Specification Analysis of Option Pricing Models Based on Time-Changed Lévy Processes*

JING-ZHI HUANG†

Penn State University and New York University

LIUREN WU‡

Fordham University

This version: March 30, 2003

*We are grateful to Rick Green (the editor), an anonymous referee, Menachem Brenner, Peter Carr, Robert Engle, Steve Figlewski, Martin Gruber, Jean Helwege, and Rangarajan Sundaram for helpful comments and discussions. We also thank seminar participants at Baruch College, the University of Notre Dame, and Washington University in St. Louis for helpful comments.

†Smeal College of Business, Penn State University, Univ Park, PA 16802; tel: (814) 863-3566; fax: (814) 865-3362; jxh56@psu.edu; www.personal.psu.edu/jxh56. Stern School of Business, New York University, New York, NY 10012; (212) 998-0925; jhuang0@stern.nyu.edu; www.stern.nyu.edu/~jhuang0.

‡Graduate School of Business, Fordham University, 113 West 60th Street, New York, NY 10023; tel: (212) 636-6117; fax: (212) 765-5573; wu@fordham.edu; www.bnet.fordham.edu/lwu.

ABSTRACT

This article analyzes the specifications of option pricing models based on time-changed Lévy processes. We classify option pricing models based on (i) the structure of the jump component in the underlying return process, (ii) the source of stochastic volatility, and (iii) the specification of the volatility process itself. Estimation of a variety of model specifications indicates that, to capture the behavior of the S&P 500 index options, one needs to incorporate a jump component with *infinite activity* and generate stochastic volatilities from two *separate* sources: the jump component and the diffusion component.

Specification Analysis of Option Pricing Models Based on Time-Changed Lévy Processes

The seminal work of Black and Scholes (1973) has spawned an enormous literature on option pricing and also played a key role in the tremendous growth of the derivatives industry. However, the model has been known to systematically misprice equity index options. While various extensions of the Black and Scholes model have been proposed and tested, researchers are still facing the challenge of finding a model that can capture both the time series and cross-sectional – across both the option strike and maturity – behavior of index options. Obviously such a model would be very valuable to participants of option markets. Perhaps equally important, this line of research would also help us understand the dynamics of the underlying return process given the cross-sectional feature of option price data. In this article we synthesize the ongoing efforts in searching for the "true" underlying return process by performing a specification analysis of option pricing models within a new general framework. We then apply this analysis to S&P 500 index options and empirically investigate some open issues regarding the specification of the index return process.

The empirical option pricing literature has documented three "anomalies" or inconsistencies with the Black and Scholes (1973) model in the data. First, the model assumes that the underlying asset return is normally distributed. However, the cross-sectional behavior of the equity index options along the strike price dimension indicates that the conditional index return distribution under the risk-neutral measure is not normally distributed. In particular, the risk-neutral distribution for the index return inferred from the options data is highly skewed to the left; see, for example, Aït-Sahalia and Lo (1998), Jackwerth and Rubinstein (1996), and Rubinstein (1994) for empirical evidence from the S&P 500 index options. To generate return non-normality and hence to reduce the mispricing of the Black-Scholes model along the strike dimension, one response of the literature is to incorporate a jump component into the underlying asset return process (e.g. Merton (1976)).

Second, the assumption of a constant return volatility made in the Black and Scholes model has also been shown to be violated in practice. For instance, empirical studies have documented so called "volatility clustering" and the "leverage effect." The former refers to the observation that while stock re-

turns are approximately uncorrelated, the return volatility exhibits strong serial dependence (e.g., Ding, Engle, and Granger (1993) and Ding and Granger (1996)). The latter stylized fact refers to observed negative correlation between stock returns and return volatilities (Black (1976)). To accommodate these stylzed facts, one direction taken in the literature is to allow return volatility to be stochastic (e.g., Heston (1993), Hull and White (1987)).

The third stylized empirical fact that cannot be explained by the Black-Scholes model is the maturity pattern of the model pricing bias along the strike dimension mentioned earlier. It has been recognized that this bias across strike (so called volatility smile/smirk) is most significant at short maturities and then flattens out as option maturity increases (e.g. Bates (1996)). More recently, Carr and Wu (2002a) document that volatility smirk in the S&P 500 index options persists even as option maturity increases up to the observable horizon of two years. This evidence implies that the conditional non-normality of the index returns does not die away with increasing horizon, in contrast to the implication of the classic central limit theorem. The literature tries to accommodate this maturity pattern by introducing jumps into the (underlying asset) return process as well as allowing for stochastic volatility with mean reversion. The rational behind this is while jumps can generate non-normal returns at very short horizons, a persistent stochastic volatility process can slow down the convergence of the return distribution to normality as maturity increases.

On balance, the consensus from the empirical option pricing literature is that in order to capture the behavior of equity index options as well as the index returns, we need stochastic volatility jump-diffusion models – models that include both stochastic return volatility and jumps in the return process.

Existing stochastic volatility jump-diffusion models of option pricing are often specified within the jump-diffusion affine framework of Duffie, Pan, and Singleton (2000). Recent examples include Bakshi, Cao, and Chen (1997), Bates (1996, 2000), Das and Sundaram (1999), Pan (2002), and Scott (1997). In these models, the underlying asset return innovation is generated by a jump-diffusion process. The diffusion component captures small and frequent market moves. The jump component, which is assumed to follow a compound Poisson process as in Merton (1976), captures the rare and large events. This is because the number of jumps within any finite time interval is assumed to be finite in the compound Poisson model. The empirical estimates for the Poisson arrival rate are usually small, averaging about one jump for every one or two years in equity indices (e.g. Andersen, Benzoni,

and Lund (2002)). This is not surprising since these models implicitly assume that the market movements can be characterized either as small diffusive moves or as rare large events. In practice, however, one often observes much more frequent discontinuous movements of different sizes in equity indices. These high frequency jumps are difficult to capture using a compound Poisson model.

Another notable feature of the existing option pricing models is that the stochastic volatility is often assumed to come solely from the diffusion component of the underlying return process. Even in models that incorporate jumps, the arrival rate of the jump events is assumed to be either a constant or a linear function of the diffusion variance. However, such specifications are mainly driven by analytical tractability. In practice, the variation in return volatility can be driven by stochastic diffusion variance as well as by variation in the arrival rates of jumps. How these two components of stochastic volatility vary over time and relatively to each other is purely an empirical issue. In this paper, we examine a sample of S&P 500 index options to determine what type of jump structure best captures the index movement. We also investigate whether arrival rates of jump events depend on the diffusion variance or depend on different factors.

The specification analysis and empirical study in this paper are based on Carr and Wu (2002b), who propose a theoretical framework of option pricing with time-changed Lévy processes. A Lévy process is a continuous time stochastic process with independent stationary increments, analogous to iid innovations in a discrete setting. In general, a Lévy process can be decomposed into a diffusion component and a jump component. In addition to the Brownian motion and the compound Poisson jump process used widely in the traditional option pricing literature, the class of Lévy processes also includes other jump processes that exhibit higher jump frequencies and hence may better capture the dynamics of equity indices than the compound Poisson process. Heuristically, a time change is a monotonic transformation of the time variable. Stochastic volatility can be generated by applying random or locally deterministic time changes to (the original time variable of) individual components of a Lévy process. In particular, stochastic volatility can be generated by applying different time changes to the diffusion and the jump components of a Lévy process. As a consequence, time-changed Lévy processes include a rich class of jump-diffusion stochastic volatility models. Furthermore, option pricing models in this new framework can have the same analytical tractability as those in the affine framework of Duffie, Pan, and Singleton (2000).

Within the class of time-changed Lévy processes, we classify model specifications into three separate but interrelated dimensions: (i) the choice of a jump component, (ii) the identification of the sources for stochastic volatility, and (iii) the specification of the volatility process itself. Such a classification scheme encompasses almost all existing option pricing models in the literature and provides a framework for future modeling efforts. Based on this framework, we design and estimate a series of models using S&P 500 index options data and test the relative goodness-of-fit of each specification. The specification analysis focuses on addressing two important questions on model design. (**Q1**) *What type of jump structure best describes the underlying price movement and the return innovation distribution?* (**Q2**) *Where does stochastic volatility come from?* To our knowledge, this paper represents the first extensive empirical study in option pricing based on the framework of time-changed Lévy processes.

The empirical analysis in this paper focuses on the performance of twelve option pricing models generated by a combination of three jump processes and four stochastic volatility specifications. The three jump processes include the standard compound Poisson jump process used in Merton (1976), the variance-gamma jump model (VG) of Madan, Carr, and Chang (1998), and the log stable model (LS) of Carr and Wu (2002a). Unlike the compound Poisson jump model (which generates a finite number of jumps within any finite time interval), both VG and LS allow an infinite number of jumps within any finite interval and hence are better suited to capture highly frequent discontinuous movements. These three different jump structures are used to answer question **Q1** posed above. The four stochastic volatility specifications considered in our empirical analysis include traditional ones such as those used in Bates (1996) and Bakshi, Cao, and Chen (1997), where the diffusion component of the total return variance is stochastic but the jump component is constant. However, we also introduce new specifications that allow stochastic volatility to be generated separately from the jump component and the diffusion component. This is motivated by question **Q2** posed earlier.

Our estimation results show that, in capturing the behavior of the S&P 500 index options, models based on VG and LS outperform those based on compound Poisson processes. This performance ranking is robust to variations in the stochastic volatility specification and holds for both in-sample and out-of-sample tests. These results suggest that the market prices index options as if there are many (actually infinite) discontinuous price movements (jumps) of different magnitudes in the S&P

500 index. This implication is in favor of incorporating high frequency jumps such as VG and LS in the underlying asset return process. The LS model is especially useful in capturing the maturity pattern of the volatility smirk in equity index options.[1]

The estimation results also indicate that variations in the index return volatility come from two *separate* sources: the instantaneous variance of the diffusion component and the arrival rate of the jump component. One implication of this finding is that the intensities of both small and large index movements vary over time and they vary separately. Furthermore, the model parameter estimates indicate that the diffusion volatility and the jump volatility behave differently (in the risk-neutral measure). In particular, while the former is more volatile, the latter exhibits much more persistence. As a result, the behavior of short term options is influenced more by the randomness from the diffusive movements, whereas the behavior of long term options is mostly influenced by the randomness in the arrival rate of jumps.

The above specification of stochastic volatility is also consistent with empirical evidence from time series that return volatilities are driven by multiple factors (e.g., Cont and da Fonseca (2002)). Using multi volatility factors obviously increases the flexibility of a model in capturing the time series behavior of the index option prices. Another implication of the above results is that a model specification with stochastic volatility driven by both diffusion and jumps can also improve the model performance cross-sectionally. The reason is as follows. Under such a specification, diffusion volatility and jump volatility – two components of the total return variance – are driven by independent random sources so their relative weight in the return variance varies along the option maturity dimension. In fact as mentioned above, our empirical evidence shows that the jump component dominates the behavior of long term options. This implies that non-normality of the (risk-neutral) return distribution will not simply reduce as the option maturity increases – since jumps are the main source of non-normality. Namely, there will be a persistent volatility smirk across the option maturity. And this is consistent with the maturity pattern of volatility smirk documented in equity index options.

To summarize, empirical results from our specification analysis of option pricing models based on Lévy processes provide further evidence for stochastic volatility jump-diffusion models. However,

[1]Notice that the central limit theorem does not apply in this model since the return variance is infinite in the log stable model. As a result, the return distribution remains to be non-normal even as the time horizon increases. Namely, the model allows a persistent deviation from normality and therefore can capture the maturity pattern of the volatility smirk.

one can improve the model by including high frequency jumps in the underlying return process and allowing the stochastic return volatility to be driven independently by diffusion and jumps.

Time change is a standard technique for generating new processes in the theory of stochastic processes. There is a growing literature on applying the technique to finance problems, which perhaps goes back to Clark (1973). He suggests that a random time change be interpreted as a cumulative measure of business activity. Ané and Geman (2000) provide empirical evidence of this interpretation. Examples of other applications include Barndorff-Nielsen and Shephard (2001), Carr, Geman, Madan, and Yor (2001), and Geman, Madan, and Yor (2001).

The remainder of this paper is organized as follows. The first section constructs option pricing models through time changing Lévy processes. Section II addresses the data and estimation issues. Section III compares the empirical performance of different model specifications. Section IV analyzes the remaining structures in the pricing errors for different models. Section V concludes with suggestions for future research.

