

Model Specification and Risk Premia: Evidence from Futures Options

MARK BROADIE, MIKHAIL CHERNOV, and MICHAEL JOHANNES*

ABSTRACT

This paper examines model specification issues and estimates diffusive and jump risk premia using S&P futures option prices from 1987 to 2003. We first develop a time series test to detect the presence of jumps in volatility, and find strong evidence in support of their presence. Next, using the cross section of option prices, we find strong evidence for jumps in prices and modest evidence for jumps in volatility based on model fit. The evidence points toward economically and statistically significant jump risk premia, which are important for understanding option returns.

THERE ARE TWO CENTRAL, RELATED, issues in empirical option pricing. The first issue is model specification, which comprises identifying and modeling the factors that jointly determine returns and option prices. Recent empirical work on index options identifies factors such as stochastic volatility, jumps in prices, and jumps in volatility. The second issue is quantifying the risk premia associated with the jump and diffusive factors using a model that passes reasonable specification hurdles.

The results in the literature regarding these issues are mixed. For example, tests using option data disagree over the importance of jumps in prices: Bakshi, Cao, and Chen (1997) (BCC) find substantial benefits from including jumps in prices, whereas Bates (2000) and others find that such benefits are economically small, if not negligible.¹ Furthermore, while studies using the time series of returns unanimously support jumps in prices, they disagree with respect to the importance of jumps in volatility. Finally, there is general disagreement regarding the magnitude and significance of volatility and jump risk premia.

*Broadie and Johannes are affiliated with the Graduate School of Business, Columbia University. Chernov is affiliated with Graduate School of Business, Columbia University and London Business School. We thank seminar participants at Columbia, Connecticut, Northwestern University, London School of Economics, London Business School, and the Western Finance Association meetings for helpful comments. David Bates, Gurdip Bakshi, Chris Jones provided especially helpful comments. We are very grateful to the anonymous referee whose comments resulted in significant improvements in the paper. We thank Tony Baer for excellent research assistance. This work was partially supported by NSF Grant #DMS-0410234.

¹ Pan (2002) finds that pricing errors decrease when jumps in prices are added for certain strike-maturity combinations, but increase for others. Eraker (2004) finds that adding jumps in returns and volatility decreases errors by only 1%. Bates (2000) finds a 10% decrease, but it falls to around 2% when time-series consistency is imposed.




Figure 1. Time series of implied volatility. This figure displays the time series of implied volatility, as measured by the VIX index, from 1987 to March 2003.

One plausible explanation for the above disparities is that most papers use data covering only short time periods. For instance, BCC and Bates (2000) use the cross section of options from 1988 to 1991 and 1988 to 1993, respectively, Pan (2002) uses two options per day from 1989 to 1996, and Eraker (2004) uses up to three options per day from 1987 to 1990. Since jumps are rare, short samples are likely to either over- or under-represent jumps and/or periods of high or low volatility, and thus could generate the disparate results. Figure 1, which displays a time-series plot of the VIX index, shows how short subsamples may be unrepresentative over the overall sample. Hence, to learn about rare jumps and stochastic volatility, and investors' attitudes toward the risks these factors embody, it is important to analyze as much data as possible.

In this paper, we use an extensive data set of S&P 500 futures options from January 1987 to March 2003 to shed light on these issues. In particular, we address three main questions. (1) Is there option-implied time-series evidence for jumps in volatility? (2) Are jumps in prices and volatility important factors in determining the cross section of option prices? (3) What is the nature of the factor risk premia embedded in the cross section of option prices?

Regarding the first question, we develop a test to detect jumps in volatility. Intuitively, volatility jumps should induce positive skewness and excess kurtosis in volatility increments. To test this conjecture, we first extract a model-based estimate of spot variance from options. We then calculate skewness and kurtosis

statistics and simulate the statistics' finite sample distribution. The tests reject a square-root stochastic volatility (SV) model and an extension with jumps in prices (the SVJ model), as these models assume that volatility increments are approximately normal. These rejections are robust to reasonable parameter variations, excluding the crash of 1987, and factor risk premia. A model with contemporaneous jumps in volatility and prices (SVCJ) easily passes these tests.

Next we turn to the information in the cross section of options prices to examine model fit and estimate risk premia. In estimating models using the cross section of option prices, we depart from the usual pure calibration approach and follow Bates (2000) by constraining certain parameters to be consistent with the time-series behavior of returns. More precisely, the volatility of volatility and the correlation between the shocks to returns and volatility should be equal under the objective and risk-neutral probability measures. We impose this constraint for both pragmatic and theoretical reasons. First, there is little disagreement in the literature over these parameter values.² Second, absolute continuity requires these parameters to be equal in the objective and risk-neutral measures. Finally, joint estimation using both options and returns is a computationally demanding task.

In terms of pricing, we find that adding price jumps to the SV model improves the cross-sectional fit by almost 50%. This is consistent with the large impact reported in BCC, but contrasts with the negligible gains documented in Bates (2000), Pan (2002), and Eraker (2004). Without any risk premium constraints, the SVJ and SVCJ models perform similarly in and out of sample. This is not surprising, as price jumps, which generate significant amounts of skewness and kurtosis, and stochastic volatility are clearly the two most important components for describing the time series of returns or for pricing options. Jumps in volatility have a lesser impact on the cross section of option prices. This does not mean volatility jumps are not important, however, as they are important for two reasons. First, volatility jumps are important for explaining the time series of returns and option prices. Second, it is dangerous to rely on risk premia estimated from a clearly misspecified model. Thus, even if the cross-sectional fit of the SVJ and SVCJ models is similar, the risk premia estimated using the SVJ model should not be trusted.

Turning to risk premia, our specification allows for the parameters that index the price and volatility jump size distributions to change across measures; we refer to the differences as "risk premia." Thus, we have a mean price jump risk premium, a volatility of price jumps risk premium, and a volatility jump risk premium. The premium associated with Brownian shocks in stochastic volatility is labeled the diffusive volatility risk premium.

The risk premia have fundamentally different sources of identification. In theory, the term structure of implied volatility primarily identifies diffusive

² As an example, the reported estimates for the volatility of volatility and correlation parameters in the SVCJ model are 0.08 and -0.48 (Eraker, Johannes, and Polson (2003)), 0.07 and -0.46 (Chernov et al. (2003)), and 0.06 and -0.46 (Eraker (2004)), respectively.

volatility premia, while the implied volatility smile identifies jump risk premia. In our sample, it is difficult to identify the diffusive volatility risk premium because most traded options are short dated and the term structure of implied volatility is flat.³ In contrast to the noisy estimates of diffusive volatility risk premia, the implied volatility smile is very informative about the risk premia associated with price jumps and volatility jumps, resulting in significant estimates.

Using the SVJ model, the mean price jump risk premia is 3% to 6%, depending on the volatility of price jumps risk premium. Mean price jump risk premia of this magnitude are significant, but not implausible, at least relative to simple equilibrium models such as Bates (1988). Using the SVCJ model, the mean price jump risk premium is smaller, about 2% to 4%, depending again on the assumptions regarding other premia. In all cases, the mean price jump risk premia are highly significant, though modest compared to previously reported estimates. We also find statistically significant volatility of price jumps and volatility jump risk premia.

Finally, to quantify the economic significance of the risk premia estimates, we consider the contribution of price jump risk to the equity risk premium and analyze how jump risk premia affect option returns. First, price jump risk premia contribute about 3% per year to an overall equity premium of 8% over our sample. Second, we use our estimates to decompose the historically high returns to put options, commonly referred to as the “put-pricing” anomaly.⁴ Based on our estimates, roughly half of the high observed returns are due to the high equity risk premium over the sample, while the other portion can be explained by modest jump risk premia. We therefore conclude that even relatively small jump risk premia can have important implications for puts. The main reason these returns appear to be puzzling is that, not surprisingly, standard linear asset pricing models have difficulty capturing jump risks.

I. Models and Methods

A. Affine Jump Diffusion Models for Option Pricing

On $(\Omega, \mathcal{F}, \mathbb{P})$, we assume that the equity index price, S_t , and its spot variance, V_t , solve

$$dS_t = S_t(r_t - \delta_t + \gamma_t)dt + S_t\sqrt{V_t}dW_t^s + d\left(\sum_{n=1}^{N_t} S_{\tau_{n-}}[e^{Z_n^s} - 1]\right) - S_t\bar{\mu}_s\lambda dt \quad (1)$$

³ On average, the slope of the term structure of implied volatility is very small. In our data set, the difference in implied volatilities between 1-month and 3- to 6-month options is less than 1% in terms of Black–Scholes implied volatility.

⁴ Bondarenko (2003), Driessen and Maenhout (2004b), and Santa-Clara and Saretto (2005) document that writing puts deliver large historical returns, about 40% per month for at-the-money puts. They argue these returns are implausibly high and anomalous, at least relative to standard asset pricing models or from a portfolio perspective.

$$dV_t = \kappa_v(\theta_v - V_t)dt + \sigma_v\sqrt{V_t}dW_t^v + d\left(\sum_{n=1}^{N_t} Z_n^v\right), \quad (2)$$

where W_t^s and W_t^v are two correlated Brownian motions ($E[W_t^s W_t^v] = \rho_t$), δ_t is the dividend yield, γ_t is equity premium, N_t is a Poisson process with intensity λ , $Z_n^s | Z_n^v \sim N(\mu_s + \rho_s Z_n^v, \sigma_s^2)$ are the jumps in prices, and $Z_n^v \sim \exp(\mu_v)$ are the jumps in volatility. The SV and SVJ models are special cases, assuming that $N_t = 0$ and $Z_n^v = 0$, respectively. The general model is given in Duffie, Pan, and Singleton (DPS) (2000).⁵

DPS specify that price jumps depend on the size of volatility jumps via ρ_s . Intuitively, ρ_s should be negative, as larger jumps in prices tend to occur with larger jumps in volatility, at least if we think of events such as the crash of 1987. Eraker, Johannes, and Polson (2003) (EJP) and Chernov, Gallant, Ghysels, and Tauchen (CGGT) (2003) find negative but insignificant estimates of ρ_s . Eraker (2004), on the other hand, finds a slightly positive but insignificant estimate. This parameter is extremely difficult to estimate, even with 15 or 20 years worth of data, because jumps are very rare events.⁶ Moreover, because ρ_s primarily affects the conditional skewness of returns, μ_s and ρ_s play a very similar role. Due to the difficulty in estimating this parameter precisely and for parsimony, we assume that the sizes of price jumps are independent of the sizes of jumps in volatility. This constraint implies that the SVCJ model has only one more parameter than the SVJ model and ensures that the SVJ and SVCJ models have the same price jump distribution, which facilitates comparisons with the existing literature. We also assume a constant intensity under $\mathbb{P}$, as CGGT and Andersen, Benzoni, and Lund (ABL, 2002) find no time-series-based evidence for a time-varying intensity, and Bates (2000) finds strong evidence for misspecification in models with state-dependent intensities.

The term $-S_t \bar{\mu}_s \lambda dt$, where $\bar{\mu}_s = \exp(\mu_s + \sigma_s^2/2) - 1$, compensates the jump component and implies that γ_t is the total equity premium. It is common to assume that the Brownian contribution to the equity premium is $\eta_s V_t$, although the evidence on the sign and magnitude of η_s is mixed (see Brandt and Kang (2004)). The jump contribution to γ_t is $\lambda \bar{\mu}_s - \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}}$, where $\mathbb{Q}$ is the risk-neutral measure. If price jumps are more negative under $\mathbb{Q}$ than $\mathbb{P}$, then $\lambda \bar{\mu}_s - \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}} > 0$. The total premium is $\gamma_t = \eta_s V_t + \lambda \bar{\mu}_s - \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}}$.

The market generated by the model in (1) and (2) is incomplete, implying that multiple equivalent martingale measures exist. We follow the literature

⁵ The earliest formal model incorporating jumps in volatility is the shot-noise model in Bookstaber and Pomerantz (1989). The empirical importance of jumps in volatility is foreshadowed in Bates (2000) and Whaley (2000), who document that there are large outliers or spikes in implied volatility increments.

