premia, their impact may be distorted. Thus, for specification analysis, returns data should lead to the same conclusion as option price data.

The main advantage of using only returns data is more pragmatic: Analyses using option price data tend to use relatively short time spans. For example, Bakshi, Cao, and Chen (1997) use data from 1988 to 1991 while Pan (2002) uses data from 1989 to 1996. This is especially important when estimating models with jumps, which we expect to occur infrequently. Longer samples spanning periods of market stress (the episodes in 1987, 1997, and 1998) will provide more accurate parameter estimates and better insights into the relative roles played by jumps and stochastic volatility.

The basis for our MCMC estimation is a time-discretization of (1)

$$\begin{aligned} Y_{(t+1)\Delta} - Y_{t\Delta} &= \mu\Delta + \sqrt{V_{t\Delta}\Delta}\varepsilon^y_{(t+1)\Delta} + \xi^y_{(t+1)\Delta}J^y_{(t+1)\Delta} \\ V_{(t+1)\Delta} - V_{t\Delta} &= \kappa(\theta - V_{t\Delta})\Delta + \sigma_v\sqrt{V_{t\Delta}\Delta}\varepsilon^v_{(t+1)\Delta} + \xi^v_{(t+1)\Delta}J^v_{(t+1)\Delta}, \end{aligned} \tag{2}$$

where $J^k_{(t+1)\Delta} = 1\,(k = y, v)$ indicates a jump arrival, $\varepsilon^y_{(t+1)\Delta}$ and $\varepsilon^v_{(t+1)\Delta}$ are standard normal random variables with correlation ρ and Δ is the time-discretization interval (one day).[12] The jump sizes retain their distributional structure and the jump times are Bernoulli random variables with constant intensities, $\lambda_y\Delta$ and $\lambda_v\Delta$. This procedure could introduce a discretization bias, although the bias is typically quite small with daily data. We provide simulations below to support this claim.

In this section, we focus on the SVCJ model, as it has the most complicated distributional structure. The posterior distribution summarizes the sample information regarding the parameters, Θ, and the latent volatility, jump times, and jump sizes:

$$p(\Theta, J, \xi^y, \xi^v, V|Y) \propto p(Y|\Theta, J, \xi^y, \xi^v, V)p(\Theta, J, \xi^y, \xi^v, V), \tag{3}$$

where J, ξ^y, ξ^v, V, and Y are vectors containing the time series of the relevant variables. The posterior combines the likelihood, $p(Y|\Theta, J, \xi^y, \xi^v, V)$, and the prior, $p(\Theta, J, \xi^y, \xi^v, V)$.

An advantage of our approach is the ability to formally incorporate prior information. The need for this is not unique to our approach, but is common in estimating models with jumps. Honore (1998) shows that without prior parameter restrictions, a time discretization of Merton's (1976) jump-diffusion model generates an unbounded likelihood function. Moreover, the prior contains information about both the parameters and the structure of the latent processes: the stochastic specifications of the jump sizes, jump times, and volatility. This reinforces the link between parameters and model specification that is often heuristically used to motivate the presence of jumps. Typically, jumps are described as large, but

[12] Our framework is not limited to the case in which the discretization interval, Δ equals the observed frequency. To reduce any bias, we could introduce additional unobserved data points between dates t and t+1 and treat them as missing data points to be included in the MCMC simulation (see Eraker (2001)).

infrequent movements in returns. This is a form of prior information as the parameters are assumed to induce infrequent but relatively large movements (low λ_y and large μ_y and/or σ_y), as opposed to frequent but small jumps.

Our priors are always consistent with the intuition that jumps are "large" and infrequent. More specifically, we choose a prior on σ_y that places low probability on the jump sizes being small, say less than one percent. For λ_y, our prior places low probability on the daily jump probability being greater than 10 percent and we place an uninformative prior on μ_y. For the other parameters, we specify extremely uninformative priors. For μ, κ, θ, and ρ_J we use mean zero normal priors with large variances and the prior on ρ is uniform over $[-1,1]$. Appendix A provides further details on the priors. It is important to note that we impose very little information through our priors.

As the posterior distribution is not known in closed form, our MCMC algorithm generates samples by iteratively drawing from the following conditional posteriors:

$$\begin{aligned}
&\textit{parameters}: p(\Theta_i|\Theta_{-i},\, J,\, \xi^y,\, \xi^v,\, V,\, Y),\; i = 1,\ldots,k\\
&\textit{jump times}: p(J_{t\Delta} = 1|\Theta,\, \xi^y,\, \xi^v,\, V,\, Y),\; t = 1,\ldots,T\\
&\textit{jump sizes}: p(\xi^y_{t\Delta}|\Theta,\, J_{t\Delta} = 1,\, \xi^v,\, V,\, Y),\; t = 1,\ldots,T\\
&\qquad\qquad\quad : p(\xi^v_{t\Delta}|\Theta,\, J_{t\Delta} = 1,\, V,\, Y),\; t = 1,\ldots,T\\
&\textit{volatility}: p(V_{t\Delta}|V_{(t+1)\Delta},\, V_{(t-1)\Delta},\, \Theta,\, J,\, \xi^y,\, \xi^v,\, Y),\; t = 1,\ldots,T,
\end{aligned}$$

where Θ_{-i} denotes the elements of the parameter vector except Θ_i. Drawing from these distributions is straightforward, with the exception of volatility, as the distribution is not of standard form (Appendix A provides details). The algorithm produces a set of draws $\{\Theta^{(g)}, J^{(g)}, (\xi^{y(g)}), (\xi^{v(g)}), V^{(g)}\}_{g=1}^{G}$ which are samples from $p(\Theta, J, \xi^y, \xi^v, V|Y)$. Johannes and Polson (2002) provide a review of the theory behind MCMC algorithms.

A. Estimating Volatility and Jumps

For specification analysis and to identify the relative importance of jumps and volatility, we require estimates of the latent volatility, jump times, and jump sizes. With continuous record observations, all are observed, but with discretely sampled observations, it is not obvious how to separate out the effects of jumps and time-varying volatility. For example, is a large movement in returns generated by a jump in returns or by high volatility? Standard latent variable estimation methods such as the Kalman filter do not apply, as our model is neither linear nor Gaussian.

Our MCMC approach provides a straightforward method to estimate the volatilities, jump times, and jump sizes by computing the posterior expectation of these variables. We, therefore, provide a Monte Carlo solution to the classical, latent variable estimation problem. The key is that our MCMC algorithm generates samples of the spot volatilities, jump times, and jump sizes, drawn from the joint posterior distribution. Given these samples, the Monte Carlo estimate of the mean of the posterior volatility distribution, for example, is

$E[V_{t\Delta}|Y] \approx \frac{1}{G}\sum_{g=1}^{G} V_{t\Delta}^{(g)}$, where $V_{t\Delta}^{(g)}$ is the variance at time t in the g^{th} iteration of the algorithm. No additional calculations are required: Latent variable estimation is just a by-product of our algorithm. Jump time and size estimates are similarly calculated.

These estimates take into account parameter uncertainty. To see this, note that we estimate $E[V_{t\Delta}|Y]$ and not $E[V_{t\Delta}|Y, \hat{\Theta}]$. The former distribution integrates out all of the parameter uncertainty. The latter distribution treats the parameter estimates as known, ignoring the fact that they are random variables.

B. Model Diagnostics and Specification Tests

The spot volatility, jump time, and jump size estimates generate a number of informal diagnostics that are useful in assessing the ability of the various models to fit the observed data. For example, consider the return residuals:

$$\frac{Y_{(t+1)\Delta} - Y_{t\Delta} - \mu\Delta - J_{t+1}^{y}\xi_{t+1}^{y}}{\sqrt{V_{t\Delta}\Delta}} = \varepsilon_{(t+1)\Delta}^{y} \approx N(0,\, 1). \tag{4}$$

It is easy to compute the posterior of these residuals using the parameter and latent variable samples. These residuals need not be exactly normally distributed (since we use a time discretization of the original model), but they should be approximately normal. Extremely large residuals suggest misspecification as the model requires abnormally large shocks to fit the observed data. Jump time and size estimates provide further diagnostic tools. For instance, any evidence of clustered jump times contradicts the i.i.d. arrival assumption. Jumps on neighboring days of opposite signs, the reversal effect of Schwert (1990), is also evidence against the i.i.d. jump arrivals and sizes.

We also compute formal specification tests. Unlike standard tests that lead to an overall evaluation of the fit of the model (e.g., omnibus chi-square tests), we compare the marginal likelihoods of the models. The advantage of this approach is that it does not rely on large sample distribution theory and provides an intuitive approach to evaluating the relative merits of competing models.

Our approach for comparing nested models is similar to that considered in Jacquier and Polson (1999). Consider testing SV versus SVJ. If we assume positive prior odds on the models, $p(SV)/p(SVJ) > 0$ (prior ignorance sets this to 1), Bayes theorem implies that the posterior odds are

$$\frac{p(SV|Y)}{p(SVJ|Y)} = \frac{p(Y|SV)}{p(Y|SVJ)}\frac{p(SV)}{p(SVJ)}, \tag{5}$$

where, $p(Y|SV)/p(Y|SVJ)$ is known as the Bayes Factor. Note that this analysis does not assume that the models are necessarily exhaustive, that is, that $p(SV) + p(SVJ) = 1$. Assuming prior ignorance, the Odds Ratio, is interpreted in the following manner: A risk neutral bookmaker would lay odds, $\frac{p(Y|SV)}{p(Y|SVJ)}$ to 1, on model SVJ versus SV. An advantage of this approach is that the test results in a single number that is the relative odds of the models given the data, and there is no appeal to approximate limiting distributions or significance levels. Appendix B derives the marginal likelihoods and Bayes Factors for the models under consideration and provides MCMC estimators.

