



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfecU.S. stock market crash risk, 1926–2010[☆]David S. Bates^{*}

University of Iowa and the National Bureau of Economic Research, Henry B. Tippie College of Business, Iowa City, IA 52242-1000, USA

ARTICLE INFO

Article history:

Received 27 January 2010

Received in revised form

31 March 2011

Accepted 25 April 2011

JEL classification:

C22

C46

G01

G13

Keywords:

Lévy processes

Time-changed Lévy processes

Stock market crashes

Option pricing

ABSTRACT

This paper examines how well alternate time-changed Lévy processes capture stochastic volatility and the substantial outliers observed in U.S. stock market returns over the past 85 years. The autocorrelation of daily stock market returns varies substantially over time, necessitating an additional state variable when analyzing historical data. I estimate various one- and two-factor stochastic volatility/Lévy models with time-varying autocorrelation via extensions of the Bates (2006) methodology that provide filtered daily estimates of volatility and autocorrelation. The paper explores option pricing implications, including for the Volatility Index (VIX) during the recent financial crisis.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

What is the risk of stock market crashes? Answering this question is complicated by two features of stock market returns: the fact that conditional volatility evolves over time, and the fat-tailed nature of daily stock market

returns. Each issue affects the other. Which returns are identified as outliers depends upon that day's assessment of conditional volatility. Conversely, estimates of current volatility from past returns can be disproportionately affected by outliers such as the 1987 crash. In standard generalized autoregressive conditional heteroskedasticity (GARCH) specifications, for instance, a 10% daily change in the stock market has one hundred times the impact on conditional variance revisions of a more typical 1% move.

This paper explores whether recently proposed continuous-time specifications of time-changed Lévy processes are a useful way to capture the twin properties of stochastic volatility and fat tails. The use of Lévy processes to capture outliers dates back at least to the Mandelbrot (1963) use of the stable Paretian distribution, and many specifications have been proposed, including the Merton (1976) jump-diffusion, the Madan and Seneta (1990) variance gamma, the Eberlein, Keller, and Prause (1998) hyperbolic Lévy, and the Carr, German, Madan, and Yor (2002) CGMY process. As all of these distributions assume identically and independently distributed (i.i.d.)

[☆] I am grateful for comments on earlier versions of the paper from the referee, from seminar participants at the universities of Iowa, Turin, London (Birkbeck), Colorado, Calgary, and Massachusetts-Amherst, and from conference participants at McGill University's 2008 Risk Management Conference, the 18th Annual Derivatives Securities and Risk Management Conference, Princeton University's 2008 Conference on Implied Volatility Models, the 2009 Western Finance Association meetings, the 2009 Conference on Liquidity, Credit Risk, and Extreme Events at the University of Chicago, the 2010 Derivatives, Volatility, and Correlation Conference at Warwick University, and the 2010 Financial Econometrics Conference at the University of Toulouse. I also thank Ken French for providing the stock turnover data from French (2008).

^{*} Tel.: +1 319 353 2288; fax: +1 319 335 3690.

E-mail address: david-bates@uiowa.edu

URL: <http://www.biz.uiowa.edu/faculty/dbates>

returns, however, they are unable to capture stochastic volatility.

More recently, Carr, German, Madan, and Yor (2003) and Carr and Wu (2004) have proposed combining Lévy processes with a subordinated time process. The idea of randomizing time dates back at least to Clark (1973). Its appeal in conjunction with Lévy processes reflects the increasing focus in finance – especially in option pricing – on representing probability distributions by their associated characteristic functions. Lévy processes have log characteristic functions that are linear in time. If the time randomization depends on underlying variables that have an analytic conditional characteristic function, then the resulting conditional characteristic function of time-changed Lévy processes is also analytic. Conditional probability densities, distributions, and option prices can then be numerically computed by Fourier inversion of simple functional transforms of this characteristic function.

Thus far, empirical research on the relevance of time-changed Lévy processes for stock market returns has largely been limited to the special cases of time-changed versions of Brownian motion and the Merton (1976) jump-diffusion. Furthermore, there has been virtually no estimation of newly proposed time-changed Lévy processes solely from time series data.¹ Papers such as Carr, German, Madan, and Yor (2003) and Carr and Wu (2004) rely on option pricing evidence to provide empirical support for their approach, instead of providing direct time series evidence. The reliance on options data is understandable. Because the state variables driving the time randomization are not directly observable, time-changed Lévy processes are hidden Markov models, creating a challenging problem in time series econometrics. Using option prices potentially identifies realizations of those latent state variables, converting the estimation problem into the substantially more tractable problem of estimating state space models with observable state variables.

While options-influenced parameter and state variable estimates should be informative under the hypothesis of correct model specification, the objective of this paper is to provide estimates of crash risk based solely upon time series analysis. Such estimates are of interest in their own right, and are useful for testing the central empirical hypothesis in option pricing: whether option prices are, in fact, compatible with the underlying time series properties of the underlying asset, after appropriate risk adjustments. Testing the compatibility hypothesis is more difficult under joint options/time series estimation approaches that are premised upon compatibility. Furthermore, option-based and joint estimation approaches are constrained by the availability of options data only since the 1980s, whereas time series estimation can exploit a longer history of extreme stock market movements.² For instance, it has been asserted

that deep out-of-the-money index put options appear overpriced, based on their surprisingly large negative returns since the 1987 crash. But all such tests require reliable estimates of downside risk; and it can be difficult to establish whether puts are overpriced based only on post-1987 data.³

Risk-adjusted time series estimates of conditional distributions can also provide useful real-time valuations of option prices, for comparison with observed option prices. At the end of the paper I compare the options-based Volatility Index (VIX) measure of volatility with time series estimates, during a 2007–2010 period spanning the recent financial crisis.

