

Model Specification and Risk Premia: Evidence from Futures Options

Mark Broadie Mikhail Chernov Michael Johannes*

September 23, 2005

Abstract

This paper examines specification issues and estimates diffusive and jump risk premia using the cross-section of S&P futures option prices from 1987 to 2003. We first test for the presence of jumps in volatility by analyzing the higher moment behavior of option implied variance. We find strong evidence supporting their presence. Based on the cross-sectional fit of option prices, we find strong evidence for jumps in prices and modest evidence for jumps in volatility. Regarding factor risk premia, we are not able to identify a statistically significant diffusive volatility risk premium, but are able to identify statistically significant, though modest, risk premia for jumps in prices and volatility. These jump risk premia are economically meaningful, as they contribute a significant component to the equity risk premium and can help explain observed put returns.

*All authors are at the Graduate School of Business, Columbia University, 3022 Broadway, New York, NY, 10027. Email addresses: mnb2@columbia.edu (Broadie); mc1365@columbia.edu (Chernov) and mj335@columbia.edu (Johannes). We thank seminar participants at Columbia, Connecticut, LSE, LBS and the WFA for helpful comments. David Bates, Gurdip Bakshi, Chris Jones provided helpful comments and clarifications. We are especially 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 Introduction

There are two central, related, issues in empirical option pricing. The first 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 these factors, using a model that passes reasonable specification hurdles. In particular, we are interested in understanding the relative importance of volatility and jump risks.

We are motivated by disparate results in the extant literature regarding these issues. For example, tests using option data disagree over the importance of jumps in prices: Bakshi, Cao, and Chen (1997) (BCC) find large benefits of jumps in prices, but Bates (2000) and others find economically-small, even negligible, benefits.¹ As we discuss in greater detail below, time-series based studies unanimously support jumps in prices but disagree over jumps in volatility. Finally, there is general disagreement regarding the magnitudes and significance of volatility and jump risk premia.

One explanation for the disparity is that most papers use data covering short time periods, often omitting periods of market stress. BCC (1997) and Bates (2000) use the entire cross-section of options from 1988 to 1991 and 1988 to 1993, respectively; Pan (2002) uses two call options per day from 1989 to 1996; and Eraker (2004) uses up to three call options per day from 1987 to 1990. Since jumps in prices or volatility are rare events, short samples will either over- or under-represent jumps and/or periods of high or low volatility, thus explaining the disparate results. Figure 1 displays a time series plot of the VIX index to show the potential problems of using short sub-samples. If the goal is to learn about

¹Pan (2002) finds that pricing errors decrease when jumps in prices are added for certain strike/maturity combinations, but increase for other combinations. Eraker (2004) finds that adding jumps in returns and jumps in volatility result in only a one percent improvement in pricing. Bates (2000) finds ten percent improvement but it falls to around two percent when time series consistency is imposed.

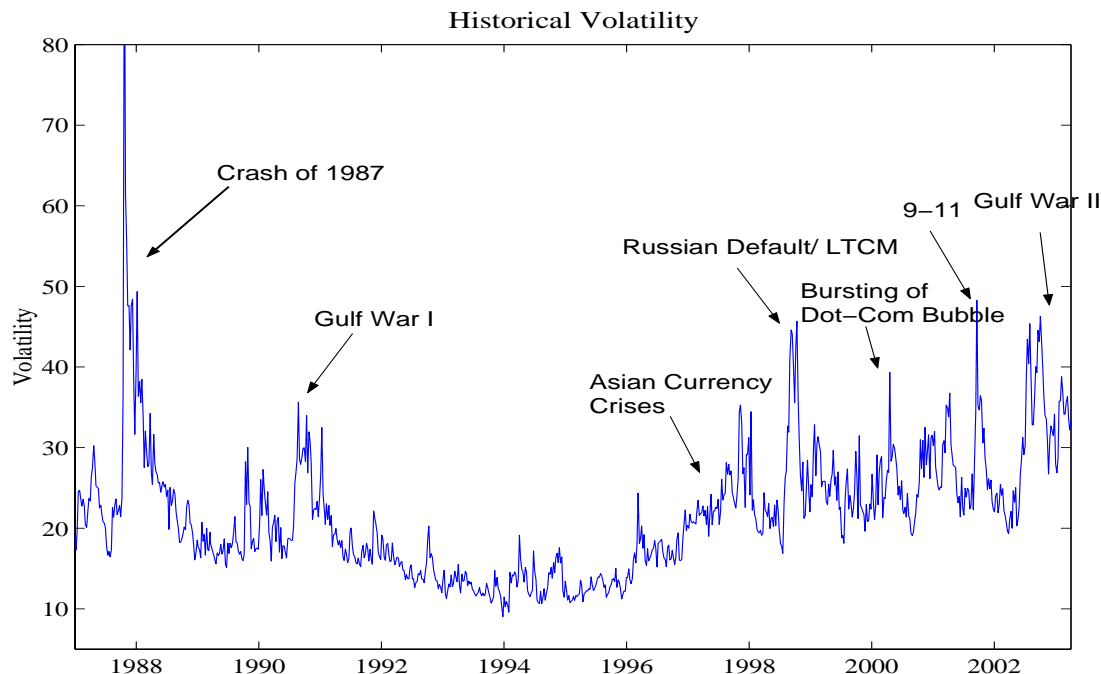


Figure 1: Time series of implied volatility (the VIX index) from 1987 to March 2003.

time-varying volatility and rare events such as jumps, it is important to include as many of these periods as is possible. It is precisely these periods of market stress that provide the best opportunity to learn about the underlying factor specification and investors' attitudes toward the risks these factors embody.

In this paper, we use an extensive data set of S&P 500 futures options from January 1987 to March 2003 to 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?; and (3) what is the nature of the factor risk premia embedded in the cross-section of option prices?

Regarding the first question, we develop tests using implied volatility to detect volatility jumps. Intuitively, positive volatility jumps induce positive skewness and excess kurtosis.

To test this, we extract a model-based estimate of spot variance from options, calculate skewness and kurtosis statistics, and simulate the statistics' finite sample distribution. A square-root stochastic volatility (SV) model and an augmentation with jumps in prices (the SVJ model) are rejected. These rejections of the SV and SVJ models are robust to reasonable variations in the volatility of volatility, to excluding the crash of 1987 and to factor risk premia. The SV and SVJ models assume that volatility increments are approximately normal, an implication which the data clearly reject. A model with contemporaneous jumps in volatility and prices (SVCJ) can easily capture these effects.

We conclude that a reasonable model requires stochastic volatility, jumps in prices, and jumps in volatility. Given this specification, we turn to the information in the cross-section of options prices over time to examine model fit and to learn about diffusive volatility and jump risk premia. We are particularly interested in understanding how risk premium estimates change when we incorporate jumps in volatility.

In estimating the models based on the cross-section of option prices, we depart from the usual routine of pure calibration 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, relying on parameter estimates previously reported in the literature (an approach similar to Benzoni 2002). The motivation for this is twofold. First, there is little disagreement over the values of these parameters, especially for the models with jumps across a number of different papers estimating these models.² Second, estimating all of the parameters using both options and returns is a computationally-demanding task.

In terms of pricing, we find that adding jumps in prices to the SV model improves

²As an example, here are reported parameter estimates for the volatility of volatility and correlation parameters in the SVCJ model: 0.08 and -0.48 (Eraker, Johannes, and Polson 2002), 0.07 and -0.46 (Chernov, Ghysels, Gallant, and Tauchen 2003) and 0.06 and -0.46 (Eraker 2004).

the cross-sectional fit by almost 50 percent. This magnitude is consistent with BCC, but contrasts with the results in Bates (2000), Pan (2002), and Eraker (2004), who find negligible gains. Without any constraints on the risk premia, the SVJ and SVCJ models perform similarly, either in or out of sample. This is not surprising, as price jumps generate significant amounts of skewness and kurtosis and, along with stochastic volatility, are clearly the two most important components for describing the time series of returns or pricing options. Jumps in volatility have a relatively small impact on the cross-section of option prices. This does not mean they are not important; in fact, the contrary is true for two reasons. First, time series tests indicate that volatility jumps are important for explaining the time series of returns and option prices. Second, it is dangerous to rely on risk premia obtained from a clearly misspecified model which implies that even if the cross-sectional fits in the SVJ and SVCJ models are similar, the risk premia in the SVJ model should not be trusted (see, e.g., Anderson, Hansen, and Sargent, 2003).

Turning to risk premia, our model specification allows for the parameters indexing the price and volatility jump size distributions to change across measures, and we label 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. We also allow for a premium associated with Brownian shocks in stochastic volatility, and this is labeled as the diffusive volatility risk premium.

Next, the different types of risk premia associated with jumps, either in prices or volatility, and diffusive volatility risk premia have fundamentally different sources of identification. In theory, the term structure of implied volatility primarily identifies diffusive 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 because the term structure of implied volatility is flat.³ In contrast to the

³On average, the slope of the term structure of implied volatility is very small. In our dataset, 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.

In the SVJ model, the risk-neutral price jump means are three to six percent lower than their objective measure counterparts, depending on whether or not we allow a volatility of price jumps risk premium. Mean price jump risk premia of this magnitude are significant, but not implausibly large, at least in the context of the equilibrium models in Bates (1988) or Naik and Lee (1990). In the SVCJ model, the mean price jump risk premium is smaller, about two or four percent, depending again on the assumptions regarding other premia. In all cases, the mean price jump risk premia are highly significant, but the sizes are modest compared to some of the prior studies. We also find a statistically significant volatility of price jumps risk premium and volatility jump risk premium.

