

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

JING-ZHI HUANG AND LIUREN WU*

*Huang is from Smeal College of Business, Penn State University. Wu is from Zicklin School of Business, Baruch College (CUNY). We are grateful to Rick Green (the editor), an anonymous referee, Menachem Brenner, Peter Carr, Robert Engle, Steve Figlewski, Martin Gruber, Jean Helwege, Marti Subrahmanyam, and Rangarajan Sundaram for helpful comments and discussions. We thank seminar participants at Baruch College, the University of Notre Dame, Salomon Smith Barney, and Washington University in St. Louis for helpful comments. We also thank Sandra Sizer Moore for her editing.

ABSTRACT

We analyze the specifications of option pricing models based on time-changed Lévy processes. We classify option pricing models based on the structure of the jump component in the underlying return process, the source of stochastic volatility, and the specification of the volatility process itself. Our estimation of a variety of model specifications indicates that to better capture the behavior of the S&P 500 index options, we need to incorporate a high frequency jump component in the return process and generate stochastic volatilities from two different sources, the jump component and the diffusion component.

Although researchers have proposed and tested various extensions of the Black and Scholes (1973) model, they have yet to find a model that can capture both the time series and cross-sectional properties of equity index options. The significance of this task is enormous. First, to price and hedge options, we need to explain the behavior of option prices across strike prices and maturities. Second, the information from the options market helps us understand the underlying asset return dynamics as the option price behavior reveals important information about the conditional risk-neutral distributions of the underlying return over different horizons.

In this paper, we present a unified framework that synthesizes the ongoing efforts to identify the “true” dynamics of the underlying return process by performing a specification analysis of option pricing models. We then apply this analysis to S&P 500 index options and empirically investigate some open issues regarding the equity index return dynamics.

The specification analysis we develop here is based on the theoretical framework of time-changed Lévy processes proposed by Carr and Wu (2003b). A Lévy process is a continuous time stochastic process with independent stationary increments. In addition to the Brownian motion and the compound Poisson jump process used in the traditional option pricing literature, the class of Lévy processes also includes other jump processes that exhibit higher jump frequencies and thus may better capture the dynamics of equity indexes. On the other hand, a time change is a monotonic transformation of the time variable. Stochastic volatility can be generated by applying random time changes to the individual components of a Lévy process. For example, we can apply different time changes separately to the diffusion component and the jump component of a Lévy process so that stochastic volatility can be generated from both components. Therefore, the framework of time-changed Lévy processes can be used to generate a wide class of jump-diffusion stochastic volatility models.

Within the class of time-changed Lévy processes, we classify model specifications into three separate but interrelated dimensions: the choice of a jump component in the asset return process, the identification of the sources for stochastic volatility, and the specification of the volatility process itself. Such a classification scheme encompasses almost all 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. Our 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) Which component, the jump component or the diffusion component, determines the time variation of return volatility?

The empirical analysis in this paper focuses on the performance of 12 option pricing models that are generated by a combination of three jump processes and four stochastic volatility specifications. The three jump processes include the standard compound Poisson jump process (MJ) 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 (2003a). The compound Poisson jump model (MJ) generates a finite number of jumps within any finite time interval. In contrast, 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 in the underlying asset return. We choose these three jump processes to address the question raised in (Q1).

The stochastic volatility specifications include the traditional specification in Bates (1996) and Bakshi, Cao, and Chen (1997), in which the instantaneous variance of the diffusion component is stochastic but the arrival rate of the jump component is constant. We also introduce new specifications that allow us to generate stochastic volatility separately from the jump and diffusion components. These new stochastic volatility specifications are designed to address the question raised in (Q2).

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 the compound Poisson process (MJ). This performance ranking is robust to variations in the stochastic volatility specification and holds for both in-sample and out-of-sample tests. Our results suggest that the market prices index options as if there are many discontinuous price movements of different magnitudes in the S&P 500 index. This implication favors incorporating high frequency jumps such as VG and LS in the underlying asset return process.

The estimation results also show 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 times and they vary separately. Furthermore, the model parameter estimates indicate that the return volatility generated from the diffusion and the jump components show different risk-neutral dynamics. The diffusion-induced volatility exhibits larger instantaneous variation, but the jump-induced return volatility shows much higher persistence. As a result, the behavior of short-term options is influenced more by the randomness from the diffusive movements, but the behavior of long-term options is mostly influenced by the randomness in the arrival rate of jumps.

The empirical results from our specification analysis provide further support for option pricing models that include both jumps and stochastic volatility. Nevertheless, our results also illustrate the importance of choosing the right jump structure and designing a good stochastic volatility specification. We find that the option pricing model performance can be significantly improved by including high-frequency jumps in the underlying return process and by generating stochastic return volatility separately from the jump and diffusion components.

The paper is organized as follows. Section I discusses the studies that form the background for our paper. Section II constructs option pricing models through time changing Lévy processes. Section III addresses the data and estimation issues. Section IV compares the empirical performance of different model specifications. Section V analyzes the remaining structures in the pricing errors for different models. Section VI concludes with suggestions for future research.

I. Background

Black and Scholes (1973) spawned an enormous literature on option pricing. Their paper has also played a key role in the growth of the derivatives industry. However, the Black-Scholes model has been known to systematically misprice equity index options, especially those that are out-of-the-money.

A key assumption underlying the Black-Scholes model is that the log return of the underlying asset is normally distributed. However, the empirical option pricing literature has documented three stylized facts that run counter to this assumption. First, for a given option maturity, the Black-Scholes implied volatilities for out-of-the-money put options are much higher than those of call options that are equally out-of-the-money.¹ This phenomenon is called the “volatility smirk.”

Second, when we plot the Black-Scholes implied volatilities against a standardized measure of moneyness, the resulting implied volatility smirk does not flatten out, but steepens slightly as the option maturity increases. This standardized moneyness measure is defined as the logarithm of the strike over the forward price of the underlying, normalized by the square root of maturity. Carr and Wu (2003a) document this phenomenon on S&P 500 index options with option maturities up to two years. More recently, Foresi and Wu (2003) find that the same maturity pattern holds for all major equity indexes in the world and for time-to-maturities up to five years.

Third, both the level of the implied volatility and the shape of the implied volatility surface across moneyness and maturity vary over time (Cont and da Fonseca (2002)).