⁶ The small sample problem is severe. Since jumps are rare (about one or two per year), samples with 15 or 20 years of data generate relatively small numbers of jumps with which to identify this parameter. For an example, using the jump parameters in Eraker, Johannes, and Polson (2003), the finite sample distribution of ρ_s , assuming price and volatility jumps are perfectly observed, results in significant mass (about 10%) greater than zero. The uncertainty is greater in reality, as price and volatility jump sizes are not perfectly observed.

by parameterizing the change of measure and estimating the risk-neutral parameters from option prices. The change of measure or density process is given by $L_t = L_t^D L_t^J$. Following Pan (2002), we assume that the diffusive prices of risk are $\Gamma_t = (\Gamma_t^s, \Gamma_t^v) = (\eta_s \sqrt{V_t}, \eta_{\kappa_v} \sigma_v^{-1} \sqrt{V_t})$ and $L_t^D = \exp(\int_0^t \Gamma_s dW_s - \frac{1}{2} \int_0^t \Gamma_s \cdot \Gamma_s ds)$. The jump component is then

$$L_t^J = \prod_{n=1}^{N_t} \left(\frac{\lambda_{\tau_n}^{\mathbb{Q}} \pi^{\mathbb{Q}}(\tau_n, Z_n)}{\lambda_{\tau_n} \pi(\tau_n, Z_n)} \right) \exp \left(\int_0^t \left\{ \int_Z [\lambda_s \pi(s, Z) - \lambda_s^{\mathbb{Q}} \pi^{\mathbb{Q}}(s, Z)] dZ \right\} ds \right), \quad (3)$$

where $Z = (Z^s, Z^v)$ are the jump sizes or marks, π and $\pi^{\mathbb{Q}}$ are the objective and risk-neutral jump size distributions, and λ_{τ_n} and $\lambda_{\tau_n}^{\mathbb{Q}}$ are the corresponding intensities. Assuming sufficient regularity (Bremaud (1981)), L_t is a $\mathbb{P}$ -martingale, $E[L_t] = 1$, and $d\mathbb{Q} = L_T d\mathbb{P}$. By Girsanov's theorem, $N_t(\mathbb{Q})$ has $\mathbb{Q}$ -intensity $\lambda_t^{\mathbb{Q}}$, $Z(\mathbb{Q})$ has joint density $\pi^{\mathbb{Q}}(s, Z)$, and $W_t^j(\mathbb{Q}) = W_t^j(\mathbb{P}) - \int_0^t \Gamma_u^j du$ for $j = s, v$ are $\mathbb{Q}$ -Brownian motions with correlation ρ .

Measure changes for jump processes are more flexible than those for diffusions. Girsanov's theorem only requires that the intensity be predictable and that the jump distributions have common support. With constant intensities and state-independent jump distributions, the only constraint is that the jump distributions be mutually absolutely continuous (see Theorem 33.1 in Sato (1999) and Corollary 1 of Cont and Tankov (2003)). We assume that $\pi^{\mathbb{Q}}(Z^v) = \exp(\mu_v^{\mathbb{Q}})$ and $\pi^{\mathbb{Q}}(Z^s) = \mathcal{N}(\mu_s^{\mathbb{Q}}, (\sigma_s^{\mathbb{Q}})^2)$, which rules out a correlation between jumps in prices and volatility under $\mathbb{Q}$. A correlation between jumps in prices and volatility would be difficult to identify under $\mathbb{Q}$ because $\mu_s^{\mathbb{Q}}$ plays the same role in the conditional distribution of returns.

Our specification allows the jump intensity and all of the jump distribution parameters to change across measures. This is more general than the specifications considered in Pan (2002) or Eraker (2004), although, we are not able to identify all of the parameters under $\mathbb{Q}$.⁷ At first glance it may seem odd that we allow $\sigma_s \neq \sigma_s^{\mathbb{Q}}$, as prior studies constrain $\sigma_s = \sigma_s^{\mathbb{Q}}$. This constraint is an implication of the Lucas economy equilibrium models in Bates (1988) and Naik and Lee (1990), which assume power utility over consumption or wealth. While the assumptions in these equilibrium models are reasonable, the arguments above imply that the absence of arbitrage does not require $\sigma_s = \sigma_s^{\mathbb{Q}}$.

Under $\mathbb{Q}$, the equity index and its variance solve

$$dS_t = S_t(r_t - \delta_t)dt + S_t \sqrt{V_t} dW_t^s(\mathbb{Q}) + d \left(\sum_{n=1}^{N_t(\mathbb{Q})} S_{\tau_n-} [e^{Z_n^s(\mathbb{Q})} - 1] \right) - S_t \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}} dt \quad (4)$$

$$dV_t = [\kappa_v^{\mathbb{Q}}(\theta_v - V_t)V_t]dt + \sigma_v \sqrt{V_t} dW_t^v(\mathbb{Q}) + d \left(\sum_{n=1}^{N_t(\mathbb{Q})} Z_n^v(\mathbb{Q}) \right), \quad (5)$$

where $\bar{\mu}_s^{\mathbb{Q}} = \exp(\mu_s^{\mathbb{Q}} + 0.5(\sigma_s^{\mathbb{Q}})^2) - 1$. For interpretation purposes, we refer to the difference between the $\mathbb{P}$ and $\mathbb{Q}$ parameters as risk premia. Specifically, we

⁷ As we discuss later, we follow Pan (2002) and Eraker (2004) and impose $\lambda^{\mathbb{Q}} = \lambda$.

let $\mu_s - \mu_s^{\mathbb{Q}}$ denote the *mean price jump risk premium*, $\sigma_s^{\mathbb{Q}} - \sigma_s$ the *volatility of price jumps risk premium*, $\mu_v^{\mathbb{Q}} - \mu_v$ the *volatility jump risk premium*, and $\eta_v = \kappa_v^{\mathbb{Q}} - \kappa_v$ the *diffusive volatility risk premium*. Below, we generally refer to $\mu_s - \mu_s^{\mathbb{Q}}$ and $\sigma_s^{\mathbb{Q}} - \sigma_s$ together as the *price jump risk premia*. We let $\Theta^{\mathbb{P}} = (\kappa_v, \theta_v, \sigma_v, \rho, \lambda, \mu_s, \sigma_s, \mu_v)$ denote the objective measure parameters and $\Theta^{\mathbb{Q}} = (\lambda, \eta_v, \mu_s^{\mathbb{Q}}, \sigma_s^{\mathbb{Q}}, \mu_v^{\mathbb{Q}})$ denote risk-neutral parameters.

It is important to note that the absolute continuity requirement implies that certain model parameters, or combinations of parameters, are the same under both measures. This is a mild but important economic restriction on the parameters. In our model, a comparison of the evolution of S_t and V_t under $\mathbb{P}$ and $\mathbb{Q}$ demonstrates that σ_v , ρ , and the product $\kappa_v \theta_v$ are the same under both measures. This implies that these parameters can be estimated using either equity index returns or option prices, but that the estimates should be the same from either data source. One way to impose this theoretical restriction is to constrain these parameters to be equal under both measures, as advocated by Bates (2000). We impose this constraint and refer to it as *time-series consistency*.

We use options on S&P 500 futures. Under $\mathbb{Q}$, the futures price F_t solves

$$dF_t = \sigma_v F_t \sqrt{V_t} dW_t^s(\mathbb{Q}) + d \left(\sum_{n=1}^{N_t(\mathbb{Q})} F_{t_n-} (e^{Z_n^s(\mathbb{Q})} - 1) \right) - \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}} F_t dt \quad (6)$$

and the volatility evolves as in equation (5). As Whaley (1986) discusses, since we do not deal with the underlying index, dividends do not impact the results. The price of a European call option on the futures is $C(F_t, V_t, \Theta, t, T, K, r) = e^{-r(T-t)} E_t^{\mathbb{Q}}[(F_T - K)^+]$, where C can be computed in closed form up to a numerical integration. Since the S&P 500 futures options are American, we use the procedure in Appendix A to account for the early exercise feature.

B. Existing Approaches and Findings

ABL, CGGT, and EJP use index returns to estimate models with stochastic volatility, jumps in prices, and in the latter two papers, jumps in volatility. Specifically, ABL use S&P 500 returns and find strong evidence for stochastic volatility and jumps in prices. They find no misspecification in the SVJ model. CGGT use Dow Jones 30 returns and find strong evidence in support of stochastic volatility and jumps in prices, but little evidence supporting jumps in volatility. In contrast, EJP use S&P 500 returns and find strong evidence for stochastic volatility, jumps in prices, and jumps in volatility. Other approaches also find evidence for jumps in prices; see, for example, Ait-Sahalia (2002), Carr and Wu (2003), and Huang and Tauchen (2005). In conclusion, these papers agree that diffusive stochastic volatility and jumps in prices are important, but they disagree over the importance of jumps in volatility.

Similar disagreement regarding specification exists among studies that use option prices. BCC calibrate the SV and SVJ models to a cross section of S&P 500 options from 1988 to 1991. They find strong evidence for both stochastic

volatility and jumps in prices, showing that adding price jumps to the SV model reduces pricing errors by 40%, but they find that the SVJ model is misspecified. Bates (2000) uses S&P 500 futures options and finds that adding price jumps to the SV model improves fit by about 10%, but only about 2% if time-series consistency is imposed, and that all models are misspecified; he suggests adding jumps in volatility. Pan (2002) uses up to two options per day and S&P 500 index returns sampled weekly from 1989 to 1996. Her tests indicate that the SVJ model outperforms the SV model in fitting returns and for certain, but not all, strike/maturity option categories. Eraker (2004) analyzes S&P 500 options from 1987 to 1991. He finds that jumps in prices and volatility improve the time-series fit, but he finds no in-sample option pricing improvement. These mixed results are surprising in the sense that the time-series evidence overwhelmingly points toward the presence of jumps in prices. One potential explanation for these inconsistent results is that the above studies use different sample periods, cross sections, and test statistics.

Regarding factor risk premia, the evidence is again inconclusive. First, theory provides no guidance regarding the sign of the diffusive volatility risk premium. Coval and Shumway (2001) and Bakshi and Kapadia (2003) find large returns to delta-hedged option positions and use this to argue for a diffusive volatility risk premium. However, these results are also consistent with price jump or volatility jump risk premia, and as Branger and Schlag (2004) note, the tests in these papers are not powerful. Moreover, the studies that formally estimate diffusive volatility risk premia obtain conflicting results, depending on the data set and the model specification used. In the SV model, Chernov and Ghysels (2000) estimate $\eta_v = -0.001$, Pan (2002) estimates $\eta_v = -0.0301$, Jones (2003) estimates $\eta_v = 0.0326$ using data from 1987 to 2003 and $\eta_v = -0.0294$ using post-1987 data, and Eraker (2004) estimates $\eta_v = -0.01$ and reports that the parameter is marginally significant. The estimates in Jones (2003), post-1987, and the estimates in Pan (2002) imply explosive volatility under the $\mathbb{Q}$ measure ($\kappa_v^{\mathbb{Q}} < 0$). Given the well-known shortcomings of the pure SV model, the extreme variation in estimates is likely due to misspecification.

In the more reasonable SVJ model, Pan (2002) argues that η_v is insignificant, and constrains it to zero.⁸ She finds an economically and statistically significant mean price jump risk premium (18%). Eraker (2004) estimates $\eta_v = -0.01$ in the SVJ and SVCJ models, but finds that the mean price jump risk premium is insignificant. Although Eraker (2004) finds marginally significant estimates of η_v , the magnitudes are extremely small. He argues (in his figure 1), that on average volatility days the presence of diffusive volatility risk premium results in an extremely small change in the term structure of implied volatility. Even on days with very high or low volatility, the difference is at most about 1% or 2% in terms of implied volatility. Finally, Driessen and Maenhout (2004a) develop a multifactor APT-style model to quantify volatility and jump risk. They find that the diffusive volatility risk premium is statistically insignificant, the price

⁸ Interestingly, Pan (2000), an earlier version of Pan (2002), reports $\eta_v = 0.029$ in the SVJ model, which is insignificant but of the opposite sign when compared to the SV model.

jump risk premia are statistically significant, and the price jump risk premia are much larger than the diffusive volatility risk premia.

II. Our Approach

A. Consistency between Returns and Option Prices

The model in equations (1) to (5) places joint restrictions on the return and volatility dynamics under $\mathbb{P}$ and $\mathbb{Q}$. This implies, for example, that the information in returns or in option prices regarding certain parameters should be consistent across measures. Specifically, $\kappa_v\theta_v$, σ_v , and ρ should be the same under $\mathbb{P}$ and $\mathbb{Q}$. Despite the fact that the parameters should be identical under both measures, option-based estimates of certain parameters, mainly σ_v and ρ , are generally inconsistent with the time series of returns and volatility, as noted in BCC and Bates (2000). These authors find that option-based estimates of σ_v are much larger and estimates of ρ are more negative than those obtained from time-series-based estimates. This inconsistency implies either the model is misspecified, or that the data source is not particularly informative about the parameters.

In principle, an efficient estimation procedure would use both returns and the cross section of option prices over time (see Chernov and Ghysels (2000), Pan (2002), and Eraker (2004)). The advantage of such an approach is that it appropriately weighs each data source, simultaneously addressing a model's ability to fit the time series of returns and the cross section of options. However, there is a crucial drawback to this approach. Computational burdens severely constrain how much and what type of data can be used. As noted earlier, Pan (2002) and Eraker (2004) use a small number of options and short data samples.

Our approach is a pragmatic compromise between the competing constraints of computational feasibility and statistical efficiency. For the parameters that are theoretically constrained to be equal across measures, we use $\mathbb{P}$ -measure parameter estimates obtained from prior time-series studies. Then, given these parameters, we use the information embedded in options to estimate volatility and the risk-neutral parameters. This two-stage approach uses the information in a long time-series of returns and the information in the entire cross section of option prices over a long time span, and is similar to the approach used in Benzoni (2002) and Duffie, Pedersen, and Singleton (2003).

