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

Bjørn Eraker*

August 5, 2001

Abstract

This paper studies the empirical performance of jump-diffusion models that allow for stochastic volatility and correlated jumps affecting both prices and volatility. The results show that the models in question provide reasonable fit to both option prices and returns data in the in-sample estimation period. This contrasts previous findings where stochastic volatility paths are found to be too smooth relative to the option implied dynamics. While the models perform well during the high volatility estimation period, they tend to overprice long dated contracts out-of-sample. This evidence points towards a too simplistic specification of the mean dynamics of volatility.

Keywords: Market crashes, jump-diffusion, Stochastic volatility, jump in volatility, affine models, Markov Chain Monte Carlo
JEL Class: C11, C15, G12

*I thank Fredrico Bandi, Lars Hansen, Nick Polson, Pietro Veronesi, and Mike Johannes, 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, and Northwestern University for helpful comments and discussions.

1 Introduction

The statistical properties of stock returns have for long been of interest to financial decision makers and academics alike. In particular, the great stock market crashes of the twentieth 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: time varying volatility models that allow for market extremes to be the outcome of normally distributed shocks that have a randomly changing variance, and 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 implausible high volatility level both prior to, and after the crash. Jump models on the other hand, can easily explain the crash of '87 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 '87 crash for example, there were large movements both prior to, and following the crash. With respect to option prices, this modelling 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 '87 crash. In addressing these issues, researchers have proposed models that incorporate both stochastic volatility and jumps components.

In particular, recent work by Bates (2000) and Pan (2000) examine 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. In conclusion, these papers point toward models that include jumps to the volatility process.

In response to these findings, Eraker, Johannes, and Polson (2000) (henceforth EJP) use returns data to investigate the performance of models with jumps in volatil-

[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.

ity as well as prices using the class of jump-in-volatility models proposed by Duffie, Pan, and Singleton (2000). 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 '87, 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 (2000). 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 contrasts studies that focus exclusively on either source of data. Secondly, the one-to-one correspondence of options to the conditional returns distribution [2] allows parameters governing the shape of this distribution to potentially be very accurately estimated from option prices. For example, EJP report fairly wide posterior standard deviations for parameters that determine 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

[2] A number of papers obtain non-parametric estimates of return distributions using options. See Rubenstein (1994), Jackwerth and Rubenstein (1996), Dumas, Fleming, and Whaley (1998) and Aït-Sahalia and Lo (1998) among others.

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 (2000) 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 Baksi, 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 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 two 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 4 (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 '87, 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 three discusses the econometric design and outlines a strategy for obtaining posterior samples by MCMC. Section four presents the data, while section five contains the empirical results. Section six summarizes the findings and suggests directions for further investigation.

2 Models and Pricing

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

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

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

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 and a component given by a Poisson counting process N_t. The specification of these jump components determines the models to be considered. We discuss the various possible specifications and their implications for option prices next.

2.1 SV model

This model was initially proposed for option pricing purposes in Heston (1993) and obtains as a special case with no jumps ($dJ_t^S = dJ_t^V = 0$). In the basic SV model, 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. σ is commonly referred to as the "volatility of volatility". Higher values of σ_V implies 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 postive (Das and Sundaram (1999)) and

[3] Attempts to fit a more elaborate drift under the objective measure was undertaken in EJP and showed to be unimportant.

increasing in both σ_V and $\|\rho\|$[4]. 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 postive/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 vise 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 when 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 moderately long time to maturity have both steeper and more concave Black-Scholes implied volatility curves across different strikes.

2.2 SVJ model

The 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

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

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 1-3) for the SV model. Noticeably however, at short time horizons, the transition density is non-gaussian.

[4]The fact that conditional kurtosis depends on σ_V corrects a typo in proposition 4 (2) of Das and Sundaram (1999).

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.

2.3 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

$$\begin{aligned} Z_t^V &\sim \exp(\mu_V) \\ Z_t^S \mid Z_t^V &\sim N(\mu_y + \rho_J Z_t^V, \sigma_S^2) \end{aligned}$$

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 in 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 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 (2000) 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.

2.4 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 an "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, Pan, and Singleton (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

$$\frac{dS_t}{S_t} = (r - \lambda\mu_y^Q)dt + \sqrt{V_t}dW(Q)_t^S + dJ(Q)_t^S \tag{3}$$

$$dV_t = (\kappa(\theta - V_t) + \eta_V V_t)dt + \sigma_V\sqrt{V_t}dW(Q)_t^V + dJ(Q)_t^V \tag{4}$$

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,

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

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 $\mu = r + \frac{\eta_J}{\sigma}$, μ is only identifiable from historical time series data. Thus, the usual problems with estimating expected rates of returns apply. A similar problem, only worse, obtains for the mean price jumps under the objective measure, μ_y. In essence, this parameter is only identified by 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.

The advantage of the model specifications outlined above is that option prices are available in semi closed 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. We therefore derive a new method for carrying out this numerical integration in an appendix. Our method takes into account the particular functional form of the integrating function to simplify and ultimately speed up the computational requirements. The method is significantly less computationally intensive than methods proposed in the past.

3 Econometric methodology

While a number of methods have been proposed for the estimation of diffusion processes, [5], the latent nature of volatility as well as the jump components, complicates estimation. In recent work Singleton (2001), Pan (2000), 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 (1999) 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 with stochastic volatility models. Extensions into multivariate models can be found in Jacquier, Polson, and Rossi (1999). Eraker (2000) propose a general approach to estimating diffusion models with arbitrary drift and diffusion functions and possibly latent, unobserved state variables

[5] See Aït-Sahalia (1998) (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.

(e.g., stochastic volatility). Finally, Johannes, Kumar, and Polson (1999) use MCMC to study jump models. Their approach was generalized in Eraker, Johannes, and Polson (2000) to the jump diffusion models considered in this paper, but estimated from returns data only.

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 :

$$P_t = F(X_t, \chi, \Theta)$$

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:

- 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.

- Observations are available at transaction times. Typically, the use of transactions data will result in unequally spaced observations at times $t_i, i = 0, ..., N$. 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 $\epsilon_{t_i,j}$

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

We wish to allow for a serial dependence in the pricing errors, ϵ_{t_i}, of each asset j. This reflects the prior belief that if an asset is miss-priced at time t, it is also likely to be miss-priced 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

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

and that $\epsilon_{t,i}$ is independent of $\epsilon_{t,j}$,for $j \neq i$. Alternatively, whenever transactions data are recorded at random transaction times t_i (case 2), the corresponding distributional assumption with respect to the pricing errors is

$$\epsilon_{t_i,j} \sim N\left(\epsilon_{t_{i-1},j} e^{\rho_j \nabla t_i}, s_j^2 \frac{1}{2\rho_j}(1 - e^{-2\rho_j \nabla t_i})\right) \tag{5}$$

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.

3.1 Joint posterior density

The specification of the error distribution, together with a specification of the full prior density 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, $Y = \{Y_{t_i,j} : i = 1, .., N, j = 1, ..., n\}$ as

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

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

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

$$\begin{aligned} p(Y, X, Z, J, \Theta) &\propto p(Y \mid X, \Theta) p(X \mid \theta) p(Z \mid J, \Theta) p(J) p(\Theta) \\ &= \prod_{i=1}^{N} p(Y_{t_i} \mid X_{t_i}, \epsilon_{t_i}, \Theta) p(X_{t_i} \mid X_{t_{i-1}}, Z_{t_i}, J_{t_i} \Theta) p(Z_{t_i} \mid J_{t_i}, \Theta) p(J_{t_i,\Theta}) p(\Theta) \end{aligned} \tag{6}$$

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

The density, $p(X_{t_i} \mid X_{t_{i-1}}, Z_{t_i}, J_{t_i}, \Theta)$, is generally unknown. There are two ways of estimating this density: 1) the fourier transform of the transition density $p(X_{t_i} \mid X_{t_{i-1}}, \theta)$ 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;

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

This is reasonable as long as the observation intervals, t_i are not too far apart. In a simulation study, Eraker, Johannes and Polson (2000) show that the discretization

bias arising from the use of daily intervals is negligable. 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 (2000)).