It is widely recognized that the implied volatility smirk is a direct result of conditional non-normality of the return on the equity index. The volatility smirk reflects asymmetry (negative skewness) and fat-tails (leptokurtosis) in the risk-neutral distribution of the underlying index return. The maturity pattern of this smirk indicates that the conditional non-normality of the return distribution does not decline with increasing horizon, as we might expect from the central limit theorem. The option pricing literature generates conditional return non-normality either by incorporating a jump component in the underlying index return process (Merton (1976)), or by allowing the return volatility to be stochastic (Heston (1993) and Hull and White (1987)). To capture the maturity pattern of the implied volatility smirk, the general consensus is that the researcher needs to incorporate both a jump component and stochastic volatility into the model. The jump component generates return non-normality over the short terms, and a persistent stochastic volatility process slows down the convergence of the return distribution to normality as the maturity increases. Incorporating stochastic volatility is also necessary to capture the time variation and dynamic behavior of the implied volatility surface.

Most option pricing models with both jumps and stochastic volatility can be 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), Eraker (2003), 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 large, rare events. The number of jumps within any given time interval is finite. Thus, these models

classify asset price moves dichotomously, as either small and diffusive or large and rare. However, in practice we observe much more frequent discontinuous movements of different sizes in equity indexes.

Another notable feature of the current option pricing literature is that researchers often assume that stochastic volatility comes 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 an affine function of the diffusion variance. Thus, the variation in return volatility is completely determined by the variation in the volatility of the diffusion component. However, such specifications of stochastic volatility are driven more by concerns of analytical tractability than by empirical evidence. In practice, the variation in return volatility can be generated either by variations in the diffusion variance, or variations in the arrival rates of jumps, or a combination of the two. The manner in which the jump component and the diffusion component contribute differently to stochastic volatility, and how these two contributions vary over time, can be determined at an empirical level, rather than by the theoretical model specification.

We adopt the time-changed Lévy process framework of Carr and Wu (2003b) to generate option pricing models. A Lévy process can accommodate not only the Brownian motion component and the compound Poisson jump component in traditional specifications, but also more recently proposed high-frequency jump processes, e.g., the normal inverse Gaussian model of Barndorff-Nielsen (1998), the generalized hyperbolic class of Eberlein, Keller, and Prause (1998), the variance-gamma (VG) model of Madan and Milne (1991) and Madan, Carr, and Chang (1998), and the log stable model of Carr and Wu (2003a).

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. This approach may go back to Clark (1973), who suggests that a random time change can be interpreted as a cumulative measure of business activity. Ané and Geman (2000) provide empirical support for this interpretation. Examples of other applications include Barndorff-Nielsen and Shephard (2001), Carr, Geman, Madan, and Yor (2003), and Geman, Madan, and Yor (2001). In this paper, we use time change to generate stochastic volatilities from different Lévy components.

II. Model Specifications

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 that we specify the following aspects: the jump component in the return process, the source for stochastic volatility, and the dynamics of the volatility process itself. We consider 12 model specifications. Under the 12 models, the characteristic function of log returns has a closed-form solution. We then convert the characteristic functions into option prices via an efficient fast Fourier transform (FFT) algorithm (Carr and Madan (1999)).

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 a risk-neutral probability measure. We specify that under this measure \mathbb{Q} , the logarithm of the underlying stock price (index level) follows a time-changed Lévy process,

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

where r denotes the instantaneous interest rate and q the dividend yield, σ is a positive constant, W is a standard Brownian motion, and J denotes a compensated pure Lévy jump martingale process. The vector $T_t \equiv [T_t^d, T_t^j]^\top$ 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 vector process with left limits satisfying the usual regularity conditions. The time change T_t is finite \mathbb{Q} -a.s. for all $t \geq 0$ and $T_t \rightarrow \infty$ as $t \rightarrow \infty$.

Although 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 . Let

$$v(t) \equiv [v^d(t), v^j(t)]^\top = \partial T_t / \partial t. \quad (2)$$

Then $v^d(t)$ is proportional to the instantaneous variance of the diffusion component and $v^j(t)$ is proportional to the arrival rate of the jump component. Following Carr and Wu (2003b), we label $v(t)$

as the instantaneous activity rate. Intuitively speaking, t is calendar time and T_t is 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.

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 by calendar time, not on business event time. Furthermore, we apply separate time changes on the diffusion and jump martingale components, which allows potentially different time-variation in the intensities (activity rates) of small and large events.

In this paper, we use “volatility” as a generic term that captures the financial activities of an asset. We do not use it as a statistical term for standard deviation. 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 by using an efficient fast Fourier inversion. Since we model the underlying asset return 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 before the time change. Then we obtain the generalized Fourier transform of the time-changed Lévy process by solving the Laplace transform of the stochastic time under an appropriate measure change.

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

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

Equation (3) decomposes the log return s_t into three components. The first component, $(r - q)t$, is from the instantaneous drift, which is determined by no-arbitrage. The second component, $(\sigma W_t - \frac{1}{2}\sigma^2 t)$, comes from the diffusion, with $\frac{1}{2}\sigma^2 t$ as the concavity adjustment. The last term, $(J_t - \xi t)$, represents the contribution from the jump component, with ξ as the analogous concavity adjustment for J_t . The generalized Fourier transform for s_t under equation (3) is given by

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

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

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

is the characteristic exponent of the diffusion component.

The characteristic exponent of the jump component, ψ_j , depends on the exact specification of the jump structure. Throughout the paper, we use a subscript (or superscript) “ d ” to denote the diffusion component and “ j ” the jump component. As a key feature of Lévy processes, neither ψ_d nor ψ_j depends on the time horizon t .² We 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.³

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

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

where $\psi \equiv [\psi_d, \psi_j]^\top$ denotes the vector of the characteristic exponents and $\mathcal{L}_{T_t}^{\mathbb{M}}(\psi)$ 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,

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

Note that equation (5) converts the issue of obtaining a generalized Fourier transform into a simpler problem of deriving the Laplace transform of the stochastic time (Carr and Wu (2003b)). The solution to this Laplace transform depends on the specification of the instantaneous activity rate $\nu(t)$ and on the characteristic exponents, the functional form of which is determined by the specification of the jump structure J_t .