This paper uses the Bates (2006) approximate maximum likelihood (AML) methodology for estimation of various time-changed Lévy processes over 1926–2006, and for out-of-sample fits over 2007–2010. AML is a filtration methodology that recursively updates conditional characteristic functions of latent variables over time given observed data. Filtered estimates of the latent variables are directly provided as a by-product, given the close link between moments and characteristic functions. The methodology's focus on characteristic functions makes it especially useful for estimating Lévy processes, which typically lack closed-form probability density functions. The paper primarily focuses on the time-changed CGMY process, which nests other Lévy processes as special cases. The approach is also compared with the stochastic volatility processes with and without normally distributed jumps previously estimated in Bates (2006).

A concern with any extended data set is the possibility that the data generating process might not be stable over time. Indeed, this paper identifies substantial instability in the autocorrelation of daily stock market returns. Autocorrelation estimates appear to be nonstationary, and peaked at the extraordinarily high level of 35% in 1971 before trending downward to the near-zero values observed since the 1980s. The instability is addressed directly, by treating autocorrelation as another latent state variable to be estimated from observed stock market returns. The paper also uses subsample estimation to test for (and find) apparent instabilities or specification issues in the one-factor volatility process used. Given these issues, I estimate a two-factor concatenated model of volatility evolution, which can be interpreted as a model of parameter drift in the unconditional mean of the one-factor variance process. Finally, I examine the sensitivity of volatility filtration and option prices to the use of different data sets and volatility models.

Overall, the time-changed CGMY process is found to be a slightly more parsimonious alternative to the Bates (2006) approach of using finite-activity stochastic-intensity jumps

(footnote continued)

movement over 1945–2010 in the Center for Research in Security Prices (CRSP) value-weighted index to exceed 10% in magnitude, whereas there were seven such movements over 1929–1932.

³ See Broadie, Chernov, and Johannes (2009) for a Monte Carlo study of unhedged one-month returns for puts on S&P 500 futures over August 1987 to June 2005. They find that their large excess return estimates often lack statistical significance, especially when volatility is stochastic.

¹ Li, Wells, and Yu (2008) use Markov chain Monte Carlo (MCMC) methods to estimate some models in which Lévy shocks are added to various stochastic volatility models. However, the additional Lévy shocks are independently and identically distributed, not time-changed.

² The –20% and +11½ movements on October 19, 1987 and October 13, 2008, respectively, were the only daily stock market

drawn from a mixture of normals, although the fits of the two approaches are very similar. Interestingly, one cannot reject the hypothesis that stock market crash risk is adequately captured by a time-changed version of the Carr and Wu (2003) log-stable process. That model's implications for upside risk, however, are strongly rejected, with the model severely underpredicting the frequency of large positive outliers.

Section 2 progressively builds up the time series model used in estimation. Section 2.1 discusses basic Lévy processes and describes the processes considered. Section 2.2 discusses time changes, the equivalence to stochastic volatility, and the leverage effect. Section 2.3 contains further modifications of the model to capture time-varying autocorrelations and day-of-the-week effects. Section 2.4 describes how the model is estimated, using the Bates (2006) AML estimation methodology for hidden Markov models.

Section 3 describes the data on excess stock market returns over 1926–2010 and presents parameter estimates, diagnostics, and filtered estimates of latent autocorrelation and volatility. Given results from the diagnostics, I develop and estimate a two-factor variance model in Section 3.7. Section 4 examines option pricing implications, and Section 5 concludes.

2. Time-changed Lévy processes

2.1. Lévy processes

A Lévy process $L(t)$ is an infinitely divisible stochastic process; i.e., one that has i.i.d. increments over nonoverlapping time intervals of equal length. The Lévy processes most commonly used in finance have been Brownian motion and the jump-diffusion process of Merton (1976), but there are many others. All Lévy processes without a Brownian motion component are pure jump processes. Such processes are characterized by their Lévy density $k(x)$, which gives the intensity (or frequency) of jumps of size x . Alternatively and equivalently, Lévy processes can be described by their generalized Fourier transform

$$F(u) \equiv Ee^{uL(t)} = \exp[tf_{dL}(u)], \quad u \in D_u \subset \mathbb{C}, \quad (1)$$

where u is a complex-valued element of the set D_u for which Eq. (1) is well-defined. If Φ is real, $F(i\Phi)$ is the characteristic function of $L(t)$, while $tf_{dL}(\Phi)$ is the cumulant generating function (CGF) of $L(t)$. Following Wu (2006), the function $f_{dL}(\Phi)$ is called the cumulant exponent of $L(t)$.⁴

The Lévy-Khintchine formula gives the mapping between jump intensities $k(x)$ and the cumulant exponent for arbitrary $u \in D_u$. Lévy processes in finance are typically specified for the log asset price and then exponentiated: $S(t) = \exp[L(t)]$. For such specifications, it is convenient to write the Lévy-Khintchine formula for pure jump processes

in the form

$$f_{dL}(u) = u\mu + \int_{\mathbb{R} - \{0\}} [e^{ux} - 1 - u(e^x - 1)]k(x)dx, \quad (2)$$

where $\mu = f_{dL}(1)$ is the continuously compounded expected return on the asset:

$$ES(t) = Ee^{L(t)} = e^{f_{dL}(1)t} = e^{\mu t}. \quad (3)$$

Pure-jump Lévy processes can be thought of as a drift term plus an infinite sum $L(t) = \int L_x dx$ of independent point processes, each drift-adjusted to make $\exp[L_x(t)]$ a martingale:

$$dL_x = x dN_x - (e^x - 1)k(x)dt, \quad (4)$$

where N_x is an integer-valued Poisson counter with intensity $k(x)$ that counts the occurrence of jumps of fixed size x . The log characteristic function of a sum of independent point processes is the sum of the log characteristic functions of the point processes, yielding Eq. (2). Exponential martingale processes of the form $L(t) = \int L_x dx$ for L_x defined in Eq. (4) are called compensated Lévy processes, as also are diffusions of the form $\sigma W_t - \frac{1}{2}\sigma^2 t$.