3.2 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 $p(V, J, Z, \Theta \mid Y)$. This accomplishes the essential estimation problem: the estimation of the marginal posterior density of the parameters, $p(\Theta \mid Y)$. Notice that we can accomplish this task by designing a sampling scheme that produces random draws, $\Theta^{(1)}, \Theta^{(2)}, ..., \Theta^{(G)}$ whose density is $p(V, J, Z, \Theta \mid Y)$ since, by a simple application of Bayes theorem, $p(\Theta \mid Y) \propto p(V, J, Z, \Theta \mid Y)$. 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 non-standard, 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 non-standard 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.

3.3 Gibbs sampling

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

For $g = 1, ...G$ simulate

For $i = 1, .., N$

$V_{t_i}^{(g)}$from $p(V_{t_i} \mid V_{t_{\setminus i}}^{(g-1)}, J^{(g-1)},\ Z^{(g-1)}, \Theta^{(g-1)}, Y)$

$J_{t_i}^{(g)}$from $p(J_{t_i} \mid V^{(g)}, J_{t_{\setminus i}}^{(g-1)},\ Z^{(g-1)}, \Theta^{(g-1)}, Y)$

$Z_{t_i}^{(g)}$from $p(Z_{t_i} \mid V^{(g)}, J^{(g)},\ Z_{t_{\setminus i}}^{(g-1)}, \Theta^{(g-1)}, Y)$

and

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

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 $p(V, J,\ Z, \Theta \mid Y)$ 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 neccessary to use the entire expression in (6). We now go through these simplifications in turn. The conditional posterior for spot volatility at time t_i can be written,

$$
\begin{aligned}
&p(V_{t_i} \mid V_{t_{\setminus i}}^{(g-1)}, J^{(g-1)}, Z^{(g-1)}, \Theta^{(g-1)}, Y) \\
&\qquad \propto p(Y_{t_i} \mid X_{t_i}, \epsilon_{t_i}, \Theta) p(X_{t_i} \mid X_{t_{i-1}}, Z_{t_i}, J_{t_i}, \Theta) p(X_{t_{i+1}} \mid X_{t_i}, Z_{t_{i+1}}, J_{t_{i+1}}, \Theta)
\end{aligned}
$$

The simplification here stems from the Markov property of the process: only the terms in (6) 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 non-standard 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 $\Pr(J_{t_i} = 1 \mid X, Z, Y, \Theta)$, which again simplifies to

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

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, Johannes and Polson (2000).

The conditional distributions for the jump sizes, $Z_{t_i} = (Z^S_{t_i}, Z^V_{t_i})$ 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, Johannes and Polson (2000) and the reader is referred to this paper for details.

Finally, we note that the conditional distribution for the parameter vector Θ is also non-standard. 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 and one 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.

4 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 Aït-Sahalia, Wang, and Yared (2000) and David and Veronesi (1999), among others. Aït-Sahalia et al. (2000) point out that the lack of timeliness between the closing prices of potentially iliquid 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 relationship to the underlying asset, but rather their theoretical relationship 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 have a large impact of these observations.

In order to facilitate estimation in a timely fashion, we need to select a subsample from the larger database. These options price observations are a random subsample of 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 as to construct a subsample for which there are at least one contract available each trading day. In total, this subsampling procedure produces 3270 call options contracts recorded over 1006 trading days, covering the period of 01/01/1987 to 12/31/1990. The remaining sampling period, 01/01/1991 to 01/03/1996, 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 in table 1.

Table 1: Descriptive Statistics for Options Data

	K/S	τ	price	IV
		period '87-'91		
mean	0.98	50.33	10.4833	20.2 %
std.dev	0.05	37.60	7.425	7.8 %
		period '91-'96		
mean	0.99	45.91	17.88	14%
std.dev	0.06	44.27	23.13	6%

Table 2: Descriptive Statistics for Returns Data

	All returns				Deleting '87 crash			
	mean	std.	skew.	kurt.	mean	std.	skew.	kurt.
Daily	0.03	1.8321	-5.06	87.50	0.05	1.15	-0.50	12.67
Weekly	0.14	11.439	-4.14	41.32	0.26	2.45	-0.70	3.99
Monthly	0.58	37.311	-1.42	7.92	1.08	4.52	0.17	4.23

5 Empirical Results

5.1 Parameter estimates

Table 3 reports the posterior means, posterior standard deviations, and the 1% and 99 % posterior percentiles of the parameters in the various models. The 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 (eg., Bates (2000), Pan (2000)).

There are several interesting features of the parameter estimates in table 3. 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 2). 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

Table 3: 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
a	0.060	0.065	0.094
	(0.034)	(0.033)	(0.030)
	[-0.021, 0.140]	[-0.013, 0.142]	[0.025, 0.162]
θ	1.933	1.652	1.353
	(0.048)	(0.053)	(0.067)
	[1.843, 2.064]	[1.523, 1.777]	[1.243, 1.560]
κ	0.019	0.019	0.023
	(0.007)	(0.006)	(0.007)
	[0.004, 0.035]	[0.005, 0.034]	[0.010, 0.038]
κ^Q	0.009	0.011	0.011
	(0.000)	(0.000)	(0.000)
	[0.009, 0.010]	[0.010, 0.011]	[0.010, 0.012]
η_V	0.010	0.009	0.013
	(0.007)	(0.006)	(0.007)
	[-0.005, 0.026]	[-0.006, 0.024]	[-0.001, 0.028]
ρ	-0.569	-0.586	-0.582
	(0.014)	(0.027)	(0.024)
	[-0.601, -0.535]	[-0.652, -0.526]	[-0.646, -0.528]
σ_V	0.220	0.203	0.163
	(0.007)	(0.007)	(0.007)
	[0.203, 0.240]	[0.187, 0.218]	[0.148, 0.181]
μ_y		-0.388	-6.062
		(3.456)	(2.274)
		[-8.560, 7.631]	[-11.566, -0.881]
μ_y^Q		-2.002	-7.508
		(1.867)	(0.932)
		[-6.079, 2.022]	[-9.725, -5.527]
η_J		1.613	1.446
		(3.789)	(2.474)
		[-7.320, 10.478]	[-4.516, 6.939]
ρ_J			-0.693
			(0.096)
			[-0.856, -0.449]
σ_y		6.634	3.630
		(1.081)	(1.106)
		[4.972, 9.508]	[1.403, 6.261]
μ_V			1.638
			(0.790)
			[0.833, 3.231]
λ		0.002	0.002
		(0.001)	(0.001)
		[0.001, 0.003]	[0.001, 0.004]

components are explaining a significant portion of the unconditional return variance.

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 1-volatility autocorrelation. 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 SVCJ model the estimate of this parameter is "significant", in the sense that its 5-percentile (but not the 1-percentile) is greater than zero. 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 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 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, Johannes and Polson (2000) 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 3 indicate that one can expect about 2-3 jumps in a stretch of 1000 trading days. According to SVJ estimates, a observed jumps will be in the range [-13.4,12.6] with

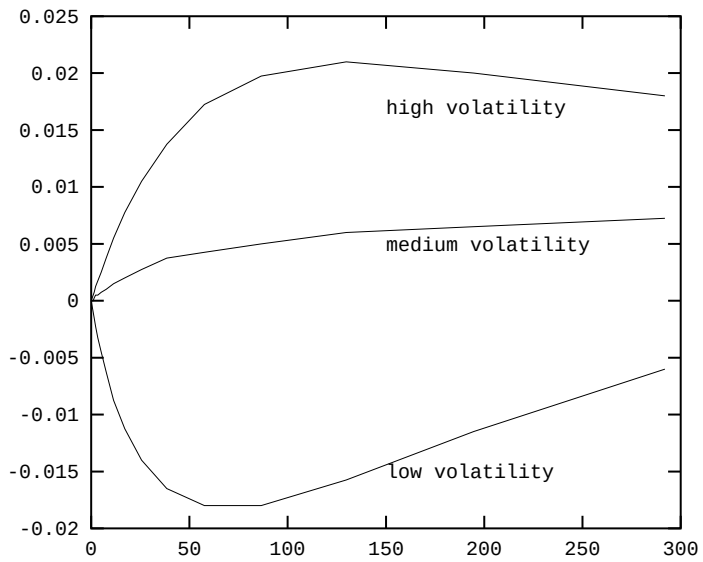

Figure 1: Volatility premium
The figure shows the increase in Black-Scholes volatility due to volatility risk premium, η_V for the SVCJ model.