C. The Jump Structure

Depending on the frequency of jump arrivals, Lévy jump processes can be classified into three categories: finite activity, infinite activity with finite variation, and infinite variation (Sato (1999)). 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, $\pi(dx)$, 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. Thus, the integral of the Lévy measure is finite:

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

Given the finiteness of this integral, 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. Under this process, the integral in equation (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

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

For all finite activity jump models, we can factor the Lévy density 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. One example of this class is the variance-gamma (VG) model of Madan and Milne (1991) and Madan, Carr, and Chang (1998). The VG process is obtained by subordinating an arithmetic Brownian motion with drift $\alpha\lambda$ and variance $\sigma_j^2\lambda$ by an independent gamma process with unit mean rate and variance rate $1/\lambda$. The Lévy measure for the VG process is given by

$$\pi_{VG}(dx) = \lambda \frac{\exp\left(-\frac{|x|}{v_{\pm}}\right)}{|x|} dx,$$

where

$$v_{\pm} = \frac{1}{2} \left(\sqrt{\alpha^2 + 2\sigma_j^2} \pm \alpha \right).$$

The parameter v_+ applies to positive jumps and v_- applies to negative jumps. The jump structure is symmetric around zero when we set $\alpha = 0$. 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, the sample paths of the VG jump process exhibit finite variation:

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

where the function $(1 \wedge |x|)$ represents the minimum of one and $|x|$. Under certain regularity conditions, the Lévy measure of large jumps always performs like a density function. Hence, 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)). There are other commonly used truncation functions for the same purpose. These include $x1_{|x|<1}$, where $1_{|x|<1}$ is an indicator function, and $x/(1+x^2)$. We can use any truncation functions, $h: \mathbb{R}^d \rightarrow \mathbb{R}^d$, which are bounded, with compact support, and satisfy $h(x) = x$ in a neighborhood of zero (Jacod and Shiryaev (1987)).

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 $\alpha \in (1, 2]$.⁴ The Lévy measure under the α -stable motion is given by

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

The process shows finite variation when $\alpha < 1$; but when $\alpha > 1$, the integral in (9) is no longer finite and the process is of infinite variation. Nevertheless, for the Lévy measure to be well-defined, the quadratic variation has to be finite:

$$\int_{\mathbb{R}^0} (1 \wedge x^2)\pi(dx) < \infty, \quad (11)$$

which requires that $\alpha \leq 2$.

The parameter α is often referred to as the tail index. 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 (2003a) for our empirical investigation. In this LS model, c_+ is set to zero in equation (10) so that only negative jumps are allowed. This restriction not only matches the asymmetric feature of the risk-neutral return distribution inferred from S&P 500 index options, but also guarantees the existence of a finite martingale measure, and thus finite option prices. Furthermore, under this model, the return has an α -stable distribution, the variance and higher moments of the asset return are infinite and hence the central limit theorem does not apply. The conditional distribution of the asset return remains non-normal as the conditioning horizon increases. This property helps explain the relatively invariant feature of the implied volatility smirk across different maturities observed for S&P 500 index options. Nevertheless, by setting c_+ to zero, the model guarantees that the conditional moments of the asset price of all positive orders are finite. This delicate balancing not only captures the salient features

of the S&P 500 index options, but also effectively addresses the criticism of Merton (1976) on using α -stable distributions to model asset returns.

As mentioned earlier, 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 we tabulate in Table I. We also include the characteristic exponent for the diffusion component for comparison. Given the Lévy measure π for a particular jump process, we can derive the corresponding characteristic exponents using the Lévy-Khintchine formula (Bertoin (1996)),

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

where b denotes a drift adjustment term.

Insert Table I About Here.

D. The Sources of Stochastic Volatility

The specification of a time-changed Lévy process given in equation (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.

D.1. SVI: 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 comes 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. The relative weight of the jump component declines as the total volatility of the return process increases.

D.2. SV2: Stochastic Volatility from Jump

Alternatively, if we apply a stochastic time change only to the jump component, i.e., $J_t \rightarrow J_{T_t^j}$, but leave the Brownian motion unchanged, stochastic volatility comes solely from the time variation in the arrival rate of jumps. Under this specification, an increase in the return volatility is due 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 (2003) are degenerate examples of this SV2 category because they apply stochastic time changes to pure jump Lévy processes.

D.3. SV3: Joint Contribution from Jump and Diffusion

To model the situation in which 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 and jump component are constant, even though the return volatility varies over time. The recent affine models in Bates (2000) and Pan (2002) are 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.

D.4. SV4: Separate Contribution from Jump and Diffusion

The general specification is to apply separate time changes to the diffusion and jump components 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 separately over time, the relative proportion of each component also varies over

time. The relative dominance of one component over the other depends on 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's specification drives both a compound Poisson jump component and a diffusion component, his model can serve 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 the intensity of the Poisson jump in the model includes both a constant term and a term proportional to the stochastic volatility factor (see also Andersen, Benzoni, and Lund (2002)). We can also see our SV4 specification as a special case of Bates's specification if we set the diffusion component to zero in one factor and the jump component to zero in the other factor. Nevertheless, our separate treatment of the jump component and the diffusion component 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. We can write the generalized Fourier transform of the log return $s_t = \ln(S_t/S_0)$ in equation (5) as

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

The complex-valued exponential martingale in equation (6) that defines the measure change becomes

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

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

Insert Table II About Here.

E. Specification of the Activity Rate Process

We finish our modeling 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} . Thus, we rewrite the Laplace transform as

$$\mathcal{L}_T^{\mathbb{M}}(\psi) = \mathbb{E}^{\mathbb{M}} \left[e^{-\psi^\top T_t} \right] = \mathbb{E}^{\mathbb{M}} \left[e^{-\int_0^t \psi^\top v(s)ds} \right]. \quad (12)$$

If we treat $\psi^\top v(t)$ as an instantaneous interest rate, equation (12) is analogous to the pricing formula for a zero coupon bond. We can then borrow from the term structure of interest rates literature 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. Multivariate activity rate processes can be modeled after affine models of Duffie and Kan (1996) and Duffie, Pan, and Singleton (2000), and the quadratic models 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,

$$dv(t) = \kappa(1 - v(t)) dt + \sigma_v \sqrt{v(t)} dZ_t, \quad (13)$$

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 equation (13) for identification purpose. For the SV4 specification, we assume that the two activity rates, $v(t) = [v^d(t), v^j(t)]^\top$, follow a vector square-root process.

Since the Laplace transform of the time change in equation (12) 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 and the drift function is adjusted to

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

We call special attention to the difference between the drift adjustment for SV2 models and that for SV1 and SV3 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 $\sqrt{v(t)}$. The two-factor SV4 model combines SV1 with SV2.

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

$$\mathcal{L}_T^{\mathbb{M}}(\psi) = \exp(-b(t)v_0 - c(t)), \quad (14)$$

where

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

with

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

For the SV2 specification, the affine structure is retained only when ρ is zero. For tractability, we restrict ρ to zero in our estimation of SV2 models. Under the SV4 model, we assume that the two activity rates are independent of each other and the activity rate for the diffusion component is correlated with the diffusion component in the return process. Then, the Laplace transform under the SV4 model becomes a product of two exponential affine forms, one for the SV1 component and the other for the SV2 component.