Finally, to understand the economic magnitudes of these premia, we analyze the contribution of price jump risk to the overall equity risk premium and study how price and volatility jump risk premia impact expected option returns. First, in the SVJ or SVCJ models, price jump risk premia contribute about three percent per year to the overall equity premium, which was about eight percent over our sample. Second, Bondarenko (2003), Driessen and Maenhout (2004b), and Santa-Clara and Saretto (2005) argue that there is a “put-pricing” anomaly, as writing puts delivers large historical returns: e.g., about 40 percent per month for at-the-money puts. They argue these returns are implausibly high, at least in the context of standard asset pricing models or from a portfolio perspective. 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 our modest price jump and volatility jump risk premia. We conclude that even relatively small premia asso-

difference in implied volatilities between one-month options and options between a three and six months is less than one percent in terms of Black-Scholes implied volatility.

ciated with jump risks can have important implications for puts and the main reason these returns appear to be puzzling is that standard models such as the CAPM or Fama-French, not surprisingly, have difficulty capturing jump risks.

2 Models and Methods

2.1 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

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

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

where, W_t^s and W_t^v are two 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 the jumps in volatility are exponentially distributed, $Z_n^v \sim \exp(\mu_v)$.⁴ The SV and SVJ models are special cases assuming that $N_t = 0$ and $Z_n^v = 0$, respectively. This model was introduced by Duffie, Pan, and Singleton (DPS) (2000).⁵ Eraker, Johannes, and Polson (2003) (EJP), Chernov, Gallant, Ghysels, and Tauchen (2003) (CGGT), and Eraker (2004) analyze these models.

DPS specified 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,

⁴All of the specifications that we consider use only finite-activity jump processes, thus we do not consider infinite-activity Levy-based models, as in Madan and Milne (1991), Madan, Carr and Chang (1998), and Carr and Wu (2003a).

⁵The earliest formal model incorporating jumps in volatility is the shot-noise model in Bookstaber and Pomerantz (1989). Although they did not adopt a formal model, 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.

at least if we think of events like the crash of 1987. Time series based estimates of ρ_s are typically negative, but always insignificant.⁶ This parameter is extremely difficult to estimate, even with 15 or 20 years worth of data, because jumps are very rare events.⁷ In addition, ρ_s primarily affects the conditional skewness of returns, and, as a result, μ_s and ρ_s play a very similar role in determining the distribution of returns. Due to the difficulty in estimating this parameter precisely and for parsimony, we assume that the price jump sizes are independent of those in volatility. This constraint implies that the SVCJ model has only one more parameter than the SVJ model and insures that the SVJ and SVCJ models have the same price jump distribution. This is convenient for comparisons to the existing literature. We also assume a constant intensity under \mathbb{P} as there is no evidence from the time series of returns that the intensity is time-varying (CGGT and ABL) and also Bates (2000) finds strong evidence for misspecification in models with state-dependent intensities. An intensity that depends on volatility is especially difficult to identify. The intuition is clear: when volatility is high, the diffusive component can easily generate large movements without jumps.

The specification includes a correction term, $-S_t \bar{\mu}_s \lambda dt$, to compensate the jump component, where $\bar{\mu}_s = \exp(\mu_s + \sigma_s^2/2) - 1$. This implies that γ_t is the total equity premium. It is commonly assumed 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.⁸ The jump contribution to the equity premium is $\lambda \bar{\mu}_s - \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}}$, where \mathbb{Q} is the risk-neutral measure. The total

⁶EJP and CGGT find negative, but insignificant estimates while Eraker (2004) find a slightly positive, but insignificant estimate based on the time series.

⁷The small sample problem is severe. Since jumps are rare (about one to two per year), samples with 15 or 20 years of data generate relatively small numbers of jumps to identify this parameter. As 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 percent) greater than zero. The uncertainty is greater in reality as price and volatility jump sizes are not perfectly observed.

⁸See, for example, the extensive discussion in Brandt and Kang (2004).

premium is $\gamma_t = \eta_s V_t + \lambda \bar{\mu}_s - \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}}$. 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 market generated by the model in (1) and (2) is incomplete, which implies there are multiple equivalent martingale measures. We follow the literature by parameterizing the change of measure and estimating the risk-neutral parameters from option prices. The change of measure or density process is $L_t = L_t^D L_t^J$. We follow Pan (2002) and assume 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 $L_t^D = \exp\left(\int_0^t \Gamma_s dW_s - \frac{1}{2} \int_0^t \Gamma_s \cdot \Gamma_s ds\right)$. The jump component is

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

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 $\lambda_{\tau_n}^{\mathbb{Q}}$ and λ_{τ_n} 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 the 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. Modulo regularity, Girsanov's theorem only constrains the intensity to be predictable and the jump distributions to have common support. With constant intensities and state independent jump distributions, the only constraint on the jump distributions is that they are 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 $\pi^{\mathbb{Q}}(Z^s) = \mathcal{N}(\mu_s^{\mathbb{Q}}, (\sigma_s^{\mathbb{Q}})^2)$.⁹

Our specification allows the jump intensity and all of the jump distribution parameters

⁹This rules out a correlation between jumps in prices and volatility under \mathbb{Q} . As in the case of the objective measure, correlation between jumps in prices and volatility would be difficult to identify because $\mu_s^{\mathbb{Q}}$ plays the same role in the conditional distribution of returns.

to change across measures. This is more general than those considered in Pan (2002) or Eraker (2004), although, as we discuss later, we are not able to identify all of the parameters under \mathbb{Q} .¹⁰ At a first glance it may seem odd that we allow $\sigma_s \neq \sigma_s^{\mathbb{Q}}$. Prior studies do not allow this, constraining $\sigma_s = \sigma_s^{\mathbb{Q}}$. This constraint is an implication of the Lucas economy equilibrium models in Bates (1988) and Naik and Lee (1990) assuming 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 constrain $\sigma_s = \sigma_s^{\mathbb{Q}}$.

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

$$dS_t = S_t(r_t - \delta_t) dt + S_t \sqrt{V_t} dW_t^s(\mathbb{Q}) + d \left(\sum_{n=1}^{N_t(\mathbb{Q})} S_{\tau_{n-}} [e^{Z_n^s(\mathbb{Q})} - 1] \right) - S_t \lambda^{\mathbb{Q}} \bar{\mu}_s^{\mathbb{Q}} dt \quad (3)$$

$$dV_t = [\kappa_v(\theta_v - V_t) - \eta_v V_t] dt + \sigma_v \sqrt{V_t} dW_t^v(\mathbb{Q}) + d \left(\sum_{n=1}^{N_t(\mathbb{Q})} Z_n^v(\mathbb{Q}) \right) \quad (4)$$

where $\bar{\mu}_s^{\mathbb{Q}} = \exp(\mu_s^{\mathbb{Q}} + 0.5(\sigma_s^{\mathbb{Q}})^2) - 1$. 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. For interpretation purposes, we label the difference between the \mathbb{P} and \mathbb{Q} parameters as risk premia. Specifically, we label $\mu_s - \mu_s^{\mathbb{Q}}$ as the *mean price jump* risk premium, $\sigma_s^{\mathbb{Q}} - \sigma_s$ is the *volatility of price jumps* risk premium, $\mu_v^{\mathbb{Q}} - \mu_v$ is the *volatility jump* risk premium, and $\eta_v = \kappa_v^{\mathbb{Q}} - \kappa_v$ is the diffusive volatility risk premium. We will often refer to $\mu_s - \mu_s^{\mathbb{Q}}$ and $\sigma_s^{\mathbb{Q}} - \sigma_s$ together as the price jump risk premia.

It is important to note that the absolute continuity requirement implies that certain model parameters, or combinations of model 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

¹⁰As we discuss later, we follow Pan (2002) and Eraker (2004) and impose $\lambda^{\mathbb{Q}} = \lambda$.

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, an approach advocated by Bates (2000). We impose this constraint and throughout we refer to this as *time-series consistency*.

We follow Bates (2000) and use options on S&P 500 futures.¹¹ It is straightforward to characterize futures prices, F_t , in our setting. Under \mathbb{Q} , the futures price, F_t , solves

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

and the volatility evolves as in equation (4). As discussed in Whaley (1986), 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. The S&P 500 futures options are American, and we use the procedure in Appendix A to account for the early exercise feature.

2.2 Existing Approaches and Findings

One approach in the literature uses a long time series of returns to estimate models and diagnose misspecification. ABL, CGGT, and EJP estimate models with stochastic volatility, jumps in prices, and in the latter two papers, jumps in volatility. ABL use S&P 500 index 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 index returns and examine models with stochastic volatility, jumps in prices, and jumps in volatility. CGGT find strong evidence in support of stochastic volatility and jumps in prices, but

¹¹The data on the SPX contract (European options on the S&P 500 index) are not available after 1996 because the Berkeley Options Database was discontinued at that time.

find little evidence supporting jumps in volatility. In contrast, EJP find strong evidence for stochastic volatility, jumps in prices, jumps in volatility. There is also evidence using other approaches for jumps in prices, see, e.g., Aït-Sahalia (2002), Carr and Wu (2003b), or Haung 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.¹²

There is a similar disagreement regarding specification among studies using option prices. BCC calibrate the SV and SVJ models to a cross-section of S&P 500 index options from 1988 to 1991. They find strong evidence for both stochastic volatility and jumps in prices and that adding jumps to the SV model reduces pricing errors by 40 percent. BCC find that the SVJ model is still misspecified. Bates (2000) finds that adding price jumps to the SV model improves the fit by about 10 percent, but only about 2 percent if time series consistency is imposed. Bates (2000) finds misspecification in all models considered and suggests adding jumps in volatility.

Pan (2002) uses up to two options per day and index returns sampled weekly from 1989 to 1996. Her tests indicate that the SVJ model outperforms the SV model in fitting returns and that SVJ model performs better than the SV model for certain strike/maturity option categories but worse for others. Eraker (2004) analyzes S&P 500 index options from 1987 to 1991. He finds that jumps in prices and volatility improve the time series fit, but finds no in-sample option pricing improvement. These mixed results are surprising in the sense that the time series evidence overwhelmingly points towards the presence of jumps in prices. One potential cause of these disparate results is the fact that the above studies use different sample periods, portions of the cross-section and test statistics.