As discussed in Carr, German, Madan, and Yor (2002), Lévy processes are finite-activity if $\int k(x)dx < \infty$ and infinite-activity otherwise. Finite-activity jumps imply there is a nonzero probability that no jumps will be observed within a time interval. Lévy processes are finite-variation if $\int |x|k(x)dx < \infty$ and infinite-variation otherwise. An infinite-variation process has sample paths of infinite length, which is also a property of Brownian motion. All Lévy processes must have finite $\int \min(x^2, 1)k(x)dx$ to be well behaved but need not have finite variance $\int x^2k(x)dx$, the stable distribution being a counterexample. A priori, all financial prices must be finite-activity processes, because price changes reflect a finite (but large) number of market transactions. However, finite-activity processes can be well approximated by infinite-activity processes, and vice versa; e.g., the Cox, Ross, and Rubinstein (1979) finite-activity binomial approximation to Brownian motion. Activity and variation therefore are treated as empirical specification issues concerned with identifying which functional form $k(x)$ for jump intensities best fits daily stock market excess returns.

I consider two particular underlying Lévy processes for log asset prices. The first is the Merton (1976) combination of Brownian motion and finite-activity normally distributed jumps:

$$d \ln S_t = \mu dt + (\sigma dW_t - \frac{1}{2}\sigma^2 dt) + (\gamma dN_t - \lambda \bar{k} dt), \quad (5)$$

where W_t is a Wiener process, N_t is a Poisson counter with intensity λ , $\gamma \sim N(\bar{\gamma}, \delta^2)$ is the normally distributed jump conditional upon a jump occurring, and $\bar{k} = e^{\bar{\gamma} + \frac{1}{2}\delta^2} - 1$ is the expected percentage jump size conditional upon a jump. The associated intensity of jumps of size x is

$$k(x) = \frac{\lambda}{\sqrt{2\pi\delta^2}} \exp\left[-\frac{(x - \bar{\gamma})^2}{2\delta^2}\right], \quad (6)$$

and the cumulant exponent takes the form

$$f_{Merton}(u) = \mu u + \frac{1}{2}\sigma^2(u^2 - u) + \lambda(e^{\bar{\gamma}u + \frac{1}{2}\delta^2 u^2} - 1 - u\bar{k}). \quad (7)$$

⁴ Carr, German, Madan, and Yor (2003) call $\psi(\Phi) = f_{dL}(i\Phi)$ the “unit time log characteristic function.” Bertoin (1996) calls $-\psi(\Phi)$ the “characteristic exponent.”

The approach can be generalized to allow alternate distributions for γ —in particular, a mixture of normals:

$$k(x) = \sum_{i=1}^2 \frac{\lambda_i}{\sqrt{2\pi\delta_i^2}} \exp\left[-\frac{(x-\bar{\gamma}_i)^2}{2\delta_i^2}\right]. \tag{8}$$

Second, I consider the generalized CGMY process of Carr, Geman, Madan, and Yor (2003), which has a Lévy density of the form

$$k(x) = \begin{cases} C_n e^{-G|x}|x|^{-1-Y_n} & \text{for } x < 0, \\ C_p e^{-M|x}|x|^{-1-Y_p} & \text{for } x > 0, \end{cases} \tag{9}$$

where $C_n, C_p, G, M \geq 0$ and $Y_p, Y_n < 2$. The associated cumulant exponent is

$$f_{CGMY}(u) = (\mu - \omega)u + V \left[w_n \frac{(G+u)^{Y_n} - G^{Y_n}}{Y_n(Y_n-1)G^{Y_n-2}} + (1-w_n) \frac{(M-u)^{Y_p} - M^{Y_p}}{Y_p(Y_p-1)M^{Y_p-2}} \right], \tag{10}$$

where ω is a mean-normalizing constant determined by $f_{CGMY}(1) = \mu$, V is variance per unit time, and w_n is the fraction of variance attributable to the downward-jump component. The corresponding intensity parameters C_n, C_p in Eq. (9) are

$$C_n = \frac{w_n V}{\Gamma(2-Y_n)G^{Y_n-2}}, \quad C_p = \frac{(1-w_n)V}{\Gamma(2-Y_p)M^{Y_p-2}}, \tag{11}$$

where $\Gamma(z)$ is the gamma function.

As discussed in Carr, Geman, Madan, and Yor (2002), the Y parameters are key in controlling jump activity near zero, in addition to their influence over tail events. The process has finite activity if $Y_p, Y_n < 0$, finite variation if $Y_p, Y_n < 1$, but infinite activity or variation if $\min(Y_p, Y_n)$ is greater or equal to zero or one, respectively. The model conveniently nests many models considered elsewhere.

For instance, $Y_n = Y_p = -1$ is the finite-activity double exponential jump model of Kou (2002), while $Y_n = Y_p = 0$ includes the variance gamma model of Madan and Seneta (1990). As Y_p and Y_n approach two, the CGMY process for fixed variance V converges to a diffusion, and the cumulant exponent converges to the corresponding quadratic form

$$f_{SV}(u) = \mu u + \frac{1}{2}V(u^2 - u). \tag{12}$$

As G and M approach zero for arbitrary Y_p, Y_n and fixed C_n, C_p , the Lévy density in Eq. (9) approaches the infinite-variance log stable process advocated by Mandelbrot (1963), with a power law property for asymptotic tail probabilities. The log-stable special case proposed by Carr and Wu (2003) is the limiting case with only negative jumps ($C_p = 0$). While infinite-variance for log returns, percentage returns have finite mean and variance under this specification. For daily stock market returns of less than 25% in magnitude, Carr and Wu's log-stable process is well approximated by a finite-variance CGMY process with minimal exponential dampening; e.g., $G = 0.001$.