I. Model Specifications

In this section, we generate candidate option pricing models by modeling the underlying asset return process as time-changed Lévy processes. Under our classification scheme, each model specification requires the specification of the following aspects: (i) the jump component in the return process; (ii) the source for stochastic volatility; and (iii) the dynamics of the volatility process itself. We consider 12 model specifications, under which the characteristic function of log returns has a closed-form solution. We then price options via an efficient fast Fourier transform (FFT) algorithm.

A. Dynamics of the Underlying Price Process

Formally, let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{Q})$ be a complete stochastic basis and $\mathbb{Q}$ be the risk-neutral probability measure. Suppose that the logarithm of the underlying stock price (index level) process, $(S_t; t \geq 0)$, follows a *time-changed Lévy process* under $\mathbb{Q}$ as the following:

$$\ln S_t = \ln S_0 + (r-q)t + \left(\sigma W_{T_t^d} - \frac{1}{2}\sigma^2 T_t^d\right) + \left(J_{T_t^j} - \xi T_t^j\right), \tag{1}$$

where r denotes the instantaneous interest rate and q the dividend yield,[2] σ is a positive constant, W is a standard Brownian motion, and J denotes a compensated pure Lévy jump martingale process, which we will elaborate later. The vector $T_t \equiv \left[T_t^d, T_t^j\right]^\top$ denotes potential *stochastic time changes* applied to the two Lévy components W_t and J_t. By definition, the time change T_t is an increasing, right-continuous process with left limits satisfying the usual regularity conditions.[3]

While stochastic time change has much wider applications, our focus here is its role in generating stochastic volatilities. For this purpose, we further restrict T_t to be continuous and differentiable with respect to t. In particular, let

$$v(t) \equiv \left[v^d(t), v^j(t)\right]^\top = \partial T_t / \partial t. \tag{2}$$

Then, $v^d(t)$ is proportional to the instantaneous variance of the diffusion component, while $v^j(t)$ is proportional to the arrival rate of the jump component. Following Carr and Wu (2002b), we label $v(t)$ as the *instantaneous activity rate*. Intuitively speaking, one can regard t as the calendar time and T_t as the business time at calendar time t. A more active business day, captured by a higher activity rate, generates higher volatility for asset returns. The randomness in business activity generates randomness in volatility.

Note that in equation (1), we apply stochastic time changes only to the diffusion and jump martingale components, but not to the instantaneous drift. The reason is that the equilibrium interest rate and dividend yield are defined on the calendar time, not on business event time. Furthermore, we apply

[2] Bakshi, Cao, and Chen (1997) also consider the role played by stochastic interest rates but find that the impact on option pricing is minimal. Here we treat both r and q as deterministic.

[3] T_t is finite $\mathbb{Q}$-a.s. for all $t \geq 0$ and that $T_t \to \infty$ as $t \to \infty$.

separate time changes on the diffusion martingale component and on the jump martingale component, allowing potentially different time-variation in the intensities (activity rates) of small and large events.

Also note that in this article, "volatility" is used as a generic term capturing the financial activities of an asset. It is not used as a statistical term for standard deviation. Just like in Heston (1993) and in many other papers, we model the stochastic "volatility" from the diffusion component by specifying a stochastic process for $v^d(t)$, which is proportional to the instantaneous *variance* of the diffusion component. In addition, we model stochastic "volatility" from the jump component by specifying a stochastic process for $v^j(t)$, which is proportional to the arrival rate of the jump component.

B. Option Pricing via Generalized Fourier Transforms

To derive the time-0 price of an option expiring at time t, we first derive the conditional generalized Fourier transform of the log return $s_t \equiv \ln(S_t/S_0)$ and then obtain the option price via an efficient fast Fourier inversion. Since the underlying asset return is modelled as a time-changed Lévy process, we derive the generalized Fourier transform of the return process in two steps. First, we derive the generalized Fourier transform of the Lévy process prior to the time change. Then, the generalized Fourier transform of the time changed Lévy process is obtained by solving the Laplace transform of the stochastic time under an appropriate measure change.

Consider first the return process before a time change. Equation (1) implies that prior to any time changes, the log return $s_t = \ln(S_t/S_0)$ follows the following Lévy process:

$$s_t = (r-q)t + \left(\sigma W_t - \frac{1}{2}\sigma^2 t\right) + (J_t - \xi t)\,. \tag{3}$$

Notice that the log return s_t is decomposed into three components in (3). On the right-hand side of (3), the first term, $(r-q)t$, is from the instantaneous drift, which is determined by no-arbitrage. The second term, $\left(\sigma W_t - \frac{1}{2}\sigma^2 t\right)$, comes from the diffusion component where $\frac{1}{2}\sigma^2 t$ is the concavity adjustment. The last term, $(J_t - \xi t)$, represents the contribution from the jump component with ξ being the analogous concavity adjustment for J_t. The generalized Fourier transform for s_t under (3) is given by

$$\phi_s(u) \equiv \mathbb{E}^{\mathbb{Q}}\left[e^{ius_t}\right] = \exp\left(iu(r-q)t - t\psi_d - t\psi_j\right), \quad u \in \mathcal{D} \in \mathbb{C}, \tag{4}$$

where $\mathbb{E}^{\mathbb{Q}}[\cdot]$ denotes the expectation operator under the risk-neutral measure $\mathbb{Q}$, $\mathcal{D}$ denotes a subset of the complex domain ($\mathbb{C}$) where the expectation is well-defined, and

$$\psi_d = \frac{1}{2}\sigma^2\left[iu + u^2\right],$$

is the *characteristic exponent* of the diffusion component. The characteristic exponent of the jump component, ψ_j, depends upon the exact specification of the jump structure.[4] As a key feature of Lévy processes (See Bertoin (1996) and Sato (1999)), neither ψ_d nor ψ_j depends upon the time horizon t. Note that $\phi_s(u)$ is essentially the characteristic function of the log return when u is real. The extension of u to the admissible complex domain is necessary for the application of the fast Fourier transform algorithm; see Titchmarsh (1975) for a comprehensive reference on generalized Fourier transforms.

Next, we apply the time change through the mapping $t \to T_t$ as defined in (1). The generalized Fourier transform of the time changed return process is given by

$$\begin{aligned}\phi_s(u) &= e^{iu(r-q)t}\mathbb{E}^{\mathbb{Q}}\left[e^{iu\left(\sigma W_{T_t^d} - \frac{1}{2}\sigma^2 T_t^d\right) + iu\left(J_{T_t^j} - \xi T_t^j\right)}\right] \\ &= e^{iu(r-q)t}\mathbb{E}^{\mathbb{M}}\left[e^{-\psi^\top T_t}\right] \equiv e^{iu(r-q)t}\mathcal{L}_T^{\mathbb{M}}(\psi), \end{aligned} \tag{5}$$

where $\psi \equiv [\psi_d, \psi_j]^\top$ denotes the vector of the characteristic exponents and $\mathcal{L}_T^{\mathbb{M}}(\psi)$ represents the Laplace transform of the stochastic time T_t under a new measure $\mathbb{M}$. The measure $\mathbb{M}$ is absolutely continuous with respect to the risk-neutral measure $\mathbb{Q}$ and is defined by a complex-valued exponential martingale,

$$\left.\frac{d\mathbb{M}}{d\mathbb{Q}}\right|_t \equiv \exp\left[iu\left(\sigma W_{T_t^d} - \frac{1}{2}\sigma^2 T_t^d\right) + iu\left(J_{T_t^j} - \xi T_t^j\right) + \psi_d T_t^d + \psi_j T_t^j\right]. \tag{6}$$

Note that in (5), the issue of obtaining a generalized Fourier transform is converted into a simpler problem, namely, one of deriving the Laplace transform of the stochastic time (see Carr and Wu (2002b)). This Laplace transform depends both on the specification of the instantaneous activity rate $v(t)$ and the characteristic exponents, the functional form of which is determined by the specification

[4] Throughout the paper, we use a subscript (or superscript) 'd' to denote the diffusion component and 'j' the jump component.

of the jump structure J_t. In what follows, we address the specification issues of the jump structure and stochastic volatility, as well as the corresponding solutions to the Laplace transform.

C. The Jump Structure

Depending upon the frequency of jump arrivals, Lévy jump processes can be classified into three categories: (1) finite activity, (2) infinite activity with finite variation, and (3) infinite variation (Sato (1999), page 65). Each jump category exhibits distinct behavior and hence results in different option pricing performance.

Formally, the structure of a Lévy jump process is captured by its *Lévy measure*, $\pi(x)$, which controls the arrival rate of jumps of size $x \in \mathbb{R}^0$ (the real line excluding zero). A *finite activity* jump process generates a finite number of jumps within any finite interval. As such, the integral of the Lévy measure is finite:

$$\int_{\mathbb{R}^0} \pi(dx) < \infty, \tag{7}$$

so that the Lévy measure has the interpretation and property of a probability density function after being normalized by this integral. A prototype example of a finite activity jump process is the *compound Poisson jump* process of Merton (1976) (MJ), which has been widely adopted by the finance literature. For such a process, the integral in (7) defines the *Poisson intensity*, λ. The MJ model assumes that conditional on one jump occurring, the jump magnitude is normally distributed with mean α and variance σ_j^2. The Lévy measure of the MJ process is given by

$$\pi_{MJ}(dx) = \lambda \frac{1}{\sqrt{2\pi\sigma_j^2}} \exp\left(-\frac{(x-\alpha)^2}{2\sigma_j^2}\right) dx. \tag{8}$$

In essence, for all finite activity jump models, one can decompose the Lévy measure into two components: a normalizing coefficient often labeled as the Poisson intensity, and a probability density function controlling the conditional distribution of the jump size.

Unlike a finite activity jump process, an *infinite activity* jump process generates an infinite number of jumps within any finite interval. The integral of the Lévy measure for such processes is no longer finite. Examples of this class include the normal inverse Gaussian model of Barndorff-Nielsen (1998),

the generalized hyperbolic class of Eberlein, Keller, and Prause (1998), and the variance-gamma (VG) model of Madan and Milne (1991) and Madan, Carr, and Chang (1998). In our empirical studies, we choose the relatively parsimonious VG model as a representative of the infinite activity jump type. The VG process is obtained by subordinating an arithmetic Brownian motion with drift α/λ and variance σ_j^2/λ by an independent gamma process with unit mean rate and variance rate $1/\lambda$. The Lévy measure for the VG process is given by

$$\pi_{VG}(dx) = \frac{\mu_\pm^2}{\nu_\pm} \frac{\exp\left(-\frac{\mu_\pm}{\nu_\pm}|x|\right)}{|x|} dx,$$

where

$$\mu_\pm = \sqrt{\frac{\alpha^2}{4\lambda^2} + \frac{\sigma_j^2}{2}} \pm \frac{\alpha}{2\lambda}, \quad \nu_\pm = \mu_\pm^2/\lambda.$$

The parameters with plus subscripts apply to positive jumps and those with minus subscripts apply to negative jumps. The jump structure is symmetric around zero when we drop the subscripts. Note that as the jump size approaches zero, the arrival rate approaches infinity. Thus, an infinite activity model incorporates infinitely many small jumps. The Lévy measure of an infinite activity jump process is singular at zero jump size.

Nevertheless, for all the above mentioned infinite activity jump models, we have

$$\int_{\mathbb{R}^0} (1 \wedge |x|)\, \pi(dx) < \infty, \tag{9}$$

so that the sample paths of the jump processes exhibit *finite variation*. The function $(1 \wedge |x|)$ here represents the minimum of 1 and $|x|$. Since, under certain regularity conditions, the Lévy measure of large jumps always performs like a density function, whether an infinite activity jump process exhibits finite or infinite variation is purely determined by its property around the singular point at zero jump size ($x = 0$). The function $(1 \wedge |x|)$ is a *truncation function* used to analyze the jump properties around the singular point of zero jump size (Bertoin (1996), page 15).[5]

[5]Other commonly used truncation functions for the same purpose include $x1_{|x|<1}$, where $1_{|x|<1}$ is an indicator function, and $x/(1+x^2)$. In essence, one can use any truncation functions, $h : \mathbb{R}^d \to \mathbb{R}^d$, which are bounded, with compact support, and satisfy $h(x) = x$ in a neighborhood of zero (Jacod and Shiryaev (1987), page 75).