Substituting the Laplace transform in equation (14) 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). We label these 12 models as ‘‘JJDSVn,’’ where $JJ \in \{MJ, VG, LS\}$ denotes the jump component, D refers to

the diffusion component, and SV_n , 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 SV_1 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?

III. Data and Estimation

We obtain from a major investment bank in New York daily closing bid and ask implied volatility quotes on the S&P 500 index options across different strikes and maturities from April 6, 1999 to May 31, 2000. The quotes are on standard European options on the S&P 500 spot index, listed at the Chicago Board of 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: that the time to maturity is greater than five business days; that the bid option price is strictly positive; and that 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.

Figure 1 shows in the left panel the histogram of moneyness of the cleaned-up option contracts, where the moneyness is defined as $k \equiv \ln(K/S)$, with K being the strike price. The observations are centered around at the money option contracts ($k = 0$). On average, there are 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 from five business days to over one year and a half, 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 of which fall 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. Thus, the contract stops trading on that expiring Thursday. To avoid potential microstructure effects, we delete from our sample the contracts that are within one week of expiry.

Insert Figure 1 About Here.

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 between -0.3988 and 0.1841 . This restriction excludes approximately 16 percent of the very deep out-of-the-money options (approximately eight percent each for calls and puts), which we deem as too illiquid to contain useful information. Note that we use an asymmetric moneyness range to reflect the fact that there are deeper out-of-the-money put option quotes than out-of-the-money call option quotes. 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 there be 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,

$$\Theta \equiv \arg \min_{\Theta} \sum_{t=1}^T \text{mse}_t, \quad (15)$$

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

$$\text{mse}_t \equiv \min_{v(t)} \frac{1}{N_t} \sum_{i=1}^{n_{t,\tau}} \sum_{j=1}^{n_{t,k}} w_{ij} e_{ij}^2, \quad (16)$$

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 Θ to minimize the sum of the daily mean squared pricing errors.⁵ To construct out-of-sample tests, we divide the data into two subsamples: 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 subsample, we fix the parameter vector Θ estimated from the first sub-sample and compute the daily mean squared pricing errors according to equation (16): At each day, we choose the activity rate levels $v(t)$ to minimize the sum of the weighted squared pricing errors on that day.

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

$$e = \begin{cases} \hat{O}(\Theta) - O_a, & \text{if } \hat{O}(\Theta) > O_a \\ 0, & \text{if } O_a \leq \hat{O}(\Theta) \leq O_b \\ \hat{O}(\Theta) - O_b, & \text{if } \hat{O}(\Theta) < O_b \end{cases} \quad (17)$$

where $\hat{O}(\Theta)$ denotes model-implied out-of-the-money 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. We set the pricing error to zero as long as the model implied price falls within the bid-ask spread of the market quote. We also normalize all prices as percentages of the underlying spot index level.

The construction of the pricing error is a delicate but important issue. For example, the pricing error can be defined with respect to implied volatility, call option price, or put option price. It can be defined as the difference in levels, in log levels, or in percentages. Here, we define the pricing error

using call option prices when $K > F$ and put option prices when $K \leq F$. This definition has become the industry standard for several reasons, one of which is that in-the-money options have positive intrinsic value that is insensitive to model specification, but can still 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 generally more reliable because they are more liquid. We refine the standard definition of the pricing error by incorporating the effects of the bid-ask spreads. Doing so reduces the potential problem of over-fitting and further accounts for the liquidity differences at different moneyness levels and maturities. Dumas, Fleming, Whaley (1998) also incorporate this bid-ask spread effect in their definition of “mean outside error.”

A. The Optimal Weighting Matrix

Like the definition of the pricing error, the construction of a “good” weighting matrix is also important in obtaining robust estimates. 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 important, it puts significantly more weight on long-term options than on short-term options. Thus, performance comparisons can be biased toward models that better capture the behavior of long-term options. Therefore, we want to estimate a weighting matrix that attaches a more balanced weighting to options at all moneyness and maturity levels, and which 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 index level. 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 assumptions that the pricing errors are independently and normally distributed, and that 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 and normally distributed, 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, the minimization problem in equation (15) also generates the maximum likelihood estimates. In principle, we can estimate the variance of the pricing errors via a two-stage procedure analogous to a two-stage least square procedure. However, the weighting obtained from such a procedure depends on 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 because for such processes the conditional return distribution over a fixed horizon does not vary over time. 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. We can then estimate the “true” option price as a percentage of the index level through a sample average, and can consider the daily deviations from such a sample average 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. Therefore, in our case the variance estimate of the option price is only an approximate measure of the variance of the pricing error. 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, to 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. The Appendix contains details for this estimation.

The left panel of Figure 2 shows 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 illustrates 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 do out-of-the-money calls ($k > 0$). This difference might be a reflection of different liquidities: For S&P 500 index options, OTM puts are more liquid and more heavily traded than OTM

calls. Given the estimated variance of the option prices, we define the optimal weight at each moneyness and maturity level as its reciprocal.

Insert Figure 2 About Here.

B. Performance Measures

We compare different models based on the sample properties of the daily mean squared pricing errors (mse_t) defined in equation (16) under the estimated model parameters. A small sample average of the daily mean squared errors for a model would indicate that on average, the model fits the option prices well. 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. We also 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:

$$t\text{-statistic} = \frac{\overline{mse}^i - \overline{mse}^j}{\text{stdev}(\overline{mse}^i - \overline{mse}^j)}, \quad (18)$$

where the overline on mse denotes the sample average and $\text{stdev}(\cdot)$ denotes the standard error of the sample mean difference. We adjust the standard error calculation for serial dependence based on Newey and West (1987), with the number of lags optimally chosen based on Andrews (1991) and an AR(1) specification.

IV. Model Performance Analysis

We analyze the parameter estimates and the sample properties of the mean squared pricing errors for each of the 12 models introduced in Section II. As mentioned earlier, our objective is to investigate

which jump type and which stochastic volatility specification perform the best in pricing S&P 500 index options. Our analysis below focuses on answering these two questions.

Tables III and IV report the parameter estimates and their standard errors for one-factor (SV1 to SV3) and two-factor stochastic volatility (SV4) models, respectively. In the tables we also report the sample average and standard deviation of the daily mean squared pricing errors, both in-sample (mse_I) and out-of-sample (mse_O). Table V reports the t -statistics defined in equation (18) for pair-wise model comparisons. With 12 models, we could have reported a 12×12 matrix of pair-wise t -tests, but to focus on 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), and Panel B compares the performance of different SV specifications for a given jump structure (MJ, VG, or LS). The table reports both in-sample and out-of-sample comparisons.