95 % probability when they occur.

A notable implication of the estimates of the mean jump sizes, μ_y and μ_y^Q in table 3, 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 a 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[6].

[6]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

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 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, but with a very sizable credibility interval for both jump models. Hence, we can not conclude that there is a significant market premium to jump risk. This conclusion contradicts findings in Pan (2000), where this premium is found to be significantly positive. Notice that since the premium is found as the difference in the mean jump sizes under the objective and risk neutral measures, its posterior distribution inherits the estimation uncertainty associated with the mean jump sizes. As such, it is particularly affected by the lack of precision in μ_y. To quantify this uncertainty, it would require more than forty years of data to uncover a "significant" jump premium given the estimates in table, since the posterior standard deviation of η_J is inversely proportional to the square root of the sample size.

In a final note on the parameter estimates in table 3, we address the results 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 estimated to be 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 to those reported

that are proportional to $1/\sqrt{\lambda}$.

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[7].

Pan (2000) and Bates (2000)'s 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 to those reported in BCC and here. Interpreting the differences with this in mind, the average jump intensity point estimates in Pan (2000) 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 (2000) obtains quite different results, with an average jump intensity of 0.005[8]. Bates also reports a jump size mean in the 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 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 4 shows parameter estimates obtained by using returns data only. These estimates were obtained using returns collected over the period Jan 1970 to Dec 1990, so comparisons with the estimates in table 3 should be done with this in mind. Not too surprisingly, the estimates i table 4 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 eighties being common. So while the sample used in EJP include the -7% returns on October 27th, 1997, and August 31st, 1998, comparably large moves did not occur in the seventies. This leads to a different estimate of jump frequency in this paper (λ= 0.004 vs. 0.0055 respectively). The other parameters are in the same ballpark as

[7]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.

[8]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.

Table 4: 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, Johannes, and Polson (2001) using 5307 time-series observations of the S&P 500 index from Jan 1970 to Dec 1990.

	SV	SVYJ	SVCJ
a	0.025	0.026	0.029
	(0.011)	(0.011)	(0.011)
	[-0.001, 0.051]	[0.000, 0.051]	[0.004, 0.055]
θ	0.886	0.836	0.569
	(0.099)	(0.126)	(0.070)
	[0.691, 1.172]	[0.610, 1.227]	[0.401, 0.734]
κ	0.017	0.011	0.018
	(0.004)	(0.004)	(0.004)
	[0.008, 0.029]	[0.004, 0.022]	[0.010, 0.027]
ρ	-0.373	-0.465	-0.499
	(0.055)	(0.059)	(0.075)
	[-0.493, -0.240]	[-0.594, -0.316]	[-0.666, -0.312]
σ_V	0.107	0.079	0.061
	(0.010)	(0.010)	(0.010)
	[0.086, 0.132]	[0.058, 0.105]	[0.037, 0.083]
μ_y		-3.676	-2.929
		(2.461)	(1.964)
		[-10.402, 1.520]	[-8.065, 1.486]
ρ_J			0.496
			(0.996)
			[-2.445, 2.984]
σ_y		6.685	4.052
		(1.672)	(1.207)
		[3.900, 12.149]	[1.668, 7.449]
μ_V			1.382
			(0.542)
			[0.593, 3.172]
λ		0.003	0.005
		(0.001)	(0.002)
		[0.001, 0.006]	[0.002, 0.010]

those reported in EJP. Notice again how the volatility-of-volatility parameter, σ_V is noticeably smaller than in table 3.

5.2 Option price fit

In this section, we discuss the empirical performance of the various models in fitting the historical option prices. Table 5 reports the posterior means of absolute option pricing errors for the different models, conditional upon moneyness and maturity. The results in table 5 are surprising indeed: the jump models do not seem to improve markedly upon the fit of the basic stochastic volatility model. In particular, the mean absolute pricing errors are about 47 cents across all moneyness and maturity categories for all three models. The exception is the shortest maturity contracts for which the SVCJ model improves about 3 cents upon the SV model.

At first glance, these results seem implausible. After all, the SVCJ model incorporates five 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 5 are mean errors over a large range of plausible (in the sense that they have positive posterior probability mass), parameter values. Hence, philosophically, the Baysian approach differs fundamentally from classical methods. Second, the SVJ and SVCJ models do not really increase the degrees of freedom by a notable amount. The number of parameters increases by five in the SVCJ model, but there are 3270 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 Baksi, Cao and Chen (1997). In their paper, since the model parameters are re-calibrated 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 "once stochastic volatility is modeled, adding other features will usually lead to second-order pricing improvements."

The lack of improvement in the pricing results in table 5, are partly difficult to reconcile with the results in Bates (2000) and Pan (2000) who both conclude

Figure 2: Model and market smiles conditional upon initial volatility and maturity

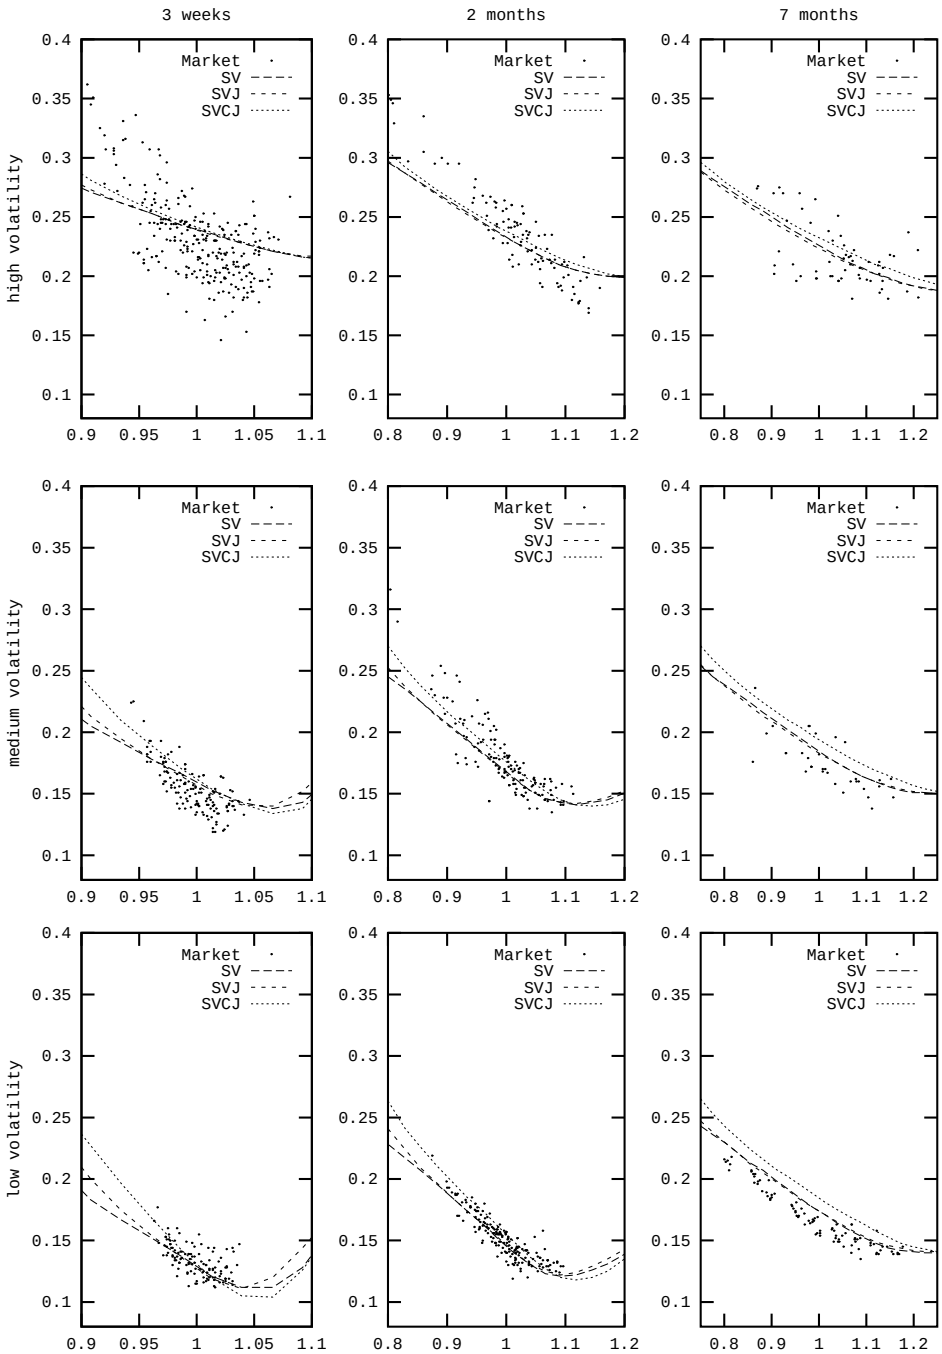

Table 5: Absolute Pricing Errors

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.