In the models that we consider, there are only three parameters that are restricted, namely, $\kappa_v\theta_v$, σ_v , and ρ . Table I summarizes the $\mathbb{P}$ -measure parameter estimates obtained by ABL, CGGT, EJP, and Eraker (2004).⁹ Although these papers use different data sets and time periods, the results are quite similar, especially for σ_v and ρ . In fact, taking into account the reported standard errors, the parameters are not statistically distinguishable. The only major difference

⁹ For the mean jump size in the SVCJ model, we use $\hat{\mu}_s = \mu_s + \rho_s\mu_v$, which is the expected jump size. EJP find that ρ_s is slightly negative, but that it is statistically insignificant. Since $\hat{\mu}_s$ does not appear under the risk-neutral measure, this does not affect our option pricing results.

Table I
Objective Measure Parameter Estimates

Objective measure parameters estimated by Eraker, Johannes, and Polson (2003), Andersen, Benzoni, and Lund (2002), Chernov et al. (2003), and Eraker (2004). The parameter values correspond to daily percentage returns. These values could be easily converted to annual decimals—another common measure—by scaling some of the parameters: for example, multiplying κ_v , and λ 252, $\sqrt{252\theta_v}/100$ gives the mean volatility, and $\sqrt{252\mu_v}/100$ gives the mean jump in volatility. In the SVCJ model, in the column labeled μ_s we report $\hat{\mu}_s = \mu_s + \rho_s\mu_v$, which is the expected jump size.

		κ_v	θ_v	σ_v	ρ	λ	μ_s (%)	σ_s (%)	μ_v
SV	EJP	0.023	0.90	0.14	−0.40	.	.	.	.
	ABL	0.016	0.66	0.08	−0.38	.	.	.	.
	CGGT	0.013	0.59	0.06	−0.27	.	.	.	.
	Eraker	0.017	0.88	0.11	−0.37	.	.	.	.
SVJ	EJP	0.013	0.81	0.10	−0.47	0.006	−2.59	4.07	.
	ABL	0.013	0.66	0.07	−0.32	0.020	0 (fixed)	1.95	.
	CGGT	0.011	0.62	0.04	−0.43	0.007	−3.01	0.70	.
	Eraker	0.012	0.83	0.08	−0.47	0.003	−3.66	6.63	.
SVCJ	EJP	0.026	0.54	0.08	−0.48	0.006	−2.63	2.89	1.48
	CGGT	0.014	0.61	0.07	−0.46	0.007	−1.52	1.73	0.72
	Eraker	0.016	0.57	0.06	−0.46	0.004	−2.84	4.91	1.25

is that CGGT and ABL's estimates of σ_v are lower, which is an expected implication given their data sets: CGGT use the Dow Jones 30 index and ABL use data from 1980 to 1996, omitting the volatile period after 1996. It is also natural to assume that there would be more variation in parameter estimates for the SV model, as it is clearly misspecified. In the case of the SVCJ model, which is the least misspecified judging by the time-series tests, the estimates of σ_v vary from 0.06 to 0.08 and the estimates of ρ vary from −0.46 to −0.48.

In our empirical implementation, we use the $\mathbb{P}$ -measure parameter estimates for $\theta_v\kappa_v$, σ_v , and ρ from EJP. First, their sample (1980 to 2000) is closest to ours (1987 to 2003). Second, they used S&P 500 returns and our options are on S&P 500 futures. Third, EJP's estimates generally have the highest σ_v and lowest ρ , which generate greater nonnormalities, and give the SV model the best chance of success. Below, we discuss the potential sensitivity of our results to the choice of $\mathbb{P}$ -measure parameter estimates.

It is easy to obtain misleading results if one ignores the theoretical restrictions that certain parameters must be consistent across measures. To see this, Figure 2 provides calibrated implied volatility curves on a representative day, August 5, 1999, placing no constraints on the parameters and minimizing the pricing errors over all strikes for the four stated maturities. This is similar to the estimation approach of BCC. For example, in the SV model, we optimize over V_t , θ_v , ρ , κ_v , and σ_v to fit the observed prices. The fits are remarkably similar across models: The root mean square errors (RMSE) of Black–Scholes implied volatility for the SV, SVJ, and SVCJ models are 1.1%, 0.6%, and 0.5%, respectively. One might be inclined to conclude that there is little, if any, benefit to the more complicated models.




Figure 2. Calibrated implied volatility curves, August 5, 1999. Parameter estimates are obtained using all four curves for each of the models, with no restrictions on the parameters. The units on the X-axis are in terms of the options' moneyness, K/F , and the units on the Y-axis are the annualized Black–Scholes implied volatility.

However, this approach ignores the fact that σ_v and ρ , should be consistent across data sources: Option-based estimates of σ_v and ρ in the SV model are grossly inconsistent with their corresponding time-series estimates. For example, the calibrated σ_v is 2.82 in the SV model, while Table I indicates that the highest reported σ_v from time-series studies is 0.14! Bates (2000), who first noted this problem, suggests constraining these parameters to be equal across measures. A simulated path using these parameters values is given in Figure 3 and shows that the option-implied parameters generate unrealistic volatility paths.

This shows that while it is possible, as a curve-fitting exercise, to make the SV model fit the market data, the resulting parameter estimates are inconsistent with the requirement of absolute continuity. Forcing a misspecified model to fit observed prices is particularly dangerous if, as is commonly the case, the fitted parameters are then used to price or hedge other derivatives. The misspecification is also important for risk premium estimation. Much of the literature documenting volatility risk premia finds these premia in the context of pure stochastic volatility models. As Anderson, Hansen, and Sargent (2003) note, model misspecification can appear in the form of a risk premium. Thus, it is important to be cautious when estimating and interpreting risk premia in poorly specified models.




Figure 3. Simulated volatility paths. This graph provides volatility paths simulated based on options ($\theta_v = 3.63$, $\kappa_v = 0.06$, $\sigma_v = 2.8$, $\rho = -0.66$), and index returns ($\theta_v = 0.90$, $\kappa_v = 0.025$, $\sigma_v = 0.15$, $\rho = -0.40$). The time corresponds to 2 years (500 trading days) and the same Brownian increments are used for both paths to allow for a direct comparison.

Figure 4 repeats the previous exercise constraining $\kappa_v\theta_v$, σ_v , and ρ to be equal to the estimates obtained in EJP (see Table I). The RMSEs for the SV, SVJ, and SVCJ models are now 8.73%, 2.97%, and 1.43%, respectively, and we see that the SV model does an extremely poor job. Also, the SVJ model has pricing errors roughly twice as large as the SVCJ model. The SV model does poorly because, once time-series consistency is imposed, it cannot generate sufficient amounts of conditional skewness and kurtosis.¹⁰

B. Time-Series Tests

Option prices are highly informative about spot volatility. In this section, we develop an intuitive test to detect volatility jumps. Our approach is similar in spirit to those implemented in Pan (2002), Johannes (2004), and Jones

¹⁰ The constraint on $\kappa_v\theta_v$ has little effect as the long-run level of volatility and the speed of mean reversion are both second-order effects on options prices and implied volatilities over the maturities for which we have data.




Figure 4. Calibrated Black-Scholes implied volatility curves, August 5, 1999. Parameter estimates are obtained using all four curves for each of the models, constraining the $\mathbb{P}$ -measure parameters to be equal to their time-series counterparts.

(2003) in that we focus on higher-moment behavior to diagnose jump-induced misspecification.

We use the following internally consistent procedure. In the first stage, given model parameters and option contract variables, we invert spot volatility from a representative at-the-money call option for every day in our sample. This provides a time series of model-implied spot variances, $\{V_t^{imp}\}_{t=1}^T$. These variances differ from Black-Scholes implied variance, as the model-based variance takes into account, for example, jumps or mean reversion in volatility. Given the implied variances, we compute the skewness and kurtosis, which are standard measures of tail behavior. As the models that we consider have state-dependent diffusion coefficients, we focus on “conditional” skewness and kurtosis using the standardized increment:¹¹

$$V_{skew} = skew^{\mathbb{P}} \left(\frac{V_{t+1} - V_t}{\sqrt{V_t}} \right) \quad \text{and} \quad V_{kurt} = kurt^{\mathbb{P}} \left(\frac{V_{t+1} - V_t}{\sqrt{V_t}} \right). \quad (7)$$

Unconditional measures of skewness and kurtosis provide the same conclusions. However, given the persistence and heteroscedasticity of volatility, the

¹¹ The motivation for the conditional measures is that if V_t follows a square root process, then $\int_t^{t+1} \sqrt{V_s} dW_s^v \approx N(0, V_t)$ and $(V_{t+1} - V_t)/\sqrt{V_t}$ is approximately normally distributed.

conditional statistics are likely to have greater power for detecting misspecification. To highlight the importance of jumps in prices, we also report the skewness and kurtosis of returns conditional on volatility, as the distribution of returns will have first-order importance on the cross section of option prices. These measures are defined as $R_{skew} = skew^{\mathbb{P}}(R_{t+1}/\sqrt{V_t})$ and $R_{kurt} = kurt^{\mathbb{P}}(R_{t+1}/\sqrt{V_t})$. We refer to these conditional measures merely as skewness or kurtosis, omitting the conditional modifier.

Pritsker (1998) and Conley, Hansen, and Liu (1997) find that asymptotic approximations are unreliable when the data are highly persistent and recommend a Monte Carlo or bootstrapping approach. Accordingly, we follow their recommendation and simulate $G = 1,000$ sample paths from the null model, $\{V_t^g\}_{t=1}^T$ for $g = 1, \dots, G$, and then compute each of the statistics for each path.

To implement this procedure, we use the $\mathbb{P}$ -measure parameters estimated from returns; specifically, we use those from EJP. We also perform sensitivity analysis by varying the parameters that control the tail behavior of the volatility process. Our conclusions regarding the misspecification of the square root volatility process hold for any set of parameters in Table I. We also document that our conclusions are not sensitive to reasonable risk premia, as we compute the statistics for the risk premia that we later estimate.

C. Estimating Pricing Errors and Risk Premia

We next focus on the information embedded in the cross section of option prices. Our goal is to understand how the misspecification manifests in option prices and to estimate risk premia. Given our flexible risk-premium specifications, μ_s , σ_s , and μ_v do not enter into the option pricing formula. Thus, the only parameters that we use from the returns-based data are λ , $\theta_v \kappa_v$, σ_v , and ρ . As we mention earlier, our conclusions regarding the relative merits of the models do not depend on the choice of the $\Theta^{\mathbb{P}}$ parameters because these parameters, once constrained to be consistent with the objective measure, have very little impact on option prices. For example, our conclusion that the SVJ and SVCJ models outperform the SV models holds for all of the parameters reported in Table I.

To estimate parameters and variances, we minimize squared differences of model and market Black–Scholes implied volatilities, that is,

$$(\hat{\Theta}^{\mathbb{Q}}, \hat{V}_t) = \arg \min \sum_{t=1}^T \sum_{n=1}^{O_t} [IV_t(K_n, \tau_n, S_t, r) - IV(V_t, \Theta^{\mathbb{Q}} | \Theta^{\mathbb{P}}, K_n, \tau_n, S_t, r)]^2, \quad (8)$$

where T is the number of days in our sample, O_t is the number of cross-sectional option prices observed on date t , $IV_t(K_n, \tau_n)$ is the market-observed Black–Scholes implied volatility for strike K_n and maturity τ_n , and $IV(V_t, \Theta^{\mathbb{Q}} | \Theta^{\mathbb{P}}, K_n, \tau_n, S_t, r)$ is the Black–Scholes implied volatility of the model price. The implied volatility metric provides an intuitive weighting of options across strikes and maturities. In contrast, minimizing squared deviations

between model and market option prices places greater weight on expensive in-the-money and long-maturity options. Indeed, others advocate discarding all in-the-money options for this reason (Huang and Wu (2004)). Christoffersen and Jacobs (2004) provide a detailed discussion of the objective function choice.

The second component in the objective function is the choice of option contracts, that is, K_n and τ_n . Since it is not possible to observe traded option prices of all strikes and maturities simultaneously, there are two ways to construct a data set, namely, to use close prices or to sample options over a window of time. We follow Bates (2000) and aggregate trades during the day. Bates (2000) chooses a 3-hour window, and we extend this window to the entire day. Since we identify diffusive volatility and price and volatility jump risk premia from longer-dated options and deep out-of-the-money (OTM) options, respectively, it is important to include as many of these as possible.

Because there are hundreds or thousands of option transactions each day, using all of them generates a number of issues. For example, the vast majority of the recorded trades in our sample involve short-maturity at-the-money (ATM) options. Equal weighting of all trades would effectively overweight the information from short-maturity ATM options, which are less interesting as all models provide similar ATM prices. As we outline in Appendix B, we take all daily transactions, fit a flexible parametric curve, and then use the interpolated curve in the objective function. It is common to perform interpolation for data reduction purposes (see also Bliss and Panigirtoglou (2004) and Huang and Wu (2004)). We view this approach as a pragmatic compromise, as it uses nearly all of the information in the cross section of option prices without, in our opinion, introducing any substantive biases.