¹²It is important recognize that the findings supporting jumps in prices are always relative to a volatility specification. Some of these specifications assume constant volatility or a linear diffusion specification. Even the models with jumps in volatility assume a simple exponential distribution. It would be interesting to see if more complicated volatility jump specifications would affect conclusions regarding jumps in prices.

Regarding factor risk premia, there are also disparate results. First, there is no theoretical argument that the diffusive volatility risk premium should be of a particular sign. 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, although this is also consistent with price jump or volatility jump risk premia. Branger and Schlag (2004) argue these tests are not powerful. The studies formally estimating diffusive volatility risk premia obtain conflicting results, depending on the dataset and the model specification. 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, finally, 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 not significant, and sets this parameter to zero.¹³ Pan (2002) does find a statistically significant mean price jump risk premium (18 percent). Eraker (2004) estimates $\eta_v = -0.01$ in the SVJ and SVCJ models, but finds the mean price jump risk premium 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 one or two percent in terms of Black-Scholes implied volatility. Finally, Driessen and Maenhout (2004a) develop a multifactor APT style model to quantify volatility and jump risk. They

¹³Interestingly, Pan (2000), in an earlier working paper version of Pan (2002), reports a value of $\eta_v = 0.029$ in the SVJ model, which is insignificant, but of the opposite sign when compared to the SV model.

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

3 Our Approach

3.1 Consistency between returns and option prices

The model in equations (1) to (4) 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). They both find that option based estimates of σ_v are much larger and estimates of ρ tend to be more negative than those obtained from the time series. This inconsistency implies either a misspecified model, or that one data source is not particularly informative about the parameter.

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). Their procedures appropriately weigh each data source, simultaneously addressing a model's ability to fit the time series of returns and the cross-section of options. Despite these advantages, there is a crucial drawback: computational burdens severely constrain how much and what type of data can be used. As noted earlier, for example, Pan (2002) and Eraker (2004) use a small number of options and short data samples. While these procedures are, in principle, efficient in their use of both options and returns data, they are inefficient as they use short time spans and/or only a few option prices. For example, it is

important to include deep out-of-the-money options to accurately estimate jump risks.

Our approach is a pragmatic compromise between the competing constraints of computational feasibility and statistical efficiency. For the parameters 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 allows us to use the information embedded in a long time series of returns and the information embedded in the entire cross-section of option prices over a long time span (1987 to 2003). The procedure reduces computational burdens and, importantly, constrains parameters that are common under both measures to be equal, thus imposing a restriction of the model. Our two-stage approach is similar to Benzoni (2002) and Duffie, Pedersen, and Singleton (2002).

In the models that we consider, there are only three parameters that are restricted, $\kappa_v\theta_v$, σ_v and ρ . Table 1 summarizes the \mathbb{P} -measure parameter estimates obtained by ABL, CGGT, EJP, and Eraker (2004).¹⁴ 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 is that CGGT and ABL's estimates of σ_v are lower, an implication of their datasets: CGGT use the Dow Jones 30 index and ABL use data from 1980-1996, 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 (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 .

¹⁴For the mean jump size in the SVCJ model, we use the value of $\hat{\mu}_s = \mu_s + \rho_s\mu_v$ reported in EJP which is the expected jump size. EJP find ρ_s to be slightly negative, but it has a large standard error. The resulting mean jump size is very close to that in the SVJ model. Essentially, in the SVCJ model, ρ_s is a nearly redundant parameter as it has a very similar impact on the conditional distribution of returns as μ_s .

Table 1: Objective measure parameters estimated by Eraker, Johannes and Polson (2003) , Andersen, Benzoni and Lund (2002), Chernov, Gallant, Ghysels, and Tauchen (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, κ_v , and λ have to be multiplied by $252\sqrt{252\theta_v}/100$ gives the mean volatility, and $\sqrt{252\mu_v}/100$ gives 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

In our empirical implementation, we use the \mathbb{P} -measure parameter estimates for $\theta_v\kappa_v$, σ_v , and ρ from EJP for a number of reasons. First, their sample (1980 to 2000) is closest to ours (1987 to 2003). Second, their dataset includes the volatile late 1990s, which implies that their estimates reflect the same volatile periods as in our sample. Third, they used S&P 500 index returns and our options are on S&P 500 index futures. Fourth, the EJP estimates generally have the highest σ_v and lowest ρ which generate greater non-normalities, giving

Since $\hat{\mu}_s$ does not appear under the risk-neutral measure, this does not affect our option pricing results.

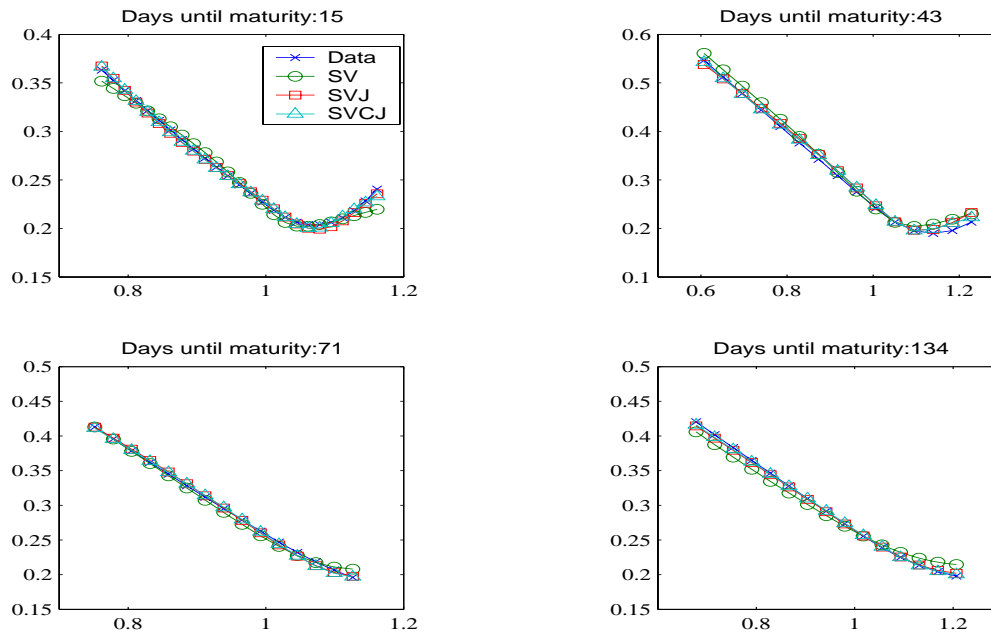


Figure 2: Calibrated implied volatility curves, August 5, 1999. Parameter estimates are obtained using all four curves for each of the models, with no restrictions on the parameters. The units on the X-axis are in terms of the options' moneyness, K/F , and the units on the Y-axis are the annualized Black-Scholes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

the SV model the best chance of success. We later 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 the four stated maturities, for all the strikes, and for each model. 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

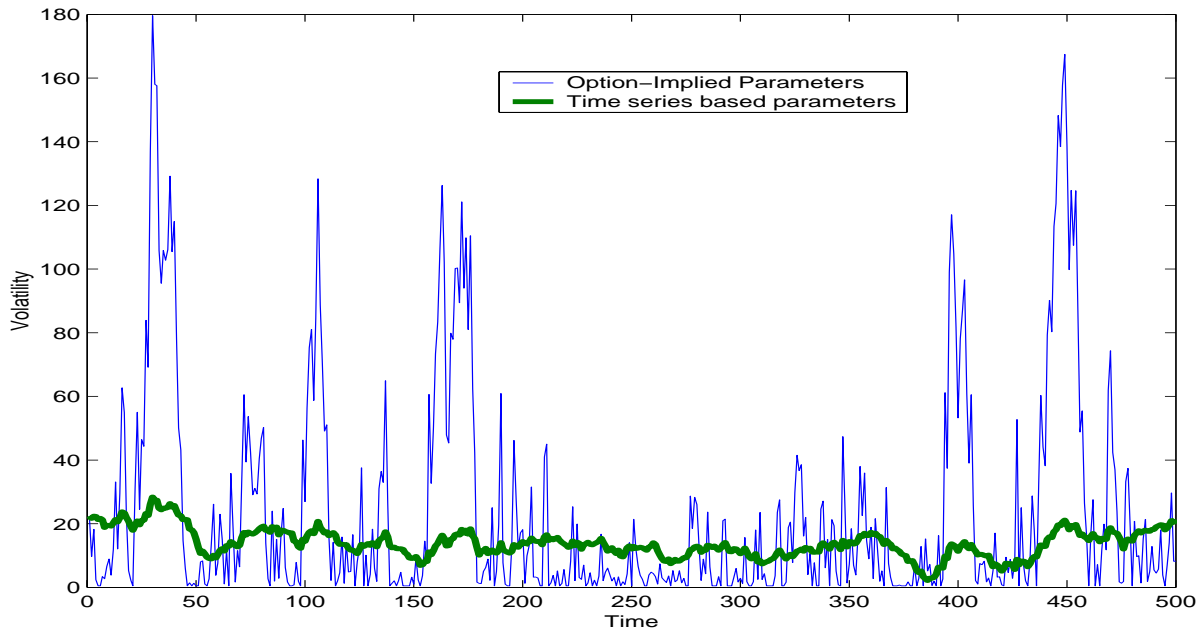


Figure 3: Volatility paths simulated based on options, ($\theta_v = 3.63$, $\kappa_v = 0.06$, $\sigma_v = 2.8$, $\rho = -.66$), and index returns ($\theta_v = 0.90$, $\kappa_v = 0.025$, $\sigma_v = 0.15$, $\rho = -.40$). The time corresponds to two years (500 trading days).

percent, 0.6 percent, and 0.5 percent, respectively. One might be inclined to conclude that there is little, if any, benefit to the more complicated models.