The fact that the models with jumps in volatility are not formally nested provides no problem for computing the Bayes Factors, for example, SVCJ versus SVIJ. We compute the Bayes Factor for the nonnested models using the fact that

$$\frac{p(SVCJ|Y)}{p(SVIJ|Y)} = \frac{p(SVCJ|Y)}{p(SVJ|Y)} \frac{p(SVJ|Y)}{p(SVIJ|Y)}. \tag{6}$$

One caveat is that the Monte Carlo standard errors may be larger than in the case of the simple nested comparisons, as the Odds Ratio is now the product of two Odds Ratios.

Following Kass and Raftery (1995), we use the following scale to interpret the Bayes Factors. Evidence against a model is positive if the log Odds Ratio is between 2 and 6, strong if it is between 6 and 10, and very strong if it is greater than 10. It is important to note that Odds Ratios do not necessarily favor more complex models, as they contain a penalty for using more parameters (due to their marginal nature). Because of this, Odds Ratios are often referred to as an "automatic Occam's razor" (see Smith and Spiegelhalter (1982)).

C. Simulation Results

We performed Monte Carlo simulations to check the reliability of our estimation approach. This is important for two reasons. First, since we time discretize the continuous-time model, it is important to check that this does not introduce any biases in parameter estimates. Second, methods for estimating multivariate jump diffusion models are not well developed and it is important to verify that we can reliably estimate the parameters for the given sample size.

Appendix C describes our simulation study, and Tables VII and VIII in that appendix summarize the results. The results indicate that our procedure provides accurate inference. Some parameters are estimated less precisely than others, for example κ, but all are close to their true values. The results also indicate that our priors are not informative, as we use the same priors for parameters common in both the SVJ and SVCJ models, even though the estimation results and true parameters differ.

III. Empirical Results

We estimate the models using S&P 500 and Nasdaq 100 index returns from January 2, 1980, to December 31, 1999, and September 24, 1985, to December 31, 1999, respectively. Excluding weekends and holidays, we have 5,054 daily observations for the S&P and 3,594 observations for the Nasdaq. Table II provides summary statistics for the continuously compounded returns, scaled by 100. In this section, we discuss the estimation results for the S&P 500, for each of the models, how the Nasdaq estimates differ from the S&P, and the role of jumps and volatility in the three periods of market stress in our sample. Tables III and IV summarize

Table II
Summary Statistics

This table provides summary statistics for daily return data on the S&P 500 from January 2, 1980, to December 31, 1999, and the Nasdaq 100 from September 24, 1985, to December 31, 1999.

	Mean	Volatility	Skewness	Kurtosis	Min	Max
S&P 500	13.0185	15.8838	– 2.6064	62.0328	– 22.8330	8.7089
Nasdaq 100	24.5841	23.2182	– 0.7287	11.9424	– 16.3405	9.7984

Table III
S&P 500 Parameter Estimates

Parameter estimates for the S&P 500 index data, January 2, 1980, to December 31, 1999. The models and parameterizations are given in Section I and the estimates correspond to percentage changes in the index value. For each parameter, we report the mean of the posterior deviation and the standard deviation of the posterior in parentheses.

	SV	SVJ	SVCJ	SVIJ
μ	0.0444 (0.0110)	0.0496 (0.0109)	0.0554 (0.0112)	0.0506 (0.0111)
θ	0.9052 (0.1077)	0.8136 (0.1244)	0.5376 (0.0539)	0.5585 (0.0811)
κ	0.0231 (0.0068)	0.0128 (0.0039)	0.0260 (0.0041)	0.0250 (0.0057)
σ_v	0.1434 (0.0128)	0.0954 (0.0104)	0.0790 (0.0074)	0.0896 (0.0115)
μ_y		– 2.5862 (1.3034)	– 1.7533 (1.5566)	– 3.0851 (3.2485)
ρ_J			– 0.6008 (0.9918)	
σ_y		4.0720 (1.7210)	2.8864 (0.5679)	2.9890 (0.7486)
μ_v			1.4832 (0.3404)	1.7980 (0.5737)
ρ	– 0.3974 (0.0516)	– 0.4668 (0.0579)	– 0.4838 (0.0623)	– 0.5040 (0.0661)
λ_y		0.0060 (0.0021)	0.0066 (0.0020)	0.0046 (0.0020)
λ_v				0.0055 (0.0032)

parameter estimation for all of the models for the S&P 500 and Nasdaq 100, respectively.[13]

A. S&P 500

The second column of Table III provides parameter posterior means and standard deviations for the SV model. The left-hand panel of Figure 1 provides spot volatility estimates over two time periods, 1987 to 1989 and 1997 to 1999. The parameter estimates are consistent with previous findings. The average annualized volatility, $\sqrt{252 \cdot \theta}$, is 15.10 percent and is close to the sample volatility of 15.89

[13] For the SV, SVJ, and SVCJ models, the MCMC algorithm appears to converge quickly. We discard the first 10,000 iterations as a "burn-in" period and use the last 90,000 to form the Monte Carlo estimates. For the SVIJ model, the algorithm appears to converge more slowly and thus we ran it for 200,000 iterations, discarding the first 10,000 as a burn-in period. For the Nasdaq 100, the algorithm converged quickly for all models, and thus we discarded the first 10,000 draws and used the last 90,000 draws.

Table IV

Nasdaq 100 Parameter Estimates

Parameter estimates for the Nasdaq 100 index data, September 24, 1985, to December 31, 1999. The models and parameterizations are given in Section I and the estimates correspond to percentage changes in the index value. For each parameter, we report the mean of the posterior deviation and the standard deviation of the posterior in parentheses.

	SV	SVJ	SVCJ	SVIJ
μ	0.1070 (0.0196)	0.1240 (0.0205)	0.1284 (0.0206)	0.1164 (0.0196)
θ	2.0242 (0.2309)	1.9067 (0.2853)	0.9249 (0.1547)	1.0593 (0.1506)
κ	0.0264 (0.0072)	0.0176 (0.0052)	0.0414 (0.0106)	0.0371 (0.0080)
σ_v	0.2155 (0.0226)	0.1692 (0.0190)	0.1216 (0.0188)	0.1395 (0.0196)
μ_y		− 2.4755 (1.0276)	− 1.8868 (0.7407)	− 2.6231 (2.5946)
ρ_J			− 0.0993 (0.1692)	
σ_y		− 2.0788 (0.4375)	1.8452 (0.3079)	2.0389 (0.4485)
μ_v			2.1054 (0.3989)	2.5227 (0.5946)
ρ	− 0.2869 (0.0553)	− 0.3366 (0.0695)	− 0.3427 (0.0975)	− 0.3910 (0.0801)
λ_y		0.0172 (0.0097)	0.0202 (0.0074)	0.0081 (0.0044)
λ_v				0.0140 (0.0058)

percent. Our estimate of ρ, -0.40, is close to the estimate obtained by Jacquier, Polson, and Rossi (2001) in a log-volatility model ($\rho = -0.39$) and Andersen, Benzoni, and Lund (2002) ($\rho = -0.38$), but is slightly smaller than those obtained from studies using option price data (Bakshi, Cao, and Chen (1997), Bates (2000), and Pan (2002) obtain estimates around -0.5).

The SV model is misspecified. This can be seen in a number of ways. First, consider the QQ or normal probability plot of the residuals in the upper left panel in Figure 2. Note the extreme nonnormality of the residuals, strong evidence of misspecification. Why does the model require such large shocks? Consider the crash in 1987. On the day of the crash, $\sqrt{V_t}$ was about three percent, which implies that an almost eight standard deviation return shock is needed to deliver the -23 percent move. This type of misspecification was also noted in stochastic volatility models by Jacquier, Polson, and Rossi (2001), Andersen, Benzoni, and Lund (2002), and Chernov et al. (2002). Second, spot volatility increased for 37 consecutive days prior to the crash in 1987. This is extremely unlikely as the mean reverting drift was exerting downward pressure on volatility throughout this period. Finally, the Bayes Factors reported in Table VI below provides evidence against the SV model in favor of the other models.

The third column in Table III provides parameter estimates for the SVJ model. Adding jumps in returns has the expected effect of reducing the demands on the volatility process. For example, average volatility falls from 15.10 percent to 14.32 percent and both σ_v and κ fall dramatically, indicating a less volatile, more persistent volatility process. Estimates of spot volatility in Figure 1 show that jumps in returns reduce spot volatility dramatically during periods of market stress, as jumps generate the largest movements. Also note that after the crash in 1987, spot volatility in the SV model remained higher than in the SVJ model for a long time.