The cumulant exponent of any finite-variance Lévy process can be written in the form

$$f_{dL}(u) = V g_{dL}(u), \tag{13}$$

where $V = f''_{dL}(0)$ is variance per unit time and $g_{dL}(u)$ is a standardized cumulant exponent with unitary variance. One can also combine Lévy processes, to nest alternative specifications within a broader specification. Any linear combination $w_1 k_1(x) + w_2 k_2(x)$ of Lévy densities for nonnegative weights that sum to one is also a valid Lévy density and generates an associated standardized weighted cumulant exponent of the form $w_1 g_1(u) + w_2 g_2(u)$, where $g_i(u)$ is the standardized cumulant exponent associated with $k_i(x)$ for $i = 1, 2$. The various $g_{dL}(u)$ specifications considered in this paper are listed in Table 1.

Table 1

Standardized cumulant exponents $g_{dL}(u) = \ln(Ee^{u dL})/dt$ for various compensated Lévy specifications that have unitary variance V per unit time and a continuously compounded expected return μ equal to zero.

Diffusion [Eq. (12)]

$$g_{SV}(u) = \frac{1}{2}(u^2 - u)$$

Normally distributed jumps

$$g_{ji}(u) = \frac{1}{\bar{\gamma}_i^2 + \delta_i^2} [e^{u\bar{\gamma}_i + \frac{1}{2}u^2\delta_i^2} - 1 - u(e^{\bar{\gamma}_i + \frac{1}{2}\delta_i^2} - 1)]$$

CGMY jumps [Eq. (10)], with ω such that $g_{CGMY}(1) = 0$

$$g_{CGMY}(u) = w_n \frac{(G-u)^{Y_n} - G^{Y_n}}{Y_n(Y_n-1)G^{Y_n-2}} + (1-w_n) \frac{(M+u)^{Y_p} - M^{Y_p}}{Y_p(Y_p-1)M^{Y_p-2}} - \omega u$$

General specification

$$g_{dL}(u) = w_{SV} g_{SV}(u) + \sum_{i=1}^2 w_{ji} g_{ji}(u) + w_{CGMY} g_{CGMY}(u)$$

Model	Weights			Parameter restrictions
	w_{SV}	w_{ji}	w_{CGMY}	
SV	1			
SVJ1	$\frac{1}{1 + \lambda_1(\bar{\gamma}_1^2 + \delta_1^2)}$	$w_{j1} = \frac{\lambda_1(\bar{\gamma}_1^2 + \delta_1^2)}{1 + \lambda_1(\bar{\gamma}_1^2 + \delta_1^2)}$		$w_{j2} = 0$
SVJ2	$\frac{1}{1 + \sum_{i=1}^2 \lambda_i(\bar{\gamma}_i^2 + \delta_i^2)}$	$\frac{\lambda_i(\bar{\gamma}_i^2 + \delta_i^2)}{1 + \sum_{i=1}^2 \lambda_i(\bar{\gamma}_i^2 + \delta_i^2)}$		
DEXP	$1 - f_{jump}$		f_{jump}	$Y_n = Y_p = -1$
VG	$1 - f_{jump}$		f_{jump}	$Y_n = Y_p = 0$
Y	0		1	$Y_n = Y_p$
YY	0		1	
YY-D	$1 - f_{jump}$		f_{jump}	
LS	0		1	$w_n = 1, G = .001$

The weights ($w_{SV}, w_{j1}, w_{j2}, w_{CGMY}$) in the general specification sum to one and give the fractions of variance attributable to the underlying Lévy processes.

2.2. Time-changed Lévy processes and stochastic volatility

Time-changed Lévy processes generate stochastic volatility by randomizing time. Because the log of $F(u)$ in Eq. (1) can be written as $\ln F(u) = g_{dl}(u)Vt$, randomizing time is fundamentally equivalent to randomizing variance. Because the connection between time changes and stochastic volatility becomes less transparent once leverage effects are added, I use explicit stochastic volatility or stochastic intensity representations of stochastic processes.

The leverage effect, or correlation between asset returns and conditional variance innovations, is captured by directly specifying shocks common to both.⁵ I initially assume that the log asset price $s_t \equiv \ln S_t$ follows a process of the form

$$ds_t = (\mu_0 + \mu_1 V_t)dt + (\rho_{sv} \sqrt{V_t} dW_t - \frac{1}{2} \rho_{sv}^2 V_t dt) + dL_t$$

$$dV_t = \beta(\theta - V_t)dt + \sigma \sqrt{V_t} dW_t. \tag{14}$$

The log increment ds_t consists of the continuously compounded expected return plus conditionally independent increments to two exponential martingales. dW_t is a Wiener increment and dL_t is the increment to a compensated Lévy process, with instantaneous variance $(1 - \rho_{sv}^2)V_t dt$. Further refinements are added below, to match properties of stock market returns more closely.

This specification has various features or implicit assumptions. First, the approach allows considerable flexibility regarding the distribution of the instantaneous shock dL_t to asset returns, which can be Wiener, compound Poisson, or any other fat-tailed distribution. Three underlying Lévy processes are considered:

1. a second diffusion that is independent of W_t , with incremental variance $(1 - \rho_{sv}^2)V_t dt$ (Heston, 1993);
2. finite-activity jumps drawn from a normal distribution or a mixture of normals; and
3. the generalized CGMY Lévy process from Eq. (9).