This approach ignores the fact that, for example, 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 time series estimates: the calibrated σ_v is 2.82 in the SV model, while Table 1 indicates that the highest reported σ_v from time series studies is 0.14! Bates (2000) first noted this problem and suggested 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

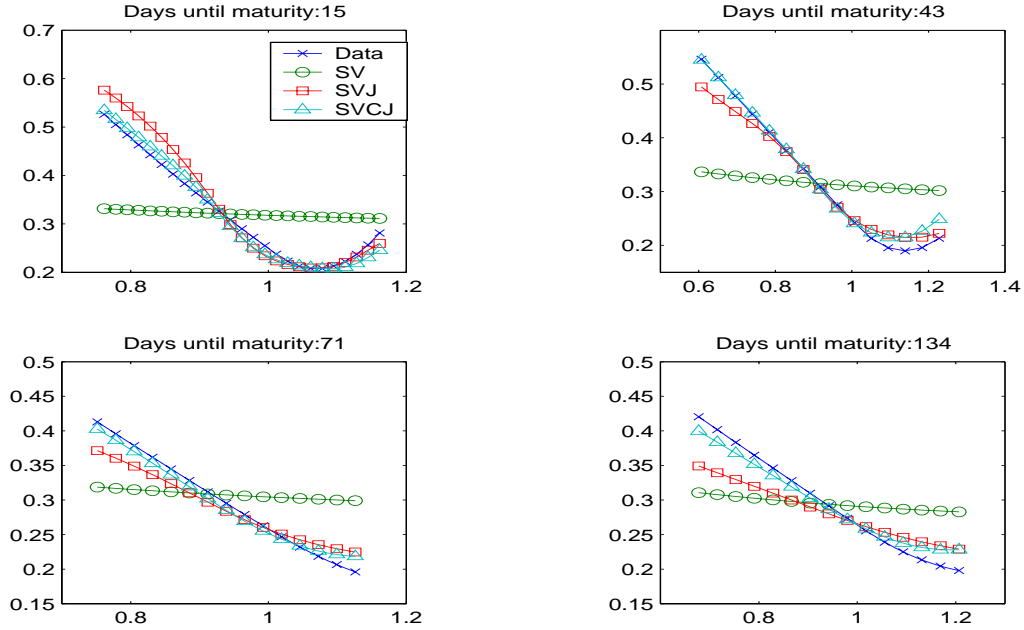


Figure 4: Calibrated Black-Scholes implied volatility curves, August 5, 1999. Parameter estimates are obtained using all four curves for each of the models, constraining the \mathbb{P} -measure parameters to be equal to their time-series counterparts.

fit the market data, the resulting parameter estimates are implausible and inconsistent with 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 noted in Anderson, Hansen, and Sargent (2000) and by others, 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 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 1). The RMSE’s for the SV, SVJ and SVCJ models are now: 8.73 percent, 2.97 percent, and 1.43 percent, 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 it cannot, once time series consistency is imposed, generate sufficient amounts of conditional skewness and kurtosis.¹⁵

3.2 Time series tests

Option prices are highly informative about spot volatility, and in this section we develop an intuitive test to detect volatility jumps. Our testing approach is similar in spirit to those implemented in Pan (2002), Johannes (2004), Jones (2003) and Eraker (2004) 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.¹⁶ Given the implied variances, we compute standard measures of tail behavior – skewness and kurtosis. As the models that we consider have state-dependent diffusion coefficients, we focus on “conditional” skewness and kurtosis using the standardized increments:¹⁷

$$V_{skew} = skew^{\mathbb{P}} \left(\frac{V_{t+1} - V_t}{\sqrt{V_t}} \right) \text{ and } V_{kurt} = kurt^{\mathbb{P}} \left(\frac{V_{t+1} - V_t}{\sqrt{V_t}} \right). \quad (5)$$

¹⁵The 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.

¹⁶Bates (1996), when discussing tests based on implied volatilities, argues that “the volatilities used should ideally be inferred using a stochastic volatility option pricing model” (p.39).

¹⁷The motivation for the conditional measures is that if V_t followed 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.

Unconditional measures of skewness and kurtosis provide the same conclusions. Given the persistence and heteroscedasticity of volatility, the conditional statistics will likely have greater power for detecting misspecification than the unconditional ones and will account for any level effects. 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 (in our case volatility) and recommend a Monte Carlo or bootstrapping approach. We follow their recommendation and simulate G samples paths from the null model, $\{V_t^g\}_{t=1}^T$ for $g = 1, \dots, G$ and then, for each path, compute each of the statistics. In practice we choose $G = 1000$.

To implement this procedure, we use the \mathbb{P} -measure parameters estimated from the time-series of returns, specifically 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 1. 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.

3.3 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 mentioned earlier, our conclusions regarding the relative

merits of the models does not depend on the choice of $\Theta^{\mathbb{P}}$ parameters and the reason is that 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 1.

We use the following objective function to estimate parameters and variances,

$$\left(\widehat{\Theta}^{\mathbb{Q}}, \widehat{V}_t\right) = \arg \min \sum_{t=1}^T \sum_{n=1}^{O_t} \left[IV_t(K_n, \tau_n, S_t, r) - IV(V_t, \Theta^{\mathbb{Q}} | \Theta^{\mathbb{P}}, K_n, \tau_n, S_t, r) \right]^2 \quad (6)$$

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.

Any estimation procedure requires assumptions on the functional form of the objective function and the number and type of option contracts used. These choices are important as they determine the weight placed on contracts with different strikes and maturities. We minimize squared differences of model and market Black-Scholes implied volatilities in (6).¹⁸ 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 heavy weight on expensive in-the-money and long maturity options. Others advocate discarding all in-the-money options for this reason (Huang and Wu 2004).

The second component in the objective function is the choice of option contracts, effectively a choice of K_n and τ_n . Since it is not possible to observe option prices of all strikes and maturities simultaneously, there are two ways to construct a dataset: use close prices or sample options over a window of time. We follow Bates (2000) and aggregate trades during the day. Bates (2000) chooses a three-hour window and we extend this window in

¹⁸See Christoffersen and Jacobs (2004) for a detailed discussion of the objective function choice.

order to capture as many transactions as possible. Since volatility and price and volatility jump risk premia are identified from longer-dated options and deep out-of-the-money (OTM) options, respectively, it is important to include as many of these as possible.

There are hundreds of option transactions each day and using all of the them generates a number of issues. First, it is computationally infeasible to use all of the data. Second, since volatility changes during the day, the recorded transactions could have different volatilities. Third, the vast majority of the recorded trades in our sample involve short maturity at-the-money (ATM) options. Equal weighting would effectively place a very high weight on short maturity ATM options. These contracts are the least interesting economically because all the models can provide similar ATM prices. As outlined in Appendix B, we take all daily transactions, fit a flexible parametric curve, and then use equally spaced observations from this 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 which 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 options. 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 3.1 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. This explains why it is important to be careful when interpreting volatility or risk premia estimates in a model which is clearly misspecified, based on, for example, time series evidence.

Another issue with cross-sectional estimation is assessing statistical significance. As

noted by Bates (2000, p.195): “A fundamental difficulty with implicit parameter estimation is the absence of an appropriate statistical theory of option pricing errors.” This, in particular, means that it is difficult to assign standard errors to parameters estimated using the cross-section. We overcome this difficulty using an extremely computationally expensive, but intuitive nonparametric bootstrapping procedure. We randomly select forty Wednesdays from 1987 to 2003 and estimate the risk premia and spot variances. We then repeat this procedure (re-estimating risk premia and spot variances) until the bootstrapped standard errors did not change appreciably. We found fifty replications to be sufficient (this is consistent with Efron and Tibshirani (1994)). We report root-mean square-errors between the model fit and our interpolated implied volatility curves across all of the fifty replications, which provides a large sample of option transactions. We also report an “out-of-sample” experiment by randomly selecting 50 days, re-estimating the spot variance (holding the risk premia estimates constant), and evaluating RMSE’s.

4 Empirical Results

4.1 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 is (a) close to maturity (to minimize the American feature), (b) at-the-money, (c) not subject to liquidity concerns, (d) is an actual transaction (not recorded at the open or the close of the market), (e) has a recorded futures transaction occurring at the same time, and (f) 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. Whaley (2000) describes

the VIX index. Although not reported, we also computed 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 follow Bates (2000) and Pan (2002) and use interpolated Treasury bill yields as a proxy for the risk-free rate.

Table 2 summarizes the implied volatilities and scaled returns for the different datasets and models. In the first panel, the first two rows labeled “VIX” provides summary statistics for the VIX index (including and excluding the crash of 1987), the rows labeled “Calls” provide statistics for our representative call option dataset, and the rows labeled “Interpolated” uses the ATM implied volatility interpolated from all of the daily transactions.¹⁹ In this first panel, the implied volatility is based on the Black-Scholes model (BSIV) and the subsequent panels report model (as opposed to Black-Scholes) based implied variances.

Although there are some *quantitative* differences across datasets, the *qualitative* nature of the results is identical: 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 datasets are due to differences in underlying indices (the VIX is based on 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). Notably, the call option dataset generally is less non-normal than the other two. In the post-crash sample (the rows labeled 1988 to 2003), the differences between the datasets are minimal.

For the formal tests, we use the call option data set. Our conclusions are robust to the use of the other datasets, although the call option data set has fewer problems (interpolation, averaging effects, stale quotes, etc.). The bottom three panels in Table 2 report

¹⁹The VIX used to be calculated from the S&P 100 index options instead of the S&P 500, although this was recently changed. As of September 22, 2003, the VIX index changed and now uses options on the S&P 500 index. Because of the new construction, it can not be interpreted as an option implied volatility anymore. In this paper we use the old VIX index (current ticker VXO).

Table 2: The table reports the summary statistics implied from market prices of option and S&P 500 returns. The first three rows provide summary statistics for variance increments and standardized returns using the VIX index, our call-option dataset generated from options sampled between 10 and 11 a.m. CST (as described in Appendix A), and a dataset based on interpolated ATM implied volatility using the entire daily cross-section (using the procedure in appendix B). In these three cases, the variance used is from the Black-Scholes model. The second third and fourth panels contain model implied variances (V_t) for the SV, SVJ and SVCJ models which include observations assuming options are priced based on the objective measure and we also include risk premiums (RP) and document the effect of increasing σ_v in the SV model.