Our approach jointly estimates V_t and Θ^Q using the cross section of option prices. Thus, if a model is poorly specified, our estimation procedure could generate implausible estimates of V_t or Θ^Q . For example, the arguments in Section II.A indicate that the SV model, once constrained to be consistent with the time series, provides a very poor fit to the cross section. Additionally, from Figure 4, it is clear that spot volatility in the SV model is higher than expected, as the estimation procedure increases spot volatility in an attempt to find the best fit including non-ATM options. This explains why it is important to be careful interpreting volatility or risk premia estimates in a model that is clearly misspecified, based on, for example, time-series evidence.

Another issue with cross-sectional estimation relates to assessing statistical significance. As Bates (2000, p. 195) notes, "A fundamental difficulty with implicit parameter estimation is the absence of an appropriate statistical theory of option pricing errors." This means, in particular, that it is difficult to assign standard errors to parameters estimated using the cross section. We overcome this difficulty using a computationally expensive, but intuitive, nonparametric bootstrapping procedure. We randomly select 40 Wednesdays from 1987 to 2003 and estimate the risk premia and spot variances. We then repeat this procedure (reestimating risk premia and spot variances) until the bootstrapped standard errors do not change appreciably. We find that 50 replications are sufficient (consistent with Efron and Tibshirani (1994)). The reported point

estimates are averages across the 50 replications. We report RMSE between the model fit and our interpolated implied volatility curves across all of the 50 replications, which provides a large sample of option transactions. We also report an “out-of-sample” experiment by randomly selecting 50 days, reestimating the spot variance (holding the risk premia estimates constant), and evaluating RMSE’s.

III. Empirical Results

A. Time-Series Specification Tests

To implement the time-series-based tests, we use a representative option price to compute a model-based estimate of V_t . We select a representative daily option price that (1) is close to maturity (to minimize the American feature), (2) is at the money, (3) is not subject to liquidity concerns, (4) is an actual transaction (not recorded at the open or the close of the market), (5) has a recorded futures transaction occurring at the same time, and (6) is a call option (to minimize the impact of the American early exercise feature). Appendix B describes the procedure in greater detail.

We also report the summary statistics using the VIX index and ATM implied volatilities extracted from daily transactions using our interpolation scheme (see Whaley (2000) for a description of the VIX index). Although not reported, we also compute all of the statistics using a sample of put options, and none of our conclusions change. We adjust the options for the American feature, as described in Appendix A. We use interpolated Treasury bill yields as a proxy for the risk-free rate.

Table II summarizes the implied volatilities and scaled returns for the different data sets and models. In the first panel, the first two rows labeled “VIX” provide summary statistics for the VIX index (including and excluding the crash of 1987), the rows labeled “Calls” provide statistics for our representative call option data set, and the rows labeled “Interpolated” use the ATM implied volatility interpolated from all of the daily transactions.¹² In this first panel, the implied volatility is based on the Black–Scholes model (BSIV) and the subsequent panels report model-based (as opposed to Black–Scholes) implied variances.

Although there are some *quantitative* differences across data sets, the *qualitative* nature of the results is unchanged: We observe large positive skewness and excess kurtosis in the variance increments and negative skewness and positive excess kurtosis in standardized returns. The minor variations across the data sets are due to differences in underlying indices (the VIX is based on the S&P 100 index) and in the timing and nature of the price quotes (the VIX is based on close prices, the calls are actual transactions in the morning, and the interpolated set averages all transactions in a given day).

¹² Formerly, the VIX was calculated from S&P 100 options instead of S&P 500 options. As of September 22, 2003, the VIX uses options on the S&P 500 index. We use the old VIX index (current ticker VXO).

Table II
Volatility and Return Summary Statistics

The first three rows provide summary statistics for variance increments and standardized returns using the VIX index, a time series of call option implied volatility (see Appendix B), and the ATM interpolated implied volatility (see Appendix B). In these three cases, the variance used is that from the Black–Scholes model. The second, third, and fourth panels contain model implied variances for the SV, SVJ, and SVCJ models assuming options are priced based on the objective measure. We also include risk premiums (RP) and document the effect of increasing σ_v in the SV model.

	Period	V_{kurt}	V_{skew}	R_{kurt}	R_{skew}
VIX (BSIV)	1987 to 2003	2996.58	50.41	13.72	-1.02
	1988 to 2003	20.93	1.74	5.69	-0.43
Calls (BSIV)	1987 to 2003	1677.16	32.78	22.99	-1.46
	1988 to 2003	15.17	1.25	5.64	-0.40
Interpolated (BSIV)	1987 to 2003	2076.58	38.21	21.04	-1.38
	1988 to 2003	25.17	1.79	5.82	-0.43
SV Model	1987 to 2003	1035.71	23.85	17.39	-1.18
	1988 to 2003	14.33	1.29	6.04	-0.41
SV Model (RP)	1987 to 2003	907.57	21.85	17.87	-1.20
	1988 to 2003	13.44	1.25	5.74	-0.40
SV Model ($\sigma_v = 0.2$)	1987 to 2003	1039.98	23.91	17.97	-1.22
	1988 to 2003	14.41	1.29	6.10	-0.43
SVJ Model	1987 to 2003	850.41	21.16	17.75	-1.20
	1988 to 2003	15.66	1.37	6.10	-0.42
SVJ Model (RP)	1987 to 2003	1048.51	24.21	15.91	-1.06
	1988 to 2003	16.01	1.40	7.09	-0.42
SVCJ Model	1987 to 2003	1015.13	23.62	16.62	-1.12
	1988 to 2003	15.16	1.34	6.31	-0.40
SVCJ Model (RP)	1987 to 2003	546.52	16.08	15.96	-1.02
	1988 to 2003	13.44	1.38	6.77	-0.35

For the formal tests, we use the call option data set. Our conclusions are the same using the other data sets, although the call option data have fewer issues (interpolation, averaging effects, stale quotes, etc.). The bottom three panels in Table II report statistics using model-based implied variances for the SV, SVJ, and SVCJ models using three sets of parameters. The first set is from EJP who, as we mention earlier, report higher σ_v estimates than other papers. Because the parameter σ_v primarily controls the kurtosis of the volatility process, this set of parameters gives the SV and SVJ models the best chance of success. For robustness, we include statistics using two additional parameter sets. The results in the rows labeled “RP” incorporate risk premia in order to gauge their impact on implied variances.¹³ The third set of results uses the SV model and $\sigma_v = 0.20$, which is roughly five standard deviations away from the point estimate in EJP in the SV model.

Table III provides quantiles of the finite sample distribution for each of the statistics using the Monte Carlo procedure described in the previous section

¹³ We use the risk premium values estimated later in the paper.

Table III
Statistics' Finite Sample Distribution

For each model and set of parameters, we report the appropriate quantiles from the statistics' finite sample distribution. The base parameters are taken from Eraker, Johannes, and Polson (2003) as reported in Table I.

	Quantile	V_{kurt}	V_{skew}	R_{kurt}	R_{skew}
SV model	0.50	3.27	0.34	3.02	-0.05
	0.95	3.51	0.41	3.14	-0.10
	0.99	3.67	0.43	3.19	-0.12
SV model $\sigma_v = 0.2$	0.50	3.55	0.48	3.05	-0.06
	0.95	3.96	0.55	3.16	-0.12
	0.99	4.26	0.60	3.23	-0.14
SVJ model	0.50	3.05	0.15	22.05	-1.48
	0.95	3.23	0.22	106.05	-5.07
	0.99	3.34	0.26	226.77	-8.66
SVCJ model	0.01	261.02	9.94	7.73	-0.63
	0.05	320.24	13.18	10.72	-0.92
	0.50	615.40	21.04	24.77	-1.91
SVCJ model $\mu_v = 0.85$ $\lambda = 0.0026$	0.95	1649.03	34.87	78.90	-4.15
	0.99	2500.51	43.73	175.67	-6.66
	0.01	3.28	0.20	3.02	0.06
SVCJ model $\mu_v = 0.85$ $\lambda = 0.0026$	0.05	13.21	1.01	3.21	-0.02
	0.50	217.70	8.98	7.13	-0.46
	0.95	1150.76	27.53	37.92	-2.20
	0.99	2012.62	39.34	94.16	-4.11

and for each of the model-parameter configurations. Note that these results are simulated under the $\mathbb{P}$ -measure. Thus, there are no separate entries for the cases incorporating risk premia, as the $\mathbb{P}$ -measure behavior does not change.

The SV model cannot generate enough positive skewness or excess kurtosis to be consistent with the observed data. For example, the model generates $V_{kurt} = 3.67$ at the 99th quantile, which is orders of magnitude lower than the value observed in the data (around 1,000). Similarly, the SV model cannot generate the large positive skewness observed in the data. We also note in passing that, not surprisingly, the SV model cannot come close to generating the observed nonnormalities in returns either.

Before concluding that the SV model is incapable of capturing the behavior of option implied volatility, it is important to document that our results are robust. To do so, we show that the results are unchanged even if we ignore the crash of 1987, we account for volatility risk premia, or we increase σ_v . The rows labeled "1988 to 2003" in Table II provide the statistics from 1988 to 2003, a period excluding the crash of 1987.¹⁴ Based on the post-1987 sample, the SV

¹⁴ Of course, we do not advocate "throwing out" data, especially outlier events in jump models. Since jumps are rare, these tail observations are invaluable for characterizing jumps. However, in this setting, the post-1987 sample highlights the severe problems with the square root process in the SV and SVJ models.

model is still incapable of generating these observed statistics, even though the parameters are estimated including the crash. If the SV model were estimated using post-1987 data, it is very likely that θ_v , κ_v , and σ_v would be lower, which implies that the model generates even less nonnormalities. The conclusion is unchanged even if we increase σ_v to 0.20.

Finally, the row labeled “SV model (RP)” in Table II indicates that the results are robust to realistic risk premia. Diffusive volatility and volatility jump risk premia change the level and speed of mean reversion in volatility, which can have a significant impact on implied variance in periods of very high volatility (e.g., October 1987). Risk premia, however, cannot explain the nonnormalities in the observed data. This is most clear in the post-crash subsample, in which risk premia have a minor impact. Thus, we can safely conclude that the SV model is incapable of capturing the observed behavior of option prices.

In the SV model, volatility increments over short time intervals,

$$V_{t+1} - V_t \approx \kappa_v(\theta_v - V_t) + \sigma_v\sqrt{V_t}(W_{t+1}^v - W_t^v), \quad (9)$$

are approximately conditionally normal (see also Table III). The data, however, are extremely nonnormal, and thus the square root diffusion specification has no chance to fit the observed data.

The following example provides the intuition by way of specific magnitudes. Consider the mini-crash in 1997: On October 27th the S&P 500 fell about 8% with Black–Scholes implied volatility increasing from 26% to 40%. In terms of variance increments, daily variance increased from 2.69 to 6.33, which translates into a standardized increment, $(V_t - V_{t-1})/\sqrt{V_t}$, of 2.22. To gauge the size of this move, it can be compared to the volatility of standardized increments over the previous 3 months, which was 0.151 (remarkably close to the value $\sigma_v = 0.14$ used above). Thus, the SV would require a roughly 16-standard deviation shock to generate this move. This example shows the fundamental incompatibility of the square root specification with the observed data. What we have here is not an issue of finding the right parameter values; rather, the model is fundamentally incapable of explaining the observed data. Whaley (2000) provides other examples of volatility “spikes.”

The SV and SVJ models share the same square root volatility process, suggesting the SVJ model is also incapable of fitting the observed data. The third panel in Table III indicates that it does generate different implied variances (due to the different volatility parameters and jumps), but it cannot generate the observed skewness or kurtosis. Subsamples or risk premia do not change the conclusion. Since the SV and SVJ models share the same volatility process, the conclusions are unchanged with $\sigma_v = 0.20$. The SVJ model can generate realistic amounts of skewness and kurtosis in returns. This should not be surprising as the jumps generate the rare, large negative returns observed in prices.

The lower panels in Tables II and III demonstrate that the SVCJ model is capable of capturing both the behavior of implied variances and standardized returns for the full sample. In Table III, the panel reports the 1st, 5th, 50th, 95th, and 99th quantiles of each of the statistics. For example, V_{kurt} based on

option prices is about 1,000, and the sample statistics fall somewhere between the 50th (about 600) and 95th quantile (1,600). The skewness generated by the model is almost identical to the value observed in the data; with the 50th quantile equal to 21.04 in simulations, compared to 21.16 in the data. The infrequent, exponentially distributed jumps in volatility naturally generate the combination of high kurtosis and positive skewness observed in the data. The model can also capture the conditional distribution of returns. The final rows in Table II indicate that the conclusions are unchanged if we include diffusive volatility and jump risk premia. The kurtosis in the SVCJ model falls in the full sample with risk premia because the jump premia (volatility and jumps in prices) alter the model-implied variances, especially during the crash of 1987.