Combinations of these processes are also considered, to nest the alternatives.

Second, the specification assumes a single underlying variance state variable V_t that follows an affine diffusion and which directly determines the variance of diffusion and jump components. This approach generalizes the stochastic jump intensity model of Bates (2000, 2006) to arbitrary Lévy processes.

Two alternate specifications are not considered, for different reasons. First, I do not consider the approach of Li, Wells, and Yu (2008), who model log-differenced asset prices as the sum of a Heston (1993) stochastic volatility process and a constant-intensity Lévy process that captures outliers. Bates (2006, Table 7) finds the stochastic-intensity jump model fits Standard & Poor's (S&P) returns better than the constant-intensity specification, when jumps are drawn from a finite-activity normal distribution or mixture of normals. Second, the diffusion assumption for V_t rules out volatility-jump models, such as the

exponential-jump model proposed by Duffie, Pan, and Singleton (2000) and estimated by Eraker, Johannes, and Polson (2003). The general issue of modeling correlated Lévy shocks to asset prices and to conditional variances is left for future research.

Define $y_T \equiv \int_{r=t}^T ds_r$ as the discrete-time return observed over horizon $\tau = T - t$, and define $f_{dl}(u) \equiv (1 - \rho_{sv}^2)V_t g_{dl}(u)$ as the cumulant exponent of dL_t conditional upon knowing V_t . By construction, $g_{dl}(u)$ is a standardized cumulant exponent, with $g_{dl}(1) = 0$ and variance $g''_{dl}(0) = 1$. A key property of affine models is the ability to compute the conditional generalized Fourier transform of (y_T, V_T) . This can be done by conditioning initially on the future variance path and iterating the expectational operator recursively backward in time:

$$F(\Phi, \psi | V_t, \tau) \equiv E[e^{\Phi \int_{r=t}^T ds_r + \psi V_T} | V_t]$$

$$= E[E[e^{\Phi \mu_0 \tau + \int_{r=t}^T \Phi[(\mu_1 - \frac{1}{2} \rho_{sv}^2)V_r dr + \rho_{sv} \sqrt{V_r} dW_r + dL_r]} \times e^{\psi V_T} | \{V_r\}_{r=t}^T] | V_t]$$

$$= E[e^{\Phi \mu_0 \tau + \int_{r=t}^T [\mu_1 \Phi + \frac{1}{2} \rho_{sv}^2 (\Phi^2 - \Phi) + (1 - \rho_{sv}^2) g_{dl}(\Phi)] V_r dr + \psi V_T} | V_t]$$

$$\equiv E[e^{\Phi \mu_0 \tau + h(\Phi) \int_{r=t}^T V_r dr + \psi V_T} | V_t] \tag{15}$$

for $h(\Phi) \equiv \mu_1 \Phi + \frac{1}{2} \rho_{sv}^2 (\Phi^2 - \Phi) + (1 - \rho_{sv}^2) g_{dl}(\Phi)$. Eq. (15) is the generalized Fourier transform of the future spot variance V_T and the average future variance $\bar{V}_{t \rightarrow T} \tau \equiv \int_{r=t}^T V_r dr$. This is a well-known problem discussed in Bakshi and Madan (2000), with an analytic solution if V_t follows an affine process. For the affine diffusion above, $F(\cdot | V_t, \tau)$ solves the Feynman-Kac partial differential equation

$$-F_\tau + \beta(\theta - V_t)F_V + \frac{1}{2} \sigma^2 V_t F_{VV} = -[\Phi \mu_0 + h(\Phi)V_t]F \tag{16}$$

subject to the boundary condition $F(\Phi, \psi | V, 0) = \exp(\psi V)$. The solution is

$$F(\Phi, \psi | V_t, \tau) = \exp[\mu_0 \tau \Phi + \theta C(\tau; \Phi, \psi) + D(\tau; \Phi, \psi)V_t] \tag{17}$$

where

$$C(\tau; \Phi, \psi) = -\frac{\beta \tau}{\sigma^2} (\rho_{sv} \sigma \Phi - \beta - \gamma)$$

$$- \frac{2\beta}{\sigma^2} \ln \left[1 + (\rho_{sv} \sigma \Phi - \beta - \gamma) \frac{(1 - e^{\gamma \tau})}{2\gamma} \right] - \frac{2\beta}{\sigma^2} \ln [1 - K(\Phi)\psi], \tag{18}$$

$$D(\tau; \Phi, \psi) = \frac{2\mu_1 \Phi + \rho_{sv}^2 (\Phi^2 - \Phi) + 2(1 - \rho_{sv}^2) g_{dl}(\Phi)}{\gamma[(e^{\gamma \tau} + 1)/(e^{\gamma \tau} - 1)] + \beta - \rho_{sv} \sigma \Phi}$$

$$+ \frac{\Lambda(\Phi)\psi}{1 - K(\Phi)\psi}, \tag{19}$$

$$\gamma = \sqrt{(\rho_{sv} \sigma \Phi - \beta)^2 - 2\sigma^2 [\mu_1 \Phi + \frac{1}{2} \rho_{sv}^2 (\Phi^2 - \Phi) + (1 - \rho_{sv}^2) g_{dl}(\Phi)]}, \tag{20}$$

$$\Lambda(\Phi) = \left[\left(\frac{e^{\gamma \tau} + 1}{e^{\gamma \tau} - 1} \right)^2 - 1 \right] / \left(\frac{e^{\gamma \tau} + 1}{e^{\gamma \tau} - 1} + \frac{\beta - \rho_{sv} \sigma \Phi}{\gamma} \right)^2, \text{ and} \tag{21}$$

$$K(\Phi) = \sigma^2 / \left(\gamma \frac{e^{\gamma \tau} + 1}{e^{\gamma \tau} - 1} + \beta - \rho_{sv} \sigma \Phi \right). \tag{22}$$

The specifications of $g_{dl}(\cdot)$ considered in this paper are listed above in Table 1.