Model/Data	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 ($\sigma_v = .2$)	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

statistics based on model based implied variances for the SV, SVJ and SVCJ model using three sets of parameters. The first set is from EJP who, as mentioned earlier, report higher σ_v estimates than other papers. The parameter σ_v primarily controls the kurtosis of the volatility process, thus this set of parameters gives the SV and SVJ models the best chance of success. For robustness, we include the statistics of implied volatility for two additional parameter sets. The results in the rows labeled ‘RP’ incorporate risk premia, to gauge their impact on implied variances.²⁰ The third set of results uses the SV model and $\sigma_v = 0.20$.²¹

Table 3 provides quantiles of the finite sample distribution for each of the statistics using the Monte Carlo procedure described in the previous section and for each of the model/parameter configurations. Note that these are results simulated under the \mathbb{P} -measure, thus there are no separate entries for the cases incorporating risk premia, as the \mathbb{P} -measure behavior is not changed.

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 99 percent quantile which is orders of magnitude lower than the value observed in the data (around 1000). Similarly, the SV model cannot generate the large positive skewness observed in the data. We also note in passing that the SV model, not surprisingly, cannot come close to generating the observed non-normalities 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 this, we show that the results are unchanged even if we ignore the crash of 1987, if we account for volatility risk premia, and if we increase σ_v . The rows labeled ‘1988 to 2003’ in Table 2 provide the statistics from 1988 to 2003 a period excluding the crash of 1987.²² Even

²⁰We use the risk premium estimates obtained later in the paper.

²¹This value is roughly five standard deviations away from the point estimate in EJP in the SV model.

²²We do not, of course, advocate “throwing out” data, especially outliers events in jump models. Since jumps are rare, these tail observations are invaluable for characterizing jumps. However, in this setting, the

Table 3: The table reports the summary statistics for the variance increments and standardized returns for the SV, SVJ and SVCJ models. 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 1.

	Quantile	V_{kurt}	V_{skew}	R_{kurt}	R_{skew}
SV simulations	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 simulations $\sigma_v = 0.2$	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 simulations	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 simulations	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
	0.95	1649.03	34.87	78.90	-4.15
	0.99	2500.51	43.73	175.67	-6.66
SVCJ simulations $\mu_v = 0.85$ $\lambda = 0.0026$	0.01	3.28	0.20	3.02	0.06
	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

based on the post 1987 sample, the SV model is still incapable of generating these values, even though the parameters used were 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 are lower which implies that the model generates even less non-normalities. The conclusion is unchanged if we again increase σ_v to 0.20.

Finally, the row labeled ‘SV model (RP)’ in Table 2 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 of volatility and this can have a significant impact on implied variance in periods of very high volatility (e.g., in October 1987). Risk premia, however, cannot explain the non-normalities in the observed data. This is most clear in the post-crash sub-sample, where 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.

The SV model is rejected because over short intervals,

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

For small time intervals, the increments are approximately conditionally normal (e.g., as shown in Table 3). The data, however, are extremely non-normal, and thus the square-root diffusion specification has no chance to fit the observed data.

The following example provides the intuition and shows specific magnitudes. Consider the mini-crash in 1997: on October 27th the S&P 500 fell about 8 percent with Black-Scholes implied volatility increasing from 26 to 40 percent. In terms of variance increments, daily variance increased from 2.69 to 6.33, which translates to 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 three months which was 0.151 (re-

post 1987 sample highlights the severe problems with the square-root process in the SV and SVJ models.

markably close to the value $\sigma_v = 0.14$ used above). Thus for the SV model to generate this move, it requires a roughly 16 standard deviation shock. Even if we doubled σ_v , it requires an eight standard deviation move. This example shows the fundamental incompatibility of the square-root specification with the observed data. It is not an issue of finding the right parameter values: the model is incapable of explaining the observed data. This is not an isolated incident, Whaley (2000) provides other examples of volatility “spikes.”

The SV and SVJ models shares the same square-root volatility process, suggesting the SVJ model is also incapable of fitting the observed data. The third panel in Table 3 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. Sub-samples 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 2 and 3 demonstrate that the SVCJ model is capable of capturing both the behavior of implied variances and standardized returns for the full sample. In Table 3, 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 1000 and the model generates these values somewhere between the 50th (about 600) and 95th quantile (1600). The skewness generated by the model is almost identical to the value observed in the data: the 50th quantile is in simulations is 21.04 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 2 indicates 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 sub-samples in Table 2 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. For example, this can be seen by comparing V_{kurt} from 1988 to 2003 in Table 2 with the 1st quantile for the SVCJ model in Table 3. This is not at all surprising since the base case parameters in EJP were estimated using data that included the crash of 1987. Jumps are rare events that generate conditional non-normalities and, therefore, if one removes these outliers, the observed data (variance increments or standardized returns) will become more normal, by construction. Of course this does not indicate that the SVCJ model is misspecified, but only that using the parameters estimated using the full sample, 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 3 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 (5,95) percent confidence band.²³ Again, we want to emphasize that we do not advocate fitting models with jumps on samples omitting large movements. We also do not suggest using these ad-hoc parameters; we use them only to illustrate the flexibility of the SVCJ model. The key point is to contrast this result with those obtained in the SV and SVJ models: these models, due to the diffusion specification, could not fit the data over *either* of the samples, even using parameters σ_v that were 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

²³This value of λ is not implausible as it is consistent with studies using 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 ($\lambda = 0.002$).

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 square-root and a constant elasticity of variance model could not explain the dynamics of implied volatility. Our procedure confirms Jones (2003) findings, and, in addition, we provide a model with jumps in volatility that is capable of capturing the observed dynamics.

4.2 Model Specification and Risk Premium Estimates

Table 4 provides risk premium estimates and overall option fit for each model, and Table 5 evaluates the significance of the model fits. We estimate the risk-neutral parameters and compute total RMSE using the procedure in Section 3.3. While the time series results indicated 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 five, ten, fifteen or forty percent.

As mentioned earlier, it is difficult to statistically identify η_v and so we report the results for versions of the models with or without a diffusive volatility risk premium. As mentioned earlier, throughout this section, we constrain $\lambda^{\mathbb{Q}}$ to be equal to λ and use the value from EJP. We originally included this variable, but it is not identified. In general, it is only possible to estimate the compensator, $\lambda^{\mathbb{Q}}\overline{\mu}_s^{\mathbb{Q}}$, and not the individual components separately. Pan (2002) and Eraker (2004) imposed the same constraint. The estimate in EJP implies about 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.

Table 4: 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 that $\sigma_s = \sigma_s^Q$.

	η_v^Q	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

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

4.2.1 Model Specification: Pricing Errors

A number of points emerge from Tables 4 and 5. 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 4, as the RMSE's of the

pricing errors are reduced by almost fifty percent, and is also true based on the bootstrapped samples. Under any of the assumptions considered on the risk premia and in all fifty bootstrapped samples, the SVJ and SVCJ models provide at least a fifteen percent pricing error improvement over the SV model and in most samples, more than 40 percent. These results are consistent with BCC (who find a 40 percent improvement) and in contrast to Bates (2000), Eraker (2004), and Pan (2002). The reason for our clear results are twofold: we impose consistency, and we use option prices spanning a long time period.

Next, consider a comparison of the overall pricing errors in the SVJ and SVCJ models in Table 4. The SVJ model, with no constraints on risk premia, has an overall pricing error of 3.48 percent. By comparison, the unconstrained SVCJ model has pricing errors of 3.31 percent, an improvement of five percent. If we impose that $\sigma_s^{\mathbb{Q}} = \sigma_s$ in the SVJ model, the SVCJ model generates a larger improvement of 18% (4.08 percent versus 3.36 percent). If $\sigma_s^{\mathbb{Q}} \neq \sigma_s$, $\sigma_s^{\mathbb{Q}}$ 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 ($\sigma_s^{\mathbb{Q}} = \sigma_s$ in the SVCJ model) shows that the pricing errors are quite similar: 3.48 percent and 3.36 percent, respectively. This result is consistent with objective measure time series results in Table 3 which show that the SVJ model can generate realistic amounts of non-normalities in returns through jumps in prices.

In Table 5, 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 $\sigma_s^{\mathbb{Q}}$, and the SVCJ model assuming $\sigma_s^{\mathbb{Q}} = \sigma_s$. 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 ($\sigma_s^{\mathbb{Q}} = \sigma_s$) model has lower pricing errors by at least five percent than the SVJ ($\sigma_s^{\mathbb{Q}} = \sigma_s$) model in 96 percent of the bootstrapped replications. Thus, the SVCJ model provides a statistically significant improvement in overall option fit. However,

Table 5: Comparison of the RMSE 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 the parantheses are out-of-sample. 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	$\sigma_s^Q \neq \sigma_s$	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	$\sigma_s^Q = \sigma_s$	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	$\sigma_s^Q = \sigma_s$	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)	

a comparison of the unconstrained SVJ model with the constrained SVCJ model indicates that the SVCJ ($\sigma_s^{\mathbb{Q}} = \sigma_s$) model outperforms by more than five percent in only six percent 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 insure the robustness of our findings, we have also evaluated the RMSE's "out-of-sample" by using the parameter values reported in Table 4 on 50 randomly selected days in our sample which were not used for risk premia estimation. For each of the 50 days, we estimated the spot variance and then computed RMSE's, 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.²⁴

We sorted option pricing errors by maturity, strike and volatility. The only major pattern that emerged is 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 provide a much better fit for all categories than the SV model, 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 as the main cross-sectional implication of jumps in volatility (the hook) can be achieved with large risk premia in a simpler model. However, as indicated by our tests in the previous section, 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 cross section.