		Moneyness (strike/spot)						
Maturity		< 0.93	0.93-0.97	0.97-1.0	1.0-1.03	1.03-1.07	> 1.07	all
	#	13	86	289	272	71	9	740
	SV	0.67	0.35	0.35	0.48	0.53	0.47	0.42
< 1 m	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
	#	13	105	248	312	257	76	1011
	SV	0.33	0.42	0.32	0.43	0.54	1.14	0.48
1-2 m	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
	#	10	74	125	188	127	79	603
	SV	0.37	0.39	0.47	0.43	0.48	0.58	0.46
2-3 m	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
	#	22	68	110	194	140	176	710
	SV	1.09	0.38	0.38	0.57	0.50	0.39	0.48
3-6 m	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
	#	5	14	23	42	40	82	206
	SV	0.74	0.64	0.39	0.63	0.74	0.45	0.56
> 6 m	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
	#	63	347	795	1008	635	422	3270
	SV	0.71	0.40	0.37	0.48	0.53	0.57	0.47
all	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

that jumps (in prices) are important in capturing systematic variations in Black-Scholes volatilities. Figure 2 suggests explanations to the surprising lack of fit for the jump models reported in table 5: 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 IV's are very similar for all three models, and on average, they tend to fit the observed data reasonably well. There are a few

exceptions:

- For the short maturity contracts (3 weeks) trading on days of high or average spot volatility, the models do not generate sufficiently steep IV's.
- For long maturity contracts trading on days with low spot volatility, all models overprice and the SVCJ model overprice the most.

The first observation is suggestive of why the models with state-dependent jump intensities in Bates and Pan outperform the SV model. Models with such state dependency will generate steeper volatility smiles in high volatility environments because the chance of observing large, negative jumps increases. Vice versa, state-dependent models generate flatter IV curves in low volatility states. It is also possible to incorporate the same effect by allowing the jump size distribution to depend on V_t.

To address the second problem, notice that we could have avoided the overpricing of long maturities in low volatility states by a lesser rate of volatility mean reversion in all models. Since the rates of mean reversion are consistent with the average term structure of implied volatility across all options, one must incorporate a specific "low volatility - low mean reversion" type effect. In other words, this observation is suggestive of a state-dependent volatility mean reversion.

To see why, notice that the IV curves for the jump models are steeper at short maturities than the SV models whenever the current market volatility is low. This has a natural interpretation: for short time invervals (say one day), the stock price distribution is approximately a mixture of normals,

$$\ln \frac{S_{t+1}}{S_t} \sim (1-\lambda)N(r, V_t) + \lambda N(\mu_y^Q, \sigma_y^2).$$

As such, the second term gains influence for lower values of V_t, and vise versa.

Given this mixture of normals interpretation, it is also quite easy to see how the low values of λ will have virtually no effect on the conditional stock price distribution except in the very far tails of the distribution. Thus, only far out or in the money option contracts are affected by the presence of jumps. This, as well as the fact that most options trade at moneyness/maturity constellations with non-negligible probability mass, suggests why the jump models do not fit the options data better than the SV model does.

The constant jump intensity is also likely to cause the problems reported for the

SVCJ model for the long maturity, low volatility states. The jumps sizes and jump probability are independent of the current level of in volatility are of magnitude and arrive independently of the current value of the process. Hence, an argument could be made similar to the one regarding price jumps: the conditional density of volatility has a relatively heavier right tail when volatility is low. This potentially explains why the SVCJ model is doing worse than SV and SVJ for long maturity options in periods of low market volatility.

It is also important to recognize that the results in table 5 could be sensitive to the dollar/cent denomination. 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 5, but also for the parameter estimates in 3. In particular, it is true that the values of the joint posterior distribution of the SVCJ model are not too different for different parameter constellations for which the importance of jumps increase 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 heteroscedastic 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 5 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. We turn to an examination of the time-series fit next.

5.3 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.

Table 6 presents the pricing errors broken down in maturity and moneyness categories. Surprisingly, the table shows that the more complicated models with jumps and jumps in volatility, perform worse than the simple stochastic volatility model. The table also show that all three models do extremely poorly in fitting the longer term options - the pricing errors are increasing from around 33 cents to more than 2 dollars on average for the longest maturity contracts.

Figure 3 share some light on the substantiative and disappointing performance of the models. This figure plots the pricing errors for less-than-one-month, two-to-three month, and more than six month contracts. As can be seen from the figure, all models are performing reasonably well in the year following the estimation period. As time progresses however, the long term contracts gradually become more and more mispriced.

There is a simple explanation for the above findings: The first period used for parameter estimation included the crash of '87, and the subsequent high volatility.

Table 6: 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 2 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.

		Moneyness (strike/spot)						
Maturity		< 0.93	0.93-0.97	0.97-1.0	1.0-1.03	1.03-1.07	> 1.07	all
	#	1213	2623	3306	3159	1398	161	11860
	SV	0.19	0.29	0.40	0.40	0.17	0.07	0.32
< 1 m	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
	#	964	1621	2830	3159	1675	244	10493
	SV	0.37	0.50	0.52	0.52	0.33	0.17	0.46
1-2 m	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
	#	552	668	1446	2027	998	303	5994
	SV	0.36	0.47	0.42	0.41	0.38	0.29	0.40
2-3 m	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
	#	635	482	861	1430	905	474	4787
	SV	0.35	0.55	0.64	0.74	0.82	0.74	0.67
3-6 m	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
	#	354	289	390	598	516	609	2756
	SV	0.93	1.59	1.66	2.37	2.53	2.66	2.10
> 6 m	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
	#	3718	5683	8833	10373	5492	1791	35890
	SV	0.36	0.46	0.52	0.60	0.59	1.18	0.56
all	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