⁵ This approach is equivalent for affine models to the change of measure approach in Carr and Wu (2004).

2.3. Autocorrelations and other refinements

That stock indexes do not follow a random walk was recognized explicitly by Lo and MacKinlay (1988) and implicitly by earlier practices in variance and covariance estimation designed to cope with autocorrelated returns; e.g., the Dimson (1979) lead/lag approach to beta estimation. The positive autocorrelations typically estimated for stock index returns are commonly attributed to stale prices in the stocks underlying the index. A standard practice in time series analysis is to prefilter the data by fitting an autoregressive moving average (ARMA) model; see, e.g., Jukivuolle (1995). Andersen, Benzoni, and Lund (2002), for instance, use a simple MA(1) specification to remove autocorrelations in S&P 500 returns over 1953–1996, a data set subsequently used by Bates (2006).

Prefiltering the data was considered unappealing in this study, for several reasons. First, the 1926–2006 interval used here is long, with considerable variation over time in trading activity and transactions costs, and structural shifts in the data generating process are probable. Indeed, Andersen, Benzoni, and Lund (2002, Table 1) find autocorrelation estimates from their full 1953–1996 sample diverge from estimates for a 1980–1996 subsample. Second, ARMA packages use a mean squared error criterion that is not robust to the fat tails observed in stock market returns. Finally, explicit consideration of autocorrelation is necessary when identifying the variance of relevance to option pricing.

Consequently, autocorrelations were treated as an additional latent variable, to be estimated as part of the overall time series model. I explore below two alternate models for daily log-differenced stock index excess returns y_t :

$$y_{t+1} = \rho_t y_t + \eta_{t+1} \quad (\text{Model 1}) \quad (23)$$

or

$$y_{t+1} = \rho_t y_t + (1 - \rho_t) \eta_{t+1} \quad (\text{Model 2}), \quad (24)$$

where

$$\eta_{t+1} \equiv \int_{r=t}^{t+\tau_t} ds_r, \quad V_{t+1} = V_t + \int_{r=t}^{t+\tau_t} dV_r, \quad \text{and} \quad (25)$$

$$\rho_{t+1} = \rho_t + \varepsilon_{t+1}, \quad \varepsilon \sim N(0, \sigma_\rho^2) \text{ and i.i.d.}$$

τ_t is the effective length of a business day, ρ_t is the daily autocorrelation, ds_t is the instantaneous intradaily shock to log asset prices, and $V_t dt \equiv \text{Var}_t(ds_t)$ is the instantaneous conditional variance of ds_t . The intradaily shocks (ds_t, dV_t) are given by Eq. (14).

Both models add an autocorrelation state variable ρ_t that captures the fact that autocorrelations of stock market returns are not constant over time. Following the literature on time-varying coefficient models, the autocorrelation is modeled as a simple random walk, to avoid constraining estimates of ρ_t . Estimation of the autocorrelation volatility parameter σ_ρ endogenously determines the appropriate degree of smoothing to use when estimating the current autocorrelation value ρ_t from past data.

The two models differ in ease of use, in their implications for the interaction between volatility and autocorrelation,

and in the pricing of risks. Model 1 assumes that the stock market excess return residual $\eta_{t+1} = y_{t+1} - \rho_t y_t$ is stationary (i.e., with a stationary conditional variance process) and that the current value of ρ_t affects only the conditional mean of y_{t+1} . Autocorrelation filtration in the model is consequently closer to standard autocorrelation estimation, and becomes identical when η_{t+1} is i.i.d. Gaussian and the autocorrelation is constant ($\sigma_\rho = 0$). Model 1 also has a semi-affine structure that permits direct estimation via the methodology of Bates (2006).

In Model 2, η_{t+1} is the permanent impact of daily shocks to stock index excess returns, and is again assumed stationary. The model assumes that infrequent trading in the component stocks (proxied by ρ_t) slows the incorporation of such shocks into the observed stock index, but that the index ultimately responds fully once all stocks have traded.⁶ Unlike Model 1, Model 2 is consistent with the LeBaron (1992) observation that annual estimates of daily stock market volatility and autocorrelation appear inversely related. Higher autocorrelations smooth shocks across periods, reducing observed market volatility. Furthermore, the model is more suitable for pricing risks; i.e., identifying the equity premium or the (affine) risk-neutral process underlying option prices. The current value of ρ_t affects both the conditional mean and higher moments of y_{t+1} , resulting in a substantially different filtration procedure for estimating ρ_t from past excess returns. The time series model is not semi-affine, but I develop below a change of variables that makes filtration and estimation as tractable as for Model 1.

Both models build upon previous time series and market microstructure research into stock market returns. For instance, the effective length τ_t of a business day is allowed to vary based upon various periodic effects. In particular, day-of-the-week effects, weekends, and holidays are accommodated by estimated time dummies that allow day-specific variation in τ_t .⁷ In addition, time dummies are estimated for the Saturday morning trading available over 1926–1952, and for the Wednesday exchange holidays in the second half of 1968 that are the focus of French and Roll (1986). Finally, the stock market closings during the Bank Holiday of March 3–15, 1933 and following the September 11, 2001 attacks are treated as 12/365- and 7/365-year returns, respectively. Treating the 1933 Bank Holiday as a 12-day interval substantially reduces the influence of its +15.5% return on parameter estimation. September 17, 2001 saw a smaller movement, of –4.7%.