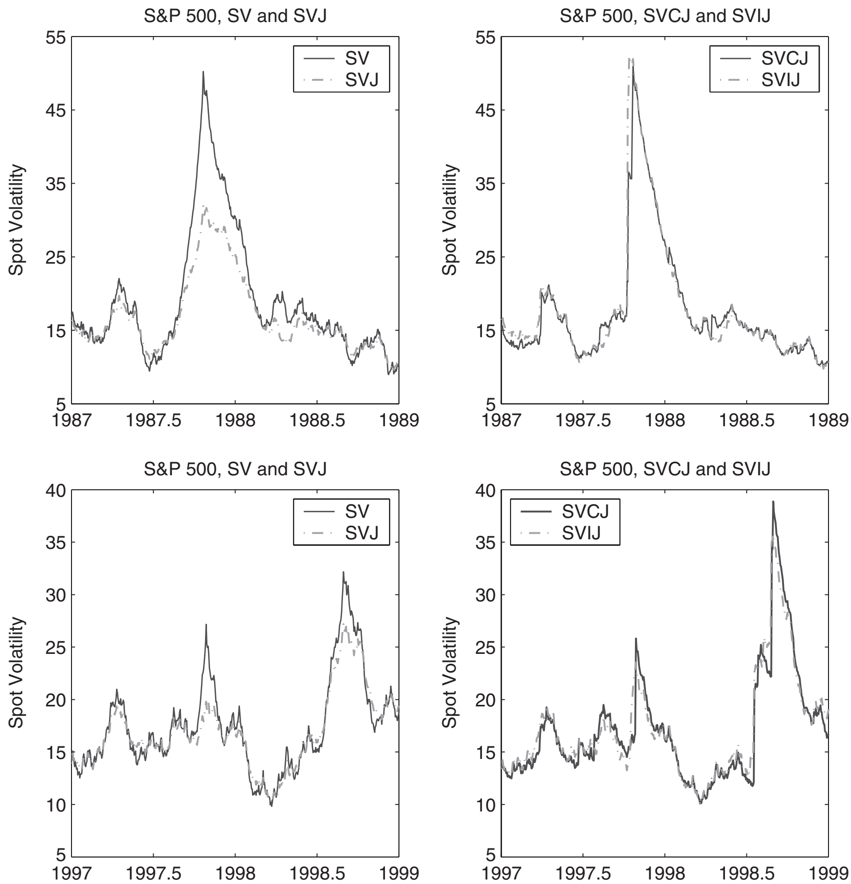

Figure 1. Estimated volatility paths for the S&P 500 index for the four models.

Thus, the effect of omitting jumps in returns lingers after the jump arrived through its impact on estimated volatility.

In the SVJ model, jumps in returns are infrequent events (about 1.5 per year), tend to be negative, and are large relative to normal day-to-day movements. A three standard deviation jump move is about -15 percent. Table V decomposes the total variance of returns into stochastic volatility and jump components. The proportion of variance due to jumps is

$$\frac{E[(\xi_t^y)^2]\lambda_y}{\bar{V} + E[(\xi_t^y)^2]\lambda_y}$$

and is 14.65 percent in the SVJ model.

The normality plots in the SVJ model are improved and the residuals now have slightly thin tails as jumps in returns capture nearly all of the large movements in returns. Figure 3, however, provides evidence that the SVJ model is misspecified.

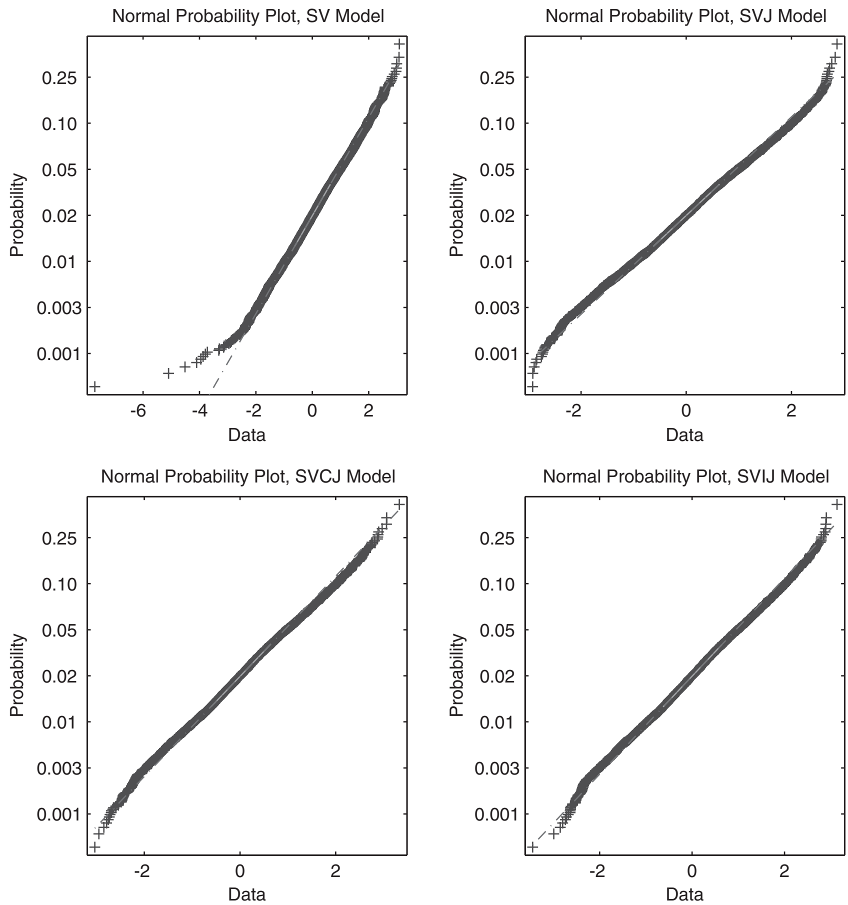

Figure 2. QQ or normality plot of the residuals for each of the models.

During the week of the crash in 1987, there were three days on which the estimated jump probabilities were extremely high, indicating a cluster of jumps. Why did this occur? During October 1987, daily volatility was always less than two percent, implying that a three standard deviation move in returns due to volatility was only six percent. Due to this, all of the moves larger than six percent were attributed to jumps in returns. Similarly, in October 1997, there were jumps estimated on neighboring days with opposite signs, the reversal effect of Schwert (1990).[14] The clustering of jump arrivals and size reversals are extremely unlikely given the i.i.d. jump time and size specifications and the infrequent nature of jumps (1.5 per year). This suggests that the SVJ model is misspecified. The Bayes

[14] Johannes, Kumar and Polson (1999) document this pattern and develop models with state-dependent arrivals and jump sizes to capture this phenomenon.

Table V

Variance Decompositions

The second and third columns give the average annualized spot stochastic volatility for each of the models and for the S&P 500 and Nasdaq 100. The fourth and fifth columns give the average annualized total volatility, which is the sum of stochastic volatility and jumps in returns components. The last two columns give the proportion of the variance due to jumps.

	Spot Volatility		Total Volatility		Return Jump	Variance (% of Total)
0	S&P	NDX	S&P	NDX	S&P	NDX
SV	15.10	22.59	15.10	22.59	–	–
SVJ	14.32	21.92	15.49	24.06	14.65	17.00
SVCJ	15.18	22.18	15.99	23.59	9.96	11.63
SVIJ	15.51	22.51	16.18	23.22	8.17	5.99

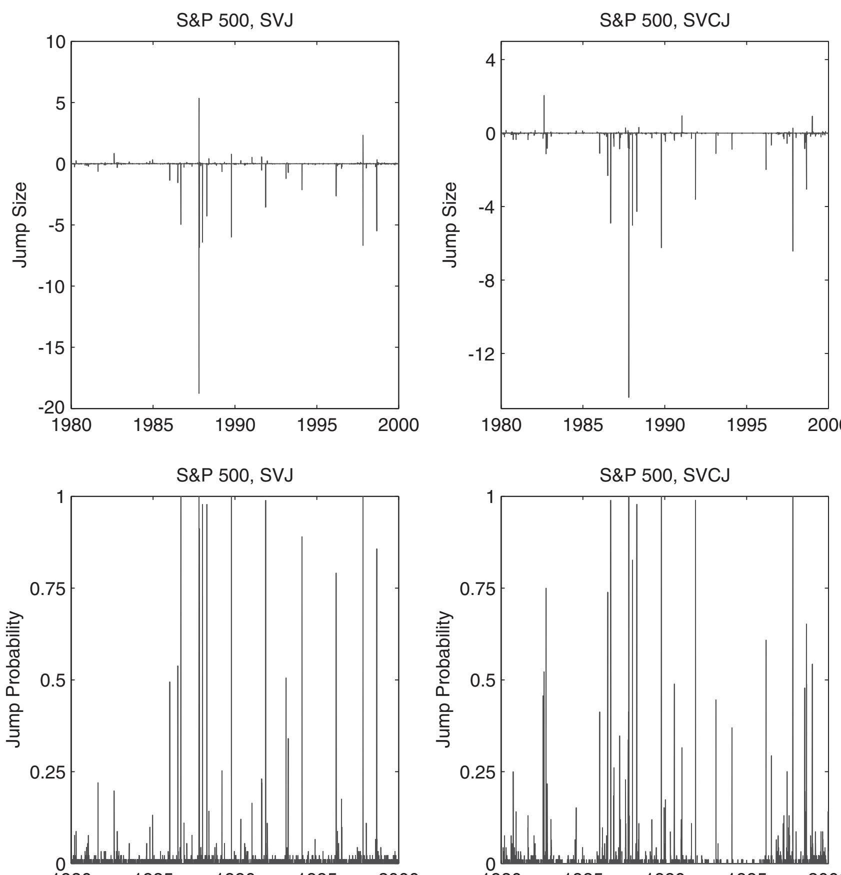

Figure 3. Estimated jump times and sizes for the SVJ and SVCJ models for the S&P 500 index.

Table VI
Log-Bayes Factors for the S&P 500 and Nasdaq 100

	SV vs. SVJ	SV vs. SVCJ	SV vs. SVIJ	SVJ vs. SVIJ	SVCJ vs. SVIJ
S&P 500	47.62	49.32	59.45	33.83	10.12
Nasdaq 100	27.16	28.95	56.03	33.19	27.08

factors in Table VI provide additional evidence against the SVJ model and in favor of the models with jumps in volatility

In conclusion, the SV and SVJ models suffer from similar problems. Periods of market stress are characterized by a short time period with multiple large movements and neither model can generate these movements. In the SV model, these large movements appear as abnormally large shocks, while in the SVJ model, they appear as clustered jumps. In both cases, diffusive volatility cannot increase rapidly enough to generate these episodes and the models are misspecified.