A comparison of the subsamples in Table II with the quantiles generated by the SVCJ models using the base parameters indicates that over the post-1987 period, the SVCJ model with parameters from EJP generates too much kurtosis and skewness. This can be seen by comparing, for example, V_{kurt} from 1988 to 2003 in Table II with the 1st quantile for the SVCJ model in Table III. This result is not at all surprising since the base case parameters in EJP are estimated using data that include the crash of 1987: Because jumps are rare events that generate conditional nonnormalities, if one removes these outliers the observed data (variance increments or standardized returns) become more normal, by construction. Of course, this does not indicate that the SVCJ model is misspecified, but rather when using the parameters estimated using the full sample, that it generates too much excess kurtosis and skewness. Naturally, if the SVCJ model were estimated omitting the crash of 1987, it is likely that the parameters governing the jump sizes would change.

To document that this is not a generic problem with the SVCJ model, the final panel in Table III shows that if one reduces the jump intensity to $\lambda = 0.0026$ and $\mu_v = 0.85$ (about two standard errors below the point estimates in EJP), the model fits both periods within a 5% to 95% confidence band.¹⁵ Note that we do not suggest using these ad hoc parameters; we simply use them to illustrate the flexibility of the SVCJ model. The key point is to contrast this result with those that obtain using the SV and SVJ models: These models, due to the diffusion specification, could not fit the data over *either* of the samples, even using a value of σ_v that is implausibly high.

We conclude that the SV and SVJ models are incapable of capturing the time-series behavior of option implied variances, while the SVCJ model can easily capture the observed behavior. Since a primary goal of this paper is to estimate risk premia, it is important to have a well-specified model as model misspecification can easily distort risk premia estimates. Finally, our results are related to Jones (2003), who shows that the square root and constant elasticity of variance models cannot explain the dynamics of implied volatility. We confirm

¹⁵ This value of λ is not implausible as it is consistent with studies that use shorter samples. For example, Pan's (2002) estimates imply about 0.3 jumps per year on average ($\lambda = 0.0012$) and Eraker (2004) finds about 0.5 jumps per year ($\lambda = 0.002$).

Table IV
Risk-Neutral Parameter Estimates

For each parameter and model, the table gives the point estimate, computed as the average parameter value across 50 bootstrapped samples, and the bootstrapped standard error. For the SVJ and SVCJ models, an entry of σ_s in the σ_s^Q column indicates that we impose the constraint $\sigma_s = \sigma_s^Q$.

	η_v	μ_s^Q (%)	σ_s^Q (%)	μ_v^Q	RMSE (%)
SV	0.005 (0.07)	—	—	—	7.18
SVJ	0.010 (0.03)	−9.97 (0.51)	σ_s	—	4.08
SVJ	0.006 (0.02)	−4.91 (0.36)	9.94 (0.41)	—	3.48
SVJ	0	−9.69 (0.58)	σ_s	—	4.09
SVJ	0	−4.82 (0.33)	9.81 (0.58)	—	3.50
SVCJ	0.030 (0.21)	−6.58 (0.53)	σ_s	10.81 (0.45)	3.36
SVCJ	0.031 (0.18)	−5.39 (0.40)	5.78 (0.70)	8.78 (0.42)	3.31
SVCJ	0	−7.25 (0.50)	σ_s	5.29 (0.18)	3.58
SVCJ	0	−5.01 (0.38)	7.51 (0.83)	3.71 (0.22)	3.39

Jones's (2003) findings, and, in addition, we provide a model with jumps in volatility that is capable of capturing the observed dynamics.

B. Model Specification and Risk Premium Estimates

Table IV provides risk premium estimates and overall option fit for each model, and Table V evaluates the significance of the model fits. We estimate the risk-neutral parameters and compute total RMSE using the procedure in Section II.C based on the cross section of option prices. While the time-series results above indicate that the SV and SVJ models are misspecified, we still report cross-sectional results for these models in order to quantify the nature of the pricing improvement in the SVCJ model and to analyze the sensitivity of the factor risk premia to model misspecification.

We discuss model specification and risk premia in turn. Our procedure generates the following intuitive metric for comparing models. We compute the number of bootstrapped samples for which the overall pricing error (measured by the relative difference between BSIVs) for one model is lower than another by 5%, 10%, 15%, or 40%.

Because it is difficult to statistically identify η_v , we report the results for versions of the models with or without a diffusive volatility risk premium. Throughout this section, we constrain λ^Q to be equal to λ and use the value from EJP. In general, it is only possible to estimate the compensator, $\lambda^Q \bar{\mu}_s^Q$, and not the individual components separately. Pan (2002) and Eraker (2004) imposed the same constraint. The estimate in EJP implies approximately 1.5 jumps per year, which is higher than most other estimates and will result in conservative price jump and volatility jump risk premia estimates.

Finally, for the jump models, we consider cases that depend on whether or not σ_s is equal to σ_s^Q . This is the first paper that allows these parameters to differ. We believe it is important to document the size of this risk premium and how it affects the estimates of the other premia.

Table V
Model Comparison Results

Comparison of the RMSEs across models. The table reads as follows: The probability that model [name in a row] is better than model [name in a column] by [number in a row]% is equal to [number in the intersection of the respective row and column]. The numbers in parentheses are out-of-sample values. For example, the probability that the RMSE of the SVCJ model is smaller than the RMSE of the SVJ ($\sigma_s^Q = \sigma_s$) model by 10% is 0.76.

		SV	SVJ $\sigma_s^Q \neq \sigma_s$	SVJ $\sigma_s^Q = \sigma_s$	SVCJ $\sigma_s^Q = \sigma_s$
SV	5%		0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
	10%		0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
	15%		0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
	40%		0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
SVJ	5%	1.00 (1.00)		1.00 (1.00)	0.26 (0.66)
	10%	1.00 (1.00)		0.98 (0.98)	0.04 (0.36)
	15%	1.00 (1.00)		0.42 (0.66)	0.00 (0.16)
	40%	1.00 (1.00)		0.00 (0.02)	0.00 (0.02)
SVJ	5%	1.00 (1.00)	0.00 (0.00)		0.00 (0.02)
	10%	1.00 (1.00)	0.00 (0.00)		0.00 (0.00)
	15%	1.00 (1.00)	0.00 (0.00)		0.00 (0.00)
	40%	0.70 (0.70)	0.00 (0.00)		0.00 (0.00)
SVCJ	5%	1.00 (1.00)	0.06 (0.00)	0.96 (0.80)	
	10%	1.00 (1.00)	0.00 (0.00)	0.76 (0.36)	
	15%	1.00 (1.00)	0.00 (0.00)	0.26 (0.00)	
	40%	0.96 (0.96)	0.00 (0.00)	0.00 (0.00)	

B.1. Model Specification: Pricing Errors

A number of points emerge from Tables IV and V. Regardless of the assumptions on the risk premium parameters, the SVJ and SVCJ models provide significant pricing improvement over the SV model. This is true in a point-wise sense in Table IV, as the RMSEs of the pricing errors are reduced by almost 50%, and is also true based on the bootstrapped samples. Under any of the risk premia assumptions and in all 50 bootstrapped samples, the SVJ and SVCJ models provide at least a 15% pricing error improvement over the SV model, and in most samples, more than a 40% improvement. These results are consistent with BCC (who find a 40% improvement) and in contrast to Bates (2000), Eraker (2004), and Pan (2002). The reasons for our clear results are twofold: We impose time-series consistency, and we use option prices that span a long time period.

Next, consider a comparison of the overall pricing errors in the SVJ and SVCJ models in Table IV. The SVJ model, with no constraints on risk premia, has an overall pricing error of 3.48%. By comparison, the unconstrained SVCJ model has pricing errors of 3.31%, which is an improvement of 5%. If we impose $\sigma_s^Q = \sigma_s$ in the SVJ model, the SVCJ model generates a larger improvement of 18% (4.08% vs. 3.36%). If $\sigma_s^Q \neq \sigma_s$, σ_s^Q increases drastically, allowing the SVJ model to generate more conditional kurtosis in returns, a role very similar to

that played by jumps in volatility. In fact, a comparison of the unconstrained SVJ model and the SVCJ model with the same number of parameters ($\sigma_s^Q = \sigma_s$ in the SVCJ model) shows that the pricing errors are quite similar: 3.48% and 3.36%, respectively. This result is consistent with objective measure time-series results in Table III, which show that the SVJ model can generate realistic amounts of nonnormalities in returns through jumps in prices.

In Table V, we test for differences for the most important variations of the models. We compare the SV model, the SVJ model with and without constraints on σ_s^Q , and the SVCJ model assuming $\sigma_s^Q = \sigma_s$. The constrained SVCJ model is of interest for model fit because the unconstrained SVJ and constrained SVCJ models have the same number of parameters. The SVCJ ($\sigma_s^Q = \sigma_s$) model has pricing errors that are at least 5% lower than those of the SVJ ($\sigma_s^Q = \sigma_s$) model in 96% of the bootstrapped replications. Thus, the SVCJ model provides a statistically significant improvement in overall option fit. However, a comparison of the unconstrained SVJ model with the constrained SVCJ model indicates that the SVCJ ($\sigma_s^Q = \sigma_s$) model outperforms by more than 5% in only 6% of the trials, an indication that SVCJ adds little to the cross section of returns when the risk premia are unconstrained in SVJ. In order to ensure the robustness of our findings, we also evaluate “out of sample” RMSEs by using the parameter values reported in Table IV on 50 randomly selected days in our sample that were not used for risk premia estimation. For each of the 50 days, we estimate the spot variance and then computed RMSEs, holding the risk premia constant. The results are qualitatively the same, with a slight improvement in the unconstrained SVJ model relative to the SVCJ model.¹⁶

We sort option pricing errors by maturity, strike, and volatility. The only major pattern that emerges is that pricing errors tend to be higher for all models in periods of high volatility. Given that our objective function focuses on absolute differences in volatility (as opposed to percentage differences), this is not a surprise. Pan (2002) finds a similar pattern. The SVJ and SVCJ models outperform the SV model in all categories, however, there is little systematic difference between the SVJ and SVCJ models.

We conclude that there is some in-sample pricing improvement by including jumps in volatility, but the effect is modest. However, as our tests in the previous section indicate, the SVJ model cannot capture the dynamics of V_t . Thus, the SVCJ is our preferred model as it is consistent simultaneously with the time series and the cross section.

B.2. Estimates of Risk Premia

Table IV summarizes the $\mathbb{Q}$ -measure parameter estimates. For each parameter and specification, the table provides estimates and bootstrapped standard errors based on 50 replications. A number of interesting findings are apparent.

¹⁶ In the interest of saving space, we do not report the unconstrained SVCJ model results. However, the unconstrained SVCJ model modestly outperforms all of the other models both in and out of sample.

Table VI
The Average Term Structure of S+P 500 Index Implied Volatility

For each year, we compute average at-the-money implied volatility from our Black–Scholes implied volatility curves as computed in Appendix B and bin the results into three categories: options that mature in under 1 month, those that mature in 1 to 2 months, and those that mature in 3 to 6 months.

Year	1 Month	2 Month	3–6 Months
1987	28.01	28.96	21.62
1988	21.67	22.30	22.39
1989	14.79	15.35	15.90
1990	19.50	19.99	20.22
1991	16.12	16.55	17.02
1992	13.14	13.67	14.62
1993	10.81	10.99	12.06
1994	11.58	11.84	12.98
1995	10.44	10.62	11.55
1996	14.34	14.37	14.64
1997	20.06	20.24	20.16
1998	21.20	22.33	23.57
1999	21.33	22.18	24.02
2000	20.63	20.32	21.36
2001	23.69	23.43	21.82
2002	25.20	24.60	23.25
2003	28.72	28.09	26.27
Mean	18.90	19.17	19.03

The diffusive volatility risk premium η_v is insignificant in every model. As mentioned earlier, there are several reasons to believe that this parameter is difficult to identify, and our finite-sample procedure generates quite large standard errors relative to the point estimate. This does not necessarily mean that $\eta_v = 0$; instead it suggests only that we cannot accurately estimate the parameter. When we constrain $\eta_v = 0$, there is virtually no change in the RMSE, which indicates that its impact on the cross section is likely to be minor.

Why is it so hard to estimate η_v , even with a long data set? The main reason is that, as shown in Table VI, the implied volatility term structure is flat, at least over the option maturities that we observe. Focussing on the whole sample, the difference between short- (less than one month) and longer-dated (3 to 6 months) implied volatility is only about 0.1%.¹⁷ In the context of our models, it is important to understand what could generate such a flat term structure. Since jumps in prices contribute a constant amount to expected average variance over different horizons, any variation in the term structure shape will arise from the stochastic volatility component.

¹⁷ According to Bates (1996), the bias arising from Jensen's inequality from the use of implied volatilities rather than implied variances is less than 0.5% for 1-month to 12-month at-the-money options. Hence, the flat term-structure observation is not likely due to Jensen's inequality.