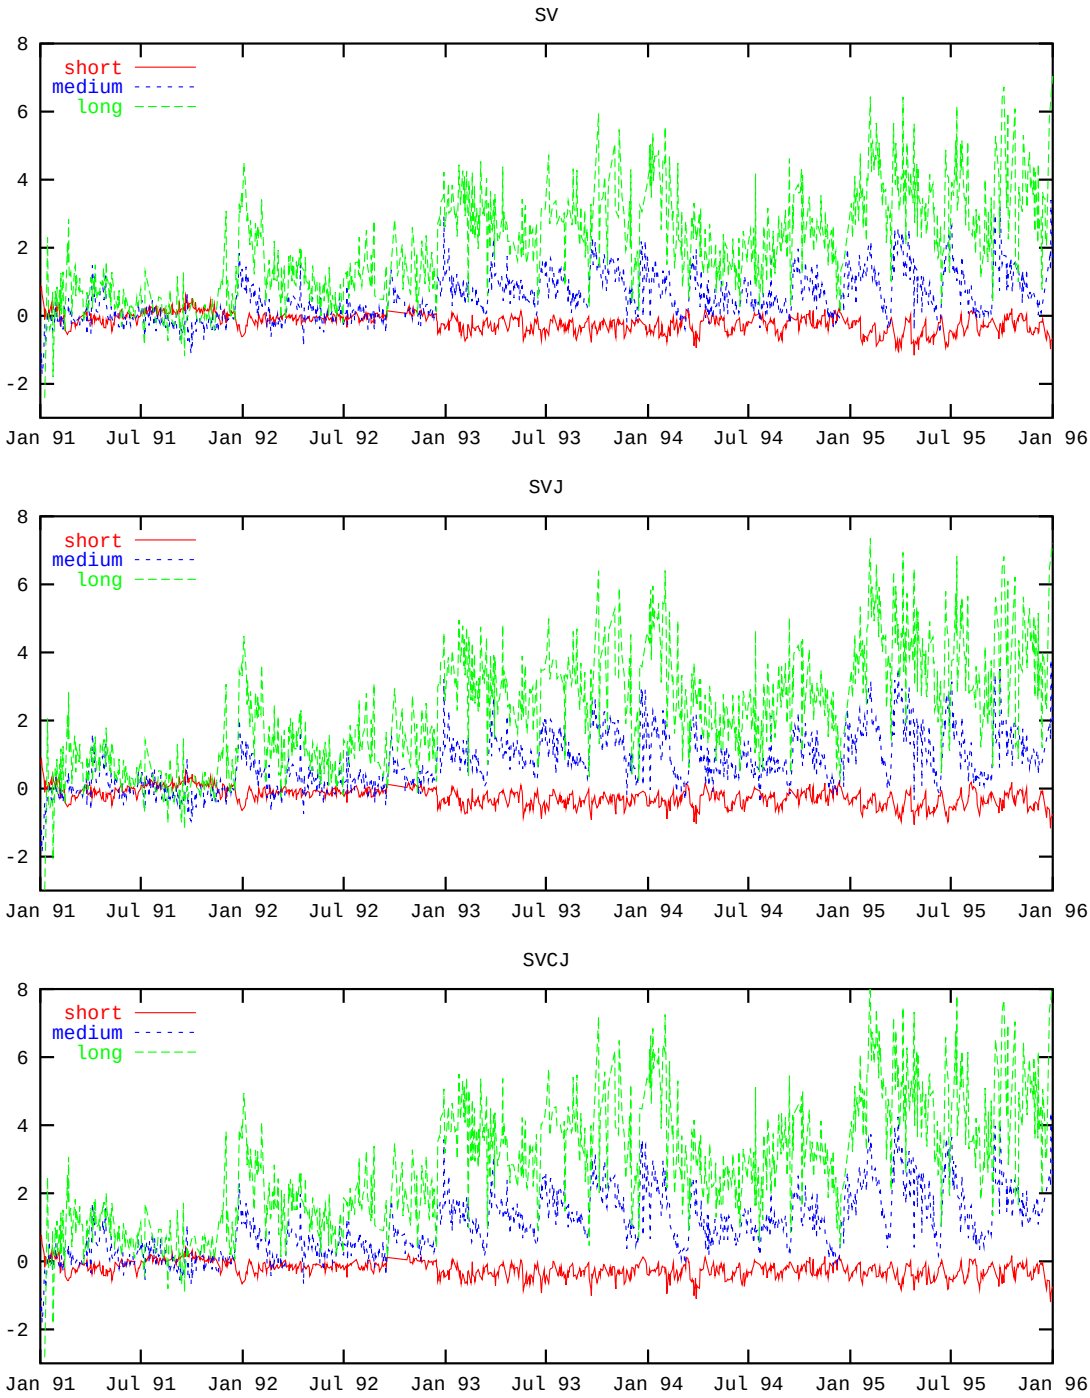

Figure 3: Out-of-sample pricing errors (model less market) for short, medium, and long maturity contracts.

This is reflected, for instance, in the relatively high parameter estimates for θ For the SV model, the estimate of θ implies an average volatility of 22%. Comparably, the annualized volatility during the out-of-sample period is 15.3% which is very close to the the historical average. 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.

Hence, the high estimated values of θ causes the overpricing of the long term contract. Notice that for the jump models, the speed of mean reversion under Q is higher than for the SV model. This implies that for these models the mispricing due to a "too high" value for θ is elevated by a more rapid mean reversion. This explains why the jump models are performing worse than the SV model out of sample.

To further shed light on how the high values of θ affects the pricing errors, table 7 presents the results of 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 strongly support the notion that jump diffusion models give better descriptions of the relative pricing of short term option contracts.

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 imply 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

Table 7: Out of sample absolute Pricing Errors - adjusted θ

The table reports mean absolute pricing errors for the pricing model from adjusting the θ parameter to reflect an average, annualized volatility of 15%.

		Moneyness (strike/spot)						
Maturity		< 0.93	0.93-0.97	0.97-1.0	1.0-1.03	1.03-1.07	> 1.07	all
	#	1213	2623	3306	3159	1398	161	11860
	SV	0.15	0.18	0.51	0.64	0.25	0.07	0.40
< 1 m	SVJ	0.12	0.15	0.39	0.48	0.21	0.06	0.31
	SVCJ	0.09	0.14	0.27	0.34	0.16	0.09	0.22
	#	964	1621	2830	3159	1675	244	10493
	SV	0.23	0.28	0.50	0.55	0.48	0.26	0.45
1-2 m	SVJ	0.22	0.26	0.39	0.44	0.42	0.26	0.37
	SVCJ	0.20	0.26	0.30	0.40	0.34	0.31	0.32
	#	552	668	1446	2027	998	303	5994
	SV	0.23	0.30	0.34	0.35	0.41	0.34	0.34
2-3 m	SVJ	0.23	0.30	0.31	0.32	0.38	0.34	0.32
	SVCJ	0.28	0.31	0.29	0.34	0.37	0.43	0.33
	#	635	482	861	1430	905	474	4787
	SV	0.31	0.51	0.58	0.71	0.65	0.42	0.57
3-6 m	SVJ	0.33	0.52	0.52	0.59	0.55	0.40	0.51
	SVCJ	0.44	0.54	0.48	0.49	0.49	0.46	0.48
	#	354	289	390	598	516	609	2756
	SV	0.56	0.88	1.45	1.38	1.60	1.22	1.24
> 6 m	SVJ	0.59	0.82	1.21	1.01	1.15	0.93	0.97
	SVCJ	0.70	0.84	0.90	0.92	0.98	0.90	0.89
	#	3718	5683	8833	10373	5492	1791	35890
	SV	0.25	0.29	0.53	0.61	0.54	0.63	0.49
all	SVJ	0.24	0.27	0.43	0.48	0.45	0.52	0.41
	SVCJ	0.26	0.26	0.33	0.41	0.39	0.55	0.36

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 determine the long-run volatility. Such models have been 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, Gallant, Ghysels, and Tauchen (99) 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.

5.4 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 occurence 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 8 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 4. As can be seen from the table, the returns-based parameter estimates produce overall pricing errors that are smaller than those in table 6, but larger than those in table 7 where 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.

Table 8: 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 4) 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.

		Moneyness (strike/spot)						
Maturity		< 0.93	0.93-0.97	0.97-1.0	1.0-1.03	1.03-1.07	> 1.07	all
	#	1213	2623	3306	3159	1398	161	11860
	SV	0.20	0.33	0.36	0.43	0.25	0.08	0.34
< 1 m	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
	#	964	1621	2830	3159	1675	244	10493
	SV	0.50	0.72	0.47	0.46	0.58	0.36	0.52
1-2 m	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
	#	552	668	1446	2027	998	303	5994
	SV	0.75	1.04	0.54	0.41	0.71	0.55	0.60
2-3 m	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
	#	635	482	861	1430	905	474	4787
	SV	1.12	1.31	0.73	0.62	0.92	0.97	0.87
3-6 m	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
	#	354	289	390	598	516	609	2756
	SV	1.61	1.55	1.32	0.94	1.23	1.53	1.33
> 6 m	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
	#	3718	5683	8833	10373	5492	1791	35890
	SV	0.65	0.67	0.51	0.49	0.63	0.93	0.58
all	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

Table 9: 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, $g = 1, .., G$. Parameter estimates correspond to a unit of time defined to be one day, and returns data scaled by 100.