Insert Table III About Here.

Insert Table IV About Here.

Insert Table V About Here.

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, to answer the question raised on jump types, we compare the performance of the three jump structures under each SV specification. If the performance ranking of the three jump structures depends crucially on the specific SV specification, then the choice of a jump structure in model design should be contingent on which SV specification we use. On the other hand, if the performance rankings are the same under

each of the four SV specifications, then we would conclude that in capturing the behavior of S&P 500 index options, the superiority of one jump structure over the others is unconditional and robust to variations in the SV specification. The empirical evidence matches the latter scenario. The infinite activity jump structures (VG and LS) outperform the traditional 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 (18). Each column in Panel A compares the performance of two jump structures under each SV specification. For example, the column “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. A t -statistics of 1.645 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 percent confidence interval. Therefore, the VG model outperforms the MJ model. A t -value of -1.645 or less implies the opposite.

The in-sample t -values in column “MJ-VG” are strongly positive under all SV specifications, so are all the in-sample t -values in the “MJ-LS” column. The out-of-sample tests reveal similar results, except for the SV4 case. Therefore, 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 result holds under all of the four SV specifications and for both in-sample and most of the out-of-sample tests.

On the other hand, the performance difference between VG and LS is much smaller and can have different signs, depending on the SV specification assumed. The t -values in the “VG-LS” column are much smaller, positive under SV1 and SV2, but negative under SV3 and SV4. Carr and Wu (2003a) 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 . Under the MJ structure, the estimates for λ fall between 0.178 under the SV4 specification (see Table IV) and 0.405 under the SV1

specification (see Table III). An estimate of 0.405 or smaller implies about one jump every two and half years, a rare event. In contrast, under the VG and LS jump structures, the number of jumps within any finite time interval is infinite. Therefore, we can expect 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, we need 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 change to different Lévy components, we can generate stochastic volatility from either the diffusion component, or the jump component, or both. Thus, it becomes a purely empirical issue as to where, exactly, the stochastic volatility comes from. We address this issue by comparing the empirical performance of four different stochastic volatility specifications in pricing the S&P 500 index options.

In Table V, Panel B 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 in the “SV1-SV2” column are all strongly negative and the in-sample t -test values in the “SV2-SV3” column are all strongly positive, which suggests that the SV2 specification is significantly outperformed by the other two one-factor SV specifications. In contrast, the in-sample t -test estimates in 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 gives similar conclusions, except under the LS jump structure where the t -statistics are much smaller.

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. The 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 tell a similar story. The volatility of volatility estimates (σ_v) are always strongly positive under SV1 specifications, slightly smaller under SV3 specifications, but 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 the MJ and LS models. These estimates indicate that 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 (2003), 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 in which instantaneous correlation is not incorporated between the return innovation and the innovation in the activity rate. Therefore, the SV2 specification cannot capture the widely documented negative correlation between stock returns and return volatilities, i.e., the “leverage effect.”⁶ Yet, under all other SV specifications, the estimates for this instantaneous correlation parameter ρ are 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. This negative correlation helps generate negative skewness in the conditional index return distribution implied by the option prices.

Consistent with our observation, Carr, Geman, Madan, and Yor (2003) also note that without the leverage effect, the performance of the SV3 specification declines to approximately the same level as that of 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, which indicates 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 implies that a high-volatility day on the market can come from either the intensified arrival of large events or the increased arrival of small, diffusive events, or both. Therefore, the exact source of high volatility is subject to further research and shall be case dependent. This result contrasts with 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.

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

Since the SV4 specification provides a framework that encompasses 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 in equation (13) for the risk-neutral activity rate dynamics, the two elements of $\sigma_v = [\sigma_v^d, \sigma_v^j]^\top$ 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 shows larger instantaneous volatility than the arrival rate of the jump component. For example, the estimates for σ_v^d are 2.646, 2.555, and 2.856 when the jump components are MJ, VG, and LS, respectively. In contrast, the corresponding estimates for σ_v^j are 2.313, 1.428, and 1.148, evidently lower than the corresponding estimates for σ_v^d .

On the other hand, the relative persistence of the activity rate dynamics is captured by the two elements of $\kappa = [\kappa^d, \kappa^j]^\top$. 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 much more persistent risk-neutral dynamics than the instantaneous variance of the diffusion component. The estimates for κ^j are 0.001, 0.63, and 0.668, when the jump components are MJ, VG, and LS, respectively, much smaller than the corresponding estimates for κ^d , which are 1.872, 1.898, and 2.119, respectively.

The parameter estimates for the SV4 specifications indicate that to match the market price behavior of S&P 500 index options, we need to derive stochastic volatilities from two separate sources, the instantaneous variance of the diffusion component and the arrival rate of the jump component. Fur-

thermore, the risk-neutral dynamics of the diffusion variance exhibits 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 and diffusion components play different roles in governing the behavior of S&P 500 options. The more volatile, but also more transient, feature of the activity rate from the diffusion component implies that the variation of the diffusion component 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 implies 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.

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-Scholes model, we need to add both a jump component and stochastic volatility. This consensus implicitly takes for granted the Brownian motion component in the Black-Scholes model. This view is not surprising since most of the jump models in the literature are variations of Merton's finite activity compound Poisson jump model. In these models, the number of jumps within a finite interval is finite. For example, under the MJDSV1 model, our estimate for the Poisson intensity is 0.405 (Table III), which implies an approximate average of one jump every two and half years. Obviously, we need to add a diffusion component to fill the gaps between the very infrequent jumps in the asset price process.

However, if we consider jump processes with infinite activity, or even infinite variation, the infinite small jumps generated from such models can fill these gaps. Carr, Geman, Madan, and Yor (2002) conclude from their empirical study that a diffusion component is no longer necessary as long as they adopt an infinite activity pure jump process. Carr and Wu (2003a) arrive at similar conclusions in their infinite variation log stable (LS) model. Carr and Wu (2003c) 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 written on this asset. They prove that a jump component, if present, dominates the short-maturity behavior of out-of-the-money options and hence can readily be 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 implies that they generate similar short-maturity behaviors for at-the-money options.

These 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 that also shows infinite variation? These questions are especially relevant here, since 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.