²⁴In the interest space, the unconstrained SVCJ model results were not included but the unconstrained SVCJ model modestly outperformed all of the model in and out of sample.

4.2.2 Estimates of Risk Premia

Table 4 summarizes the \mathbb{Q} -measure parameter estimates. For each parameter and specification, the table provides estimates and bootstrapped standard errors based on 50 replications. There are a number of interesting findings.

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

Why is it so hard to estimate η_v even with a long dataset? The main reason is that, as shown in Table 6, the implied volatility term structure is flat, at least over the option maturities that we observe.²⁵ Focussing on the whole sample, the difference between short (less than one month) and longer-dated (three to six months) implied volatility is only about 0.1 percent. 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.

The stochastic volatility model could generate a flat average term structure via two conduits. First, η_v could be small which implies that the term structure is flat on average. Second, η_v could be large (of either sign) but the term structure will 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

²⁵We thank Chris Jones for useful discussions regarding the term structure of implied volatilities.

could only distinguish these two competing explanations if we had 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. While true, joint estimation still is confronted with the fact that the term structure is flat. This implies that merely using a different estimation procedure will not likely alleviate the problem that η_v is insignificant, unless, the procedure incorporates long-dated options. This, in large part, explains why the existing literature using options and returns gets either 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 percent (imposing $\sigma_s^{\mathbb{Q}} = \sigma_s$) to about -5 percent ($\sigma_s^{\mathbb{Q}} \neq \sigma_s$). In the SVJ and SVCJ models, estimates of μ_s are generally around -2 to -3 percent based on the time series of returns, which implies a modest mean price jump risk premium of about 2 percent to 6 percent. In the SVCJ model, the estimates of $\mu_s^{\mathbb{Q}}$ are again significant and generate a risk premium of similar magnitude, about 2 percent to 3 percent when ($\sigma_s^{\mathbb{Q}} \neq \sigma_s$) and 4 percent to 5 percent 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 percent to high of -4.82 percent. 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.

Third, in both the SVJ and SVCJ models, there is strong evidence that $\sigma_s^{\mathbb{Q}} \neq \sigma_s$. This premium attached to the volatility of price jumps is an effect not previously documented. Moreover, as noted in the previous paragraph, this has important implications for the magnitudes of premium attached the mean price jump size. Estimates of σ_s in the SVJ model are around 4 percent, while estimates of $\sigma_s^{\mathbb{Q}}$ are more than 9 percent. However, it

Table 6: The term structure of S&P 500 index 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 one month, from one to two months and from three to sixth 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

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 percent (with a standard error of 0.7) in the unconstrained SVCJ model and around 7.5 percent (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.

Fourth, for every variant of the SVCJ model, we find 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: 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

$$E^{\mathbb{Q}}[V_t] = \frac{\kappa_v \theta_v + \mu_v^{\mathbb{Q}} \lambda}{\kappa_v + \eta_v}. \quad (7)$$

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 percent, an average sized jump increases annualized volatility to 25 percent under \mathbb{P} and 34 percent under \mathbb{Q} . Even if $\sigma_s^{\mathbb{Q}} = \sigma_s$, the estimate

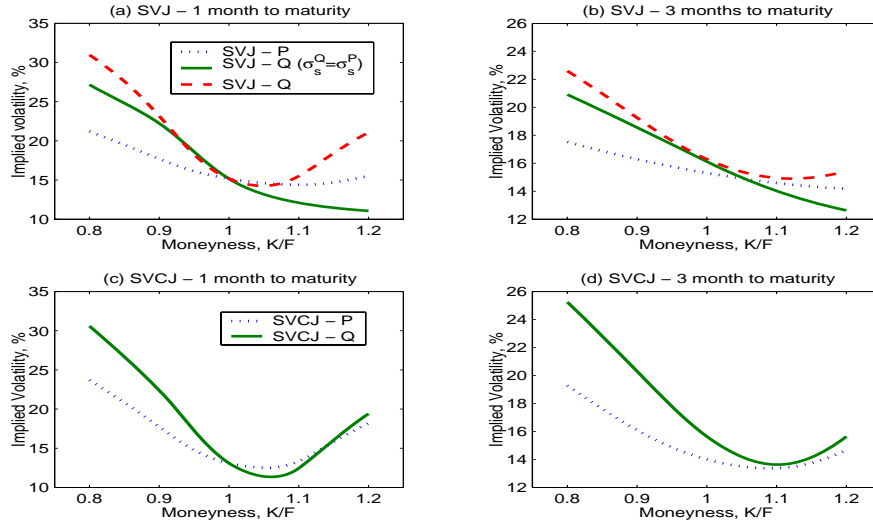


Figure 5: Black-Scholes implied volatility curves for the SVJ and SVCJ models based on \mathbb{P} and \mathbb{Q} measure parameters.

of 5.29 implies that an average sized, risk-neutral jump increases volatility to 39 percent, 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 one and three 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, but 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 where $\sigma_s^{\mathbb{Q}} = \sigma_s$ 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 with using \mathbb{P} -measure parameters or with $\sigma_s^{\mathbb{Q}} = \sigma_s$ generate monotonically declining implied volatility curves (the smirk). These models generate very little risk neutral conditional kurtosis and they therefore miss the hook, or 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 a sufficiently flexible risk premia, one cannot distinguish different models based on options cross-section only. A simpler model with elaborate risk premia will likely be favored. 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.

4.2.3 Interpreting the risk premia

In this section, we examine the “reasonableness” of these premia and document their economic implications. To do this, we discuss the mean price jump size premia in the context of some simple equilibrium models, characterize the contribution of jump risk to the overall equity premium, and examine the impact of price and volatility jump risk premia on option returns.

First, consider Bates’ (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 risk-aversion parameter in a power utility and $\delta_{S,W}$ is the correlation between the jumps in stock market and total wealth. It is reasonable to assume that jumps are highly correlated with wealth (at least financial wealth) and therefore even small values of A can easily generate the two to six percent wedge that we find. Levels of risk aversion under ten are generally considered to be reasonable (Campbell, Lo and MacKinlay, 1997). Bates (2000) estimates $\mu_s^{\mathbb{Q}}$ to be around nine percent and $\sigma_s^{\mathbb{Q}}$ to be around 10 or 11 percent 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 2.1, the contribution of price jump risk premia is 2.7 and 2.9 percent per annum in the SVJ and SVCJ models (with $\eta_v = 0$), respectively. Over our sample, the equity premium was about eight percent, implying that jumps generate roughly one-third of total premium. As a benchmark, time series studies find that jumps in prices explain about 10 to 15 percent 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 which identifies a “put-pricing puzzle” (see, e.g., Bondarenko 2003, Driessen and Maenhout 2004b, Jones 2005, or Santa-Clara and Saretto 2005). These authors document, using data similar to ours (Bondarenko uses S&P 500 futures options from 1987 to 2000), that average monthly returns of ATM and OTM puts range from -40 to -95 percent, 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 are also puzzling from a portfolio perspective (Driessen and Maenhout 2004b) or based on a nonlinear factor model (Jones 2005).

Our models and risk premium estimates provide a natural setting to explore a risk-based explanation of the “put-pricing puzzle.” To examine this issue, we simulate from the SVCJ model and calculate put returns with and without risk premia and compare them to results previously reported. Table 7 provides these average values and the (5,95) percent bootstrapped confidence bounds reported by Bondarenko (2004). To compute model based

Table 7: The impact of risk premia on options returns in the SVCJ model. We compare out-of-the-money (OTM) average options 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- P) and using the estimated risk premia (SVCJ- Q). The dagger (\dagger) denotes returns outside the confidence intervals.

		Moneyess	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- P		-20.70 \dagger	-21.91 \dagger	-21.78 \dagger	-19.96 \dagger
	SVCJ- $Q-\mu_s^Q$		-68.46 \dagger	-58.36	-45.21	-32.18
	SVCJ- $Q-\mu_v^Q$		-75.78 \dagger	-64.67	-50.01	-35.27

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 one 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, and the case SVCJ- $\mathbb{Q}-\mu_s^Q$ considers only the effect of a mean price jump risk premium and SVCJ- $\mathbb{Q}-\mu_v^Q$ adds the volatility jump risk premium. In each case, we match the overall average risk premium over the sample (eight percent).

Table 7 provides a number of interesting implications. The SVCJ- \mathbb{P} results indicate that a model without any risk premia generates about -20 percent 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 percent return to the writer. Thus, a large component of the put pricing puzzle is due to the (also puzzling) 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 for the deep OTM puts, and here the returns are economically close.²⁶

We conclude that a 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 non-normal 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 have an important economic implications impact, we are agnostic about the exact sources of these jump risk premia. There are a number of potential sources: asymmetric utility of gains and losses, which leads investors to care more about large negative returns; the presence of heterogeneous investors with more risk averse investors buying put options from the less risk-averse; or alternatively, institutional explanations. Some of these are already being explored. Pan, Liu, and Wang (2005) argue that an aversion to parameter uncertainty could generate the premia. Bollen and Whaley (2004) show that there are

²⁶The 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 moneyness category (see Table 2 in Bondarenko).

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 heterogenous investors. Disentangling the sources of the premia and their relationship to investor demand appears to be a fruitful avenue for future research.

5 Conclusions

In this paper, we have used the time series and cross-section of option prices to address a number of important option pricing issues. Using the time series, we find extremely 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. Based on 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 jumps in prices is allowed to have volatility of price jumps risk premium. In this case, the SVJ requires a relative large premium attached to the volatility of jumps in prices to generate the substantial amounts of risk-neutral kurtosis required to fit the cross-section.

We find that estimates of risk-neutral mean price jumps are consistent across models, and are on the order of five to seven percent and are highly statistically significant, implying a mean price jump risk premium of about two to five percent. 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 three percent to the total equity premium, which was eight percent 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 in the 1990s. We find that a large proportion of the puzzle can be explained by the high returns on the underlying index and the remaining proportion can be generated by our modest price and volatility jump risk premia.