	SV	SVJ	SVCJ
mean	0.202	0.198	0.151
std.	0.005	0.004	0.005

5.5 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 (2000) 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 9 reports estimates of σ_V from the simulated values of the historical volailities, $V_t^g, t = 1, .., T, g = 1, .., G$ analogous to those in BCC and Bates. The numbers match almost exactly those reported in table 3. The point estimates in table 9 are slightly smaller than the posterior means reported in table 3, however the posterior credibility intervals in table 3 do indeed overlap with the point estimates in table 9. Hence, the mismatch between 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 dy-

namics. 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 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 $\triangle W_t^S = W_t^i - W_{t-1}^i$ for $i = \{S, V\}$, the continuously arriving shocks to prices and volatility. These can be interpreted as the model standardized "residuals" [9]. Figures 4 and 5 plot these residual for returns and volatility, respectively. It is evident from these plots that the SVCJ model fares far better in explaining large stock price movements than the other models. For example, the -22% crash of 10/19/87 and the 6% drop on 10/13/89 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 model, this changes because of 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 shows the estimated spot volatilities for each model, along with the residual plots corresponding to the shocks in the volatility process. As is well known, the infamous crash of '87 had a huge impact on option prices across all maturities, especially in the days following the crash. The effect on the estimated, latent spot

[9] The analysis of residuals in Bayesian analysis is not as straightforward as it is with classical methods. For example, consider "residuals" obtained from $y_i \sim N(0, s^2)$ where the parameter of interest, s, has an conjugate inverse Gamma posterior distribution. Then it is easy to see that the posterior distribution of the residual term y_i/s is a t-distribution(!), and not the normal distribution assumed under the null hypothesis. See Zellner (1975) for details. Zellner recommends that an 'optimal' value of the parameter is chosen for residual diagnostics.