Tables III and IV show that under all the tested models with infinite-activity jump components (VG or LS), the estimates for the diffusion parameter σ are all significantly different from zero. This finding indicates 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 (2003a) is that we incorporate stochastic volatility, but they consider pure Lévy processes without stochastic volatility. Thus, our identification of the diffusion component comes from its role in generating stochastic volatility. The separate specification of the two activity rate processes under SV4 implies that the relative proportion of small (diffusive) movements and large (jump) movements can vary over time. Their different risk-neutral dynamics further imply 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 must be incorporated by using a diffusion component in the return process, because 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. We can conceivably incorporate a jump component in the activity rate process, as in Chernov, Gallant, Ghysels, and Tauchen (1999) and Eraker,

Johannes, and Polson (2003). Doing so would allow us to 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 be reduced.

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 we calibrate all models 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 shown in Table IV, the estimate of σ (the diffusion parameter) is 0.312 for MJDSV4, 0.310 for VGDSV4, but 0.308 for LSDSV4. We observe similar declines 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). However, the most dramatic decline 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. SV2 differs from all other SV specifications in generating stochastic volatility from the jump component only and by not accommodating a leverage effect. Thus, consistent with our discussion above, without a role in either generating stochastic volatility or accommodating a leverage effect, the diffusion component is hardly needed when the jump component also shows infinite variation, as in the case of LSDSV2.

Combining all the evidence, we conclude that as the frequency of jump arrival increases from MJ to VG and then to LS, the need for a diffusion 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 providing a separate source of stochastic volatility and 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.

V. 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, we 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 show no obvious structures along both the moneyness and the maturity dimensions.

Since an option's time-to-maturity and moneyness change every day, we estimate the pricing error at fixed moneyness and maturity by using nonparametric smoothing (See the Appendix). It follows from the pricing error definition in equation (17) that a positive pricing error indicates model overpricing compared to market data and a negative pricing error indicates model 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 show 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 which corresponds to a particular model specification. The four rows, 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).

Insert Figure 3 About Here.

For comparison, we use the same scale for all panels except for the second row, where we use a larger scale to accommodate the larger pricing errors from the SV2 models. Figure 3 shows that the three SV2 models exhibit large mean pricing errors along both the maturity and the moneyness dimensions. 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. We 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. The two-factor SV4 specifications show much better performance, except at very long maturities (one-year, dotted line), where we still observe some moneyness bias.

In the option pricing literature, it is 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-Scholes implied volatilities for out-of-the-money puts are much higher than those for equally out-of-the-money calls. This volatility smirk is a direct result of conditional non-normality in asset returns. The downward slope of the smirk reflects asymmetry (negative skewness) in the risk-neutral distribution. The positive curvature of the smirk reflects the fat-tails (leptokurtosis) of this distribution. Furthermore, the central limit theorem implies that under very general conditions, the conditional return distribution should converge to normality as the maturity increases. As a result, the volatility smirk should flatten out. However, Carr and Wu (2003a) find that the implied volatility smirk in fact steepens slightly as maturity increases. 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. The first approach incorporates a persistent stochastic volatility process to slow down the convergence to normality. The second approach, proposed by Carr and Wu (2003a), adopts an α -stable process 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. The SV2 models imply that as maturity increases, volatility smirk flattens out faster than observed in the data. Thus, the remaining structure in the mean pricing error for SV2 models indicates that the SV2 specification fails to account for the maturity pattern of the volatility smirk for S&P 500 index options. This failure implies that not just any persistent stochastic volatility model will work. The better performance of other SV specifications also suggests that for successful model design, we need also to 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 look promising in generating a persistent volatility smirk across the maturity horizon. The better performance of the LSDSV4 model can be traced 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 (2003a) 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. 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 also improves the performance by generating variations in the relative proportion of the jump and diffusion components along the option maturity dimension. Since the estimated risk-neutral dynamics for the jump-induced return volatility is more persistent than the estimated dynamics for the diffusion-induced volatility ($\kappa^j = 0.668$ compared to $\kappa^d = 2.119$ under LSDSV4, see Table IV), the impact of the more persistent jump component dominates the behavior of long-term options and 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 the central limit theorem and helps maintain a relatively stable and slightly steepening maturity pattern for the implied volatility smirk.

VI. Conclusion

Within the general framework of time-changed Lévy processes, we analyze the specification of option pricing models based on three separate but interrelated dimensions: the specification of the jump component in the return process, the source for stochastic volatility, and 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 several model design issues that have not been answered in the literature. We focus on answering two questions: What type of jump structure best captures the behavior of the S&P 500 index options? Where does stochastic volatility of the S&P 500 index return come from?

We find that a high-frequency jump structure always outperforms the traditional low-frequency compound Poisson jump specification. 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 occurrences rather than rare events. We also find that stochastic volatility comes from two separate sources, the instan-

taneous variance of a diffusion component and the arrival rate of a jump component. The risk-neutral dynamics of the diffusion variance is more volatile, but the risk-neutral dynamics of the arrival rate of the jump component shows much more persistence. As a result, the stochastic volatility from the diffusion component dominates the behavior of short-term option prices, and the stochastic volatility from jumps dominates that of long-term option prices. In short, our empirical results for the S&P 500 index options demonstrate that a good model for the underlying index return dynamics should include a high-frequency jump component and should allow stochastic volatility to be generated separately from the diffusion component and the jump component.

To maintain the parsimony of the model, 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 inquiry for 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 both jumps into the activity rate process and jumps in the asset return processes. Nevertheless, all these studies consider only compound Poisson jumps with potentially time-varying arrival rates. In light of our findings in this paper on the superior performance of high frequency jumps, it would be intriguing to see whether incorporating such high-frequency jumps in the volatility process will also deliver superior performance over specifications of Poisson jumps in volatility used in the current literature.

Since the objective of our 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. A line for future research is to 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 (2003) and Pan (2002).

Appendix. 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

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

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 Ait-Sahalia and Lo (1998) in choosing independent Gaussian kernels and setting the bandwidths according to

$$h_j = c_j \sigma_j \mathcal{N}^{-1/6}, \quad j = \tau, k, \quad (\text{A2})$$

where σ_j is the standard deviation of the regressor Z_j , and c_j is a constant that is typically of order of magnitude one. The larger the coefficient c_j , 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.
- Andrews, Donald, 1991, Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* 59, 817–858.
- Ané, Thierry, 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).
- 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, 2002, The fine structure of asset returns: An empirical investigation, *Journal of Business* 75, 305–332.
- Carr, Peter, Hélyette Geman, Dilip Madan, and Marc Yor, 2003, Stochastic volatility for Lévy processes, *Mathematical Finance* 13, 345–382.
- Carr, Peter, and Dilip Madan, 1999, Option valuation using the fast Fourier transform, *Journal of Computational Finance* 2, 61–73.
- Carr, Peter, and Liuren Wu, 2003a, Finite moment log stable process and option pricing, *Journal of Finance* 58, 753–777.
- Carr, Peter, and Liuren Wu, 2003b, Time-changed Lévy processes and option pricing, *Journal of Financial Economics*, forthcoming.
- Carr, Peter, and Liuren Wu, 2003c, What type of process underlies options? A simple robust test, *Journal of Finance* 58, 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, Pennsylvania 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 surface, *Quantitative Finance* 2, 361–377.
- 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 Quantitative Analysis* 34, 211–240.
- 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, Bjørn, 2003, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *Journal of Finance*, forthcoming.
- Eraker, Bjørn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in equity index volatility and returns, *Journal of Finance* 58, 1269–1300.
- Foresi, Silverio, and Liuren Wu, 2003, Crash-o-phobia: A domestic fear or a worldwide concern? Manuscript, Fordham University.
- 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 Frank 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.
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- 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, Edward Charles, 1986, *Introduction to the Theory of Fourier Integrals*, 3rd ed. (Chelsea Pub. Co., New York).