When the integral in (9) is no longer finite, the sample path of the process exhibits *infinite variation*. A typical example is an α-stable motion with $\alpha \in (1,2]$; see two monographs, Samorodnitsky and Taqqu (1994) and Janicki and Weron (1994), on such a process. The Lévy measure under the α-stable process is given by

$$\pi(dx) = c_{\pm}|x|^{-\alpha-1}dx. \tag{10}$$

The process exhibits finite variation when $\alpha < 1$; but when $\alpha > 1$, the integral in (9) is no longer finite and the process is of infinite variation.[6] The parameter α is often referred to as the *tail index* while the parameters $c_{\pm}$ control both the scale and the asymmetry of the process. Within this category, we choose the *finite moment log stable* (LS) process of Carr and Wu (2002a) in our empirical investigation. In this LS model, c_{+} is set to zero in (10) so that only negative jumps are allowed. This restriction guarantees the existence of a finite martingale measure (and hence finite option prices) and ensures that the conditional moments of the asset price of all positive orders are finite. This latter feature allows the model to explain the slow decay of the implied volatility smirk across different maturities observed for S&P 500 index options. The reason is that the central limit theorem does not apply in this model (since the return has an α-stable distribution, and the variance and higher moments of the asset return are infinite) and, as a result, conditional distribution of the asset return *remains* non-normal as the conditioning horizon increases. Note that this LS model also addresses the criticism of Merton (1976) on using α-stable distributions to model asset returns.

As mentioned earlier, in order to calculate option prices via equation (5), we need to know the characteristic exponents of the specified jump process. The three jump processes considered here, MJ, VG, and LS, all have analytical characteristic exponents, which are tabulated in Table I. We also include the characteristic exponent for the diffusion component for ease of comparison. Given the Lévy measure π for a particular jump process, the corresponding characteristic exponents can be derived using the well-known *Lévy-Khintchine formula* (Bertoin (1996), page 12),

$$\psi_j(u) \equiv -iub + \int_{\mathbb{R}^0} \left(1 - e^{iux} + iux1_{|x|<1}\right) \pi(dx), \tag{11}$$

[6] Nevertheless, for the Lévy measure to be well-defined, the quadratic variation has to be finite:

$$\int_{\mathbb{R}^0} (1 \wedge x^2)\pi(dx) < \infty,$$

which requires that $\alpha \leq 2$.

Table I
Characteristic Exponent of the Lévy Components in the Asset Return Process

Component	$\psi_d(u)$ or $\psi_j(u)$
Diffusion	$\frac{1}{2}\sigma^2\left[iu-(iu)^2\right)$
Poisson Jump (MJ)	$\lambda\left[iu\left(e^{\alpha+\frac{1}{2}\sigma_j^2}-1\right)-\left(e^{iu\alpha-\frac{1}{2}u^2\sigma_j^2}-1\right)\right]$
Variance Gamma (VG)	$\lambda\left[-iu\ln\left(1-\alpha-\frac{1}{2}\sigma_j^2\right)+\ln\left(1-iu\alpha+\frac{1}{2}\sigma_j^2u^2\right)\right]$
Log Stable (LS)	$\lambda\left(iu-(iu)^\alpha\right)$

where b denotes a drift adjustment term.

D. The Sources of Stochastic Volatility

The specification of a time-changed Lévy process given in (1) makes it transparent that stochastic volatility can come either from the instantaneous variance of the diffusion component or from the arrival rate of the jump component, or both. We consider four cases that exhaust the potential sources of stochastic volatility.

SV1: Stochastic volatility from diffusion: If we apply a stochastic time change to the Brownian motion only, i.e., $W_t \rightarrow W_{T_t^d}$, and leave the jump component J_t unchanged, stochastic volatility will come solely from the diffusion component. The arrival rate of jumps remains constant. Examples using this specification include Bakshi, Cao, and Chen (1997), and Bates (1996). Under this specification, whenever the asset price movement becomes more volatile, it is due to an increase in the diffusive movements in the asset price. The frequency of large events remains constant. Thus, the *relative* weight of the diffusion and jump components in the return process varies over time. In particular, the relative weight of the jump component declines as the total volatility of the return process increases.

SV2: Stochastic volatility from jump: If, instead, we apply a stochastic time change to the jump component only, i.e., $J_t \rightarrow J_{T_t^j}$, but leave the Brownian motion unchanged, stochastic volatility will come solely from the time variation in the arrival rate of jumps. Under this specification, an increase in the return volatility is attributed solely to an increase in the discontinuous movements (jumps) in the

asset price. Hence, the relative weight of the jump component increases with the return volatility. The models proposed in Carr, Geman, Madan, and Yor (2001) can be regarded as degenerate examples of this SV2 category as they apply stochastic time changes to pure jump Lévy processes.

SV3: Joint contribution from jump and diffusion: To model the situation where stochastic volatility comes simultaneously from both the diffusion and jump components, we can apply the same stochastic time change T_t (a scalar process) to both W_t and J_t. In this case, the instantaneous variance of the diffusion and the arrival rate of jumps vary synchronously over time. Under SV3, the relative proportions of the diffusion component and the jump component are constant even though the return volatility varies over time. The recent affine models in Bates (2000) and Pan (2002) can be regarded as variations of this category. In these models, both the arrival rate of the Poisson jump and the instantaneous variance of the diffusion component are driven by one stochastic process.

SV4: Separate contribution from jump and diffusion: The most general specification is to apply *separate* time changes to the diffusion component and the jump component so that the time change T_t is a bivariate process. Under this specification, the instantaneous variance of the diffusion component and the arrival rate of the jump component follow separate stochastic processes. Hence, variation in the return volatility can come from either or both of the two components. Since the two components vary over time separately, the relative proportion of each component also varies over time. The relative dominance of one component over the other depends upon the exact dynamics of the two activity rates. Specification SV4 encompasses all the previous three specifications (SV1-SV3) as special cases.

Under the affine framework of Duffie, Pan, and Singleton (2000), Bates (2000) also specifies a two-factor stochastic volatility process. Since each of the two volatility factors in Bates (2000) drives both a compound Poisson jump component and a diffusion component, his model can be regarded as a two-factor extension of our SV3 model. Alternatively, his model can also be regarded as a mixture of SV1 and SV3 specifications since in the model the intensity of the Poisson jump includes both a constant term and a term proportional to the stochastic volatility factor. One can also regard our SV4 specification as a special case of Bates (2000) by setting the diffusion component to zero in one factor and the jump component to zero in the other factor. Nevertheless, the approach in this paper that treats

Table II
Generalized Fourier Transforms of Log Returns under Different SV Specifications

x_t denotes the time changed component and y_t denotes the unchanged component in the log return $s_t = \ln(S_t/S_0)$. J_t denotes a compensated pure jump martingale component, and ξ its concavity adjustment.

Model	x_t	y_t	$\phi_s(u)$
SV1	$\sigma W_t - \frac{1}{2}\sigma^2 t$	$J_t - \xi t$	$e^{iu(r-q)t - t\psi_j}\mathcal{L}_T^{\mathbb{M}}(\psi_d)$
SV2	$J_t - \xi t$	$\sigma W_t - \frac{1}{2}\sigma^2 t$	$e^{iu(r-q)t - t\psi_d}\mathcal{L}_T^{\mathbb{M}}(\psi_j)$
SV3	$\sigma W_t - \frac{1}{2}\sigma^2 t + J_t - \xi t$	0	$e^{iu(r-q)t}\mathcal{L}_T^{\mathbb{M}}(\psi_d + \psi_j)$
SV4	$\left[\sigma W_t - \frac{1}{2}\sigma^2 t, J_t - \xi t\right]^\top$	0	$e^{iu(r-q)t}\mathcal{L}_T^{\mathbb{M}}(\left[\psi_d, \psi_j\right]^\top)$

the jump and the diffusion components separately makes it easier to identify the different roles played by the two components.

We now derive the generalized Fourier transform of the log return for each of the four SV specifications. Let x denote the time-changed component and y the unchanged component in the log return, and ψ_x and ψ_y denote their respective characteristic exponents. The generalized Fourier transform of the log return $s_t = \ln(S_t/S_0)$ in (5) can be rewritten as

$$\phi_s(u) = \mathbb{E}^{\mathbb{Q}}\left[e^{iu(r-q)t + y_t + x_{T_t}}\right] = e^{iu(r-q)t - t\psi_y}\,\mathbb{E}^{\mathbb{M}}\left[e^{-\psi_x T(t)}\right] = e^{iu(r-q)t - t\psi_y}\mathcal{L}_T^{\mathbb{M}}(\psi_x). \tag{12}$$

The complex-valued exponential martingale in (6) that defines the measure change can be rewritten as

$$\left.\frac{d\mathbb{M}}{d\mathbb{Q}}\right|_t = \exp\left(iuy_t + iux_{T_t} + \psi_y t + \psi_x T_t\right). \tag{13}$$

Table II summarizes the x and y components, as well as the generalized Fourier transform of the log return, for each of the four SV specifications.

E. Specification of the Activity Rate Process

We close the modeling effort by specifying an activity rate process $v(t)$ and deriving the Laplace transform of the stochastic time $T_t = \int_0^t v(s)ds$ under the new measure $\mathbb{M}$. For this purpose, we rewrite the Laplace transform as

$$\mathcal{L}_T^{\mathbb{M}}(\psi) = \mathbb{E}^{\mathbb{M}}\left[e^{-\psi^\top T_t}\right] = \mathbb{E}^{\mathbb{M}}\left[e^{-\int_0^t \psi^\top v(s)ds}\right], \tag{14}$$

which is analogous to the pricing formula for a zero coupon bond if we treat $\psi^\top v(t)$ as an instantaneous interest rate. We can thus borrow the abundant literature in term structure models for the modeling of the activity rate. For example, we can model the activity rate of a Brownian motion after the term structure model of Cox, Ingersoll, and Ross (1985) and, in fact, recover the Heston (1993) stochastic volatility model.[7] Multivariate activity rate processes can be modeled after, among others, affine models of Duffie and Kan (1996) and Duffie, Pan, and Singleton (2000) and quadratic ones of Leippold and Wu (2002).

Despite the large pool of candidate processes for the activity rate modeling, we leave the specification analysis of different activity rate models for future research. For the empirical work in this paper, we focus on one activity rate process, i.e., the Heston (1993) model. Under the risk-neutral measure $\mathbb{Q}$, the activity rate process satisfies the following stochastic differential equation,

$$dv(t) = \kappa(1 - v(t))\,dt + \sigma_v\sqrt{v(t)}dZ_t, \tag{15}$$

where Z_t denotes a standard Brownian motion under $\mathbb{Q}$, which can be correlated with the standard Brownian motion W_t in the return process by: $\rho dt = \mathbb{E}^{\mathbb{Q}}[dW_t dZ_t]$. Note that the long run mean of the activity rate is normalized to unity in (15) for identification purpose. For the SV4 specification, we assume that the two activity rates, $v(t) = \left[v^d(t), v^j(t)\right]^\top$, follow a vector square-root process.

[7] To obtain the Heston (1993) model, we can apply a stochastic time change to the Brownian motion in the stock return process in the Black-Scholes model, use the square-root process of Cox, Ingersoll, and Ross (1985) to model the activity rate, and allow the activity rate and the stock return to be correlated.

Since the Laplace transform of the time change in (14) is defined under measure $\mathbb{M}$, we need to obtain the activity rate process under $\mathbb{M}$. By Girsanov's Theorem, under measure $\mathbb{M}$, the diffusion function of $v(t)$ remains unchanged while the drift function is adjusted to

$$\mu^{\mathbb{M}} = \begin{cases} \kappa(1-v(t)) + iu\sigma\sigma_v\rho\, v(t), & \text{for } SV1, SV3, \text{ and } SV4; \\ \kappa(1-v(t)) + iu\sigma\sigma_v\rho\sqrt{v(t)}, & \text{for } SV2. \end{cases}$$

Note the difference between the drift adjustment for SV2 models and that for all other models. This difference occurs because the diffusion component in the return process is time changed under all SV specifications except for the SV2 specification. Therefore, given that $dW_{T_t} = \sqrt{v(t)}dW_t$ holds in probability, the drift adjustment term for SV2 models is different from the drift adjustment term for all other SV specifications by a scaling of $\sqrt{v(t)}$.