The stochastic volatility model could generate a flat average term structure via two conduits. First, η_v could be small, implying that the term structure is flat on average. Second, η_v could be large (of either sign), but the term structure could still be flat over short horizons if risk-neutral volatility is very persistent. Given the near unit root behavior of volatility under $\mathbb{P}$, volatility will also be very persistent under $\mathbb{Q}$ for a wide range of plausible η_v s, generating a flat term structure of implied volatility. This implies that we can only distinguish between these two competing explanations if we have long-dated options.

It is also the case that using a more efficient estimation procedure, such as one including returns and option prices, would improve the accuracy of the parameter and risk premia estimates. However, joint estimation must still confront the fact that the term structure is flat, which implies that merely using a different estimation procedure is not likely to alleviate the problem that η_v is insignificant, unless the procedure incorporates long-dated options. In large part, this explains why the existing literature using options and returns finds unstable, insignificant, or economically small estimates.

We find evidence for modest but highly significant jump risk premia in the SVJ and SVCJ models, as the information in the volatility smile allows us to accurately estimate $\mu_s^{\mathbb{Q}}$, $\sigma_s^{\mathbb{Q}}$, and $\mu_v^{\mathbb{Q}}$. In the SVJ model, $\mu_s^{\mathbb{Q}}$ ranges from about -9% (imposing $\sigma_s^{\mathbb{Q}} = \sigma_s$) to about -5% ($\sigma_s^{\mathbb{Q}} \neq \sigma_s$). In the SVJ and SVCJ models, estimates of μ_s are generally around -2% to -3% based on the time-series of returns, which implies a modest mean price jump risk premium of about 2% to 6% . In the SVCJ model, the estimates of $\mu_s^{\mathbb{Q}}$ are again significant and generate a risk premium of similar magnitude, about 2% to 3% when ($\sigma_s^{\mathbb{Q}} \neq \sigma_s$) and 4% to 5% when ($\sigma_s^{\mathbb{Q}} = \sigma_s$). When $\sigma_s^{\mathbb{Q}} \neq \sigma_s$, the SVJ and SVCJ models deliver remarkably consistent results: The estimates of $\mu_s^{\mathbb{Q}}$ vary from a low of -5.39% to a high of -4.82% . The risk premium estimates do not appear to depend on whether or not jumps in volatility are present, and thus are a robust finding. A significant mean price jump risk premium should not be a surprise since jumps cannot be perfectly hedged with a finite number of instruments.

In both the SVJ and SVCJ models, there is strong evidence that $\sigma_s^{\mathbb{Q}} \neq \sigma_s$, an effect that has not previously been documented. Moreover, as we note in the previous paragraph, this has important implications for the magnitudes of the premium attached to the mean price jump size. Estimates of σ_s in the SVJ model are around 4% , while estimates of $\sigma_s^{\mathbb{Q}}$ are more than 9% . However, it appears that this premium is largely driven by specification. As mentioned earlier, the time-series tests indicate the presence of jumps in volatility. These jumps generate large amounts of excess kurtosis in the distribution of returns. Since the SVJ model does not allow volatility to jump, it can only fit observed option prices by drastically increasing $\sigma_s^{\mathbb{Q}}$ to create a large amount of risk-neutral kurtosis. When jumps in volatility are allowed in the SVCJ model, estimates of $\sigma_s^{\mathbb{Q}}$ fall to about 6% (with a standard error of 0.7%) in the unconstrained SVCJ model and to around 7.5% (with a standard error of 0.83%) when η_v is constrained to be zero. Unlike the significant estimates of $\mu_s^{\mathbb{Q}}$ and the insignificant estimates of η_v in all models, the very large risk premium attached to σ_s

in the SVJ model appears to be driven by model specification, although, even with jumps in volatility, there is evidence for a modest premium.

For every variant of the SVCJ model, there is strong evidence for a volatility jump risk premium, that is, $\mu_v^{\mathbb{Q}} > \mu_v$. In the SVCJ model, $\mu_v^{\mathbb{Q}}$ plays two roles; namely, it generates conditional kurtosis in returns and it influences the long-run risk-neutral mean of V_t . The fact that $\mu_v^{\mathbb{Q}} > \mu_v$ shows the need for greater risk-neutral kurtosis to fit the volatility smile. As in the other models, η_v is insignificant. However, this parameter impacts estimates of $\mu_v^{\mathbb{Q}}$: By constraining $\eta_v = 0$, $\mu_v^{\mathbb{Q}}$ estimates fall drastically. This occurs because the long-run mean of volatility is

$$E^{\mathbb{Q}}[V_t] = \frac{\kappa_v \theta_v + \mu_v^{\mathbb{Q}} \lambda}{\kappa_v + \eta_v}. \quad (10)$$

A large value of $\mu_v^{\mathbb{Q}}$, while generating a large conditional kurtosis of returns, has little impact on the long-run mean of the variance because its impact is largely nullified by η_v . When $\eta_v = 0$, $\mu_v^{\mathbb{Q}}$ estimates are more reasonable and more precise.

To interpret the magnitude of $\mu_v^{\mathbb{Q}}$, recall that EJP report an estimate of μ_v of about 1.5 (with a standard error of 0.34). With no constraints on $\sigma_s^{\mathbb{Q}}$, we estimate $\mu_v^{\mathbb{Q}}$ to be 3.71. If average annualized volatility is 15%, an average sized jump increases annualized volatility to 25% under $\mathbb{P}$ and 34% under $\mathbb{Q}$. Even if $\sigma_s^{\mathbb{Q}} = \sigma_s$, the estimate of 5.29 implies that an average-sized, risk-neutral jump increases volatility to 39%, which is plausible given the large increases observed historically.

Finally, to see the impact of risk premia on option prices, Figure 5 displays the Black–Scholes implied volatility curves for the SVJ and SVCJ models for 1 and 3 months to expiration. For the SVJ model, the solid line displays the implied volatility curve based on $\mathbb{P}$ -measure parameters, the dotted line includes risk premia and constrains $\sigma_s^{\mathbb{Q}} = \sigma_s$, and the dashed line allows $\sigma_s^{\mathbb{Q}} \neq \sigma_s$. For the SVCJ model, we consider the case in which $\sigma_s^{\mathbb{Q}} = \sigma_s$, as the implied volatility curves when $\sigma_s^{\mathbb{Q}} \neq \sigma_s$ are qualitatively similar.

Note first from the two upper panels that the SVJ model using the $\mathbb{P}$ -measure parameters or with $\sigma_s^{\mathbb{Q}} = \sigma_s$ generates monotonically declining implied volatility curves (the smirk). These models generate very little risk-neutral conditional kurtosis and they therefore miss the hook, that is, the increase in implied volatility for ITM puts or OTM calls (see, DPS, Pan (2002), and EJP). When the constraint on $\sigma_s^{\mathbb{Q}}$ is relaxed, the model generates more risk-neutral kurtosis (note the difference between the dashed line and the solid and dotted lines). Risk premia play a lesser role in the SVCJ model as this specification can generate a reasonable smile effect at both short- and intermediate-term maturities even under the $\mathbb{P}$ measure. It is also clear from the Figure 5 that when $\sigma_s^{\mathbb{Q}}$ is unconstrained, the SVJ model generates implied volatility curves that are quite similar to those of the SVCJ model.

This discussion illustrates that, given sufficiently flexible risk premia, one cannot distinguish among different models based on options' cross section only.




Figure 5. Black-Scholes implied volatility curves for the SVJ and SVCJ models based on $\mathbb{P}$ - and $\mathbb{Q}$ -measure parameters. This figure plots the implied volatility curves generated from the SVJ and SVCJ models using estimated parameters under various assumptions on the risk premia. A “ $\mathbb{P}$ ” after a model indicates that the figures are computed with no risk premia. SVJ- $\mathbb{Q}$ ($\sigma_s^{\mathbb{Q}} = \sigma_s^{\mathbb{P}}$) indicates that the price jump volatility is held constant across measures, but other risk premia are included using the values from Table IV. SVJ- $\mathbb{Q}$ and SVCJ- $\mathbb{Q}$ are implied volatility curves computed using the risk premia in Table IV.

However, a good option pricing model must be able to fit both cross-sectional and time-series properties. In our case, this means that SVCJ is the only model capable of successfully addressing all aspects of the data.

B.3. Interpreting the Risk Premia

In this section, we examine the degree to which the risk premia are reasonable, and we assess their economic significance. To do so, we examine the mean price jump size premia in the context of simple equilibrium models, the contribution of jump risk to the overall equity premium, and the impact of price and volatility jump risk premia on option returns.

First, consider Bates’s (1988) constant volatility jump diffusion equilibrium model. Bates finds that $\mu_s^{\mathbb{Q}} = \mu_s^{\mathbb{P}} - A\delta_{S,W}$, where A is the power utility

risk-aversion parameter and $\delta_{S,W}$ is the covariance between the jumps in the stock market and those in total wealth. It is reasonable to assume that stock price jumps are highly correlated with wealth (at least financial wealth) and therefore reasonable values of A can easily generate the 2% to 6% wedge that we find. Levels of risk aversion under 10 are generally considered to be reasonable. Bates (2000) estimates μ_s^Q to be around 9% and σ_s^Q to be around 10% to 11%, and notes that these estimates give “little reason to believe that the jump risk premia introduce a substantial wedge between the ‘risk-neutral’ parameters implicit in option prices and the true parameters” (p. 193). Additionally, the $\mathbb{P}$ -measure parameters are measured with more noise than the risk-neutral parameters, which indicates that the wedge between the two could be even smaller in a statistical sense.

Next, consider the jump risk contribution to the equity premium. Using $\mathbb{P}$ -measure parameter estimates from EJP and the decomposition of the equity risk premium from Section I.A, the contribution of the price jump risk premia is 2.7% and 2.9% per annum in the SVJ and SVCJ models (with $\eta_v = 0$), respectively. Over our sample, the equity premium is about 8%, implying that jumps generate roughly one third of the total premium. As a benchmark, time-series studies find that jumps in prices explain about 10% to 15% of overall equity volatility (EJP or Huang and Tauchen (2005)). It appears that jumps generate a relatively larger share of the overall equity premium. While significant, it is difficult to argue that these premia are unreasonable.

Finally, we consider our risk premium estimates in the context of a rapidly growing literature that identifies a “put-pricing puzzle” (see e.g., Bondarenko (2003), Driessen and Maenhout (2004b), Jones (2006), or Santa-Clara and Saretto (2005)). Using data similar to ours (Bondarenko uses S&P 500 futures options from 1987 to 2000), these authors document that average monthly returns of ATM and OTM puts are approximately -40% to -95% , respectively, and have high Sharpe ratios. Naturally, average returns of this magnitude are difficult to explain using standard risk-based asset pricing models such as the CAPM or Fama–French three-factor model, and they are also puzzling from a portfolio perspective (Driessen and Maenhout (2004b)) or based on a nonlinear factor model (Jones (2006)).

Our models and risk premium estimates provide a natural setting in which to explore a risk-based explanation of the put-pricing anomaly. To examine this issue, we simulate the SVCJ model, calculate put returns with and without risk premia, and compare the returns to results previously reported. Table VII provides these average values and the 5% to 95% bootstrapped confidence bounds reported by Bondarenko (2003). To compute model-based returns, we estimate the population values of average options returns by simulating 20,000 monthly index and option returns from the SVCJ model. We are careful to precisely follow the empirical design of Bondarenko and compute the holding period returns of options with 1 month left to maturity. We consider three scenarios with respect to the values of the risk premia. The case SVCJ- $\mathbb{P}$ corresponds to the zero risk premia, the case SVCJ- $Q-\mu_s^Q$ considers only the effect of a mean price jump risk premium, the case SVCJ- $Q-\mu_v^Q$ adds the volatility jump risk premium. In

Table VII
The Impact of Risk Premia on Option Returns in the SVCJ Model

We compare out-of-the-money (OTM) average option returns (measured in percent) and their bootstrapped percentiles reported by Bondarenko (2003) (Average, 5%, 95%) to population average options returns implied by the SVCJ model assuming zero risk premia (SVCJ- $\mathbb{P}$) and using the estimated risk premia (SVCJ- $\mathbb{Q}$). The dagger (†) denotes returns outside the confidence intervals.

	Moneyiness	6%	4%	2%	0%
Data	Average	−95.00	−58.00	−54.00	−39.00
	5%	−99.00	−80.00	−72.00	−54.00
	95%	−89.00	−35.00	−36.00	−24.00
Model	SVCJ- $\mathbb{P}$	−20.70 †	−21.91 †	−21.78 †	−19.96 †
	SVCJ- $\mathbb{Q}$ - $\mu_s^{\mathbb{Q}}$	−68.46 †	−58.36	−45.21	−32.18
	SVCJ- $\mathbb{Q}$ - $\mu_v^{\mathbb{Q}}$	−75.78 †	−64.67	−50.01	−35.27

each case, we match the overall average equity risk premium over the sample, which is about 8%.