NOTES

¹See, for example, Aït-Sahalia and Lo (1998), Jackwerth and Rubinstein (1996), and Rubinstein (1994) for empirical documentation of this phenomenon in S&P 500 index options.

²See Bertoin (1996) and Sato (1999).

³See Titchmarsh (1986) for a comprehensive reference on generalized Fourier transforms.

⁴See Samorodnitsky and Taqqu (1994) and Janicki and Weron (1994).

⁵We thank an anonymous referee for suggesting this estimation procedure.

⁶Black (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).

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

Component	$\psi(u)$
Diffusion	$\frac{1}{2}\sigma^2 [iu + u^2]$
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)	$\lambda \left[-iu \ln \left(1 - \alpha - \frac{1}{2}\sigma_j^2 \right) + \ln \left(1 - iu\alpha + \frac{1}{2}\sigma_j^2 u^2 \right) \right]$
Log Stable (LS)	$\lambda (iu - (iu)^\alpha)$

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

The term x_t denotes the time changed component and y_t denotes the unchanged component in the log return $s_t = \ln(S_t/S_0)$. The term 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	$[\sigma W_t - \frac{1}{2}\sigma^2 t, J_t - \xi t]^\top$	0	$e^{iu(r-q)t} \mathcal{L}_T^{\mathbb{M}}([\Psi_d, \Psi_j]^\top)$

Table III
Parameter Estimates of One-Factor SV Models

We estimate model parameters by minimizing the sum of daily mean squared errors. Given model parameters, we obtain the daily mean squared errors 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. We also report the sample average and standard deviation of the daily mean squared error for both the in-sample period (mse_I , the first 139 days) and the out-of-sample period (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 (0.288)	0.173 (0.010)	0.301 (0.031)	0.318 (0.037)	0.157 (0.006)	0.272 (0.030)	0.309 (0.037)	0.044 (0.001)	0.175 (0.029)
λ	0.405 (0.459)	0.364 (0.135)	0.223 (0.190)	0.253 (0.194)	0.593 (0.068)	0.985 (1.115)	0.028 (0.003)	0.077 (0.007)	0.053 (0.007)
α	-0.091 (0.052)	-0.393 (0.100)	-0.408 (0.149)	-0.247 (0.119)	-0.391 (0.029)	-0.244 (0.122)	1.673 (0.044)	1.578 (0.028)	1.738 (0.066)
σ_j	0.175 (0.113)	0.235 (0.035)	0.000 (0.000)	0.264 (0.108)	0.013 (0.001)	0.003 (0.001)	---	---	---
κ	1.039 (0.320)	2.070 (0.000)	1.110 (0.346)	0.813 (0.322)	2.054 (0.237)	0.974 (0.327)	0.795 (0.293)	0.867 (0.000)	1.304 (0.324)
σ_v	2.574 (0.620)	0.000 (0.000)	1.983 (0.361)	2.136 (0.431)	0.001 (0.000)	1.745 (0.321)	2.253 (0.358)	0.000 (0.000)	1.839 (0.344)
ρ	-0.704 (0.073)	---	-0.648 (0.092)	-0.692 (0.075)	---	-0.662 (0.101)	-1.000 (0.000)	---	-1.000 (0.000)
mse_I	0.334 (0.254)	1.159 (0.314)	0.307 (0.247)	0.302 (0.244)	0.927 (0.272)	0.279 (0.236)	0.256 (0.218)	0.859 (0.262)	0.279 (0.237)
mse_O	2.105 (1.123)	2.599 (1.127)	1.752 (0.968)	1.868 (0.982)	2.339 (1.046)	1.531 (0.874)	1.813 (1.030)	1.610 (0.734)	1.739 (0.968)

Table IV
Parameter Estimates of SV4 Models

We estimate model parameters by minimizing the sum of daily mean squared errors. Given model parameters, we obtain the daily mean squared errors 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. We also report the sample average and standard deviation of the daily mean squared error for both the in-sample period (mse_I , the first 139 days) and the out-of-sample period (mse_O , the last 150 days). The pricing error is defined in percentages of the spot price.

Θ	MJDSV4		VGDSV4		LSDSV4	
σ	0.312	(0.020)	0.310	(0.010)	0.308	(0.0099)
λ	0.178	(3.450)	0.001	(0.000)	0.001	(0.0001)
α	-0.486	(0.513)	-0.534	(0.076)	1.296	(0.0739)
σ_j	0.281	(0.439)	0.440	(0.025)	---	(0.0000)
κ^d	1.872	(0.498)	1.898	(0.254)	2.119	(0.3884)
κ^j	0.001	(0.014)	0.630	(0.102)	0.668	(0.3262)
σ_v^d	2.646	(0.358)	2.555	(0.201)	2.856	(0.3466)
σ_v^j	2.313	(22.018)	1.428	(0.160)	1.148	(0.6824)
ρ^d	-0.701	(0.063)	-0.730	(0.036)	-0.719	(0.0421)
mse_I	0.140	(0.151)	0.128	(0.150)	0.128	(0.1491)
mse_O	0.705	(0.647)	0.677	(0.551)	0.743	(0.5868)

Table V
Pair-Wise t -statistics for Model Comparisons

Entries report the t -statistics defined in equation (18). 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 1.645 implies that the mean squared pricing error from model i is significantly larger than the mean squared error from model j . Hence, model j outperforms model i , at 95 percent confidence interval. A t -value less than -1.645 implies the opposite. We base in-sample tests on the first 139 days of option price data and out-of-sample tests on the last 150 days of data, given parameter estimates from the first subsample.