For Model 1, the cumulant generating function of future returns and state variable realizations conditional upon current values is analytic and of the semi-affine

⁶ Jukivuolle (1995) distinguishes between the observed and true stock index when trading is infrequent, and proposes using a standard Beveridge-Nelson decomposition to identify the latter. This paper differs in assuming that the parameters of the autoregressive integrated moving average (ARIMA) process for the observed stock index are not constant.

⁷ Gallant, Rossi, and Tauchen (1992) use a similar approach, and also estimate monthly seasonals.

form

$$\begin{aligned} \ln F(\Phi, \xi, \psi | y_t, \rho_t, V_t) &\equiv \ln E[e^{\Phi y_{t+1} + \xi \rho_{t+1} + \psi V_{t+1}} | y_t, \rho_t, V_t] \\ &= \mu_0 \tau \Phi + \theta C(\tau; \Phi, \psi) + \frac{1}{2} \sigma_\rho^2 \xi^2 + (\xi + \Phi y_t) \rho_t + D(\tau; \Phi, \psi) V_t \\ &\equiv c(\tau; \Phi, \psi, \xi) + (\xi + \Phi y_t) \rho_t + D(\tau; \Phi, \psi) V_t, \end{aligned} \quad (26)$$

where $C(\tau; \Phi, \psi)$ and $D(\tau; \Phi, \psi)$ are given in Eqs. (18) and (19). For Model 2, the conditional cumulant generating function is of the non-affine form

$$\begin{aligned} \ln F(\Phi, \xi, \psi | y_t, \rho_t, V_t) &= \mu_0 \tau (1 - \rho_t) \Phi + \theta C(\tau; (1 - \rho_t) \Phi, \psi) + \frac{1}{2} \sigma_\rho^2 \xi^2 \\ &\quad + (\xi + \Phi y_t) \rho_t + D(\tau; (1 - \rho_t) \Phi, \psi) V_t \end{aligned} \quad (27)$$

given the shocks to y_{t+1} are scaled by $(1 - \rho_t)$.⁸

2.4. Filtration and maximum likelihood estimation

If the state variables (ρ_t, V_t) were observed along with returns, it would in principle be possible to evaluate the joint transition densities of the data and the state variable evolution by Fourier inversion of the joint conditional characteristic function $F(i\Phi, i\xi, i\psi | y_t, \rho_t, V_t)$, and to use this in a maximum likelihood procedure to estimate the parameters of the stochastic process. However, because (ρ_t, V_t) are latent, this is a hidden Markov model that must be estimated by other means.

For Model 1, the assumption that the cumulant generating function in Eq. (26) is affine in the latent state variables (ρ_t, V_t) implies that the hidden Markov model can be filtered and estimated using the approximate maximum likelihood (AML) methodology of Bates (2006). The AML procedure is a filtration methodology that recursively updates the conditional characteristic functions of the latent variables and future data conditional upon the latest datum. Define $\mathbf{Y}_t \equiv \{y_1, y_2, \dots, y_t\}$ as the data observed up through period t , and define

$$G_{t|t}(i\xi, i\psi) \equiv E[e^{i\xi \rho_t + i\psi V_t} | \mathbf{Y}_t] \quad (28)$$

as the joint conditional characteristic function that summarizes what is known at time t about (ρ_t, V_t) . The probability density of the latest observation y_{t+1} conditional upon \mathbf{Y}_t can be computed by Fourier inversion of its conditional characteristic function:

$$p(y_{t+1} | \mathbf{Y}_t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} G_{t|t}[i\Phi y_t, D(\tau; i\Phi, 0)] e^{c(\tau; i\Phi, 0, 0) - i\Phi y_{t+1}} d\Phi. \quad (29)$$

Conversely, the joint conditional characteristic function $G_{t+1|t+1}(i\xi, i\psi)$ needed for the next observation can be updated given y_{t+1} by the characteristic-function equivalent of Bayes' rule:

$$G_{t+1|t+1}(i\xi, i\psi) = \frac{1}{2\pi p(y_{t+1} | \mathbf{Y}_t)}$$

$$\times \int_{-\infty}^{\infty} G_{t|t}[i\xi + i\Phi y_t, D(\tau; i\Phi, i\psi)] e^{c(\tau; i\Phi, i\xi, i\psi) - i\Phi y_{t+1}} d\Phi. \quad (30)$$

The algorithm begins with an initial joint characteristic function $G_{1|1}(\xi)$ and proceeds recursively through the entire data set, generating the log likelihood function $\sum \ln p(y_{t+1} | \mathbf{Y}_t)$ used in maximum likelihood estimation. Filtered estimates of the latent variables can be computed from derivatives of the joint conditional moment generating function, as can higher conditional moments:

$$E[\rho_{t+1}^m V_{t+1}^n | \mathbf{Y}_{t+1}] = \left. \frac{\partial^{m+n} G_{t+1|t+1}(\xi, \psi)}{\partial \xi^m \partial \psi^n} \right|_{\xi = \psi = 0}. \quad (31)$$

The above procedure would permit exact maximum likelihood function estimation of parameters if implementable. However, the procedure requires storing and updating each entire function $G_{t|t}(\cdot)$ based on point-by-point univariate numerical integrations. As such a procedure would be slow, the AML methodology instead approximates $G_{t|t}(\cdot)$ on each date t by a moment-matching joint characteristic function, and updates the approximation based upon updated estimates of the moments of the latent variables. Given an approximate prior $\hat{G}_{t|t}(\cdot)$ and a datum y_{t+1} , Eqs. (30) and (31) are used to compute the posterior moments of (ρ_{t+1}, V_{t+1}) , which are then used to create an approximate $G_{t+1|t+1}(\cdot)$. The overall procedure is analogous to the Kalman filtration procedure of updating conditional means and variances of latent variables based upon observed data, under the assumption that those variables and the data have a conditional normal distribution. However, Eqs. (29) through (31) identify the optimal nonlinear moment updating rules for a given prior $G_{t|t}(\cdot)$, whereas standard Kalman filtration uses linear rules. I show below that this modification in filtration rules is important when estimating latent autocorrelations and volatilities under fat-tailed Lévy processes. Furthermore, Bates (2006) proves that the iterative AML filtration is numerically stable and shows that it performs well in estimating parameters and latent variable realizations.