Figure 4: Return residuals

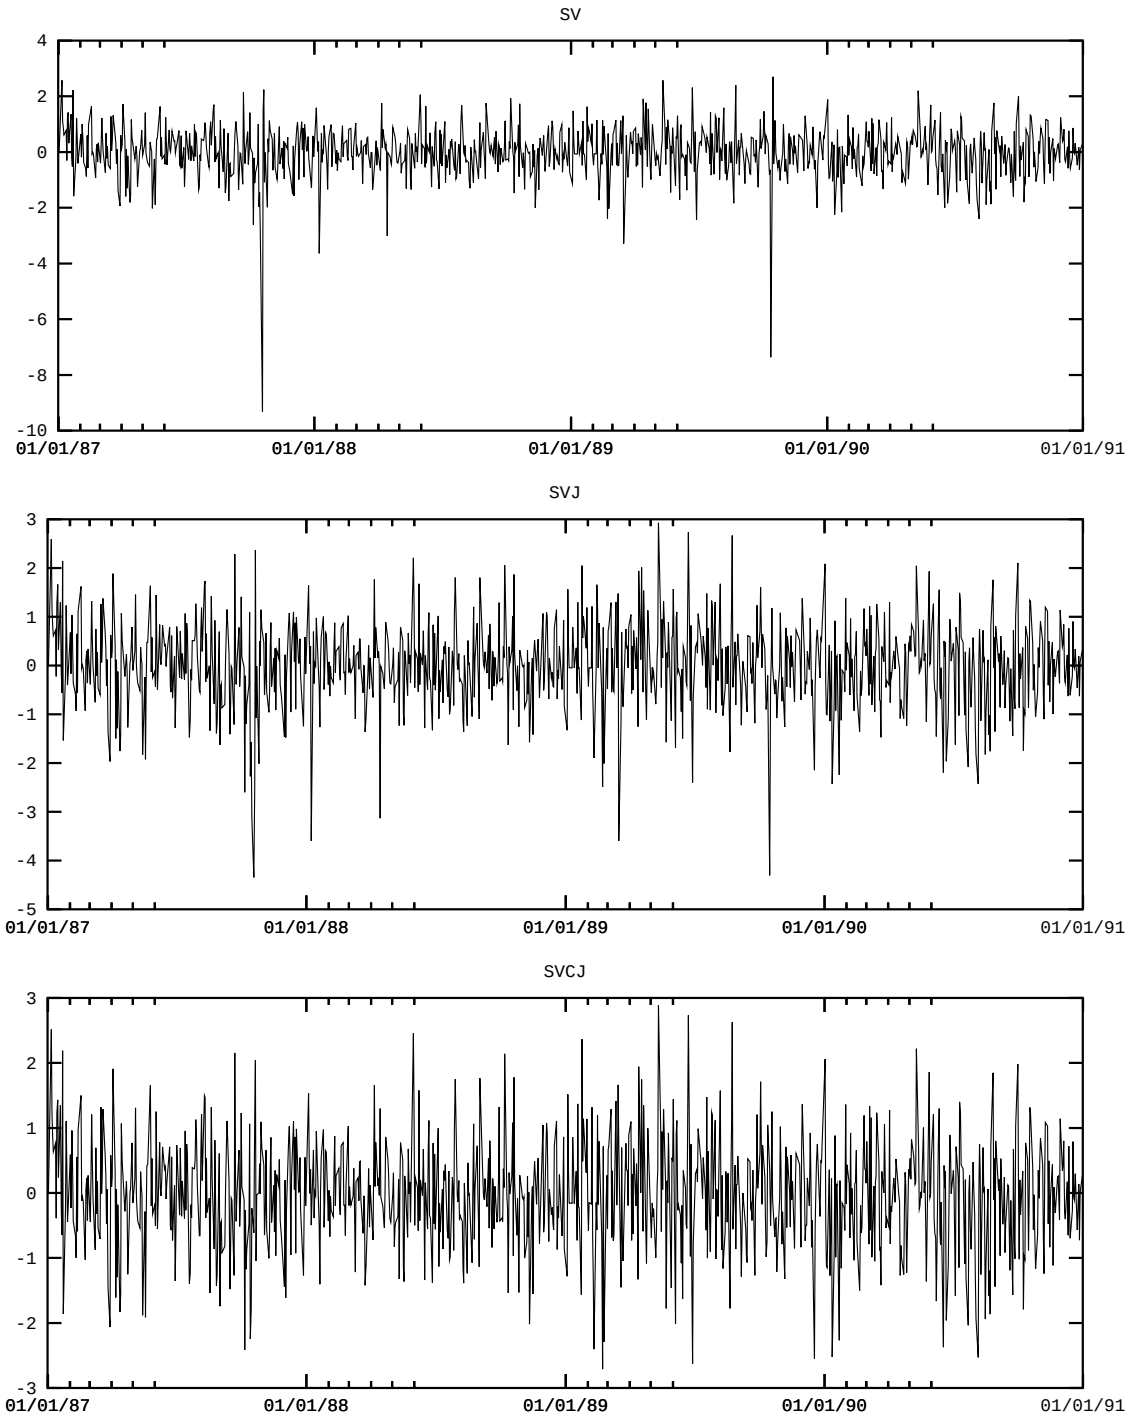

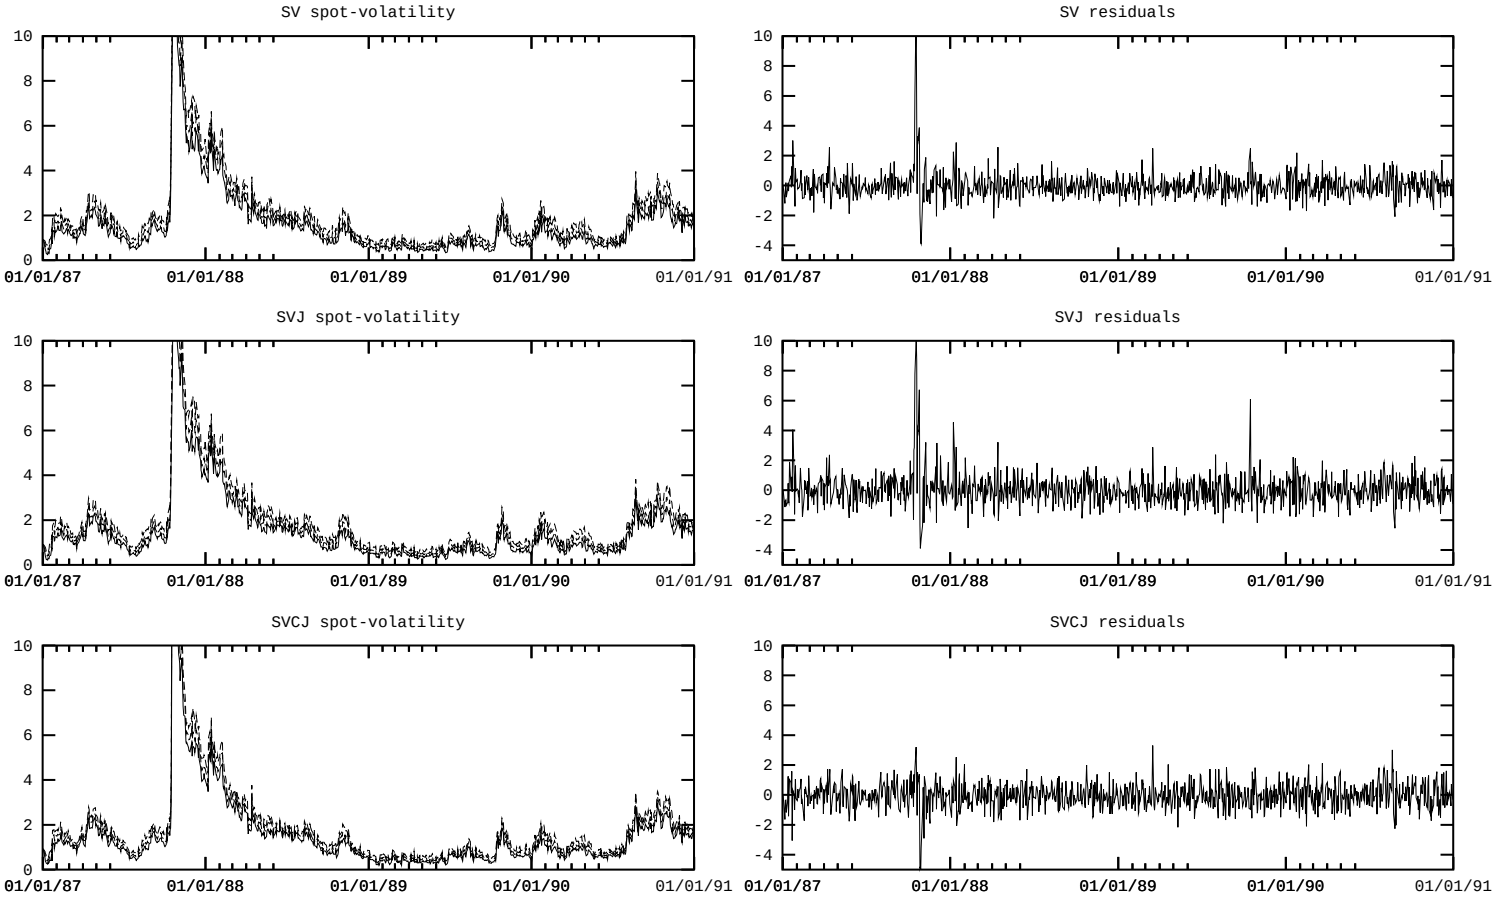

Figure 5: Posterior volatility percentiles and residuals

volatility is a huge jump in volatility. Regardless of whether such jumps are specified as part of the model (SVCJ), or not, the estimated spot volatility increase dramatically on the day of the crash. For the SV model, the posterior median of V_t is 5.7 on the day before the crash (10/16/87), 11.5 on the day of the crash (10/19/87) and 15 on the following day (10/20/87). The corresponding numbers for the SVJ model are 5.5,9.8 and 14.5, while the spot volatility in the SVCJ model is estimated to have increased from 4.1 on the Friday before the crash, to 18.8 and 18.6 on the day of, and the day following the crash. Hence, the dynamic restrictions implicit in the joint likelihood function for returns and volatility *do* have an impact on the posterior distribution of the spot volatilities. However, the restrictions are not sufficient to produce residuals consistent with the brownian motion assumption.

The upper plots in figure 5 show the standardized residual plots for the models. There are huge outliers in the innovations to the volatility series estimated for the SV and SVJ models with values greater than 10 on the day of the crash - an event which is extremely unlikely under the assumed $N(0,1)$ distribution. The SVCJ model is an exception, and this model essentially explains the increase in spot volatility on this day as a jump to volatility.

Interestingly, all three models generate too large *negative* movements in the days following the 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 in the SVCJ model are assumed exponentially distributed, and hence positive.

Finally, table 10 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 are 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 1% percentile of 3.56 and 4.33 respectively.

Table 10: 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
		Return residuals	
Skewness	-1.891	-0.692	-0.129
	(0.088)	(0.294)	(0.060)
	[-2.078, -1.709]	[-1.633, -0.364]	[-0.240, -0.016]
Kurtosis	20.352	7.223	3.959
	(1.121)	(3.282)	(0.267)
	[18.083, 22.615]	[4.525, 18.630]	[3.563, 4.620]
		Volatility residuals	
Skewness	2.487	2.045	-0.300
	(0.290)	(0.249)	(0.134)
	[1.940, 3.194]	[1.475, 2.638]	[-0.564, -0.068]
Kurtosis	27.758	19.256	6.017
	(4.515)	(3.076)	(0.693)
	[19.876, 38.507]	[13.687, 28.204]	[4.337, 7.700]

6 Concluding remarks

This paper has presented estimates of parameters governing the dynamics of jump diffusion models of stock price data from joint returns and option price data. For the purpose of estimation, we have discussed a novel implementation of Markov Chain Monte Carlo (MCMC) which allows the posterior distributions of parameters, latent volatility, jump times and jump sizes to be estimated. The results show several interesting features:

The parameter estimates show a surprisingly low arrival frequency of jumps. Although jumps occur infrequently, our parameter estimates imply sizable negative movements in stock prices, reminiscent of the crash of '87. Other parameter estimates are comparable to those reported elsewhere.

The in-sample performance measures show surprisingly little improvement in the fit of jump diffusion models to the option price data. The diagnostics show that the anticipated increase in implied skewness at short maturity contracts, is really only present whenever the spot volatility is high. Conversely, in low volatility states the stochastic volatility model seems adequate to capture the steepness in implied Black-Scholes volatilities across moneyness categories. Since the arrival intensity of jumps is constant and does not depend on current market volatility in the model examined here, it seems impossible for these models to capture the variations in steepness in different volatility states. This finding is consistent with models that allow for a state-dependent jump arrival intensity, as in Bates (2000). All models seem to perform adequately for long maturity contracts in the sense that they capture the skewness in the conditional returns distributions implied by option prices.

Our time-series diagnostics indicate that the MCMC estimator employed does not produce a volatility of volatility that is too high relative to the variation in the volatility trajectories simulated from the posterior distribution. This finding is at odds with previous evidence presented in Baksi, Cao and Chen (1997), Bates (2000) and Pan (2000). As such, it represents an important re-appraisal of all the models considered. In part, the reconciliation of the joint options/returns estimate with that of the latent volatility series is likely to be caused by a binding constraint in the time-series based part of the likelihood function. This also implies that the crash of '87, and its subsequent impact on spot-volatility, are important factors influencing, and possibly increasing, the estimate of volatility of volatility. Since this parameter

increases the kurtosis in the conditional returns distribution, the role of jumps in capturing option prices, is likely to be aversely affected.

With respect to the time-series dimension of the problem, the SVCJ model that allows for simultaneous jumps to prices and volatility, performs markedly better than the pure stochastic volatility (SV), and the combined price-jump, stochastic volatility (SVJ) models.

Overall, the results point toward models that either allow for multiple factors, or models that incorporate an increased level of complexity through the jump distributions. As such, models that incorporate state (volatility) dependence in either the jump arrival intensity, or the jump size distribution, are likely candidates. While these suggestions possibly explain the option prices, they are unlikely to be sufficient to explain the time-series of both returns and latent volatility. The main problem with the estimated V_t process is that it is reverting from its high following the '87 crash far to rapidly to be consistent with the assumption of a linear drift. There are two potential model explanations that can address this. First, conceivably the volatility jumps must be allowed to be negative. This poses tractability issues in the resulting option-pricing formula[10]. A second conceivable way of modeling the sharp decline in volatility following the crash is to increase the mean reversion. As the mean reversion generally tends to be very low, this would hence call for a state dependent volatility reversion. Two factor volatility models, along the lines of those studied in Bates (2000) for option pricing, accomplish this. His analysis, however, also indicates outliers in the volatility innovations of these two factor volatility models. Moreover, the implied Black-Scholes volatilities at long maturities are, during periods of low volatility, also shown to be too low to be consistent with with the models. Together, these observations seem to imply that the mean reversion imposed by the linear drift specification, is to low in high volatility states, and too high in low volatility states. It is conceivable that both observations indicating state dependent volatility reversion would call for a non-linearity in the drift of the volatility process.