The fourth column of Table III provides parameter estimates for the SVCJ model and the results indicate that jumps in volatility play an important role. When a jump arrives, volatility increases from $\sqrt{V_{t-}}$ to $\sqrt{V_{t-} + \xi_v}$, and the parameter estimates point toward significant increases. For example, when volatility is 15 percent, an average size jump increases volatility to 24 percent. The solid line in the right hand panels of Figure 1 shows the estimated volatility path for the SVCJ model, and it is very different from those in the SV and SVJ models.

As jump-driven high volatility generates many of the large moves in the volatility process, σ_v is smaller than in the SVJ model. In the SVCJ model, volatility mean reverts faster, as κ is almost double its value in the SVJ model. Jumps still arrive at the rate of about 1.5 per year, but jumps in returns are smaller. Jumps in returns play a lesser role in the SVCJ model than in the SVJ model, as the proportion of total variance coming from jumps in returns is now only 9.96 percent. The jump size estimates in the upper right hand panel of Figure 3 indicate that nearly all of the jumps are negative, as the increased volatility after a jump generates the reversal effect. The estimate of ρ falls again, which is consistent with Duffie, Pan, and Singleton (2000).

The normal residual plot in Figure 3 gives no indication of misspecification in the SVCJ model. As jumps in returns are smaller, they explain less of the large movements and the Brownian increments are almost perfectly normally distributed. Although it is difficult to see in Figure 2, there was only a single jump during the week of the crash in 1987 on October 19, when returns had a negative jump of -14 percent and volatility jumped upward from 40 percent to just over 50 percent. The Bayes factors in Table VI support SVCJ over the SV and SVJ models.

Figure 1 shows that the SVIJ and SVCJ models exhibit very similar behavior, as spot volatility in the two models is almost identical. The jump times and sizes are also similar, and are therefore not reported. Allowing volatility to jump independently of returns provides additional flexibility over the SVCJ model, although the model is harder to estimate as jumps in volatility are not signalled by a jump in returns. This added flexibility allows volatility to play an even more important role than in the SVCJ model, as the proportion of variance due to jumps in returns is less than in the SVJ or SVCJ models. Diagnostics do not provide any evidence of misspecification in the SVIJ model. The Bayes Factor comparing SVCJ and SVIJ is over 10, evidence that favors the SVIJ model. A caveat is in order: The SVIJ model is more difficult to estimate, as the jumps do not occur simultaneously and this results in greater posterior parameter uncertainty for all of the jump size parameters.

B. Nasdaq 100

Table IV provides the Nasdaq parameters estimates, using the same priors that were used for the S&P, and Figures 4 and 5 provide volatility and jump estimates. The biggest difference between the Nasdaq and S&P returns is that Nasdaq's

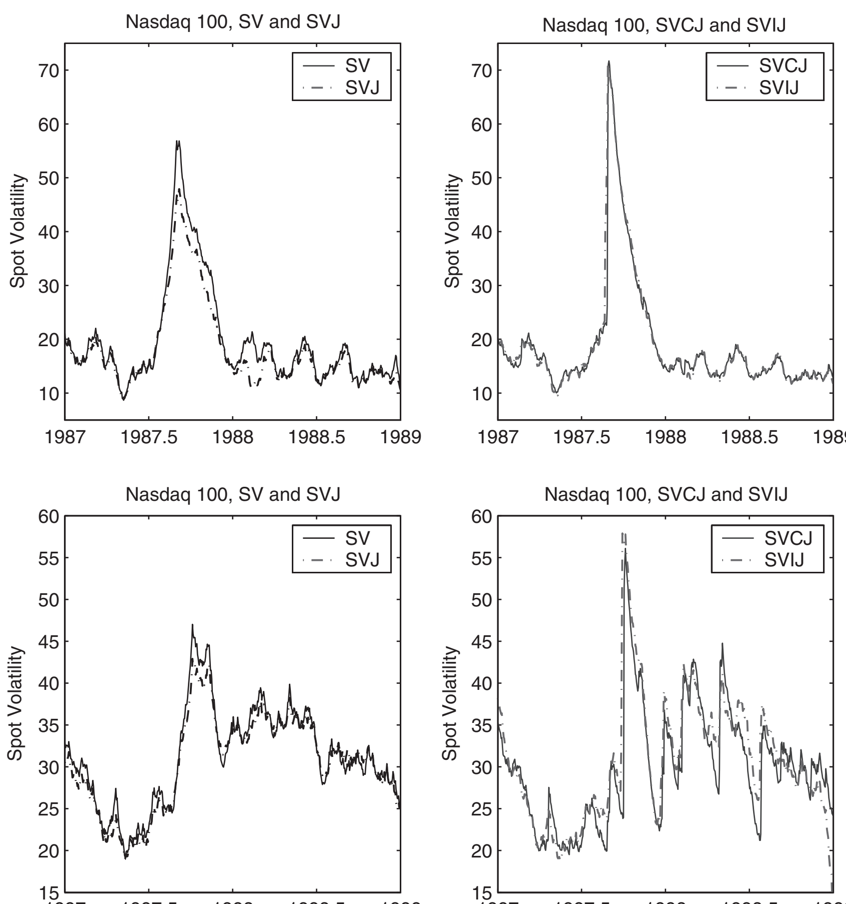

Figure 4. Estimated volatility paths for the Nasdaq 100 for the four models.

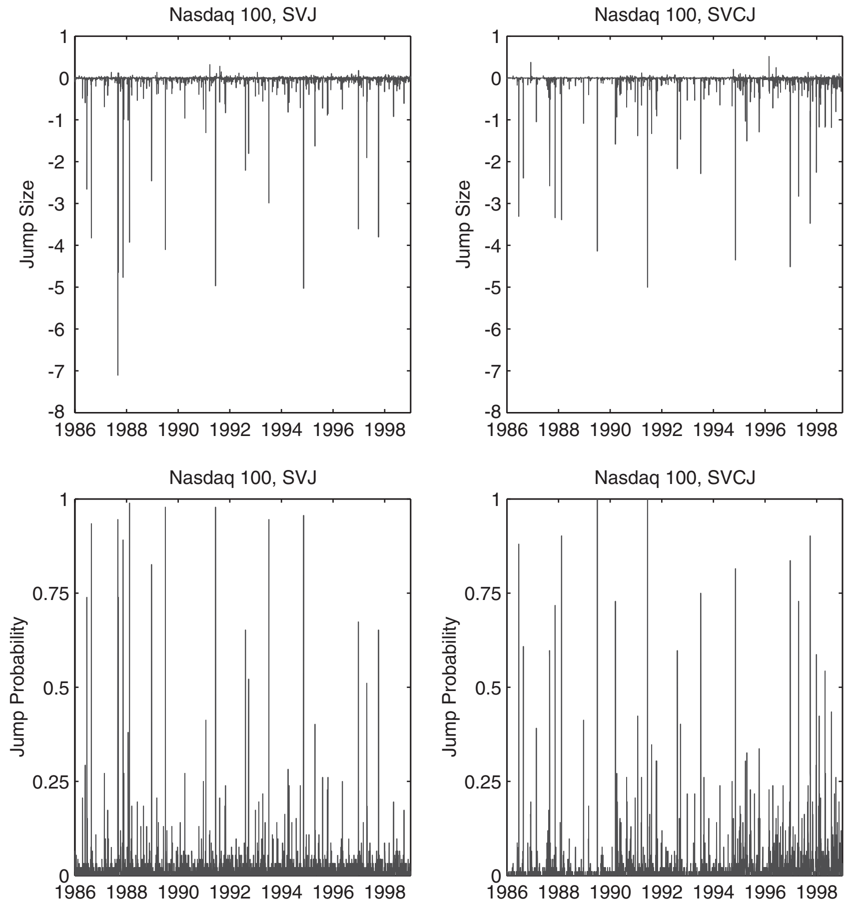

Figure 5. Estimated jump times and sizes for the SVJ and SVCJ models for the Nasdaq 100 index.

volatility is much higher and more volatile than S&P volatility. For example, in the SVJ model, the average daily variance for the Nasdaq is 1.91, more than double the average for the S&P, 0.91. For each of the models, σ_v is roughly 50 percent higher for the Nasdaq than for the S&P and the speed of volatility mean reversion is also higher for the Nasdaq.

Although volatility is systematically higher, this does not imply that jumps play a lesser role; in fact, the contrary is true. Jumps arrive about three times more often, although the sizes are typically smaller. Since the largest single day move in the Nasdaq is smaller than in the S&P (-16 percent compared to -23 percent), jumps in returns need not generate as many extremely large moves (see Figure 5). The combination of the slightly lower jump size volatility with the higher arrival rate results in the proportion of total variance due to jumps increasing slightly, except for the SVIJ model (see Table V). An easier way to see how the volatility structure of the S&P and Nasdaq indices differs is to compare spot volatility estimates in Figure 6. The volatility of the two indices was similar prior to 1990, but since then, the volatility of the Nasdaq has been higher, often three times higher than the volatility of the S&P.

The leverage effect, ρ, is less pronounced for the Nasdaq. This is consistent with the sample statistics, as Nasdaq's skewness is smaller than that of the S&P. As in the case of the S&P, ρ increases, in absolute value, from SV to SVIJ. The average size of a jump in volatility in the Nasdaq is larger than in the S&P. However, as Nasdaq volatility is higher than the volatility of the S&P, the proportional increase in volatility due to an average-sized jump is smaller for the Nasdaq. The specification diagnostics, both formal and informal, indicate that models with jumps in volatility are preferred over those without jumps in volatility.