As the drift $\mu^{\mathbb{M}}$ remains affine for models SV1 and SV3 for any $\rho \in [-1,1]$, the arrival rate process belongs to the affine class. The Laplace transform of T_t is then exponential-affine in v_0 (the current level of the arrival rate), and is given by

$$\mathcal{L}_T^{\mathbb{M}}(\psi) = \exp\left(-b(t)v_0 - c(t)\right), \tag{16}$$

where

$$\begin{aligned} b(t) &= \frac{2\psi(1-e^{-\eta t})}{2\eta - (\eta-\kappa^*)(1-e^{-\eta t})}; \\ c(t) &= \frac{\kappa}{\sigma_v^2}\left[2\ln\left(1 - \frac{\eta-\kappa^*}{2\eta}\left(1-e^{-\eta t}\right)\right) + (\eta-\kappa^*)t\right], \end{aligned}$$

with

$$\eta = \sqrt{(\kappa^*)^2 + 2\sigma_v^2\psi}, \quad \kappa^* = \kappa - iu\rho\sigma\sigma_v.$$

The SV4 model also satisfies the affine structure in a vector form. For tractability, we assume that the two arrival rates are independent and separately correlated with the return process. Then, the above solutions for $b(t)$ and $c(t)$ can be regarded as solutions to the coefficients for each of the two activity rates. For the SV2 specification, the affine structure is retained only when $\rho = 0$. For tractability, we restrict $\rho = 0$ in our estimation of SV2 models.

Substituting the Laplace transform in (16) into the generalized Fourier transforms in Table II, we can derive in analytical forms the generalized Fourier transforms for all 12 models: three jump specifications (MJ, VG, and LS) multiplied by four stochastic volatility specifications (SV1-SV4). These 12 models are labeled as "JJDSVn," where JJ $\in \{\text{MJ}, \text{VG}, \text{LS}\}$ denotes the jump component, D refers to the diffusion component, and SVn, with $n = 1, 2, 3, 4$, denotes a particular stochastic volatility specification. For example, when the Merton jump diffusion model (MJD) is coupled with the SV1 specification, we have the model labeled as "MJDSV1." This is the same specification as the one considered in Bakshi, Cao, and Chen (1997) and Bates (1996). Taken together, the 12 models are designed to answer two important questions: (1) what type of jump process performs best in capturing the behavior of S&P 500 index options? (2) where does stochastic volatility come from?

II. Data and Estimation

We have daily closing bid and ask implied volatility quotes on the S&P 500 index options across different strikes and maturities from April 6th, 1999 to May 31st, 2000, obtained from a major investment bank in New York. The quotes are on standard European options on the S&P 500 spot index, listed at the Chicago Board Options Exchange (CBOE). The implied volatility quotes are derived from out-of-the-money (OTM) option prices. The same data set also contains matching forward prices F, spot prices (index levels) S, and interest rates r corresponding to each option quote, compiled by the same bank. We apply the following filters to the data: (1) the time to maturity is greater than five business days; (2) the bid option price is strictly positive; (3) the ask price is no less than the bid price. After applying these filters, we also plot the mid implied volatility quote for each day and maturity against strike prices to visually check for obvious outliers. After removing these outliers, we have 62,950 option quotes over a period of 290 business days.

The left panel of Figure 1 depicts the histogram of moneyness of the cleaned up option contracts, where the moneyness is defined as $k \equiv \ln(K/S)$. The observations are centered around at the money options ($k = 0$). On average, we have more OTM put option quotes ($k < 0$) than OTM call option quotes ($k > 0$), reflecting the difference in their respective trading activities. The right panel of Figure 1 plots the histogram of the time-to-maturity for the option contracts. The maturities of the option

contracts range between five business days and over one and a half years, with the number of option quotes declining almost monotonically as the time-to-maturity increases. These exchange-traded index options have fixed expiry dates, all on the Saturday following the third Friday of a month. The terminal payoff at expiry is computed based on the opening index level on that Friday. The contract hence stops trading on that expiring Thursday. We delete from our sample contracts that are within one week of expiry to avoid potential microstructure effects.

Since the FFT algorithm that we use returns option prices at fixed moneyness with equal intervals, we linearly interpolate across moneyness to obtain option prices at fixed moneyness. We also restrict our attentions to the more liquid options with moneyness $k = \ln(K/S)$ between -0.3988 and 0.1841, where K denotes the strike price and S the spot index level. This restriction excludes approximately 16% very deep out-of-money options (approximately 8% calls and 8% puts) which we deem as too illiquid to contain useful information. Note that the moneyness range is asymmetric to reflect the fact that there are deeper out-of-the-money put options quotes than out-of-the-money call options. Within this range, we sample options with a fixed moneyness interval of $\Delta k = 0.03068$ (a maximum of 20 strike points at each maturity). For the interpolation to work with sufficient precision, we require that at each day and maturity, we have at least five option quotes. We also refrain from extrapolating: we only retain option prices at fixed moneyness intervals that are within the data range. Visual inspection indicates that at each date and maturity, the quotes are so close to each other along the moneyness line that interpolation can be done with little error, irrespective of the interpolation method. We delete one inactive day from the sample when the number of sample points is less than 20. The number of sample points in the other active 289 days ranges from 92 to 144, with an average of 118 sample points per day. In total, we have 34,361 sample data points for estimation.

We estimate the vector of model parameters, Θ, by minimizing the weighted sum of squared pricing errors as follows,

$$\Theta \equiv \arg\min_{\Theta} \sum_{t=1}^{\mathcal{T}} \mathtt{mse}_t, \tag{17}$$

where $\mathcal{T}$ denotes the total number of days and $\texttt{mse}_t$ denotes the mean squared pricing error at date t, defined as

$$\texttt{mse}_t \equiv \min_{v(t)} \frac{1}{N_t} \sum_{i=1}^{n_{t,\tau}} \sum_{j=1}^{n_{t,k}} w_{ij} e_{ij}^2, \tag{18}$$

where $n_{t,\tau}$ and $n_{t,k}$ denote respectively the number of maturities and the number of moneyness levels per each maturity at date t, N_t denotes the total number of observations at date t, w_{ij} denotes an optimal weight, and e_{ij} represents the pricing error at maturity i and moneyness j. Note that there are two layers of estimation involved. First, given the set of model parameters, Θ, we identify the instantaneous activity rates level $v(t)$ at each date t by minimizing the weighted mean squared pricing errors on that day. Next, we choose model parameters Θ to minimize the sum of the daily mean squared pricing errors.[8] To construct out-of-sample tests, we divide the data into two sub-samples: we use the first 139 days of data to estimate the model parameters and then the remaining 150 days of data to test the models' out-of-sample performance. To evaluate out-of-sample performance on the second sub-sample, we fix the parameter vector Θ estimated from the first sub-sample and compute the daily mean squared pricing errors according to (18) by minimizing the squared pricing errors each day with respect to the activity rate levels $v(t)$.

The pricing error matrix $e = (e_{ij})$ is defined as follows,

$$e = \begin{cases} \widehat{O}(\Theta) - O_a, & \text{if} \quad \widehat{O}(\Theta) > O_a \\ 0, & \text{if} \quad O_a \leq \widehat{O}(\Theta) \leq O_b \\ \widehat{O}(\Theta) - O_b, & \text{if} \quad \widehat{O}(\Theta) < O_b \end{cases} \tag{19}$$

where $\widehat{O}(\Theta)$ denotes model implied out-of-the-money (OTM) option prices (put prices when $K \leq F$ and call prices when $K > F$) as a function of the parameter vector Θ, and O_a and O_b denote, respectively, the ask and bid prices observed from the market. The pricing error is assumed to be zero as long as the model implied price falls within the bid-ask spread of the market quote. All prices are normalized as percentages of the underlying spot price.

The construction of the pricing error is a delicate but important issue. For example, the pricing error can be defined on implied volatility, call option price, or put option price. It can be defined as

[8] We thank an anonymous referee for suggesting this estimation procedure.

the difference in levels, in log levels, or in percentages. Here, we define the pricing error using call option prices when $K > F$ and using put option prices when $K \leq F$. Such a definition has become the industry standard for several reasons. One reason is that in-the-money options have positive intrinsic value which is insensitive to model specification and yet can be the dominant component of the total option value. Another reason is that when there is a discrepancy between the market quotes on out-of-the-money options and their in-the-money counterparts, the former quotes are in general more reliable as they are more liquid, probably because in the presence of transactions costs, out-of-money options represent a cheaper way to speculate on or hedge against changes in future volatility. We refine the standard definition of the pricing error by incorporating the effects of the bid-ask spreads. This reduces the potential problem of over-fitting and further accounts for the liquidity differences at different moneyness levels and maturities. Dumas, Fleming, and Whaley (1998) also incorporate this bid-ask spread effect in their definition of "mean outside error."

A. The Optimal Weighting Matrix

Similar to the definition of the pricing error, the construction of a "good" weighting matrix is also important in obtaining robust estimates. Existing empirical studies often use identity weighting matrix. Under our definition of the pricing error, an identity weighting matrix puts more weight on near-the-money options than on deep out-of-the-money options. More importantly, it puts significantly more weight on long term options than on short term options. Thus, performance comparisons may be biased toward models that better capture the behavior of long term options. In this section, we seek to estimate a weighting matrix which (a) attaches a more balanced weighting to options at all moneyness and maturity levels and (b) can be applied to the estimation and comparison of all relevant models.

One way to achieve this is to estimate an optimal weighting matrix based on the variance of the option prices, normalized as percentages of the underlying spot price. Specifically, we estimate the variance of the percentage option prices at each moneyness and maturity level via nonparametric regression and use its reciprocal as the weighting for the pricing error at that moneyness and maturity. This weighting matrix is optimal in the sense of maximum likelihood under the following assumptions: (a) the pricing errors are independently normally distributed and (b) the variance of the pricing error is well approximated by the variance of the corresponding option prices as percentages of the index level.

When the pricing errors are independently normally distributed, the minimization problem in (17) also generates the maximum likelihood estimates if we set the weighting at each moneyness and maturity level to the reciprocal of the variance estimate of the pricing error at that moneyness and maturity. In principle, the variance of the pricing errors can be estimated via a two-stage procedure analogous to a two-stage least square procedure. However, the weighting obtained from such a procedure depends upon the exact model being estimated. We use the variance of the option price (as a percentage of the index level) as an approximate measure for the variance of the pricing error. This approximation is *exact* when the return to the underlying stock index follows a Lévy process without stochastic volatility. This is because the conditional return distribution over a fixed horizon does not vary over time in such processes. As a result, for a given option maturity and moneyness, the option price normalized by the underlying index level does not vary with time either. The "true" option price as a percentage of the index level can then be estimated through a sample average and the daily deviations from such a sample average can be regarded as the pricing error. Therefore, the variance of the pricing error is equivalent to the variance of the option prices normalized by the index level.

However, all our model specifications incorporate some type of stochastic volatility. Thus, the variance of the option prices includes both the variance of the pricing error and the variation induced by stochastic volatility. The variance estimate of the option price is therefore only an approximate measure of the variance of the pricing error in our case. Nevertheless, our posterior analysis of the pricing errors confirms that such a choice of weighting matrix is reasonable. The idea of choosing a common metric, upon which different and potentially non-nested models can be compared, is also used in the distance metric proposed by Hansen and Jagannathan (1997) for evaluating different stochastic discount factor models.

Since the moneyness and maturity of the options vary every day, we estimate the mean option value and the option price variance as percentages of the index level at fixed moneyness and maturities through a nonparametric smoothing method. Refer to Appendix A for details. The left panel of Figure 2 portrays the smoothed mean surface of out-of-the-money option prices. As expected, option prices are the highest for at the money options and they also increase with maturities. The right panel portrays the variance estimates of the option prices. Overall, the variance increases with the maturity of the option. For the same maturity, out-of-the-money puts ($k < 0$) have a smaller variation than out-the-money calls

($k > 0$). This might be a reflection of different liquidities: OTM puts are more liquid and more heavily traded than OTM calls for S&P 500 index options. Given the estimated variance of the option prices, the optimal weight at each moneyness and maturity level is defined as its reciprocal.

B. Performance Measures

Different models are compared based on the sample properties of the daily mean squared pricing errors ($\texttt{mse}_t$) defined in equation (18), under the estimated model parameters. A small sample average of the daily mean squared errors for a model would indicate that the model fits the option prices well on average. A small standard deviation for a model would further indicate that the model is capable of capturing different cross-sectional properties of the option prices at different dates. Our analysis is based on both the in-sample mean squared errors of the first 139 days and the out-of-sample mean squared errors of the last 150 days. In addition, we gauge the statistical significance of the performance difference between any two models *i* and *j* based on the following *t*-statistic of the sample differences in daily mean squared errors,

$$\texttt{t-statistic} = \frac{\overline{\texttt{mse}}^i - \overline{\texttt{mse}}^j}{\texttt{stdev}(\texttt{mse}_t^i - \texttt{mse}_t^j)/\sqrt{\mathcal{T}}}, \tag{20}$$

where the overline on $\texttt{mse}$ denotes the sample average and $\texttt{stdev}(\cdot)$ denotes the standard deviation.