Interestingly, this interpretation is also supported by the (independent) out-of-sample evidence. This evidence shows that estimates of the mean-volatility parameter, θ, obtained using the volatile '87-'91 period, are too high to be consistent with the

[10]It is possible to show that the fourier transform of the jump size distribution is no longer affine if the model allows for negative volatility jumps while maintaining positivity of the process. An affine transform is necessary for the model to remain in the class discussed in Duffie, Pan, and Singleton (2000)

following ('91-'96) low volatility period. This raises the possibility that a model for which the volatility reversion would be lower in the low volatility period, would perform better than the models considered here.

References

Aït-Sahalia, Y. (1996). Testing continuous-time models of the spot interest rate. *Review of Financial Studies*, *9*, 385–426.

Aït-Sahalia, Y. (1998). Maximum Likelihood Estimation of Discretely Sampled Diffusions: A Closed-Form Approach. *Working paper, Princeton.*

Aït-Sahalia, Y. and Lo, A. (1998). Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices. *Journal of Finance*, *53*, 499–547.

Aït-Sahalia, Y., Wang, Y., and Yared, F. (2000). Do Options Markets Correctly Assess the Probabilities of Movement in the Underlying Asset?. *Journal of Econometrics*, *forthcoming.*

Amin, K. I. and Ng, V. (1993). Option valuation with systematic stochastic volatility. *Journal of Finance*, *48*, 881–910.

Andersen, T. G., Benzoni, L., and Lund, J. (1999). Estimating Jump-Diffusions for Equity Returns. *Working paper, Northwestern University.*

Andersen, T. G. and Lund, J. (1997). Estimating continuous-time stochastic volatility models of the short-term interest rate. *Journal of Econometrics*, *7*, 343–377.

Baksi, G., Cao, C., and Chen, Z. (1997). Empirical Performance of Alternative Option Pricing Models. *Journal of Finance*, *52*, 2003–2049.

Baksi, G. and Madan, D. (1999). Crash Discovery in Stock and Options Markets. *Working paper, University of Maryland.*

Ball, C. and Torous, W. (1985). On Jumps in Common Stock Prices and Their Impact on Call Option Pricing. *Journal of Finance*, *40*, 155–173.

Bandi, F. M. and Philips, P. C. B. (1999). Econometric Analysis of Diffusion Models. *Working paper, University of Chicago.*

Bandi, F. M. and Philips, P. C. B. (2000). Semiparametric Methods for diffusion model estimation. *Working paper, University of Chicago.*

Bates, D. (1996). Jump and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies*, *9*, 69–107.

Bates, D. (2000). Post-'87 Crash fears in S&P 500 Futures Options. *Journal of Econometrics*, *94*, 181–238.

Benzoni, L. (1998). Pricing Options under Stochastic Volatility: An Econometric Analysis. *Working paper*, *University of Minnesota.*

Black, F. (1976). Studies in Stock Price Volatility Changes. *Proceedings of the 1976 meeting of the Business and Economic Statistics Section, American Statistical Association*, 177–181.

Carlin, B. P., Polson, N. G., and Stoffer, D. S. (1992). A Monte Carlo Approach to Nonnormal and Nonlinear State-Space Modeling. *Journal of the American Statistical Association*, *87*, 493–500.

Chernov, M., Gallant, A. R., Ghysels, E., and Tauchen, G. (99). A New Class of Stochastic Volatility Models with Jumps: Theory and Estimation. *Working Paper, Duke University.*

Chernov, M. and Ghysels, E. (2000). Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for te Purpose of Options Valuation. *Journal of Financial Economics*, *forthcoming.*

Chesney, M. and Scott, L. (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*, 267–84.

Conley, T., Hansen, L. P., Luttmer, E., and Scheinkman, J. (1997). Estimating Subordinated Diffusions from Discrete Time Data. *Review of Financial Studies*, *10*, 525–577.

Dai, J. and Singleton, K. S. (2000). Specification Analysis of Affine Term Structure Models. *Journal of Finance*, *LV*, 1943–1978.

Das, S. and Sundaram, R. (1999). Of Smiles and Smirks: A Term-Structure Perspective. *Journal of Financial and Quantitative Analysis*, *34(2)*, 211–239.

David, A. and Veronesi, P. (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, D. and Glynn, P. (1996). Estimation of Continuous-Time Markov Processes Sampled at Random Time Intervals. *Manuscript, Stanford University.*

Duffie, D., Pan, J., and Singleton, K. J. (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica*, *forthcoming.*

Duffie, D. and Singleton, K. (1993). Simulated Moments Estimation of Markov Models of Asset Prices. *Econometrica*, *61*, 929–952.

Dumas, B., Fleming, J., and Whaley, R. (1998). Implied Volatility Functions: Empirical Tests. *Journal of Finance*, *53*, 2059–2106.

Eraker, B. (2000). MCMC Analysis of Diffusion Models with Application to Finance. *Journal of Business and Economic Statistics*, *forthcoming.*

Eraker, B., Johannes, M. J., and Polson, N. G. (2000). The Impact of Jumps in Returns and Volatility. *Working paper, University of Chicago, Columbia University.*

Gallant, A. R. and Tauchen, G. (1996). Which Moments to Match?. *Econometric Theory*, *12*, 657–681.

Gourieroux, C., Monfort, A., and Renault, E. (1993). Indirect Inference. *Journal of Applied Econometrics*, *8*, S85–S118.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, *50*, 1029–1054.

Hansen, L. P. and Scheinkman, J. (1995). Back to the Future: Generating Moment Implications for Continuous-Time Markov Processes. *Econometrica*, *63*, 767–804.

Heston, S. (1993). Closed-Form Solution of Options with Stochastic Volatility with Application to Bond and Currency Options. *Review of Financial Studies*, *6*, 327–343.

Hull, J. and White, A. (1987). The pricing of options with stochastic volatilities. *Journal of Finance*, *42*, 281–300.

Jackwerth, J. C. and Rubenstein, M. (1996). Recovering Probability Distributions from Option Prices. *Journal of Finance*, *51*, 1611–1631.

Jacquier, E. and Jarrow, R. (2000). Bayesian Analysis of Contingent Claim Modell Error. *Journal of Econometrics*, *94*, 145–180.

Jacquier, E., Polson, N. G., and Rossi, P. (1994). Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics*, *12*, 371–389.

Jacquier, E., Polson, N. G., and Rossi, P. (1999). Stochastic Volatility: Univariate and Multivariate Extensions. *Working paper, University of Chicago.*

Jiang, G. J. and Knight, J. L. (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, M. S., Kumar, R., and Polson, N. G. (1999). State Dependent Jump Models: How do US Equity Indices Jump?. *Working paper, University of Chicago.*

Melino, A. and Thornbull, S. (1990). Pricing Foreign Currency options with Systematic Stochastic Volatility. *Journal of Econometrics*, *45*, 239–265.

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

Pan, J. (2000). Integrated Time-Series Analysis of Spot and Option Prices. *Journal of Financial Economics*, *forthcoming.*

Pastorello, S., Renault, E., and Touzi, N. (1998). Statistical Inference for Random Variance Option Pricing. *Working Paper, Universita di Bologna, Italy.*

Rubenstein, M. (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, M. (1994). Implied Binomial Trees. *Journal of Finance*, *49*, 771–818.

Scott, L. O. (1987). Option Pricing when the Variance Changes Randomly: Theory, Estimation, and an Application. *Journal of Financial and Quantitative Analysis*, *22*, 419–483.

Singleton, K, J. (2001). Estimation of Affine Asset Pricing Models Using the Empirical Characteristic Function. *Journal of Econometrics*, *forthcoming*.

Stanton, R. (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 Stein, J. C. (1991). Stock Price Distributions with Stochastic Volatility: An Analytic Approach. *Review of Financial Studies*, *4*, 727–752.

Viceira, L. M. and Chacko, G. (1999). Spectral GMM Estimation of Continuous-Time Processes. *Working paper, Harvard University*.

Wiggins, J. B. (1987). Option values under stochastic volatility: Theory and empirical estimates. *Journal of Financial Economics*, *19*, 351–372.

Zellner, A. (1975). Bayesian Analysis of Regression Error Terms. *Journal of the American Statistical Association*, *70*, 138–144.