C. Periods of Markets Stress: Jumps and Volatility

The estimates of spot volatility and jumps provide a means to evaluate the contribution of these factors to the three periods of market stress in our sample:

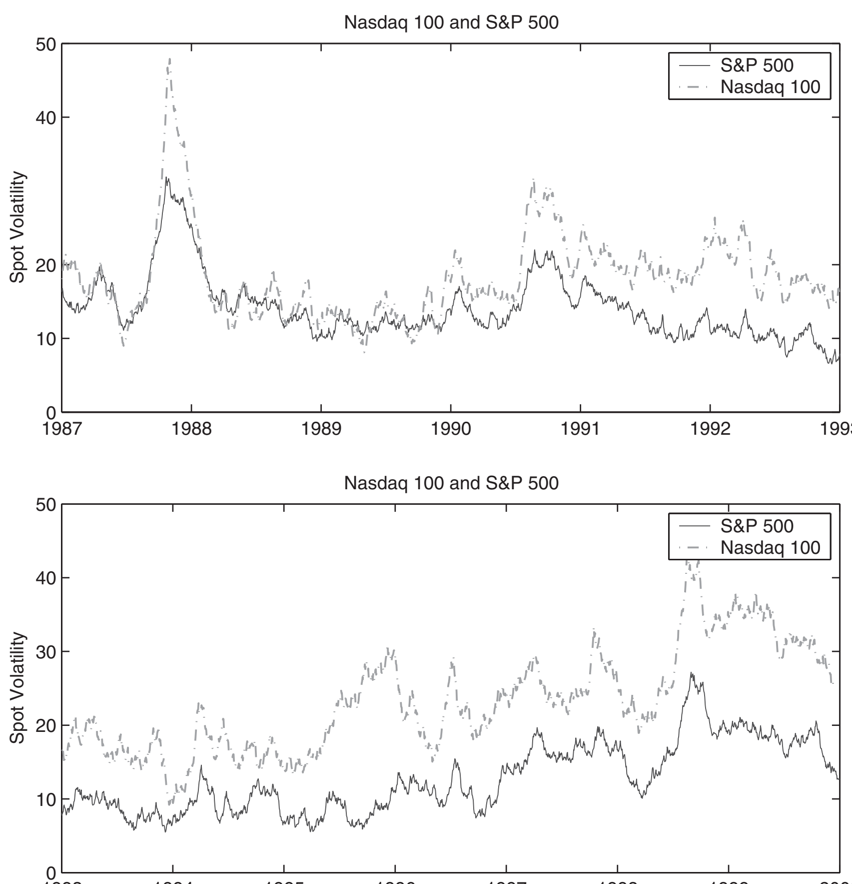

Figure 6. Implied volatility paths for the S&P 500 and Nasdaq 100 for the SVJ model.

October 1987, October 1997, and late Summer–Fall 1998. A priori, it is unclear if jumps in returns or excess volatility were responsible for these periods, as both can generate large movements. Understanding how these factors contribute to these periods is important, because it is precisely the extreme movements during periods of market stress that may cause investors to demand large premia to carry these risks. We focus on the behavior of volatility and jumps in the SVJ and SVCJ models for the S&P 500 index.

In October 1987, the SVJ model attributes most of the large movements to jumps in returns. Of the 23 percent decline on the day of the crash, a jump in returns generated 18.8 percent of the move and there were also two other jumps during the week. However, when allowed to jump, volatility plays a more prominent role in generating market stress. In the SVCJ model, there was a jump about two weeks before the crash, which increased volatility from 21 percent to 36 percent, and on October 16, volatility jumped from 35 to 42 percent. Another jump arrived on the day of the crash, increasing volatility to 50 percent and delivering a -14 percent jump in returns. This implies that in the SVCJ model, jumps in returns generated more than half of the crash. The large movements after the crash (+8 percent, -4 percent, and -8 percent) are attributed to Brownian shocks combined with the high volatility, in contrast to the SVJ model. The SVIJ model delivers similar results, although the exact timing of jumps in volatility and returns is slightly different.

In October 1997, there was a -7 percent move on the 27th and a $+5$ percent move on the 28th. The SVJ model attributes both moves to jumps in returns and volatility stays constant at about 19 percent. SVCJ estimates indicate that there was a single jump on the 27th that lowered returns by -6 percent and increased volatility from 16 percent to 25 percent. Thus the first move was due to a jump in returns, while the second was due to high volatility, a more plausible explanation than back-to-back jumps in returns. These periods indicate that both jumps in volatility and returns are important components of periods of market stress.

Late Summer and early Fall 1998 were periods with a number of relatively large moves in the three to five percent range. Our estimates indicate that most of the moves were generated by high volatility (as opposed to jumps in returns), which, in turn, was generated by jumps in volatility. The first, on July 21, 1998, occurred coincidental to Alan Greenspan's comments on the economy and one day after the *Wall Street Journal* first reported the large losses at Long Term Capital Management. Although the move in the S&P 500 was only -1.6 percent, in the SVCJ model, our algorithm estimated that a jump in volatility occurred with relatively high probability (about 50 percent), increasing volatility from 12 percent to just over 20 percent. Jumps in returns and volatility also occurred at the end of August, corresponding to the halt in trading of the Russian ruble. Over time, volatility gradually mean-reverted back to relatively normal levels.

These three periods indicate that jumps in volatility and returns are the primary components that generated the periods of market stress, while diffusive stochastic volatility plays a secondary role. This suggests that it may be more appropriate to assign risk premia to jumps in returns and volatility components, rather than diffusive stochastic volatility.

IV. Option Pricing Implications

This section analyzes two option pricing implications: how the models induce different option prices and the effect of parameter and spot volatility estimation risk on option prices. Eraker (2002) provides a more detailed analysis of the option pricing implications.

A. Differences across Models

The time-series results indicate that jumps in volatility and returns play an important role in determining the dynamics of returns. To evaluate their economic significance, we examine the cross section and term structure of implied volatilities for the different models, conditional on estimated parameters and spot volatility. These results are comparable to those in Duffie, Pan, and Singleton (2000) with one important difference. Duffie, Pan, and Singleton calibrate the parameters and spot volatility to minimize option pricing errors on a given day, while ours use only the information contained in historical returns.

Figure 7 displays implied volatility (IV) curves for the SV, SVJ, and SVCJ models for three maturities. The SVCJ and SVIJ models deliver similar curves, so we omitted the SVIJ curves. The IV curves were computed using call option prices conditional on the posterior mean of spot volatility and parameters for various strikes. To frame the results, we also plot market implied volatilities for options traded on October 31, 1994, a randomly selected average volatility day in our sample.[15]

Results indicate that jumps in returns and volatility induce important differences in the shape of the IV curves, especially for short maturities. First, and most apparent, the addition of jumps in returns and jumps in volatility significantly increases the curvature of the IV curves. For short maturity options, the difference between the SV, SVJ, and SVCJ IV curves for far in-the-money (ITM) or out-of-the money (OTM) options is quite large. This should not be a surprise, as jumps in returns and volatility increase the conditional nonnormalities of the underlying distribution. Our results are different from those in Andersen, Benzoni, and Lund (2002) and Pan (2002) who require significant jump risk premia to generate differences between the IV curves generated by the SV and SVJ model.

Note the substantial increase in IV for ITM options for the SVCJ model. This hook or "tipping at the end" effect (Duffie, Pan, and Singleton (2000) and Pan (2002)) is not present in the SVJ model. It appears to be an important feature of market implied volatilities, although the SVCJ model, without volatility jump risk premia, does not generate a sharp enough hook to match observed IV curves. Third, the SV model generates very flat IV curves, as it does not generate any substantive conditional departures from normality. Last, we find a significant flattening out effect as time to maturity increases for all of the models. This occurs because as maturity increases, the fat-tails and asymmetries in the conditional distribution are driven to a larger extent by diffusive volatility, rather than jumps.

[15] We thank the referee for suggesting this.

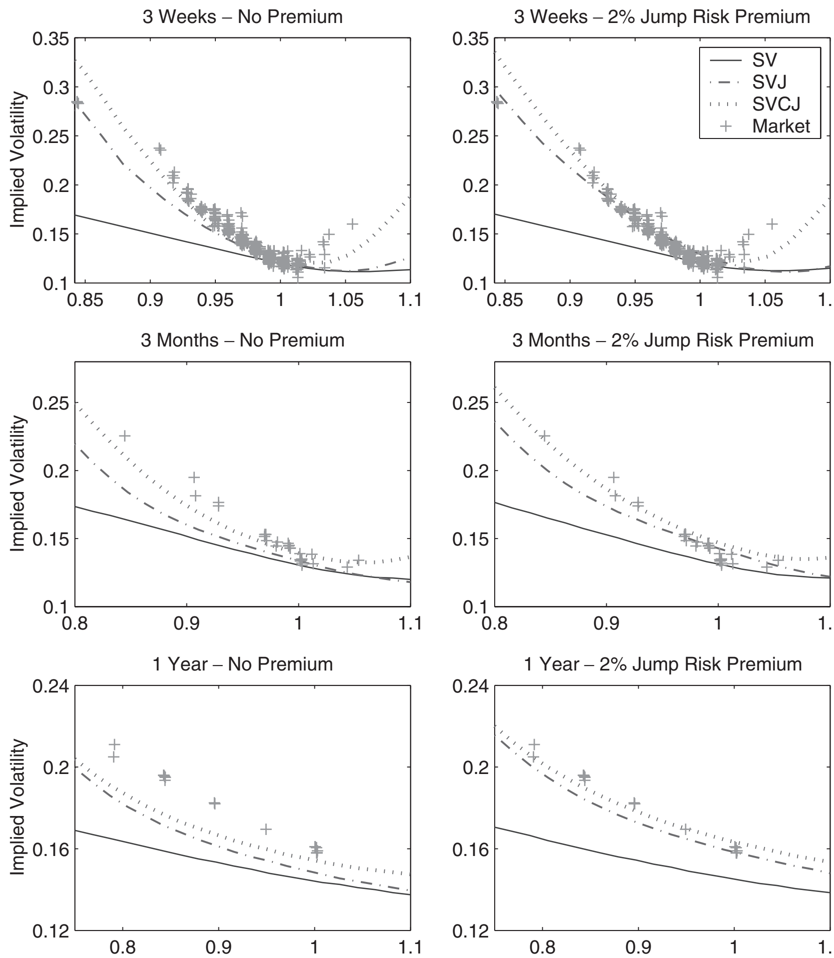

Figure 7. Implied volatility curves for the SV, SVJ, and SVCJ models and market implied volatility curves for a randomly selected day in our sample, October 31, 1994.