<i>Panel A. Testing Which Jump Structure Performs the Best</i>						
Cases/Tests	<i>MJ – VG</i>	<i>MJ – LS</i>	<i>VG – LS</i>	<i>MJ – VG</i>	<i>MJ – LS</i>	<i>VG – LS</i>
	In-sample comparison			Out-of-sample comparison		
SV1	2.57	6.24	3.98	5.73	7.75	2.07
SV2	15.51	7.27	1.98	7.37	6.95	6.09
SV3	4.28	5.91	–0.00	6.48	2.70	–6.42
SV4	4.65	4.26	0.09	0.67	–0.93	–5.87
<i>Panel B. Testing Where Stochastic Volatility Comes From</i>						
Cases/Tests	<i>SV1 – SV2</i>	<i>SV1 – SV3</i>	<i>SV2 – SV3</i>	<i>SV4 – SV1</i>	<i>SV4 – SV2</i>	<i>SV4 – SV3</i>
	In-sample comparison					
MJ	–15.03	3.48	15.01	–4.42	–25.48	–4.05
VG	–13.52	2.88	15.01	–4.44	–24.70	–4.38
LS	–17.64	–3.55	14.54	–4.39	–25.44	–4.23
	Out-of-sample comparison					
MJ	–7.07	6.24	11.95	–8.12	–14.56	–7.58
VG	–8.31	6.84	13.23	–6.66	–11.38	–5.06
LS	1.50	1.87	–1.17	–7.42	–5.65	–5.43




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 6, 1999 to May 31, 2000, comprising 290 business days and 62,950 option contracts.




Figure 2. Mean and variance surface of option prices. Using independent Gaussian kernels, we estimate nonparametrically 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). The sample of S&P 500 index options is daily from April 6, 1999 to May 31, 2000, comprising 290 business days and 62,950 option quotes.

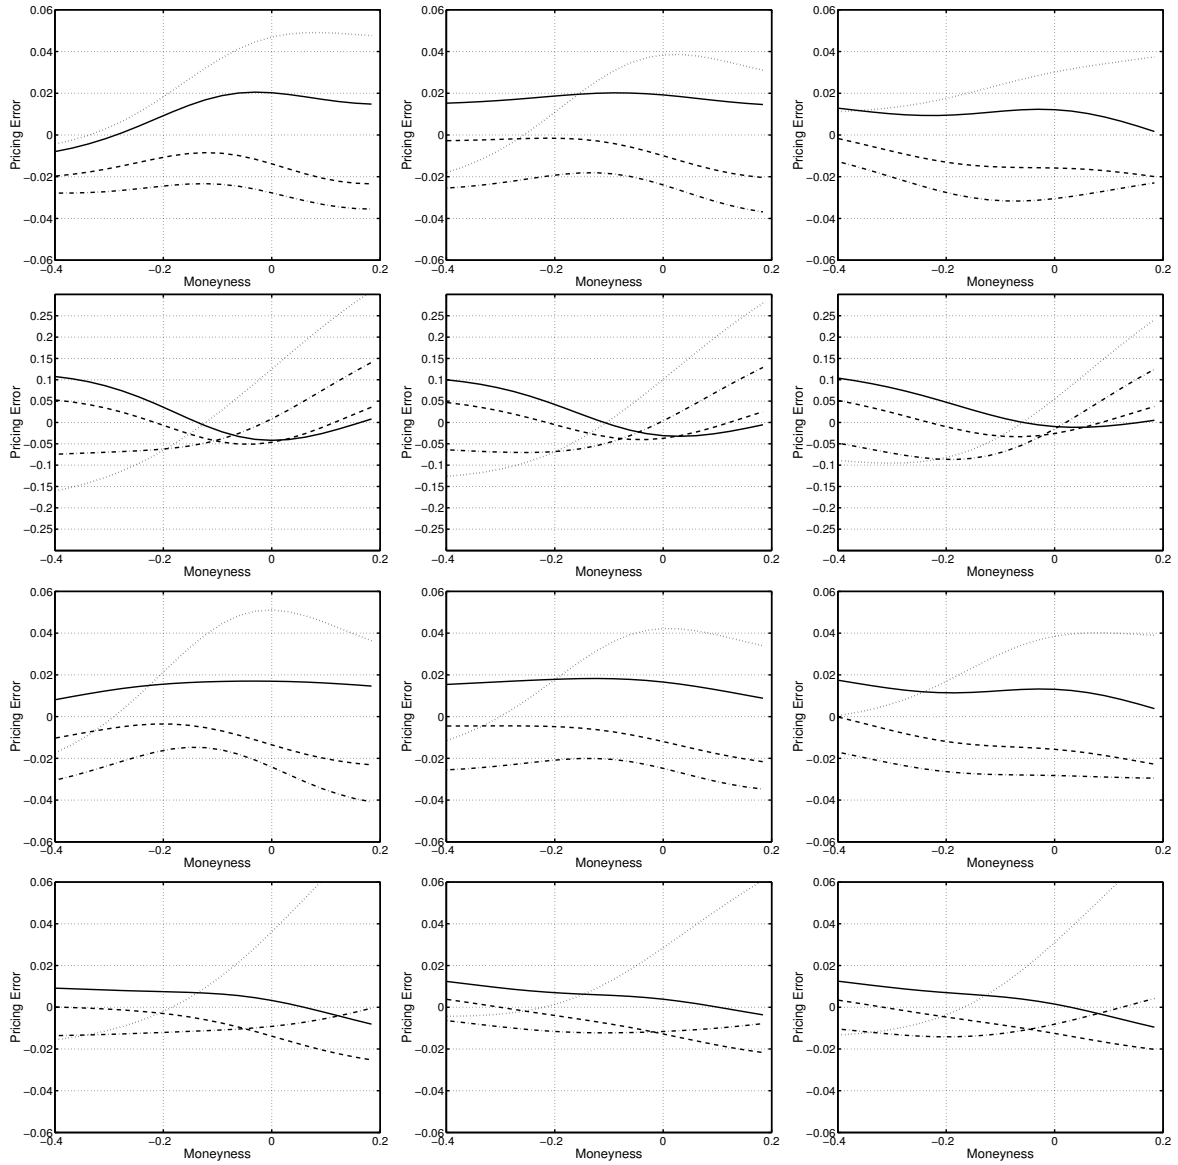


Figure 3. In-sample mean pricing errors. We define pricing errors as the difference between the model-implied option price and the market-observed price, outside the bid-ask spread and in percentages of the underlying spot level. Using independent Gaussian kernels, we estimate the mean pricing errors at fixed moneyness ($k = \ln(K/S)$) and maturities. 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 we use a much larger scale to accommodate the much larger pricing errors of SV2 models.