III. Model Performance Analysis

We now analyze the parameter estimates and the sample properties of the mean squared pricing errors for each of the 12 models introduced in Section I. As mentioned earlier, our objective is to investigate which jump type and which stochastic volatility specification deliver the best performance in pricing S&P 500 index options. Our analysis below is focused on answering these two questions.

Tables III and IV report the parameter estimates and their standard errors for one-factor (SV1-SV3) and two-factor stochastic volatility (SV4) models, respectively. We also report in the tables the sample average and standard deviation of the daily mean squared pricing errors, both in sample ($\texttt{mse}_I$) and out of sample ($\texttt{mse}_O$). These parameter estimates for each model are then used to calculate the

t-statistics for pair-wise model comparisons using (20). The results for the comparison are reported in Table V. With 12 models, we could have reported a 12×12 matrix of pair-wise t-tests; but to focus on addressing the two questions raised above, we report the t-tests in two panels. Panel A compares the performance of different jump structures under each stochastic volatility specification (SV1 to SV4); Panel B compares the performance of different SV specifications for a given jump structure (MJ, VG, or LS). Both in-sample and out-of-sample comparisons are reported.

A. What Jump Structure Best Captures the Behavior of S&P 500 Index Options?

Since our 12 models are combinations of three jump structures and four SV specifications, we compare the performance of the three jump structures, MJ, VG, and LS, under each SV specification to answer the question on jump types. If the performance ranking of the three jump structures depends crucially on the specific SV specification, the choice of a jump structure in model design should be contingent on the SV specification to be used. On the other hand, if the performance rankings are the same under each of the four SV specifications, we would conclude that the superiority of one jump structure over the others in capturing the behavior of S&P 500 index options is unconditional and robust to perturbations in SV specifications. The empirical evidence favors the latter: the infinite activity jump structures (VG and LS) outperform the classic finite activity compound Poisson (MJ) jump structure under all four SV specifications.

Panel A of Table V addresses the question based on the t-statistics defined in equation (20). Each column in Panel A compares the performance of two jump structures under each SV specification. For example, the column under "$MJ-VG$" compares the performance of the Merton jump model (MJ) against the performance of the variance-gamma model (VG), under each of the four SV specifications. In particular, a t-statistics of 1.96 or higher implies that the pricing error from the MJ model is significantly larger than the pricing error from the VG model under a 95% confidence interval, and hence, the VG model outperforms the MJ model. A t-value of -1.96 or less implies the opposite.

The t-values under column $MJ-VG$ are strongly positive under all SV specifications, both in sample and out of sample. The same is also observed for all t-tests under the $MJ-LS$ column. Thus, our test results indicate that out of the three jump structures, the most commonly used compound Poisson

jump structure of Merton (1976), performs significantly worse than both the VG and the LS jump structures. This results holds under all of the four SV specifications and for both in-sample and out-of-sample tests. The performance difference between VG and LS, on the other hand, is much smaller and can have different signs depending upon the SV specification assumed. The t-values under the $VG-LS$ column are much smaller, positive under SV1, SV2, and SV4, but negative under SV3. Carr and Wu (2002a) obtain similar performance rankings for the three jump structures without incorporating any stochastic volatilities. Our results show that this ranking remains unchanged in the presence of stochastic volatility.

The key structural difference between the Merton jump model and the other two types of jump structures lies in the jump frequency specification. Within any finite time interval, the number of jumps under MJ is finite and is captured by the jump intensity measure λ. The estimates for λ under the MJ structure fall between 0.086 under the SV4 specification (see Table IV) and 0.405 under the SV1 specification (see Table III). Specifically, an estimate of 0.405 or smaller implies that on average, one observes one jump every two and half years or so, a rare event. In contrast, under the VG and LS jump structures, the number of jumps under any finite time interval is infinite. One thus expects to observe much more frequent jumps of different magnitudes than in the Merton jump case. Our estimation results indicate that, to capture the behavior of S&P 500 index options, one needs to incorporate a much more frequent jump structure in the underlying return process than the classic Merton model allows.

B. Where Does Stochastic Volatility Come From?

By applying stochastic time changes to different Lévy components, one can generate stochastic volatility from either the diffusion component, or the jump component, or both. It thus becomes a purely empirical issue as to where exactly the stochastic volatility comes from. We address this issue by comparing the empirical performances of four different stochastic volatility specifications in pricing the S&P 500 index options.

Panel B of Table V compares the performance of the four stochastic volatility specifications under each of the three jump structures. We first look at the three one-factor SV specifications: SV1, SV2, and

SV3. We find that the in-sample t-test values under the "$SV1 - SV2$" column are all strongly negative and that the in-sample t-test values under the "$SV2 - SV3$" column are all strongly positive, suggesting that the SV2 specification is significantly outperformed by the other two one-factor SV specifications. In contrast, the in-sample t-test estimates under the "$SV1 - SV3$" column are much smaller and have different signs under different jump specifications: positive under MJ and VG, negative under LS. The out-of-sample performance comparison delivers similar conclusions, except under the LS jump structure, where the t-statistics are much smaller.

Recall that under the SV2 specification, the instantaneous variance of the diffusion component is constant and all stochastic volatilities are attributed to the time variation in the arrival rate of jumps. Inferior performance of SV2, as compared to SV1 and SV3, indicates that the instantaneous variance of the diffusion component should be stochastic. The parameter estimates of the three one-factor SV specifications in Table III also tell a similar story. The volatility of volatility estimates (σ_v) are always strongly positive under SV1 specifications, slightly smaller under SV3 specifications, but are close to zero under SV2 specifications, when only the arrival rate of the jump component is allowed to be stochastic. For example, the estimate of σ_v is 2.136 under VGDSV1, 1.745 under VGDSV3, but a mere 0.001 under VGDSV2. Similar results hold for MJ and LS models. These estimates indicate that, overall, the arrival rate of the jump component is not as volatile as the instantaneous variance of the diffusion component. This evidence supports traditional stochastic volatility specifications but casts doubt on the performance of the stochastic volatility models of Carr, Geman, Madan, and Yor (2001), which generate stochastic volatility from pure jump models.

Another important structural difference between the SV2 specification and the other SV specifications is that SV2 is the only specification where instantaneous correlation is *not* incorporated between the return innovation and the innovation in the arrival rate. Hence, the SV2 specification cannot capture the widely documented negative correlation between stock returns and return volatilities, i.e., the "leverage effect."[9] Yet, under all other SV specifications, the estimates for this instantaneous correlation parameter, ρ, are all strongly negative (see Table III), suggesting the importance of incorporating such a leverage effect in capturing the behavior of S&P 500 index option prices. In particular, this

[9]Black (1976) first documented this phenomenon and attributed it to the "leverage effect;" however, various other explanations have also been proposed in the literature, e.g., Haugen, Talmor, and Torous (1991), Campbell and Hentschel (1992), Campbell and Kyle (1993), and Bekaert and Wu (2000).

negative correlation helps in generating negative skewness in the conditional index return distribution implied by the option prices.

Consistent with this observation, Carr, Geman, Madan, and Yor (2001) also note that, without the leverage effect, the performance of the SV3 specification declines to approximately the same level as the SV2 specification. Therefore, this lack of negative correlation under SV2 constitutes another key reason for its significantly worse performance compared to other one-factor SV specifications.

In contrast to the three one-factor SV specifications, the SV4 specification allows the instantaneous variance of the diffusion component and the arrival rate of the jump component to vary separately. The *t*-statistics in Table V indicate that this extra flexibility significantly improves the model performance. The *t*-tests for performance comparisons between SV4 and all the one-factor SV specifications are strongly negative, both in sample and out of sample, indicating that the two-factor SV4 models perform much better than all the one-factor SV models. This superior performance of the SV4 models indicates that stochastic volatility actually comes from two *separate* sources: the instantaneous variance of the diffusion component and the arrival rate of the jump component.

The superior performance of the SV4 models has important implications in practice. First, it implies that a high volatility day on the market can be due to either intensified arrival of large events or increased arrival of small, diffusive events, or both. The exact source of high volatility is hence subject to further research and shall be case dependent. This result is in contrast to the implication of earlier option pricing models, e.g. Bates (1996) and Bakshi, Cao, and Chen (1997), both of which assume that variations in volatility can only come from variations in the diffusive volatility.

Furthermore, the superior performance of SV4 models also indicate that, of the four SV specifications, SV4 models suffer the least from model misspecification. Hence, comparisons of different jump structures should be the least biased when the comparison is based upon the SV4 framework. The ranking of the three jump structures under SV4 specifications is, from worst to best, $MJ < VG < LS$, with the difference between any pair being statistically significant based on the *t*-statistics. Recall that the jump frequency increases from MJ to VG and to LS. The performance ranking is in line with this ranking of jump frequency for different jump structures. Therefore, we conclude that the market prices

the S&P 500 index options as if the discontinuous index level movements are frequent occurrences and not rare events.

C. How Do the Risk-Neutral Dynamics of the Two Activity Rates Differ?

Since the SV4 specification provides an encompassing framework for all the one-factor SV specifications, we can learn more about the risk-neutral dynamics of the activity rates by investigating the relevant parameter estimates of the SV4 models, which are reported in Table IV.

Based on the square-root specification for the risk-neutral activity rate dynamics, the two elements of $\sigma_v = \left[\sigma_v^d, \sigma_v^j\right]^\top$ capture the instantaneous volatility of the two activity rate processes, with σ_v^d capturing the instantaneous volatility of the diffusion variance and σ_v^j the instantaneous volatility of the jump arrival rate. The estimates indicate that the variance of the diffusion component exhibits larger instantaneous volatility than the arrival rate of the jump component. For example, the estimates for σ_v^d are 2.417, 2.600, and 4.697 when the jump components are MJ, VG, and LS, respectively. In contrast, the corresponding estimates for σ_v^j are 1.644, 1.433, and 2.582, about half the magnitude for σ_v^d.

On the other hand, the relative persistence of the activity rate dynamics is captured by the two elements of $\kappa = \left[\kappa^d, \kappa^j\right]^\top$. A smaller value for κ implies a more persistent process. The estimates reported in Table IV indicate that the arrival rate of the jump component exhibits a much more persistent risk-neutral dynamics than the instantaneous variance of the diffusion component. Specifically, the estimates for κ^j are 0.002, 0.001, and 0.096, when the jump components are MJ, VG, and LS, respectively, much smaller than the corresponding estimates for κ^d, which are 2.949, 3.045, and 3.466, respectively.

The parameter estimates for the SV4 specifications indicate that, to match the market price behavior of S&P 500 index options, one needs to derive stochastic volatilities from two separate sources: the instantaneous variance of the diffusion component and the arrival rate of the jump component. Furthermore, the risk-neutral dynamics of the diffusion variance needs to exhibit higher instantaneous volatility and much less persistence than the risk-neutral dynamics of the jump arrival rate. Such different risk-neutral dynamics for the two activity rate processes dictate that the jump component and the diffusion component play different roles in governing the behavior of S&P 500 options. In particular,

the more volatile but also more transient feature of the activity rate from the diffusion component implies that it is more likely to dominate the price behavior of the short term options. On the other hand, although the activity rate from the jump component is not as volatile, its highly persistent nature dictates that its impact is more likely to last longer and hence dominate the behavior of long term options. These different impacts generate potentially testable implications on the time series behavior of S&P 500 index options. This is left for future research.

D. Shall We Take the Diffusion Component for Granted?

One consensus in the option pricing literature is that to account for the pricing biases in the Black and Scholes (1973) model, one needs to *add* both a jump component and stochastic volatility. This consensus implicitly takes the Brownian motion component in the Black-Scholes model for granted. This is not surprising given that most of the jump models in the literature are variations of the *finite activity* compound Poisson jump model of Merton (1976). In these models, the number of jumps within a finite interval is finite. For example, under the MJDSV1 model, which is also estimated in Bates (1996) and Bakshi, Cao, and Chen (1997), our estimate for the Poisson intensity is 0.405, which implies approximately an average of one jump every two and half years. Obviously, one needs to add a diffusion component to fill the "gaps" between the very infrequent jumps.