To see the impact of adding a mean jump size risk premium, the right hand panel of Figure 7 displays the IV curves with a modest two percent risk premium in the mean jump size in returns. The premium increases the OTM slope of the IV curves, but the effect is small. Casual observation indicates that the risk premium improves the models' ability to fit option prices, especially for longer maturities. The fact that this small risk premium brings model implied prices close to market prices is in contrast to Pan (2002), who estimates the risk premium to be 18 percent.

B. Estimation Risk and Option Prices

Our MCMC approach, through the posterior distribution, quantifies estimation risk: the uncertainty inherent in estimating parameters and spot volatility.

This section examines how this uncertainty impacts option prices. If Θ and V_t were known, the price of a call option struck at K, with τ days to maturity, conditional on the current spot index level, S_t, is the usual option pricing function, $C(V_t, \Theta, S_t, K, \tau) = E_t^{\mathbb{Q}}[e^{-r(\tau)} \max(S_{t+\tau} - K, 0)]$, where $\mathbb{Q}$ is the pricing measure.

This ignores that Θ and V_t are unknown and must be estimated. In our setting, $p(\Theta, V_t|Y)$, which our MCMC algorithm computes, quantifies the uncertainty in estimating V_t and Θ. From the econometrician's perspective, with uncertain spot volatility and parameters, the price of a call option is given by

$$C_t(S_t, K, \tau) = \int C(V_t, \Theta, S_t, K, \tau) p(\Theta, V_t|Y) d\Theta dV_t,$$

which integrates out all posterior uncertainty. This assumes that the agents pricing the options know the true volatility and parameters. We do not address the difficult issue of evaluating the impact of parameter uncertainty on equilibrium prices (see, e.g., Anderson, Hansen, and Sargent (2000) for one approach to this issue).

To quantify the impact of estimation risk, we compute the posterior distribution of option prices when both Θ and V_t are uncertain (Case 1), when only V_t is unknown (Case 2), and when only Θ is uncertain (Case 3). We do this by integrating out the uncertainty in Θ and/or V_t, as summarized by the posterior distribution. In the case where Θ or V_t are uncertain, we condition on the posterior mean for the other quantity. Figure 8 displays the posterior mean and a (10 percent, 90 percent) posterior coverage interval for each case, for three maturities for the SVCJ model, which is representative of the other jump models.

A comparison of the graphs in the upper and lower left panels indicates that nearly all of the ATM uncertainty is generated by uncertainty in V_t. The opposite effect occurs for OTM options. For this case, the middle left panel indicates that parameter uncertainty dominates. Why does this occur? Note that while the posterior means of μ_y and σ_y are -1.7 and 2.9, a one standard deviation symmetric coverage interval is (0.2 percent, -3.3 percent) for μ_y and (2.3 percent, 3.5 percent) for σ_y. Since the tails in the conditional distribution over short intervals load heavily on these parameters, any uncertainty over these parameters has a large impact on OTM options. A similar argument holds for λ_y and μ_v.

At medium time horizons, V_t and Θ uncertainty is still slightly greater for ATM and OTM options, but the effect is smaller. At long horizons, a different effect occurs: Parameter uncertainty dominates and spot volatility uncertainty has a minimal effect. The intuition for this is clear. The conditional distribution of the index value over a year is largely determined by the long-run behavior of the model, which in turn is determined by the parameters driving volatility: θ, κ, and σ_v.

V. Conclusions

This paper analyzed models with jumps in returns and in volatility. For both the S&P 500 and Nasdaq 100 index, results indicate that both of these jump com-

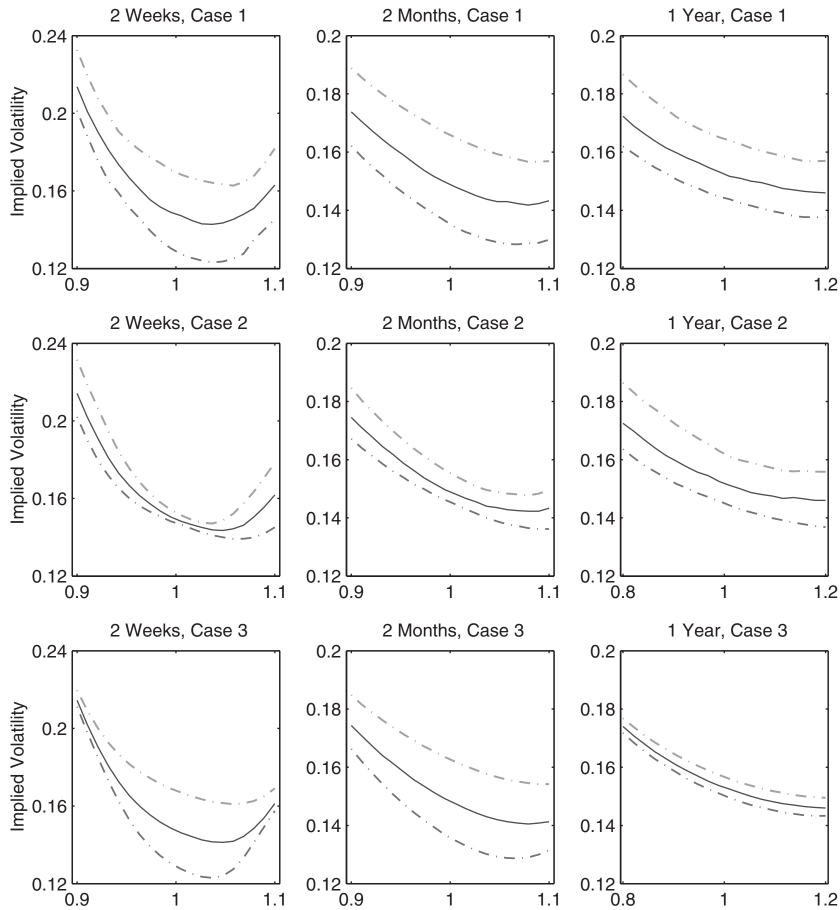

Figure 8. The impact of estimation risk. The top panels integrate out both volatility and parameter uncertainty. The middle panels integrate out only parameter uncertainty (conditional on estimated volatility). The bottom panels integrate out only volatility uncertainty (conditional on posterior mean of parameters).

ponents are important, and models without jumps in volatility are misspecified. Models with only diffusive stochastic volatility and jumps in returns are misspecified, because they do not have a component driving the conditional volatility of returns, which is rapidly moving.

The information in the time series of returns indicates that jumps in returns and jumps in volatility have a strong impact on option prices. Compared to a model with only jumps in returns, models with jumps in volatility result in a significant increase in IV for deep in-the-money or out-of-the-money options. The fact that jumps in volatility have such a large impact on IV is somewhat surprising, as the original motivation for including jumps in volatility was to improve on the dynamics of spot volatility and not necessarily to generate more realistic IV curves, although, of course, these two goals are clearly related.

Parameter and volatility estimation risk also has an important impact on option prices. While parameter uncertainty might result in some option price uncertainty for short maturity at-the-money options, the impact can be as large as ± 2 implied volatilities for out-of-the-money options. Volatility uncertainty for at-the-money short maturity options leads to prices that differ by ± 2 implied volatilities, far greater than the bid–ask spread. Although the results on parameter uncertainty in option pricing are new, the importance of volatility uncertainty is not new. Merton (1980) recognized this as an important feature and states that the "most important direction is to develop accurate variance estimation models which take into account of the errors in variance estimates" (p. 355). Our MCMC algorithm provides such a method.

Our results indicate that parameter and spot volatility estimation risk is substantial. To obtain more accurate estimates, option price data may be extremely useful. Chernov and Ghysels (2000), Eraker (2002), and Pan (2002) all use spot and option price data to estimate various models. However, while the addition of option prices will aid in the estimation of the spot volatility, it is unclear if it will significantly reduce the parameter uncertainty. This occurs because the risk premia embedded in option prices introduce additional parameters, which are typically difficult to estimate. Unless the risk premia are restricted as in Pan (2002), it is not clear if the aggregate parameter uncertainty will increase or decrease with the addition of option price data.