Table VII provides a number of interesting implications. The SVCJ- $\mathbb{P}$ results indicate that a model without any risk premia generates about −20% per month for average put returns. These large negative put returns are solely driven by the very high S&P 500 returns: A short put position has a high “beta” on the index (around −25 to −30 for ATM puts, see Coval and Shumway (2001)). Intuitively, if the equity premium is high, puts often end up OTM, with a 100% return to the writer. Thus, a large component of the put pricing puzzle is due to the high equity premium over the 1990s. The SVCJ- $\mathbb{Q}$ - $\mu_s^{\mathbb{Q}}$ results indicate that adding a small mean price jump risk premium alone generates returns that are inside the confidence bands for most strikes. The final row indicates that adding the volatility jump risk premium generates option returns that are very close to the historical sample means. The only case that is not within the bounds is that of deep OTM puts; here, the returns are economically close.¹⁸

We conclude that jump risk premia provide an attractive risk-based explanation for the put-pricing puzzle. Although the returns are extremely large, these option positions are highly levered. Since our intuition and models (CAPM) are often based on normal distributional assumptions, the high returns (and Sharpe ratios) seem puzzling. However, these models are not well suited for understanding nonnormal risks such as those embedded in jumps. In our model with jumps, once we allow for even modest jump risk premia, the returns on these option strategies are not necessarily puzzling.

Although this section argues that these jump premia are not unreasonably large and that they have important economic implications, we are agnostic

¹⁸ The average return of 6% OTM puts is outside the confidence band, but the bands could be unreliable because of the small number of observations in this moneyiness category (see table 2 in Bondarenko (2003)).

about their exact sources. These jump risk premia may arise, for example, from standard utility functions; from asymmetric utility of gains and losses, which leads investors to care more about large negative returns; from the presence of heterogeneous investors with more risk-averse investors buying put options from the less risk-averse investors; from an inability to hedge jump risks; or from institutional explanations. Several papers explore some of these potential explanations. Liu, Pan, and Wang (2005) argue that an aversion to parameter uncertainty could generate the premia. Bollen and Whaley (2004) show that there are demand effects, in the sense that changes in implied volatility are related to the demand for options, and Garleanu, Pedersen, and Poteshman (2005) model demand effects with heterogeneous investors. Disentangling the sources of the jump risk premia and characterizing their relationship to investor demand appears to be a fruitful avenue for future research.

IV. Conclusions

In this paper, we use the time series and cross section of option prices to address a number of important option pricing issues. Using the time series, we find strong evidence for jumps in volatility, which, in conjunction with prior work, implies that stochastic volatility, jumps in prices, and jumps in volatility are all important components for option pricing. Using the cross section of option prices, we find that models with jumps in prices (with or without jumps in volatility) drastically improve overall pricing performance. Jumps in volatility offer a significant pricing improvement in the cross section unless a model with only price jumps is allowed to have a premia attached to the volatility of price jumps. In this case, the SVJ model requires that a relatively large premium be attached to the volatility of price jumps in order to generate the substantial amount of risk-neutral kurtosis observed in the cross section.

We find that estimates of risk-neutral mean price jumps are consistent across models, are on the order of 5% to 7%, are highly statistically significant, and imply a mean price jump risk premium of about 2% to 5%. We also find evidence for volatility of price jumps and volatility jump risk premia. The premia are economically plausible. The mean price jump risk premium is consistent with a modest level of risk aversion in simple equilibrium models. Jump risks contribute just under 3% to the total equity premium of 8% over our sample. Our jump risk premia also have important implications for the so-called “put-pricing puzzle,” which refers to the extremely high returns to writing put options that were observed in the 1990s. We find that a large proportion of the puzzle can be explained by the high returns on the underlying index, and that the remaining proportion can be generated by our modest price and volatility jump risk premia.

While our results resolve a number of existing issues in the literature, we conclude by mentioning three topics that represent promising avenues for future research. First, joint estimation and new types of data (variance swaps, volatility futures, etc.) would improve parameter estimates, especially η_v and ρ_s , if options of multiple strikes and maturities were used. Second, it is certainly

the case that our preferred model has shortcomings. For example, our model (as well as all others estimated in the literature) assumes that the long-run mean of volatility is constant. Casual observation indicates that this may be a tenuous assumption, as there are long periods during which volatility is higher or lower than its unconditional long-run mean. DPS and Pan (2002) suggest a model with a time-varying central tendency factor, which could be identified from longer-dated options. Such a model, estimated using more efficient methods and longer-dated options, might also resolve the issues surrounding the diffusive risk premium estimates. Similarly, more flexible variance jump distributions are of interest. Finally, we find preliminary evidence that the risk-neutral jump parameters vary over time, increasing in periods of market stress and decreasing during other periods. Santa-Clara and Yan (2005) also find evidence for time-varying jump risk premia. It would be interesting to develop diagnostics based on, for example, the slope of the implied volatility curve, to identify these time-varying premia, and to further examine their implications.

Appendix A: Adjusting for the Early Exercise Premium

Given that the S&P 500 futures options are American, this complicates the parameter estimation procedure because of the considerable additional time needed to compute model American prices versus European prices. We circumvent this computational difficulty by transforming American prices to European prices, and then estimating model parameters based on European prices. This approach has two main advantages. First, the computational savings is very significant, rendering the parameter estimation procedure feasible. Second, it eliminates the need to develop analytical approximations for American option values under each model.

We now quantify the magnitude of the approximation error introduced by this procedure. Assume that market prices are generated by a particular model, for example, the SV model. Let $SV^A(\tilde{\Theta})$ denote the American option price under the SV model with parameters $\tilde{\Theta} = (\Theta, K, \tau, S_t, V_t, r, \delta)$, and let $SV^E(\tilde{\Theta})$ denote the corresponding European option price with the same parameters. Suppose that the market American price C is $SV^A(\tilde{\Theta})$. According to our model assumptions, a European option would trade for $SV^E(\tilde{\Theta})$. Using the observed price C we compute an American Black–Scholes implied volatility, that is, a value σ^{BS} such that $C = BS^A(\sigma^{BS}, \tilde{\Theta})$, where BS^A denotes the Black–Scholes American option price. We then estimate that an equivalent European option would trade in the market at a price $BS^E(\sigma^{BS}, \tilde{\Theta})$, where BS^E denotes the Black–Scholes European option price.

We test the accuracy of this procedure by computing absolute, $BS^E(\sigma^{BS}, \tilde{\Theta}) - SV^E(\tilde{\Theta})$, and relative, $(BS^E(\sigma^{BS}, \tilde{\Theta}) - SV^E(\tilde{\Theta}))/SV^E(\tilde{\Theta})$, errors over a large set of model parameters chosen to be representative of those estimated based on our data set. Similar definitions apply for the SVJ model. For the SV model, we test all combinations of the parameters that straddle the estimated values in Table I: $\kappa_v \in \{0.008, 0.032\}$, $\sigma_v \in \{0.1, 0.25\}$, $\rho \in \{-0.1, -0.7\}$, $\theta_v \in \{0.57, 1.59\}$,

Table AI
European Price Approximation Errors

We test the accuracy of the approximation procedure by computing absolute and relative errors over a large set of model parameters that are consistent with our data set. The relative error $(BS^E(\sigma^{BS}, \tilde{\Theta}) - M^E(\tilde{\Theta})) / M^E(\tilde{\Theta})$, where M is either SV or SVJ, is reported for options with prices above \$0.5. The absolute error $BS^E(\sigma^{BS}, \tilde{\Theta}) - M^E(\tilde{\Theta})$ is relevant for lower priced options. We summarize the errors by the root mean square (RMS) and maximum absolute (Max) errors.

Model Error			SV				SVJ			
			Relative		Absolute		Relative		Absolute	
			RMS	Max	RMS	Max	RMS	Max	RMS	Max
Call	1.5 months		0.17%	0.26%	0.0006	0.0017	0.28%	0.71%	0.0014	0.0044
	6 months		0.26%	1.08%	0.0029	0.0065	0.37%	0.85%	0.0020	0.0035
Put	1.5 months		0.05%	0.11%	0.0008	0.0034	0.10%	0.42%	0.0013	0.0060
	6 months		0.30%	1.26%	0.0020	0.0045	0.36%	1.37%	0.0023	0.0064

$V_t \in \{0.57, 1.59\}$, $r \in \{0.008\%, 0.03\%\}$, and $\tau \in \{1.5, 6\}$ (months). We choose rather long maturities as the exercise premium is increasing in maturity. Given the initial stock price $S_t = 100$, we test strikes from the set $K \in \{85, 90, 105\}$ for put options and $K \in \{95, 100, 115\}$ for call options. We test a total of 768 combinations of SV model and option parameters. For the SVJ model, we test all combinations of the previous parameters, together with all combinations of the jump parameters: $\lambda \in \{0.004, 0.008\}$, $\mu_s \in \{-1.00\%, -5.00\%\}$, and $\sigma_s \in \{2.00\%, 8.00\%\}$. All of the parameters are in daily units.¹⁹ We test a total of 6,144 combinations of SVJ model and option parameters.

We determine accurate American prices using two-dimensional finite difference routines. These finite difference routines are much slower than the Fourier inversion methods that we use for pricing European options. The computation time required for American options makes calibration to a very large set of options impractical.

Let ε_i denote the error (either absolute or relative) for the i th set of option parameters. We summarize the errors by the RMS error measure and the worst case, or maximum, absolute error. For low-priced options, absolute error is the relevant error measure, while for higher-priced options, relative error is more relevant. We choose a price of \$0.50 as the separator between low-priced and high-priced options, consistent with previous studies (e.g., Broadie and Detemple (1996)).

Table AI provides summary results for the SV and SVJ models. The results show that approximating European prices under the SV (or SVJ) model by using the market American prices and subtracting the BS early exercise premium leads to very small approximation errors. RMS relative errors for high-priced options are less than 0.4% while the maximum absolute relative error is 1.4%.

¹⁹ In terms of annual decimal parameters, the values are: $\kappa \in \{2, 8\}$, $\sigma_v \in \{0.1, 0.25\}$, $\rho \in \{-0.1, -0.7\}$, $\theta \in \{(0.12)^2, (0.2)^2\}$, $V_t \in \{(0.12)^2, (0.2)^2\}$, $r \in \{0.02, 0.08\}$, $\lambda = \{1, 2\}$, and $\tau \in \{30/252, 120/252\}$. The jump distribution parameters are unchanged.

Table AII
Exercise Premium in the SVJ Model

We illustrate the magnitudes of the exercise premium and the approximation error on the example of individual put options under the SVJ model. We evaluate the put option prices assuming $S_t = \$100$, $V_t = 1.59$, $\sigma_v = 0.25$, $\rho = -0.7$, $\lambda = 0.008$, $\mu_s = -5.0\%$, $\sigma_s = 8.0\%$, and $r = 3.0\%$.

Strike	Maturity	κ_v	θ_v	SVJ^A	SVJ^E	BS^E	Absolute ε
85	1.5 months	0.008	1.59	0.201	0.201	0.201	0.000
90	1.5 months	0.008	1.59	0.561	0.560	0.561	0.000
105	1.5 months	0.008	1.59	6.154	6.139	6.139	-0.001
85	6 months	0.008	1.59	1.583	1.574	1.576	0.003
90	6 months	0.008	1.59	2.658	2.638	2.643	0.005
105	6 months	0.008	1.59	8.994	8.884	8.904	0.020
85	1.5 months	0.032	0.57	0.161	0.161	0.161	0.000
90	1.5 months	0.032	.057	0.460	0.459	0.460	0.001
105	1.5 months	0.032	.057	5.909	5.894	5.894	0.000
85	6 months	0.032	.057	0.953	0.945	0.949	0.004
90	6 months	0.032	.057	1.804	1.787	1.795	0.008
105	6 months	0.032	.057	6.963	6.846	6.883	0.038

To put these numbers in perspective, a \$10 option with the largest observed relative error of 1.4% has an absolute price error of only \$0.14. For low-priced options RMS absolute errors are less than \$0.003 while the maximum absolute error is only \$0.007. These errors are far smaller than typical bid-ask spreads. Furthermore, the largest errors occur for extreme option parameters that do not obtain in our data set. For example, the largest errors are for long maturity in-the-money options with a large difference between initial and long-run volatilities, while the most actively traded options in our data set are those with short maturities and out-of-the-money strikes.

In order to better understand these summary statistics, results for individual put options in the SVJ model are given in Table AII. It is clear that this procedure works well because the early exercise premia (i.e., American minus European option values) are small to begin with, and approximating the SV or SVJ early exercise premium by the corresponding early exercise premium in the BS model reduces the error even further.

To recap, our calibration procedure begins by converting American market prices to equivalent European market prices. This allows us to use computationally efficient European pricing routines that render the large-scale calibration procedure feasible. The results in this Appendix show that the approximation error using this procedure is minimal.