However, if one considers jump processes with infinite activity, or even infinite variation, the infinitely many small jumps generated from such models can be imagined to fill these gaps. In particular, Carr, Geman, Madan, and Yor (2002) conclude from their empirical study that a diffusion component is no longer necessary as long as one adopts an infinite activity pure jump process. Carr and Wu (2002a) arrive at similar conclusions in their infinite variation log stable (LS) model. Most recently, Carr and Wu (2002c) propose a method to identify the presence of jump and diffusion components in the underlying asset price process by investigating the short maturity behavior of at-the-money and out-of-the-money options underlying such an asset. In particular, they prove that a jump component, if present, dominates the short maturity behavior of out-of-the-money options and hence can be readily identified. A diffusion component, if present, usually dominates the short maturity behavior of at-the-money options. Nevertheless, they find that, in theory, an infinite variation jump component can also generate the same short maturity behavior for at-the-money options as does a diffusion process. The

same infinite variation feature for both a Brownian motion and an infinite variation pure jump process dictates that they generate similar short maturity behaviors for at-the-money options.

The above empirical and theoretical findings lead us to ask questions beyond the traditional framework of thinking: Do we really need a diffusion component if we include an infinite activity jump component in the option pricing model? Can we separately identify a diffusion component from an infinite activity jump component, especially one also with infinite variation? These questions are especially relevant here as our estimation results strongly favor the infinite activity jump components, and the infinite variation LS jump component in particular, over the more traditional finite activity compound Poisson MJ jump specification.

As can be seen from Tables III and IV, under all the tested models with infinite activity jump components (VG or LS), the estimates for the diffusion component, σ, are all significantly different from zero, indicating that the diffusion component is both identifiable and needed. The key difference between our models and those estimated in Carr, Geman, Madan, and Yor (2002) and Carr and Wu (2002a) is that we have incorporated stochastic volatility while they consider pure Lévy processes without stochastic volatility. Thus, our identification of the diffusion component comes from its role in generating stochastic volatility. In particular, the separate specification of the two activity rate processes under SV4 implies that the relative proportion of small (diffusive) movements and large (jump) movement can vary over time. Their different risk-neutral dynamics further implies that the two components can separately dominate the price behaviors of options at different maturities.

Furthermore, our empirical work focuses on a purely diffusive specification for the activity rate process, i.e., the Heston (1993) model. Under such a specification, any instantaneous negative correlation between the activity rate process and the return innovation has to be incorporated via a diffusion component in the return process because, after all, a pure jump component is by definition orthogonal to any diffusion components. Thus, under our specification, the diffusion component in the return process is not only important in providing a separate source of stochastic volatility, but also indispensable in providing a vehicle to accommodate the "leverage effect." Conceivably, one can incorporate a jump component in the activity rate process as in Chernov, Gallant, Ghysels, and Tauchen (1999) and Eraker, Johannes, and Polson (2003), and thus accommodate the leverage effect via a correlation between the jump component in the return process and the jump component in the activity rate processes. When

these two jump components exhibit infinite variation, the need for a separate diffusion component could potentially be reduced. We leave this issue for future research.

Indeed, even under our diffusive activity rate specification, the model parameter estimates indicate that the relative proportion of the diffusion component declines as the jump specification goes from finite activity (MJ) to infinite activity but finite variation (VG) and to infinite variation (LS). Given that all models are calibrated to the same data set, the estimate of the diffusion parameter σ represents the relative weight of the diffusion component compared to the jump component. The decline in the relative weight of the diffusion component holds for all SV specifications. For instance, among the SV4 models, as shown in Table IV, the estimate of σ (the diffusion component) is 0.279 for MJDSV4, 0.276 for VGDSV4, but 0.262 for LSDSV4. Similar declines are also observed under SV1 specifications (from 0.352, to 0.318, and then to 0.309) and SV3 specifications (from 0.301, to 0.272, and then to 0.175). The most dramatic decline, however, comes under the SV2 specification: The estimate for σ is 0.173 under MJDSV2, 0.157 under VGDSV2, but a meager 0.044 under LSDSV2. Recall that SV2 differs from all other SV specifications in (a) generating stochastic volatility from the jump component only, and (b) not accommodating a leverage effect. Thus, consistent with our above discussion, without a role in either generating stochastic volatility or accommodating a leverage effect, the diffusion component is hardly needed when the jump component also exhibits infinite variation as in the case of LSDSV2.

Putting all the evidence together, we conclude that as the frequency of jump arrival increases from MJ, to VG, and then to LS, the need for a diffusive component declines. The many small jumps in infinite variation jump components can partially replace the role played by a diffusion component. Nevertheless, under our specifications, the diffusion component plays important roles in (i) providing a separate source of stochastic volatility and (ii) accommodating the leverage effect between the return innovation and the activity rate process. Therefore, under our specifications, the diffusion component cannot be totally replaced by the jump component, even if the jump component exhibits infinite variation.

IV. Pricing Error Analysis

Another way to investigate the robustness and performance of different model specifications is to check for remaining structures in the pricing errors of these models. If a model is specified reasonably well, one should find minimal structure in the pricing errors on the S&P 500 index options. We check for remaining structures in the mean pricing error at each moneyness and maturity. The mean pricing error of a good model should be close to zero and exhibit no obvious structures along both the moneyness and the maturity dimensions.

Since an option's time-to-maturity and moneyness change everyday, we estimate the pricing error at fixed moneyness and maturity via nonparametric smoothing (Appendix A). The pricing error is defined as the difference between the model implied price and the observed market price, as a percentage of the underlying spot level. Thus, a positive pricing error implies that the model overprices and a negative pricing errors implies underpricing. Figure 3 reports the smoothed in-sample pricing errors at different moneyness and maturities under each of the 12 model specifications. The mean out-of-sample pricing errors exhibit similar structures and are not reported for the sake of brevity. Twelve panels are shown as a four-by-three matrix in Figure 3, each of them corresponding to a particular model specification. The four rows runs, from top to bottom, correspond to the four SV specifications: SV1, SV2, SV3, and SV4. The three columns, from left to right, correspond to the three jump structures: MJ, VG, and LS. Thus, the panel at the top left corner denotes mean pricing errors from model MJDSV1, the panel at the bottom right corner denotes mean pricing errors from model LSDSV4, and so on. Within each panel, the four lines represent pricing errors for four maturities: 0.1 (solid), 0.5 (dashed), 1.0 (dot-dashed), and 1.5 years (dotted).

For comparison, we use the same scale for all panels except for the second row, where a larger scale is used to accommodate the larger pricing errors from the SV2 models. As can be seen from the figure, the three SV2 models exhibit large mean pricing errors along both the maturity and the moneyness dimensions. In particular, at short maturities, SV2 models overprice out-of-the-money put options ($k < 0$) relative to out-of-the-money call options ($k > 0$). At long maturities, the pattern is reversed. Out-of-the-money put options are underpriced relative to out-of-the-money call options. One can also see from the figure that SV1 and SV3 models perform well along the moneyness dimension but

not as well along the maturity dimension. In contrast, the three two-factor SV4 specifications exhibit much better performance. As shown in the bottom row in Figure 3, the in-sample pricing errors for SV4 models are much smaller than other SV specifications. In particular, very little structure is left in the pricing errors of the LSDSV4 model (the bottom right panel in Figure 3).

In the option pricing literature, it has become a standard practice to document the option price behavior in terms of the Black-Scholes implied volatility. For S&P 500 index options, at a given maturity level, the Black and Scholes (1973) implied volatilities for out-of-the-money puts are much higher than those for out-of-the-money calls. (See empirical documentations in, for example, Aït-Sahalia and Lo (1998), Jackwerth and Rubinstein (1996), and Rubinstein (1994).) This phenomenon is commonly referred to as the "volatility smirk." It is widely accepted that the implied volatility smirk is a direct result of conditional non-normality in asset returns. In particular, the downward slope of the smirk reflects asymmetry (negative skewness) in the risk-neutral distribution, while the positive curvature of the smirk reflects the fat-tails (leptokurtosis) of this distribution. Yet, the central limit theorem implies that under fairly general conditions, the conditional return distribution should converge to normality as the maturity increases. As a result, the volatility smirk should flatten out accordingly. However, Carr and Wu (2002a) find that when implied volatilities are graphed against a standard measure of moneyness, the resulting implied volatility smirk does not flatten out as maturity increases up to the maximum observable horizon of two years. Such a maturity pattern seems to run against the implications of the central limit theorem and presents challenges for option pricing modeling. The literature has used two approaches to account for this maturity pattern of the volatility smirk: (1) incorporating a persistent stochastic volatility process to slow down the convergence to normality, and (2) adopting an α-stable process as in Carr and Wu (2002a) so that the traditional central limit theorem does not hold and return non-normality does not disappear with aggregation.

The bias of SV2 models as shown in Figure 3 implies that the SV2 models generate steeper implied volatility smirks than observed in the data at short maturities and flatter ones than observed in the data at long maturities. Taken together, the SV2 model implies that volatility smirk flattens out faster than observed in the data as maturity increases. Thus, the remaining structure in the mean pricing error for SV2 models indicates that the SV2 specification fails to meet the challenge of accounting for the maturity pattern of the volatility smirk for S&P 500 index options. This observed failure implies

that not just any persistent stochastic volatility model will work. The better performance of other SV specifications further suggests that, for successful model design, it is imperative to also address the issue of how the stochastic volatility is incorporated into the return process.

Both Figure 3 and Tables IV and V show that SV4 type models are promising in generating a persistent volatility smirk across the maturity horizon. In particular, the best performance of the LSDSV4 model can be attributed to a combination of two attributes: the LS jump structure and the SV4 specification. The LS jump structure is specifically designed by Carr and Wu (2002a) to capture the maturity pattern of the implied volatility smirk for S&P 500 index options. Under this LS model, the central limit theorem does not apply and conditional non-normality remains as maturity increases so that the model can generate a relatively stable maturity pattern for the implied volatility smirk. The SV4 specification further improves the performance by generating variations in the relative proportion of the jump component and the diffusion component along the option maturity dimension. In particular, since the estimated stochastic jump volatility is more persistent than the estimated diffusive volatility under the risk-neutral measure ($\kappa^j = 0.096$ versus $\kappa^d = 3.466$ under LSDSV4, see Table IV), the impact of the more persistent jump component dominates the behavior of long term options while the more transient diffusion component contributes more to short term options. Since non-normality is mainly generated from the jump component, the progressively increasing proportion of the jump component with increasing maturities counteracts with the central limit theorem and helps further in maintaining a relatively stable, and slightly steepening, maturity pattern for the implied volatility smirk.

V. Concluding Remarks

We classify option pricing models based on time-changed Lévy processes. Specifically, we consider candidate (underlying asset) return processes that are generated by applying stochastic time changes to Lévy processes - which can have both diffusion and jump components. We then classify option pricing models by the following features: (i) the specification of the jump component in the return process; (ii) the source for stochastic volatility, namely, if it is generated from stochastic diffusive volatility or jump volatility, or both; and (iii) the dynamics of the volatility process itself. Based on this classification scheme, we propose and test a variety of new option pricing models and address a

few model design issues which have not been answered in the literature. In particular, we focus on answering two questions: (i) what type of jump structure best captures the behavior of the S&P 500 index options? (ii) where does stochastic volatility come from?

We find that a high frequency jump structure always outperforms the traditional low frequency jump specification - the compound Poisson model. The implication of this finding is that the market prices the S&P 500 index options as if discrete movements in the index level are frequent events instead of rare events. We also find that stochastic volatility comes from two separate sources: the instantaneous variance of a diffusion component and the arrival rate of a jump component. While the risk-neutral dynamics of the diffusion variance is more volatile, the risk-neutral dynamics of the arrival rate of the jump component exhibits much more persistence. As a result, stochastic volatility from diffusion dominates the behavior of short term option prices while stochastic volatility from jumps dominates that of long term option prices.

In summary, our empirical results with the S&P 500 index option data indicate that a model of the underlying index returns should include a high frequency jump component in the return process and allow the stochastic return volatility to be driven independently by diffusion and jumps.

To maintain the scope of the paper to a manageable level, we consider only one activity rate specification in our empirical study, which is the square-root model of Heston (1993). Yet, the framework proposed here provides fertile ground for extensions and future research. One potential line of future research is to investigate the relative performance of different activity rate specifications. In particular, a series of recent studies, e.g., Chernov, Gallant, Ghysels, and Tauchen (1999) and Eraker, Johannes, and Polson (2003), incorporate jumps into the activity rate process, in addition to jumps in the asset return processes. Nevertheless, all these studies consider only compound Poisson jumps with potentially time varying arrival rates. In light of the findings in this article on the better performance of models which include a high frequency jump component in the return process, it is intriguing to see whether incorporating such jumps in the stochastic volatility process will also deliver superior performance over specifications of Poisson jumps in volatility used in the existing literature.