Appendix A. Posterior Distributions for Jumps and Volatility

The conditional posteriors for the jump sizes and jump times are new and are derived as follows. We use the SVCJ model, because it is the most complex, in terms of distributional structure. Recall the prior structure for the jump sizes in SVCJ, $\xi_t^v \sim \exp(\mu_V)$ and $\xi_t^y | \xi_t^v \sim N(\mu_y + \rho_J \xi_t^v, \sigma_y^2)$. This specification allows us to use the Gibbs sampler to exploit the conditional independence of the jumps in volatility, as it is easy to draw from the conditional posterior. The conditional posterior for the jump sizes to volatility is $p(\xi_{(t+1)\Delta}^v | J_{(t+1)\Delta} = 1, \Theta, \xi^y, V, Y)$ which, by Bayes rule is proportional to

$$p(Y_{(t+1)\Delta}, V_{(t+1)\Delta} | V_{t\Delta}, J_{(t+1)\Delta} = 1, \Theta, \xi_{(t+1)\Delta}^y, \xi_{(t+1)\Delta}^v) \cdot p(\xi_{(t+1)\Delta}^v | J_{(t+1)\Delta} = 1, \Theta, \xi_{(t+1)\Delta}^y)$$

where the first term, $p(Y_{(t+1)\Delta}, V_{(t+1)\Delta} | V_{t\Delta}, J_{(t+1)\Delta} = 1, \Theta, \xi_{(t+1)\Delta}^y, \xi_{(t+1)\Delta}^v)$, is a bivariate normal distribution and the second term is, by Bayes rule, proportional to the product of $p(\xi_{(t+1)\Delta}^y | \xi_{(t+1)\Delta}^v, J_{(t+1)\Delta} = 1, \Theta)$ and $p(\xi_{(t+1)\Delta}^v | J_{(t+1)\Delta} = 1, \Theta)$. This is a product of a normal and an exponential. To find the full conditional posterior, completing the square for all three terms as a function of $\xi_{(t+1)\Delta}^v$ leads to a truncated normal

$$p(\xi_{(t+1)\Delta}^v | J_{(t+1)\Delta} = 1, \Theta, \xi^y, V, Y) \propto 1_{[\xi_{(t+1)\Delta}^v > 0]} N(\alpha_{(t+1)\Delta}^v, \omega_{(t+1)\Delta}^v),$$

where $\alpha^v_{(t+1)\Delta}$ and $\omega^v_{(t+1)\Delta}$ are straightforward to compute. The conditional posterior given a jump is therefore a truncated normal distribution. When $J_{(t+1)\Delta} = 0$, the conditional posterior is $\xi^v_t \sim \exp(\mu_v)$, as the data provides no information about the jump size.

The posterior for the jumps in returns is similarly derived. Bayes Rule implies that

$$p(\xi^y_{(t+1)\Delta}|\xi^v_{(t+1)\Delta}, J_{(t+1)\Delta} = 1, \Theta, V_{t\Delta}, Y_{(t+1)\Delta}) \propto$$
$$p(Y_{(t+1)\Delta}|\xi^y_{(t+1)\Delta}, \xi^v_{(t+1)\Delta}, J_{(t+1)\Delta} = 1, \Theta, V_{t\Delta})p(\xi^y_{(t+1)\Delta}|\xi^v_{(t+1)\Delta}, \Theta).$$

Since both of the densities are Gaussian, we have that

$$\xi^y_{(t+1)\Delta}|\xi^v_{(t+1)\Delta}, J_{(t+1)\Delta} = 1, \Theta, V_{t\Delta}, Y_{t+1)\Delta} \propto N(\alpha^y_{(t+1)\Delta}, \omega^y_{(t+1)\Delta}) \tag{7}$$

where $\alpha^y_{(t+1)\Delta}$ and $\omega^y_{(t+1)\Delta}$ are easy to compute.

For the jump times, which are assumed to arrive contemporaneously, the posterior combines information from the returns and from the volatility. As $J_{(t+1)\Delta}$ can take only two values, its posterior is Bernoulli, $Ber(\lambda_{(t+1)\Delta})$. To compute the Bernoulli probability, we use the conditional independence of increments to volatility and returns to get that

$$p(J_{(t+1)\Delta} = 1|V_{(t+1)\Delta}, V_{t\Delta}, Y_{(t+1)\Delta}, \xi^y_{(t+1)\Delta}, \xi^v_{(t+1)\Delta}, \Theta)$$
$$\propto \lambda\Delta \cdot p(Y_{(t+1)\Delta}, V_{(t+1)\Delta}|V_{t\Delta}, J_{(t+1)\Delta} = 1, \xi^y_{(t+1)\Delta}, \xi^v_{(t+1)\Delta}, \Theta),$$

which, again, is easy to calculate, since $p(Y_{(t+1)\Delta}, V_{(t+1)\Delta}|V_{t\Delta}, J_{(t+1)\Delta} = 1, \xi^y_{(t+1)\Delta}, \xi^v_{(t+1)\Delta}, \Theta)$ is a bivariate Gaussian density. Computing the conditional posterior for $J_{(t+1)\Delta} = 0$ proceeds similarly, which gives the Bernoulli probability.

The conditional posterior for volatility, $p(V_{t\Delta} \mid Y_{t\Delta}, V_{(t+1)\Delta}, V_{(t-1)\Delta}, \xi^y_{t\Delta}, \xi^v_{t\Delta}, J_{t\Delta}, \Theta)$, is not a known distribution. To sample from it, we use a random-walk Metropolis algorithm (see Johannes and Polson (2002)). The conditional posteriors for the parameters are standard. Given our conjugate priors, the conditional posteriors for $\{\alpha, \kappa, \theta, \sigma_v, \lambda_y, \lambda_v, \mu_y, \sigma^2_y, \mu_v\}$ are all standard distributions, and we omit the derivations as they can be found in standard texts. For ρ we use an independence Metropolis algorithm with a proposal density centered at the sample correlation between the Brownian increments.

We now discuss our choices of prior distributions and parameters. Wherever possible, we choose standard conjugate priors, which allows us to directly draw from the conditional posteriors. Our prior distributions for the parameters are: $\alpha \sim N(1, 25)$, $\kappa\theta \sim \mathrm{N}(0,1)$, $\kappa \sim N(0,1)$, $\sigma^2_v \sim IG(2.5, 0.1)$, $\rho \sim \mathcal{U}(-1, 1)$, $\lambda_y = \lambda_v \sim \beta(2,40)$, $\mu_y \sim N(0, 100)$, $\sigma^2_y \sim \mathcal{IG}(5.0, 20)$, $\mu_v \sim \mathcal{G}(20, 10)$, and $\rho_J \sim N(0, 4)$ where $\mathcal{G}$ refers to a Gamma distribution, $\mathcal{IG}$ refers to the Inverse Gamma distribution, and $\mathcal{U}$ a standard uniform distribution. All of the prior distributions are uninformative with the exception of σ_y and λ_y. The simulation results given below demonstrate that the information imposed by these priors is minor relative to the information in the likelihood function. This is further demonstrated by the fact that we use the same priors for the Nasdaq and S&P data, which deliver drastically different parameter estimates.

Appendix B. Bayes Factor Calculations

Suppose that you wish to compare SV and SVJ. Let $\Omega = \{\Theta, \{V_{t\Delta}, J_{t\Delta}, \xi_{t\Delta}\}_{t=1}^{T}\}$ be a matrix of latent variables and parameters, which gives the marginal likelihoods as

$$p(Y|SV) = \int p(Y|\Omega, SV)p(\Omega|SV)d\Omega$$

$$p(Y|SVJ) = \int p(Y|\Omega, SVJ)p(\Omega|SVJ)d\Omega.$$

Now, as the SVJ model embeds the simple SV model when the entire vector of jump times is zero ($J = 0$), we have the identity that $p(Y|\Omega, SV) = p(Y|\Omega, J = 0, SVJ)$. Moreover, if we assume in the priors that $p(\Omega|SV) = p(\Omega|J = 0, SVJ)$ (parameters in common have the same priors), then

$$p(Y|SV) = \int p(Y|\Omega, J = 0, SVJ)p(\Omega|J = 0, SVJ)d\Omega = p(Y|J = 0, SVJ).$$

Bayes rule also implies that

$$P(J = 0|Y, SVJ) = \frac{P(Y|J = 0, SVJ)}{P(Y|SVJ)}P(J = 0|SVJ), \tag{8}$$

and we have the important identity that we use to develop an MCMC estimator,

$$odds(sv : svj) = \frac{P(J = 0|Y, SVJ)}{P(J = 0|SVJ)}. \tag{9}$$

This is simply a ratio of prior ordinate to posterior ordinate, and the key observation is that this can be directly computed from the MCMC output, whereas the marginal likelihoods above are not available (see Jacquier and Polson (1999)).