Our results resolve a number of different existing issues in the literature, but, to conclude, we mention three topics which are promising avenues for future research. First, joint estimation and new types of data (variance swaps, volatility futures, etc) would improve parameter estimates, especially for 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 observations indicates that this may be a tenuous assumption, as there are long periods when 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. This 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 have found some 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.

A Adjusting for the Early Exercise Premium

The CME S&P 500 futures options are American and this complicates the parameter estimation procedure because of the considerable additional time to compute model American prices versus model European prices. We circumvent this extreme computational difficulty by transforming market prices from American to European, and then estimating model parameters on European prices. This has two main advantages. First, the savings in computation time are very significant and make the parameter estimation procedure feasible. Second, it eliminates the need to develop analytical approximations for American option values under each model. Next we quantify the magnitude of the approximation error introduced by this procedure.

We assume market prices are generated by a particular model, e.g., the SV model. We 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, i.e., a value σ^{BS} so 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 tested 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 found in our data set. Similar definitions apply for the SVJ model. For the SV model, we tested all combinations of the parameters which straddle the estimated values in Table 8: $\kappa_v \in \{0.008, 0.032\}$, $\sigma_v \in \{0.1, 0.25\}$, $\rho \in \{-0.1, -0.7\}$, $\theta_v \in \{0.57, 1.59\}$, $V_t \in \{0.57, 1.59\}$, $r \in \{0.008\%, 0.03\%\}$, and $\tau \in \{1.5, 6\}$

(months). We chose rather long maturities as the exercise premium is increasing in maturity. Given the initial stock price $S_t = 100$, for put options the strikes tested are from the set $K \in \{85, 90, 105\}$ and for call options $K \in \{95, 100, 115\}$. A total of 768 combinations of SV model and option parameters were tested. For the SVJ model, we tested 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.²⁷ A total of 6144 combinations of SVJ model and option parameters were tested.

Accurate American prices were determined with two-dimensional finite difference routines. These finite difference routines are hundreds of times slower than the Fourier inversion methods we used 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, a relevant error measure is absolute error, while for higher priced options, relative error is more relevant. We choose a price of \$0.5 as the separator between low-priced and high-priced options, as was done in previous studies (e.g., Broadie and Detemple, 1996).

Table 8 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 range are less than 0.4 percent while the maximum absolute relative error was 1.4 percent. To put these numbers in perspective, a \$10 option with the largest observed relative error of 1.4 percent relative error has an

²⁷In terms of annual decimal parameters, the values are: $\kappa \in \{2, 8\}$, $\sigma_v \in \{0.1, 0.25\}$, $\rho \in \{-0.1, -0.7\}$, $\theta \in \{(0.12)^2, (0.2)^2\}$, $V_t \in \{(0.12)^2, (0.2)^2\}$, $r \in \{0.02, 0.08\}$, $\lambda = \{1, 2\}$, $\tau \in \{30/252, 120/252\}$. The parameters of the jump distribution are unchanged.

Table 8: European Price Approximation Errors. We tested the accuracy of the procedure by computing absolute and relative errors over a large set of model parameters chosen to be representative of those found in 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		SV				SVJ			
Error		Relative		Absolute		Relative		Absolute	
		RMS	Max	RMS	Max	RMS	Max	RMS	Max
Call	1.5 mo	0.17%	0.26%	0.0006	0.0017	0.28%	0.71%	0.0014	0.0044
	6 mo	0.26%	1.08%	0.0029	0.0065	0.37%	0.85%	0.0020	0.0035
Put	1.5 mo	0.05%	0.11%	0.0008	0.0034	0.10%	0.42%	0.0013	0.0060
	6 mo	0.30%	1.26%	0.0020	0.0045	0.36%	1.37%	0.0023	0.0064

absolute price error of only \$0.14, which is smaller than typical bid-ask spreads. For low-priced options RMS absolute errors are less than \$0.003 while the maximum absolute error was only 0.007. These errors are far smaller than typical bid-ask spreads. Furthermore, the largest errors occur for extreme option parameters not often found 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 dataset 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 9. 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 that approximating an 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 market American prices to equivalent European market prices. This allows us to use computationally efficient European pricing routines which make the large-scale calibration procedure feasible. The results

Table 9: 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 $S_t = \$100, V_t = 1.59, \sigma_v = 0.25, \rho = -0.7, \lambda = 0.008, \mu_s = -5.00\%, \sigma_s = 8.00\%, r = 0.03\%$.

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

in this Appendix show that the approximation error in this procedure are minimal.

B Data Issues

We obtained the daily trade files for the S&P 500 futures and options on the S&P 500 futures from the Chicago Mercantile Exchange (CME). The files record transactions and bid-ask quotes for both the futures and the options. We ignore bid-ask quotes and use transaction prices. Beginning with 2,246,426 option transactions, we were able to find a matching futures transaction within five minutes for 2,081,727 transactions. There were 947,635 call and 1,134,092 put transactions. By maturity, 1,090,462 transactions were under 30 days-to-maturity, 803,971 were between 30 and 89 days to maturity and 187,294 were 90 days or more to maturity. By strike, 412,327 had $K/F < 0.95$, 1,474,517 had

$0.95 < K/F < 1.05$, and 194,883 had $K/F > 1.05$.

S&P 500 futures options do not have the wildcard feature that is present in the S&P 100 options contract. A wildcard feature arises because the trading day in options ends at 3:15 CST even though that the stocks underlying the index cease trading at 3:00 CST. Moreover, the S&P 100 options are American and holders of options can exercise their options until 3:20. If exercised, the options are *cash-settled* based on the index value at 3:00 CST. This generates an additional valuable option for option holders, commonly referred to as a wildcard option (see Fleming and Whaley 1994).

For S&P 500 futures options, Chapter 351A of the CME Rulebook states that “an option may be exercised by the buyer on any business day that the option is traded. Exercise of an option is accomplished by the clearing member representing the buyer presenting a exercise notice to the Clearing House by 7:00 p.m. on the day of exercise.” This seems to indicate that there is a four hour wildcard option. There is not, however, a wildcard feature because the option is not cash settled (according to Chapter 351A section 02.B). Exercised options receive a long or short position in the futures marked at the bid/ask price at the time of exercise. 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 for empirical implementation: a time series of “representative” option prices for the time series specification tests and also the cross-section of option prices for each maturity on each date. For the “representative” option data set we used the following four selection criteria: (1) Select all option-futures transaction pairs that were traded between 9:30 a.m. and 10:30 a.m., had a time difference of less than a minute, and were out-of-the-money by 2 percent or less. If there were 3 or more pairs matching these conditions, the pair that had the median volatility was selected; (2) if less than 3 records were selected in criterion 1, then more records were added to the running total 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 were still in effect. If there were 3 or more pairs matching both criteria 1 and 2, then the 3 records that best satisfy criteria were selected for median computation. Of these three best records, the record with the median volatility was selected; (3) if less than 3 records were selected in criteria 1 and 2, then more records were added to the running total by allowing pairs with any moneyness (ordered by closeness to $K/S = 1$). The time difference condition was still in effect. If there were 3 or more pairs matching criteria 1, 2, and 3, then the 3 records that best satisfy criteria were selected for median computation. Of these three best records, the record with the median volatility was selected; (4) if less than 3 records were selected in criteria 1, 2, and 3, then more records were added to the running total by allowing pairs with any time difference. In this case, records were ordered so that time difference was the lowest (in 5 minute blocks, so that a time difference of 4 minutes and 59 seconds was just as good as a time difference of 1 minute and 1 second). Any ties in the ordering of the 5 minute blocks was broken by choosing the pair closest to the money. If there were 3 or more pairs matching criteria 1,2, 3, and 4, then the 3 records that best satisfy criteria were selected for median computation. Of these three best records, the record with the median volatility was selected. If there were less than 3 pairs, then the ‘best’ record was selected.

For risk premium estimation, we constructed a data set using transactions throughout each trading day using every put and call transaction that can be time-matched within five minutes to a futures transaction. This procedure typically produced hundreds of matched options-futures transactions. We then used the following curve-fitting procedure to combine all of the 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 .²⁸ For each day

²⁸This is the volatility parameter which equates the market option price to the American option price under the Black-Scholes model. These values are easily computed using an iterative solver together with a

and maturity, we fit a piecewise quadratic function to the implied volatilities:

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

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 estimated with least squares. the knot point was allowed to vary on the nearest maturity, but for longer maturities it was fixed at 1. If the maturity had 10 or fewer option transactions, then the piecewise quadratic was replaced with a linear function. Figure 6 shows a representative day, August 6, 1999 with four fitted curves for each of the maturities that traded. The maturities were 14 days, 42 days, 70 days and 133 days and we had 302, 134, 17 and 24 put and call transactions, respectively, for these maturities.

We experimented with a number of other specifications, including piecewise cubic functions, fixed knot points for all maturities, linear and piecewise functions, etc. We evaluated these alternative specifications by a cross-validation procedure. For each maturity on each day, the set of options was divided in half. For each half, parametric curves were fit and the quality (fitting error) was assessed using the other half of the data. The overall fitting error was the average of the two out-of-sample results. For maturities with a small number, say m , data points, $m - 1$ points were used to fit a curve and the quality was assessed on the m^{th} point. The overall fitting error was the average of the m results each time leaving out one data point. The piecewise quadratic approach was the best when there were more than 10 data points: piecewise cubic functions were found to overfit the data (performed relatively poorly on the cross-validation test) while linear functions tended to underfit the data. For 10 data points or fewer, the linear function approach was found to be the best, because quadratic and piecewise quadratic functions tended to overfit the data. We exper-

binomial or finite difference American option pricing routine.