Finally, since the objective of this paper is to analyze the relative performance of different models in pricing options on a daily basis, we focus on the model specifications under the risk-neutral measure

and calibrate the models solely to the options data. One can potentially analyze the properties of the risk premia on the jump component, the diffusion component, and the stochastic activity rates through an integrated analysis of the time series of both the option prices and the underlying spot prices along the lines of Eraker (2001) and Pan (2002).

Appendix A. Nonparametric Estimation of Weighting Matrix

Since the moneyness and maturity of option contracts vary over time, we estimate the variance of the pricing error at fixed moneyness and maturity via nonparametric kernel regression.

Let τ denote time to maturity and $k = \ln(K/S)$ denote moneyness. Define the information set $\mathbf{Z} \equiv (\tau, k)$. Suppose that there are total $\mathcal{N}$ observations. Given a kernel function $\mathcal{K}(\cdot)$ and the *bandwidth* matrix H, the kernel estimate of the variance $V(\mathbf{Z})$ is given by

$$\widehat{V}(\mathbf{Z}) = \frac{\sum_{i=1}^{\mathcal{N}} \mathcal{K}\left(\frac{|\mathbf{Z}-\mathbf{Z}_i|}{H}\right)(e_i)^2}{\sum_{i=1}^{\mathcal{N}} \mathcal{K}\left(\frac{|\mathbf{Z}-\mathbf{Z}_i|}{H}\right)} - [\widehat{\mu}(\mathbf{Z})]^2; \quad \widehat{\mu}(\mathbf{Z}) = \frac{\sum_{i=1}^{\mathcal{N}} \mathcal{K}\left(\frac{|\mathbf{Z}-\mathbf{Z}_i|}{H}\right) e_i}{\sum_{i=1}^{\mathcal{N}} \mathcal{K}\left(\frac{|\mathbf{Z}-\mathbf{Z}_i|}{H}\right)}. \tag{A1}$$

There are a variety of choices of kernels and bandwidth in the literature. We refer to the monograph by Simonoff (1996) on this issue. In our analysis, we follow Aït-Sahalia and Lo (1998) in choosing independent Gaussian kernels and setting the bandwidths according to:

$$h_j = c_j \sigma_j \mathcal{N}^{-1/6}, \qquad j = \tau, k, \tag{A2}$$

where σ_j is the standard deviation of the regressor Z_j and c_j is a constant which is typically of order of magnitude one. The larger the coefficient c_j is, the smoother the estimates are across moneyness and maturities. In our application, we set $c_\tau = c_k = 4$.

References

Aït-Sahalia, Yacine, and Andrew Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.

Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2002, An empirical investigation of continuous-time equity return models, *Journal of Finance* 57, 1239–1284.

Ané, T., and Hélyette Geman, 2000, Order flow, transaction clock and normality of asset returns, *Journal of Finance* 55, 2259–2284.

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

Barndorff-Nielsen, Ole E., 1998, Processes of normal inverse gaussian type, *Finance and Stochastics* 2, 41–68.

Barndorff-Nielsen, Ole E., and Neil Shephard, 2001, Non-Gaussian Ornstein-Uhlenbeck based models and some of their uses in financial economics, *Journal of the Royal Statistical Society-Series B* 63, 167–241.

Bates, David, 1996, Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options, *Review of Financial Studies* 9, 69–107.

Bates, David, 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.

Bekaert, Geert, and Guojun Wu, 2000, Asymmetric volatilities and risk in equity markets, *Review of Financial Studies* 13, 1–42.

Bertoin, Jean, 1996, *Lévy Processes*. (Cambridge University Press Cambridge).

Black, Fisher, 1976, Studies of stock price volatility changes, in *Proceedings of the 1976 American Statistical Association, Business and Economical Statistics Section* (American Statistical Association, Alexandria, VA).

Black, Fisher, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Campbell, John Y., and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Review of Economic Studies* 31, 281–318.

Campbell, John Y., and Albert S. Kyle, 1993, Smart money, noise trading and stock price behavior, *Review of Economic Studies* 60, 1–34.

Carr, Peter, Hélyette Geman, Dilip Madan, and Marc Yor, 2001, Stochastic volatility for Lévy processes, *Mathematical Finance* forthcoming.

Carr, Peter, Hélyette Geman, Dilip Madan, and Marc Yor, 2002, The fine structure of asset returns: An empirical investigation, *Journal of Business* 75, 305–332.

Carr, Peter, and Liuren Wu, 2002a, Finite moment log stable process and option pricing, *Journal of Finance* forthcoming.

Carr, Peter, and Liuren Wu, 2002b, Time-changed Lévy processes and option pricing, *Journal of Financial Economics* forthcoming.

Carr, Peter, and Liuren Wu, 2002c, What type of process underlies options? A simple robust test, *Journal of Finance* forthcoming.

Chernov, Mikhail, A. Ronald Gallant, Eric Ghysels, and George Tauchen, 1999, A new class of stochastic volatility models with jumps: Theory and estimation, Cirano working paper 99s-48, Pennsylvia State University.

Clark, Peter K., 1973, A subordinated stochastic process with finite variance for speculative prices, *Econometrica* 41, 135–155.

Cont, Rama, and Jose da Fonseca, 2002, Dynamics of implied volatility surfaces, *Quantitative Finance* 2, 45–60.

Cox, John C., Jonathan E. Ingersoll, and Stephen R. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–408.

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

Ding, Zhuanxin, Robert F. Engle, and Clive W. J. Granger, 1993, A long memory property of stock returns and a new model, *Journal of Empirical Finance* 1, 83–106.

Ding, Zhuanxin, and Clive W. J. Granger, 1996, Modeling volatility persistence of speculative returns: A new approach, *Journal of Econometrics* 73, 185–215.

Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379–406.

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

Dumas, Bernard, Jeff Fleming, and Robert E. Whaley, 1998, Implied volatility functions: Empirical tests, *Journal of Finance* 53, 2059–2106.

Eberlein, Ernst, Ulrich Keller, and Karsten Prause, 1998, New insights into smile, mispricing, and value at risk: The hyperbolic model, *Journal of Business* 71, 371–406.

Eraker, Bjorn, 2001, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, manuscript, Duke University.

Eraker, Bjorn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in equity index volatility and returns, *Journal of Finance* forthcoming.

Geman, Hélyette, Dilip Madan, and Marc Yor, 2001, Time changes for Lévy processes, *Mathematical Finance* 11, 79–96.

Hansen, Lars Peter, and Ravi Jagannathan, 1997, Assessing specification errors in stochastic discount factor models, *Journal of Finance* 52, 557–590.

Haugen, Robert A., Eli Talmor, and Walter N. Torous, 1991, The effect of volatility changes on the level of stock prices and subsequent expected returns, *Journal of Finance* 46, 985–1007.

Heston, Stephen, 1993, Closed-form solution for options with stochastic volatility, with application to bond and currency options, *Review of Financial Studies* 6, 327–343.

Hull, John, and Alan White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281–300.

Jackwerth, Jens Carsten, and Mark Rubinstein, 1996, Recovering probability distributions from contemporary security prices, *Journal of Finance* 51, 347–369.

Jacod, Jean, and Albert N. Shiryaev, 1987, *Limit Theorems for Stochastic Processes*. (Springer-Verlag Berlin).

Janicki, Aleksander, and Aleksander Weron, 1994, *Simulation and Chaotic Behavior of* α*-Stable Stochastic Processes*. (Marcel Dekker New York).

Leippold, Markus, and Liuren Wu, 2002, Asset pricing under the quadratic class, *Journal of Financial and Quantitative Analysis* 37, 271–295.

Madan, Dilip, and F. Milne, 1991, Option pricing with VG martingale components, *Mathematical Finance* 1, 39–56.

Madan, Dilip B., Peter P. Carr, and Eric C. Chang, 1998, The variance gamma process and option pricing, *European Finance Review* 2, 79–105.

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

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

Rubinstein, Mark, 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.

Samorodnitsky, Gennady, and Murad S. Taqqu, 1994, *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. (Chapman & Hall New York).

Sato, Ken-Iti, 1999, *Lévy Processes and Infinitely Divisible Distributions*. (Cambridge University Press Cambridge).

Scott, Louis O., 1997, Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: Applications of fourier inversion methods, *Mathematical Finance* 7, 413–426.

Simonoff, Jeffrey S., 1996, *Smoothing Methods in Statistics*. (Springer-Verlag New York).

Titchmarsh, E. C., 1975, *Theory of Fourier Integrals*. (Oxford University Press London) 2nd edn.

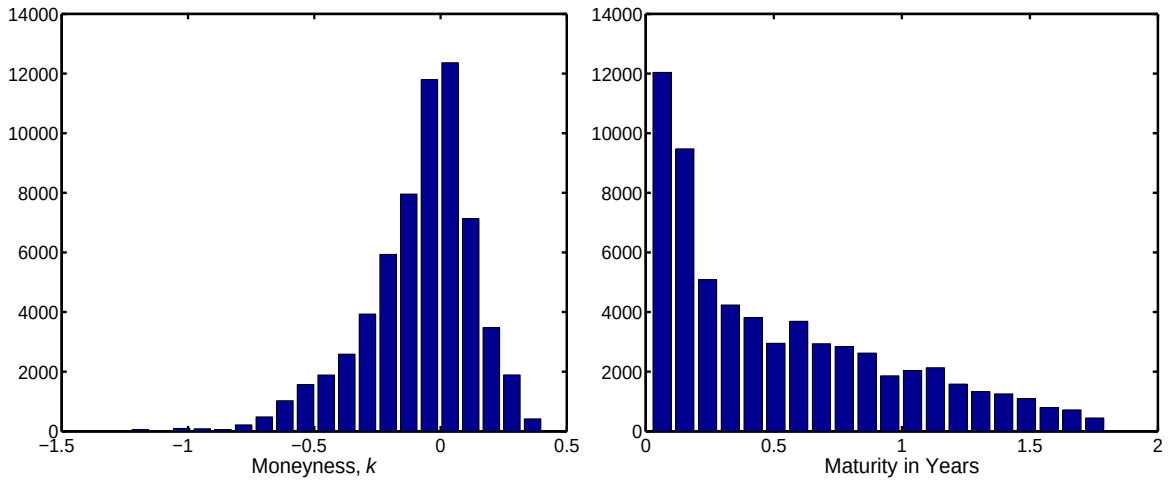

Figure 1. Histograms of OTM Option Prices

The left panel depicts the histogram of the moneyness $k = \ln(K/S)$ and the right panel depicts the histogram of the maturities for all available S&P 500 index options in our cleaned sample. The sample is daily from April 6th, 1999 to May 31st, 2000, with 290 business days and 62,950 option contracts.

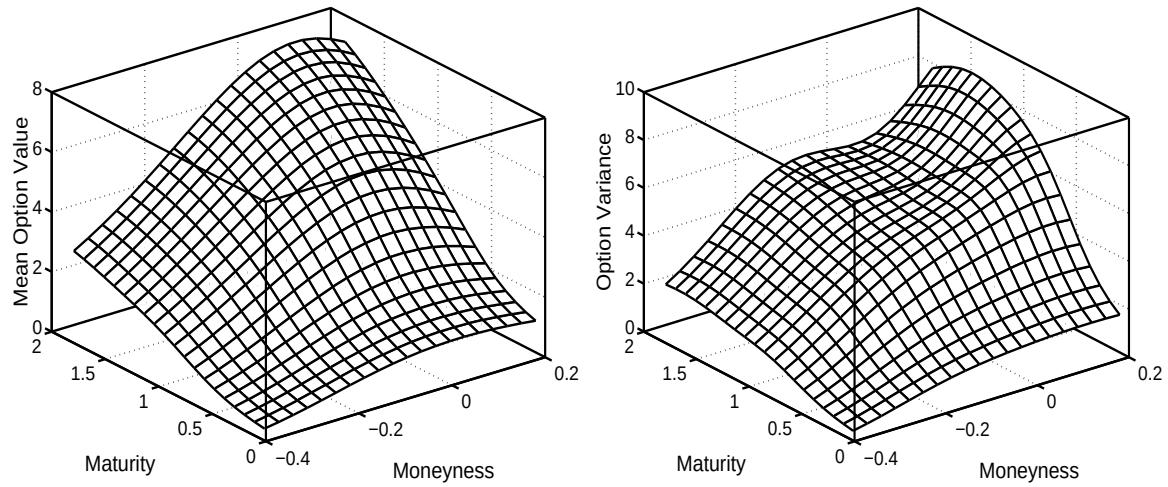

Figure 2. Mean and Variance Surface of Option Prices

The mean (left panel) and variance (right panel) of S&P 500 index option prices as percentages of the index level at each moneyness $k = \ln(K/S)$ and maturity (in years) are estimated nonparametrically using independent Gaussian kernels. The sample of S&P 500 index options is daily from April 6th, 1999 to May 31st, 2000 with 290 business days and 62,950 option quotes.