In the SVJ model, we can compute

$$\begin{aligned} p(J = 0|SVJ) &= \int_0^1 p(J = 0|\lambda, SVJ)p(\lambda|SVJ)d\lambda \\ &= \int_0^1 (1-\lambda)^T \frac{\lambda^{\alpha_0-1}(1-\lambda)^{\beta_0-1}}{\mathcal{B}(\alpha_0, \beta_0)}d\lambda = \frac{\mathcal{B}(\alpha_0, T+\beta_0)}{\mathcal{B}(\alpha_0, \beta_0)}. \end{aligned}$$

For the other portion, a similar argument gives

$$p(J = 0|Y, SVJ) = \int_0^1 p(J = 0|\lambda, Y, SVJ)p(\lambda|Y, SVJ)d\lambda, \tag{10}$$

which a straightforward computation yields

$$p(J = 0|Y, SVJ) = \int_J \left[\frac{B(\alpha_0 + \sum_{t=0}^{T} J_t, \beta_0 + 2T - \sum_{t=0}^{T} J_t)}{B(\alpha_0 + \sum_{t=1}^{T} J_t, \beta_0 + T - \sum_{t=1}^{T} J_t)}\right]p(J|Y)dJ$$

since

$$p(\lambda|Y,\, SVJ) = \int_J p(\lambda|J,\, Y,\, SVJ)p(J|Y)dJ = \int_J p(\lambda|J,\, SVJ)p(J|Y)dJ$$
$$= \int_J \frac{\lambda^{\alpha_0+\sum_{t=0}^{T} J_t - 1}(1-\lambda)^{\beta_0+T-\sum_{t=1}^{T} J_t - 1}}{\mathcal{B}(\alpha_0 + \sum_{t=1}^{T} J_t,\, \beta_0 + T - \sum_{t=1}^{T} J_t)} p(J|Y)dJ.$$

Computing this integral with the Monte Carlo samples gives

$$p(J=0|Y,\, SVJ) = \frac{1}{G}\sum_{g=1}^{G} \frac{\mathcal{B}(\alpha_0 + \sum_{t=0}^{T} J_t^g, \beta_0 + 2T - \sum_{t=0}^{T} J_t^g)}{\mathcal{B}(\alpha_0 + \sum_{t=1}^{T} J_t^g, \beta_0 + T - \sum_{t=1}^{T} J_t^g)}, \tag{11}$$

and we therefore have an MCMC estimator:

$$odds(sv : svj) = \frac{\mathcal{B}(\alpha_0, \beta_0)}{\mathcal{B}(\alpha_0,\, T+\beta_0)} \frac{1}{G}\sum_{g=1}^{G} \frac{\mathcal{B}(\alpha_0 + \sum_{t=0}^{T} J_t^g,\, \beta_0 + 2T - \sum_{t=0}^{T} J_t^g)}{\mathcal{B}(\alpha_0 + \sum_{t=1}^{T} J_t^g,\, \beta_0 + T - \sum_{t=1}^{T} J_t^g)}. \tag{12}$$

The computation of the other Odds Ratios is similar. In following the literature, we compute log odds ratios.

Appendix C. Simulation Experiments

Our simulations used 100 artificial data sets consisting of 4,000 data points. The data were generated using the Euler discretization of the continuous-time model with $\Delta = 1/20$. Hence, our estimation method is measured against the artificial data generated by the true continuous time process. We used a posterior sample size of 50,000 for each of the MCMC runs.

The number of data points and the true parameters are suggestive of what one can expect from daily equity price data, although we tried to stack the deck against our methodology by making the jumps extremely unlikely. An intensity of 0.006 indicates that there are only 1.5 jumps per year. With 4,000 daily observations, this implies that we essentially only have 26 jumps per sample.

Tables VII and VIII report simulations for the SVJ and SVCJ models. The tables report the means and the root mean squared error (RMSE) and the results indicate that the algorithm delivers extremely accurate estimates for most of the parameters, with the exception of κ, which appears to be slightly biased. This is not necessarily surprising: This parameter governs the speed of mean reversion of stochastic volatility and is difficult to estimate. Our results are similar to those found reported in Jacquier et al. (1994) for the stochastic volatility parameters.

To see how much of the error in estimating jump components was due to the extremely rare jumps in returns, Table VIII reports the results when we increased the arrival intensity to 0.015 (3.6 jumps per year). As expected, increasing the arrival rate of jumps improves the ability of the algorithm to estimate the parameters of the jump distribution, although the increased noise makes estimation of the stochastic volatility components slightly less accurate. The results for the

Table VII
Simulation Results for the SVJ Model

This table provides a summary of the simulation results for the SVJ model. One hundred sample paths were simulated off the true parameters and we report the average across the simulations, as well as the root mean squared error (RMSE). We consider high and low jump intensity cases.

	α	θ	κ	σ_v	μ_y	σ_y	ρ	λ
				Low Jump Intensity				
True	0.0500	0.8000	0.015	0.1000	– 3.0000	3.5000	– 0.4000	0.0060
Mean	0.0528	0.8212	0.0209	0.1117	– 3.7220	3.0424	– 0.4228	0.0068
RMSE	0.0129	0.1008	0.0060	0.0089	1.6786	0.5773	0.0612	0.0023
				High Jump Intensity				
True	0.0500	0.8000	0.015	0.1000	– 3.0000	3.5000	– 0.4000	0.0150
Mean	0.0537	0.8739	0.0214	0.1169	– 3.3374	3.5955	– 0.4129	0.0154
RMSE	0.0135	0.3849	0.0060	0.0245	0.8584	0.4732	0.0660	0.0031

Table VIII
Simulation Results for the SVCJ Model

This table provides a summary of the simulation results for the SVCJ model. One hundred sample paths were simulated off the true parameters and we report the average across the simulations as well as the root mean squared error (RMSE).

	α	θ	κ	σ_v	ρ
True	0.0500	0.5000	0.0300	0.1000	– 0.5000
Mean	0.0507	0.4755	0.0330	0.0817	– 0.4921
RMSE	0.0114	0.0601	0.0062	0.0076	0.0711
	μ_y	σ_y	ρ_J	μ_v	λ
True	– 2.0000	3.5000	– 0.4000	1.0000	0.0080
Mean	– 2.3140	2.9116	– 0.0599	1.4272	0.0093
RMSE	1.1087	0.4259	0.4369	0.3622	0.0027

SVCJ model are similar, although we are not able to reliably estimate the ρ_J parameter. The variable ρ_J is essentially a regression parameter between the jumps in returns and volatility, $\xi^y = \mu_y + \rho_J \xi^v + \varepsilon$ where $\varepsilon \sim \mathcal{N}(0, \sigma_y^2)$. Since jumps are latent and jumps are rare events, this is not surprising.

References

Alizadeh, Sassan, Michael Brandt, and Francis X. Diebold, 2002, Range-based estimation of stochastic volatility models, *Journal of Finance* 57, 1047–1091.

Andersen, Torben, Luca Benzoni, and Jesper Lund, 2002, Towards an empirical foundation for continuous-time equity return models, *Journal of Finance* 57, 1239–1284.

Andersen, Torben, Tim Bollerslev, Francis Diebold, and Heiko Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* 61, 43–76.

Andersen, Torben, Hyung-Jin Chung, and Bent Sorensen, 1999, Efficient method of moments estimation of a stochastic volatility model: A Monte Carlo study, *Journal of Econometrics* 91, 61–87.

Anderson, Evan, Lars Hansen, and Thomas Sargent, 2000, Robustness, detection and the price of risk, Working paper, University of Chicago.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003–2049.

Bates, David, 1996, Jumps 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.

Brandt, Michael, and Pedro Santa-Clara, 2002, Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets, *Journal of Financial Economics* 63, 161–210.

Chacko, George, and Luis Viceira, 2001, Spectral GMM estimation of continuous-time processes, Working paper, Harvard 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.

Chernov, Mikhail, Eric Ghysels, A. Ronald Gallant, and George Tauchen, 2002, Alternative models for stock price dynamics, Working paper, Columbia University.

Das, Sanjiv, and Randarajan Sundaram, 1999, Of smiles and smirks: A term structure perspective, *Journal of Financial and Quantitative Analysis* 34, 211–240.

Duffie, Darrell, and Jun Pan, 1997, An overview of value at risk, *Journal of Derivatives*, Spring, 7–49.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Durham, Garland, and A. Ronald Gallant, 2001, Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes, Working paper, University of North Carolina.

Eraker, Bjørn, 2001, MCMC analysis of diffusion models with applications to finance, *Journal of Business and Economic Statistics* 19–2, 177–191.

Eraker, Bjørn, 2002, Do equity prices and volatility jump? Reconciling evidence from spot and option prices, Working paper, Duke University.

Heston, Steven, 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.

Honore, Peter, 1998, Pitfalls in estimating jump-diffusion models, Working paper, University of Aarhus.

Jacquier, Eric, and Nicholas Polson, 1999, Odds Ratios in MCMC: An application to stochastic volatility, Working paper, University of Chicago.

Jacquier, Eric, Nicholas Polson, and Peter Rossi, 1994, Bayesian analysis of stochastic volatility models, *Journal of Economics and Business Statistics* 12, 371–389.

Jacquier, Eric, Nicholas Polson, and Peter Rossi, 2002, Bayesian analysis of stochastic volatility models with fat-tailed and correlated errors, *Journal of Econometrics*, forthcoming.

Johannes, Michael, Rohit Kumar, and Nicholas Polson, 1999, State dependent jump models: How do U.S. equity indices jump?, Working paper, Columbia University.

Johannes, Michael, and Nicholas Polson, 2002, MCMC methods for continuous-time financial econometrics, in Yacine Aït-Sahalia and Lars Hansen, eds.: *Handbook of Financial Econometrics* (Elsevier, New York).

Jones, Christopher, 2002, The dynamics of stochastic volatility: Evidence from underlying and options markets, Working paper, University of Rochester.

Kass, Robert, and Adrian Raftery, 1995, Bayes factors, *Journal of the American Statistical Association* 90, 773–795.

Liu, Jun, Francis Longstaff, and Jun Pan, 2003, Dynamic asset allocation with event risk, *Journal of Finance* 58, 231–259.

Merton, Robert, 1976, Option pricing when the underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

Merton, Robert, 1980, Estimating the expected return on the market, *Journal of Financial Economics* 8, 323–363.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Piazzesi, Monika, 2001, Macroeconomic jump effects and the yield curve, working paper, UCLA.

Robert, Christian, and George Casella, 1999, *Monte Carlo Statistical Methods* (Springer-Verlag).

Schwert, William, 1990, Stock market volatility, *Review of Financial Studies* 3, 77–102.

Smith, Adrian, and David Spiegelhalter, 1982, Bayes factors and choice criteria for linear models, *Journal of the Royal Statistical Association, Series B* 42, 213–220.