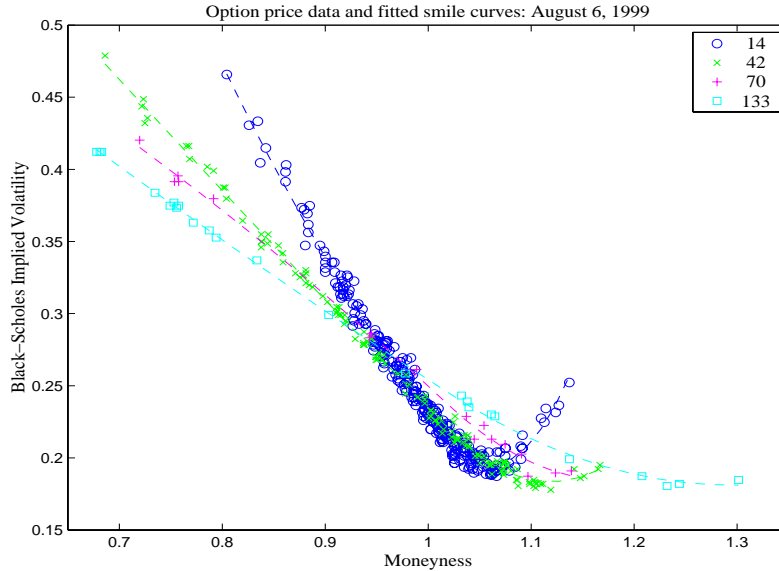


Figure 6: Observed call and put prices for various maturities and the resulting piecewise quadratic fits.

plemented with a number of other cases and specifications. For example, x_0 was allowed to vary within the range of x -values in the data for each maturity and x_0 was fixed at $x_0 = 1$ for all maturities. These generated only minor differences in the implied volatility fits. If the optimizer did not converge on an x_0 inside the range of x -values in the data then x_0 was fixed at 1. If the maturity had less than 10 option transactions, then the piecewise quadratic was replaced with a linear function.

Given the resulting smile curves, option models are calibrated by minimizing squared differences between market and model Black-Scholes implied volatilities using equation (6) and a discrete set of points from these curves. For each day, we chose 25 equally-spaced points for the shortest option maturity and 15 equally-spaced points for longer maturities.²⁹

²⁹Since implied volatility curves tend to flatten at longer maturities (Bates 2000 or Pan 2002). More points were used for shorter maturities to ensure that any increased curvature would be captured. The points vary day by day as the range of strikes and maturities varies throughout the sample.

In effect, option models are calibrated in a two-stage approach. The first stage fits implied volatility curves to transactions data. The second stage finds model parameters which best fit the implied volatility curves. This two-stage approach dramatically reduces the computational requirements without sacrificing accuracy.

The optimization was done using the nonlinear programming package CFSQP (Lawrence, Zhou and Tits 1997). Because the objective function is not globally convex, each optimization problem was solved from multiple diverse starting points to ensure convergence to the global optimum. The starting points were chosen by scaling points from the low-discrepancy Sobol' sequence. Fitting model parameters $\Theta^{\mathbb{Q}}$ and state variables V_t to daily data from 1987 through 2003 was not feasible because of the large number of decision variables and the computation time required to evaluate the objective function.^{30,31} We therefore chose to fit models to subsamples of days throughout the data set. To ensure that the results using a subsample of data did not differ significantly from results that would have been obtained with the full sample, we experimented with different numbers of dates and with random and equally spaced data. We found that our estimates of the optimal model parameters using forty equally spaced dates was adequate, i.e., our estimates did not change significantly with larger subsamples.

³⁰Evaluating the objective function for n days of data requires roughly $70n$ implied volatility computations, each of which requires computing model option prices that involve numerical integration to evaluate inverse Fourier transforms.

³¹Huang and Wu (2003) use a two-stage estimation procedure, alternating between optimizing over model parameters and state variables (i.e., daily volatilities). While significantly reducing the dimension of the search, we found this procedure could converge to local minima that gave substantively different results. In our procedure we jointly optimize over $(V_t, \Theta^{\mathbb{Q}})$ 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*, v.1, 68-123.
- Bakshi, Gurdip, and Charles Cao, 2003, Risk Neutral Kurtosis, Jumps, and Option Pricing: Evidence from 100 Most Active Firms on the CBOE, Working paper, University of Maryland.
- 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, 2002, 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 Maddala, G.S. and C.R. Rao (Eds), *Handbook of Statistics, Vol. 14, Statistical Methods in Finance*. Elsevier, Amsterdam, 567-611.
- Bates, David, 2000, Post-'87 Crash Fears in S&P 500 Futures Options, *Journal of Econometrics* 94, 181-238.
- Bates, David, 2003, Empirical Option Pricing: A Retrospection, *Journal of Econometrics* 116, 387-404.

- Bollen, Nick and Robert Whaley, Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?, *Journal of Finance* 59, 711-753, 2004.
- Bondarenko, Oleg, 2003, Why are Put Options So Expensive? Working paper, University of Illinois, Chicago.
- Benzoni, Luca, 2002, Pricing Options under Stochastic Volatility: An Empirical Investigation, Working Paper, University of Minnesota.
- Bookstaber, Richard and Steven Pomerantz, 1989, An Information-Based Model of Market Volatility, *Financial Analysts Journal* 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.
- Campbell, John, Andrew Lo, and Craig MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press.
- Carr, Peter and Liuren Wu, 2003a, Finite Moment Log Stable Process and Option Pricing *Journal of Finance* 58, 753-777.
- Carr, Peter and Liuren Wu, 2003b, What Type of Process Underlies Options? A Simple Robust Test, *Journal of Finance* 58, 2581-2610.
- 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.
- Chernov, Mikhail, Eric Ghysels, A. Ronald Gallant and George Tauchen, 1999, A New Class of Stochastic Volatility Models with Jumps: Theory and Estimation, Working paper, Pennsylvania State University.
- Chernov, Mikhail, Eric Ghysels, A. Ronald Gallant and George Tauchen, 2003, Alternative Models for Stock Price Dynamics, *Journal of Econometrics* 116, 225-257.
- Conley, Timothy, Lars Peter Hansen, and W. F. 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.
- Christoffersen, Peter and Kris Jacobs, 2004, The Importance of the Loss Function in Option Valuation, *Journal of Financial Economics* 72, 291-318.
- Dai, Qiang, Kenneth Singleton, and Wei Yang, 2003, Regime Shifts in a Dynamic Term Structure Model of U.S. Treasury Bond Yields, Working paper, Stanford University.
- Das, Sanjiv and Rangarajan Sundaram, 1999, Of Smiles and Smirks: A Term Structure Perspective, *Journal of Financial and Quantitative Analysis* 34, 211-239.
- Duarte, Jefferson, 2004, Evaluating an Alternative Risk Preference in Affine Term Structure Models, *Review of Financial Studies* 17, 379-404.
- 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.
- Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform Analysis and Asset Pricing for Affine Jump–Diffusions, *Econometrica* 68, 1343-1376.

- 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.
- Efron, Bradley and Robert Tibshirani, 1994, *An Introduction to the Bootstrap*, CRC Press.
- Eraker, Bjørn, 2001, MCMC Analysis of Diffusion Models with Applications to Finance, *Journal of Business and Economic Statistics* 19-2, 177-191.
- Eraker, Bjørn, 2003, 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.
- Heston, Steven, 1993, A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Review of Financial Studies* 6, 327-343.
- Hedges, L.V. and I Olkin, 1985, *Statistical Methods for Meta-Analysis*, Orlando, Academic Press.
- Honore, Peter, 1998, Pitfalls in Estimating Jump–Diffusion Models, Working Paper, University of Aarhus.
- Huang, Xin and George Tauchen, 2005, The Relative Contribution of Jumps to Total Price Variance, working paper, Duke University.
- 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.

- Jacod, Jean and Albert Shiryaev, 2003, *Limit Theorems for Stochastic Processes, Second edition*, Springer, New York.
- Jones, Christopher, 2003, The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets, *Journal of Econometrics* 116, 181-224.
- Jones, Christopher, 2005, A Nonlinear Factor Analysis of S&P 500 Index Option Returns, forthcoming, *Journal of Finance*.
- Johannes, Michael, 2004, The Statistical and Economic Role of Jumps in Interest Rates, *Journal of Finance* 59, 227-260
- Lawrence, C.T., J. L. Zhou and A.L. Tits, 1997, User's Guide for CFSQP Version 2.5: A C Code for Solving (Large Scale) Constrained Nonlinear (Minimax) Optimization Problems, Generating Iterates Satisfying All Inequality Constraints, Institute for Systems Research, University of Maryland, Technical Report TR-94-16r1, College Park, MD.
- Liu, Jun, Francis Longstaff and Jun Pan, 2003, Dynamic Asset Allocation with Event Risk, *Journal of Finance* 58, 231-259.
- Liu, Jun, Jun Pan and Tan Wang, 2005, An Equilibrium Model of Rare-Event Premia, *Review of Financial Studies* 18, 131-164.
- Merton, Robert, 1976, Option Pricing when the Underlying Stock Returns are Discontinuous, *Journal of Financial Economics* 3, 1235-144.
- Madan, Dilip, Peter Carr, and Erik Chang, 1998, The Variance Gamma Process and Option Pricing, *European Finance Review*, 2, 79-105.
- Madan, Dilip and Frank Milne, 1991, Option Pricing with VG Martingale Components, *Mathematical Finance* 1, 39-56.
- 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.
- Pan, Jun, Jun Liu, and Tan Wang, 2005, An Equilibrium Model of Rare-Event Premia and Its Implication for Option Smirks, *Review of Financial Studies* 18, 131–164.
- Pritsker, Matthew, 1998, Nonparametric Density Estimation and Tests of Continuous Time Interest rate Models, *Review of Financial Studies* 11, 449-487.
- Renyi, Alfred, 1961, On measures of entropy and information, *Proc. 4th Berkeley Symp. Math. Stat. and Prob.* 1, 547-561.
- 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, volume 68, Cambridge University Press, Cambridge, UK.
- Scott, Louis, 1997, Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Stochastic Interest Rates: Application of Fourier Inversion Methods. *Mathematical Finance* 7, 413-426.
- 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.