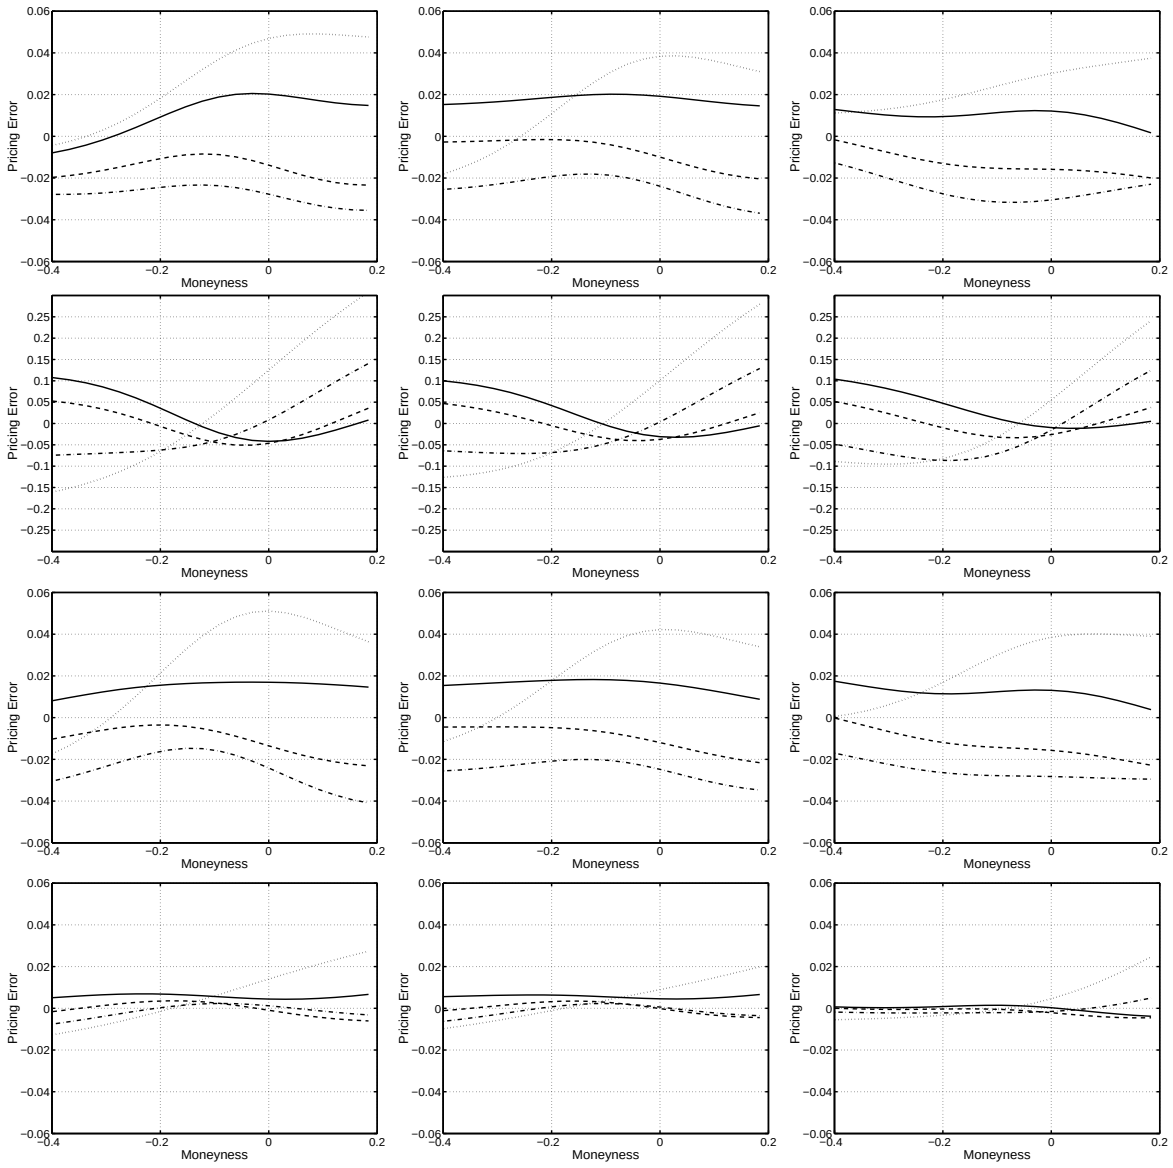

Figure 3. In-Sample Mean Pricing Errors

Pricing errors are defined as the difference between the model implied option price and the market observed price, in percentages of the underlying spot level. Mean pricing errors at fixed moneyness ($k = \ln(K/S)$) and maturity are estimated nonparametrically using independent Gaussian kernels. Each panel denotes one model. The jump component of the model is, from left to right, MJ, VG, and LS. The stochastic volatility specification is, from top to bottom, SV1, SV2, SV3, and SV4. The four lines in each panel denote for maturities: 0.1 (solid line), 0.5 (dashed line), 1.0 (dash-dotted line), and 1.5 years (dotted line). For ease of comparison, we use the same scale for all panels except for the second row, where a much larger scale is used to accommodate the much larger pricing errors of SV2 models.

Table III
Parameter Estimates of One-Factor SV Models

Model parameters are estimated by minimizing the sum of daily mean squared errors. Given model parameters, the daily mean squared errors are obtained by choosing the activity rate level at that day to minimize the sum of the weighted squared pricing errors on that day. Entries report the parameter estimates and standard errors (in parentheses), based on the first 139 days of data. Also reported are the sample average and standard deviation of the daily mean squared error for both the in-sample period ($\texttt{mse}_I$, the first 139 days) and the out-of-sample period ($\texttt{mse}_O$, the last 150 days). The pricing error is defined in percentages of the spot price.

	Model specifications								
Θ	MJDSV1	MJDSV2	MJDSV3	VGDSV1	VGDSV2	VGDSV3	LSDSV1	LSDSV2	LSDSV3
σ	0.352	0.173	0.301	0.318	0.157	0.272	0.309	0.044	0.175
	(0.288)	(0.010)	(0.031)	(0.037)	(0.006)	(0.030)	(0.037)	(0.001)	(0.029)
λ	0.405	0.364	0.223	0.253	0.593	0.985	0.028	0.077	0.053
	(0.459)	(0.135)	(0.190)	(0.194)	(0.068)	(1.115)	(0.003)	(0.007)	(0.007)
α	−0.091	−0.393	−0.408	−0.247	−0.391	−0.244	1.673	1.578	1.738
	(0.052)	(0.100)	(0.149)	(0.119)	(0.029)	(0.122)	(0.044)	(0.028)	(0.066)
σ_j	0.175	0.235	0.000	0.264	0.013	0.003	− − −	− − −	− − −
	(0.113)	(0.035)	(0.000)	(0.108)	(0.001)	(0.001)	− − −	− − −	− − −
κ	1.039	2.070	1.110	0.813	2.054	0.974	0.795	0.867	1.304
	(0.320)	(0.000)	(0.346)	(0.322)	(0.237)	(0.327)	(0.293)	(0.000)	(0.324)
σ_v	2.574	0.000	1.983	2.136	0.001	1.745	2.253	0.000	1.839
	(0.620)	(0.000)	(0.361)	(0.431)	(0.000)	(0.321)	(0.358)	(0.000)	(0.344)
ρ	−0.704	− − −	−0.648	−0.692	− − −	−0.662	−1.000	− − −	−1.000
	(0.073)	− − −	(0.092)	(0.075)	− − −	(0.101)	(0.000)	− − −	(0.000)
mse_I	0.334	1.159	0.307	0.302	0.927	0.279	0.256	0.859	0.279
	(0.254)	(0.314)	(0.247)	(0.244)	(0.272)	(0.236)	(0.218)	(0.262)	(0.237)
mse_O	2.105	2.599	1.752	1.868	2.339	1.531	1.813	1.610	1.739
	(1.123)	(1.127)	(0.968)	(0.982)	(1.046)	(0.874)	(1.030)	(0.734)	(0.968)

Table IV
Parameter Estimates of SV4 Models

Model parameters are estimated by minimizing the sum of daily mean squared errors. Given model parameters, the daily mean squared errors are obtained by choosing the activity rate level at that day to minimize the sum of the weighted squared pricing errors on that day. Entries report the parameter estimates and standard errors (in parentheses), based on the first 139 days of data. The superscripts d and j on a parameter denote respectively the diffusion and jump components of that parameter vector. Also reported are the sample average and standard deviation of the daily mean squared error for both the in-sample period ($\texttt{mse}_I$, the first 139 days) and the out-of-sample period ($\texttt{mse}_O$, the last 150 days). The pricing error is defined in percentages of the spot price.

Θ	MJDSV4		VGDSV4		LSDSV4	
σ	0.279	(0.009)	0.276	(0.008)	0.262	(0.006)
λ	0.086	(0.617)	0.003	(0.022)	0.032	(0.009)
α	−0.119	(0.041)	−0.184	(0.057)	1.833	(0.034)
σ_j	0.276	(0.056)	0.298	(0.076)	− − −	(0.000)
κ^d	2.949	(0.354)	3.045	(0.337)	3.466	(0.480)
κ^j	0.002	(0.088)	0.001	(0.008)	0.096	(0.045)
σ_v^d	2.417	(0.326)	2.600	(0.341)	4.697	(0.750)
σ_v^j	1.644	(4.346)	1.433	(0.237)	2.582	(0.701)
ρ^d	−0.788	(0.085)	−0.707	(0.080)	−0.522	(0.081)
ρ^j	−0.931	(2.464)	−0.999	(0.018)	−0.645	(0.140)
mse_I	0.096	(0.144)	0.089	(0.144)	0.074	(0.133)
mse_O	0.666	(0.562)	0.625	(0.528)	0.216	(0.202)

Table V
Pair-Wise t-statistics for Model Comparisons

Entries report the t-statistics defined in equation (20). Tests in Panel A compare the performance of different jump structures under each stochastic volatility specification (SV1 to SV4). Tests in Panel B compare the performance of different SV specifications given a fixed jump structure. For each test (Model i -Model j), a t-value greater than 1.96 implies that the mean squared pricing error from model i is significantly larger than the mean squared error from model j and hence model j outperforms model i, at 95% confidence interval. A t-value less than -1.96 implies the opposite. In-sample tests are based on the first 139 days of option price data while the out-of-sample tests are based on the last 150 days of data, given parameter estimates from the first sub-sample.

Panel A. Testing Which Jump Structure Performs the Best

Cases/Tests	$MJ-VG$	$MJ-LS$	$VG-LS$	$MJ-VG$	$MJ-LS$	$VG-LS$
	In sample comparison			Out of sample comparison		
SV1	6.70	12.45	7.31	15.98	13.74	2.69
SV2	31.73	15.75	4.75	14.95	16.47	15.89
SV3	11.95	13.27	−0.01	19.12	5.05	−19.18
SV4	9.88	5.78	3.69	10.12	12.04	11.89

Panel B. Testing Where Stochastic Volatility Comes From

Cases/Tests	$SV1-SV2$	$SV1-SV3$	$SV2-SV3$	$SV4-SV1$	$SV4-SV2$	$SV4-SV3$
			In sample comparison			
MJ	−29.53	6.74	31.16	−15.43	−44.55	−14.34
VG	−27.67	6.49	30.93	−14.09	−43.01	−13.44
LS	−35.67	−6.00	33.02	−14.94	−43.63	−14.91
			Out of sample comparison			
MJ	−10.12	18.10	17.74	−21.41	−33.46	−17.98
VG	−11.67	19.09	21.30	−20.07	−29.75	−15.54
LS	3.88	3.43	−3.09	−21.17	−27.66	−21.43