Appendix B: Data Issues

We obtain daily time and sales files for the S&P 500 futures and futures options from the Chicago Mercantile Exchange (CME). The files record transactions and bid-ask quotes for both the futures and the options. We only use

transaction prices. Beginning with 2,246,426 option transactions, we are able to find a matching futures transaction within five minutes for 2,081,727 transactions, which consists of 947,635 call and 1,134,092 put transactions. By maturity, 1,090,462 transactions are under 30 days to maturity, 803,971 are between 30 and 89 days to maturity, and 187,294 are 90 days or more to maturity. By strike, 412,327 have $K/F < 0.95$, 1,474,517 have $0.95 < K/F < 1.05$, and 194,883 have $K/F > 1.05$.

S&P 500 futures options do not have the wild card feature that is present in S&P 100 options. A wildcard feature arises because the trading day in options ends at 3:15 CST even though the stocks underlying the index cease trading at 3:00 CST. The S&P 100 options are American and holders of options can exercise their options until 3:20. If exercised, the options are settled in cash based on the index value at 3:00 CST. This generates an additional valuable option for option holders, which is commonly called the wildcard option (see Fleming and Whaley (1994)). For S&P 500 futures options, there is not a wildcard feature as exercised options receive a long or short position in the futures marked at the bid or ask price at the time of exercise (see Chapter 351A, Section 02.B of the CME rulebook). Since the futures trade after hours, the futures contract will take into account any news and there is no wildcard effect.

We use two different data sets, namely, a time series of “representative” option prices for the specification tests and the cross section of option prices for each maturity on each date. For the representative option data set we use the following four selection criteria. (1) Select all option-futures pairs that are traded between 9:30 a.m. and 10:30 a.m., have a time difference of less than 1 minute, and for which $|K/S| < 1.02$. If three or more pairs match these conditions, the pair that has the median volatility is selected. (2) If less than three records are selected using criterion 1, then we add more records by allowing pairs traded during any time of the day (ordered by closeness to a trade time of 10:00 a.m.). The time difference and strike conditions are still in effect. If three or more pairs match both criteria 1 and 2, then the three records that best satisfy the criteria are selected for median computation. (3) If less than three records are selected in criteria 1 and 2, then we add more records by allowing pairs with any moneyness (ordered by closeness to $K/S = 1$). The time difference condition is still in effect. If three or more pairs match criteria 1, 2, and 3, then we select the three records that best satisfy the criteria for median computation. (4) If less than three records are selected in criteria 1, 2, and 3, then we add more records by allowing pairs with any time difference. In this case, records are ordered so that time difference is the lowest (in 5 minute blocks, so that a time difference of 4 minutes and 59 seconds is just as good as a time difference of 1 minute and 1 second). Any ties in the ordering of the 5 minute blocks is broken by choosing the pair closest to the money. If three or more pairs match criteria 1, 2, 3, and 4, then we select the three records that best satisfy the criteria for median computation. If there are less than three pairs, then we select the “best” record.

For risk premium estimation, we construct a data set for each trading day using every option transaction that can be time-matched within 5 minutes to a futures transaction. This typically produces hundreds of matched options-futures transactions per day. We then use the following curve-fitting procedure to combine all matched transactions into a representative curve for each option maturity on each trading day. First, for each option price, we compute its American option-implied volatility under the Black–Scholes model, $\sigma_t^{i,j}$ for each strike K_j and time to maturity τ_i . We compute $\sigma_t^{i,j}$ using an iterative solver together with a binomial or finite difference American option pricing routine. Then, for each day and maturity, we fit a piecewise quadratic function to the implied volatilities, that is,

$$y = 1_{[x \leq x_0]} [a_2(x - x_0)^2 + a_1(x - x_0) + a_0] \\ + 1_{[x > x_0]} [b_2(x - x_0)^2 + a_1(x - x_0) + a_0] + \varepsilon,$$

where y is the American Black–Scholes implied volatility, x is the moneyness (K/S), x_0 is the knot point of the piecewise quadratic, and the coefficients are least-square estimates. The knot point is allowed to vary on the nearest maturity, but for longer maturities it is fixed at 1. If the maturity has 10 or fewer option transactions, then a linear function replaces the piecewise quadratic function. Figure B1 shows a representative day, August 6, 1999, with four fitted curves for each of the traded maturities. The maturities are 14 days, 42 days, 70 days, and 133 days and we have 302, 134, 17, and 24 put and call transactions, respectively, for these maturities.

We experiment with a number of other specifications, including piecewise cubic functions, fixed knot points for all maturities, linear and piecewise functions, etc., and evaluate these alternative specifications by a cross-validation procedure. For each maturity on each day, we divide the set of options in half. For each half, we fit parametric curves and assess the fitting error using the other half of the data. The overall fitting error is the average of the two out-of-sample results. For maturities with a small number of data points, say m , we use $m - 1$ points to fit the curve and assess the error in fitting the m th point. The overall fitting error is the average of the m results, each time leaving out one data point. The piecewise quadratic approach is the best when there were more than 10 data points: Piecewise cubic functions overfit the data (perform relatively poorly on the cross-validation test), while linear functions tend to underfit the data. For 10 data points or fewer, we find the linear function approach is the best, as quadratic and piecewise quadratic functions tend to overfit the data. We experiment with a number of other cases and specifications. For example, we allow x_0 to vary within the range of x -values in the data for each maturity, and we fix x_0 at $x_0 = 1$ for all maturities. These generate only minor differences in the implied volatility fits. If the optimizer does not converge to an x_0 inside the range of x -values in the data then we fix x_0 at 1.




Figure B1. Market observed implied volatilities and fitted implied volatility curves. This figure displays observed call and put prices for various maturities and the resulting piecewise quadratic fits.

Given the resulting smile curves, we calibrate option models by minimizing squared differences between market and model Black–Scholes implied volatilities using equation (8) and a discrete set of points from these curves. In effect, we calibrate option models in a two-stage approach. The first stage fits implied volatility curves to transactions data. The second stage finds model parameters that best fit the implied volatility curves. This two-stage approach dramatically reduces the computational requirements without sacrificing accuracy. Because the objective function is not globally convex, each optimization problem is solved from multiple diverse starting points to ensure convergence to the global optimum.²⁰ Other details of the optimization procedure are available on request.

²⁰ Huang and Wu (2003) use a two-stage estimation procedure, alternating between optimizing over parameters and daily volatilities. While significantly reducing the search dimension, we find that this procedure can converge to local minima that given substantively different results. In our procedure, we jointly optimize over (V_t, Θ^Q) even though this procedure is more computationally intensive.

REFERENCES

- Ait-Sahalia, Yacine, 2002, Telling from discrete data whether the underlying continuous-time model is a diffusion, *Journal of Finance* 57, 2075–2112.
- Andersen, Torben, Luca Benzoni, and Jesper Lund, 2002, Towards an empirical foundation for continuous-time equity return models, *Journal of Finance* 57, 1239–1284.
- Anderson, Evan, Lars Hansen, and Thomas Sargent, 2003, A quartet of semi-groups for model specification, robustness, price of risk, and model detection, *Journal of the European Economic Association* 1, 68–123.
- Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003–2049.
- Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta hedged gains and the pricing of volatility risk, *Review of Financial Studies* 16, 527–566.
- Bates, David, 1988, Pricing options under jump-diffusion processes, Working paper, University of Pennsylvania.
- Bates, David, 1996, Testing option pricing models, in G. S. Maddala and C. R. Rao, eds. *Handbook of Statistics, Vol. 14, Statistical Methods in Finance* (Elsevier, Amsterdam).
- Bates, David, 2000, Post-'87 crash fears in S&P 500 futures options, *Journal of Econometrics* 94, 181–238.
- Benzoni, Luca, 2002, Pricing options under stochastic volatility: An empirical investigation, Working paper, University of Minnesota.
- Bliss, Robert, and Nikolaos Panigirtzoglou, 2004, Option-implied risk aversion estimates, *Journal of Finance* 59, 407–446.
- Bollen, Nicholas, and Robert Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* 59, 711–753.
- Bondarenko, Oleg, 2003, Why are put options so expensive? Working paper, University of Illinois, Chicago.
- Bookstaber, Richard, and Steven Pomerantz, 1989, An information-based model of market volatility, *Financial Analysts Journal* 45, 37–46.
- Brandt, Michael, and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economics* 72, 217–257.
- Branger, Nicole, and Christian Schlag, 2004, Can tests based on option hedging errors correctly identify volatility risk premia? Working paper, Goethe University.
- Bremaud, Pierre, 1981, *Point Processes and Queues: Martingale Dynamics* (Springer-Verlag, New York).
- Broadie, Mark, and Jerome Detemple, 1996, American option valuation: New bounds, approximations, and a comparison of existing methods, *Review of Financial Studies* 9, 1211–1250.
- Carr, Peter, and Liuren Wu, 2003, What type of process underlies options? A simple robust test, *Journal of Finance* 58, 2581–2610.
- Chernov, Mikhail, A. Ronald Gallant, Eric Ghysels, and George Tauchen, 2003, Alternative models for stock price dynamics, *Journal of Econometrics* 116, 225–257.
- Chernov, Mikhail, and Eric Ghysels, 2000, A study towards a unified approach to the joint estimation of objective and risk-neutral measures for the purpose of options valuation, *Journal of Financial Economics* 56, 407–458.
- Christoffersen, Peter, and Kris Jacobs, 2004, The importance of the loss function in option valuation, *Journal of Financial Economics* 72, 291–318.
- Conley, Timothy, Lars Peter Hansen, and Wen Fang Liu, 1997, Bootstrapping the long run, *Macroeconomic Dynamics* 1, 279–311.
- Cont, Rama, and Peter Tankov, 2003, Non-parametric calibration of jump-diffusion option pricing models, *Journal of Computational Finance* 7, 1–49.
- Coval, Joshua, and Tyler Shumway, 2001, Expected option returns, *Journal of Finance* 56, 983–1009.
- Driesen, Joost, and Pascal Maenhout, 2004a, The world price of jump and volatility risk, Working paper, Insead and University of Amsterdam.

- Driesen, Joost, and Pascal Maenhout, 2004b, A portfolio perspective on option pricing anomalies, Working paper, Insead and University of Amsterdam.
- Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.
- Duffie, Darrell, Lasse Heje Pedersen, and Kenneth Singleton, 2003, Modeling sovereign yield spreads: A case study of Russian debt, *Journal of Finance* 55, 119–159.
- Efron, Bradley, and Robert Tibshirani, 1994, *An Introduction to the Bootstrap* (CRC Press, New York).
- Eraker, Bjørn, 2004, Do equity prices and volatility jump? Reconciling evidence from spot and option prices, *Journal of Finance* 59, 1367–1403.
- Eraker, Bjørn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in equity index volatility and returns, *Journal of Finance* 58, 1269–1300.
- Fleming, Jeff, and Robert Whaley, 1994, The value of the wildcard option, *Journal of Finance* 49, 215–236.
- Garleanu, Nicolae, Lasse Pedersen, and Allen Poteshman, 2005, Demand based option pricing, Working paper, University of Pennsylvania.
- Huang, Jing-Zhi, and Liuren Wu, 2004, Specification analysis of option pricing models based on time-changed Levy processes, *Journal of Finance* 59, 1405–1439.
- Huang, Xin, and George Tauchen, 2005, The relative contribution of jumps to total price variance, *Journal of Financial Econometrics* 3, 456–499.
- Johannes, Michael, 2004, The statistical and economic role of jumps in interest rates, *Journal of Finance* 59, 227–260.
- Jones, Christopher, 2003, The dynamics of stochastic volatility: Evidence from underlying and options markets, *Journal of Econometrics* 116, 181–224.
- Jones, Christopher, 2006, A nonlinear factor analysis of S&P 500 index option returns, *Journal of Finance* 41, 2325–2363.
- Liu, Jun, Jun Pan, and Tan Wang, 2005, An equilibrium model of rare-event premia and its implication for option smirks, *Review of Financial Studies* 18, 131–164.
- Naik, Vasanttilak, and Moon Lee, 1990, General equilibrium pricing of options on the market portfolio with discontinuous returns, *Review of Financial Studies* 3, 493–521.
- Pan, Jun, 2000, Integrated time-series analysis of spot and option prices, Working paper, MIT Sloan School of Management.
- Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.
- Pritsker, Matthew, 1998, Nonparametric density estimation and tests of continuous-time interest rate models, *Review of Financial Studies* 11, 449–487.
- Santa-Clara, Pedro, and Alessio Saretto, 2005, Option strategies: Good deals and margin calls, Working paper, UCLA.
- Santa-Clara, Pedro, and Shu Yan, 2005, Crashes, volatility, and the equity premium: Lessons from S&P 500 options, Working paper, UCLA.
- Sato, Ken-Iti, 1999, Levy process and infinitely divisible distributions, *Cambridge Studies in Advanced Mathematics*, Vol. 68 (Cambridge University Press, Cambridge, UK).
- Whaley, Robert, 1986, Valuation of American futures options: Theory and empirical tests, *Journal of Finance* 41, 127–150.
- Whaley, Robert, 2000, The investor fear gauge, *Journal of Portfolio Management* 26, 12–17.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.