Autocorrelations can be negative or positive, while conditional variance must be positive. Consequently, different two-parameter distributions were used to summarize conditional distributions of the two latent variables at each point in time: Gaussian for autocorrelations (in Model 1), gamma for variances. Furthermore, because volatility estimates mean-revert within months while autocorrelation estimates evolve over years, realizations of the two latent variables were assumed independent conditional upon past data. These assumptions resulted in an approximate cumulant generating function (CGF) for Model 1 of the form

$$\ln \hat{G}_{t|t}(\xi, \psi) = [\hat{\rho}_t \xi + \frac{1}{2} W_{t|t} \xi^2] - v_t \ln(1 - \kappa_t \psi). \quad (32)$$

To initiate the AML filtration, initial spot variance V_1 is assumed drawn from its unconditional gamma distribution, using the parameter values (κ_1, v_1) given in Table 2. Because autocorrelations are assumed nonstationary, no unconditional distribution exists. Consequently, the AML algorithm for Model 1 is initiated using a relatively diffuse conditional distribution for the initial autocorrelation ρ_1 that spans a much wider range than the plausible $(-1, +1)$ domain. The distributional assumptions for latent

⁸ Dilip Madan informs me that practitioners distinguish between time-scaled and space-scaled models of time-varying volatility. GARCH models are typically space-scaled, whereas Model 1 is a time-scaled model of stochastic volatility. Model 2 contains both (stationary) time scaling via V_t and the time dummies, and (nonstationary) space scaling via $1 - \rho_t$.

Table 2

Distributional approximations of latent state variables conditional upon past data \mathbf{Y}_t . All latent state variables are assumed conditionally independent of each other.

State variable	Conditional distribution	Distributional parameters	Initial parameters at time $t = 1$
Spot variance V_t (Models 1 and 2)	Gamma	(κ_t, ν_t) where $E[V_t \mathbf{Y}_t] \equiv \hat{V}_{t t} = \kappa_t \nu_t$ and $\text{Var}[V_t \mathbf{Y}_t] \equiv P_{t t} = \kappa_t^2 \nu_t$	$(\kappa_1, \nu_1) = \left(\frac{\sigma^2}{2\beta}, \frac{2\beta\theta}{\sigma^2}\right)$
Autocorrelation ρ_t (Model 1)	Gaussian	$\hat{\rho}_{t t} \equiv E[\rho_t \mathbf{Y}_t] \quad W_{t t} \equiv \text{Var}[\rho_t \mathbf{Y}_t]$	$(\hat{\rho}_{1 1}, W_{1 1}) = (0, 10^2)$
$1 - \rho_t$ (Model 2)	Inverse Gaussian	$\hat{\rho}_{t t} \equiv E[\rho_t \mathbf{Y}_t] \quad W_{t t} \equiv \text{Var}[\rho_t \mathbf{Y}_t]$	$(\hat{\rho}_{1 1}, W_{1 1}) = (0, 0.5^2)$

state variables under Models 1 and 2 are summarized in Table 2.

The parameters $(\hat{\rho}_{t|t}, W_{t|t}; \kappa_t, \nu_t)$, or equivalently the moments $(\hat{\rho}_{t|t}, W_{t|t}; \hat{V}_{t|t}, P_{t|t})$, summarize what is known about the latent variables at time t . These were updated daily using the latest observation y_{t+1} and Eqs. (29)–(31). For each day, five univariate integrations were required: one for the density evaluation in Eq. (29), and four for the mean and variance evaluations in Eq. (31). An upper Φ_{\max} was computed for each integral for which upper truncation error would be less than 10^{-10} in magnitude. The integrands were then integrated over $(-\Phi_{\max}, \Phi_{\max})$ to a relative accuracy of 10^{-9} , using the double-precision adaptive Gauss-Legendre quadrature routine DQDAG in the International Mathematics and Statistics Library (IMSL) and exploiting the fact that the integrands for negative Φ are the complex conjugates of the integrands evaluated at positive Φ . On average between 234 and 448 evaluations of the integrand were required for each integration.⁹

The nonaffine specification $y_{t+1} = \rho_t y_t + (1 - \rho_t) \eta_{t+1}$ in Model 2 necessitates additional restrictions upon the distribution of latent ρ_t . In particular, it is desirable that the scaling factor $1 - \rho_t$ be nonnegative, so that the lower tail properties of η_{t+1} originating in the underlying Lévy specifications do not influence the upper tail properties of y_{t+1} . Consequently, the conditional distribution of latent $1 - \rho_t$ for Model 2 is modeled as inverse Gaussian – a two-parameter unimodal distribution with conditional mean $1 - \hat{\rho}_{t|t}$ and conditional variance $W_{t|t}$. Appendix A derives the filtration procedure for this model, exploiting a useful change of variables procedure. The filtration is initiated at $\rho_1 \sim (0, 0.5^2)$, and again assumes that ρ_t and V_t are conditionally independent for all t .

3. U.S. stock market returns, 1926–2010

3.1. Data

Two value-weighted measures of the U.S. stock market are readily available: the Center for Research in Security Prices (CRSP) value-weighted index, and the S&P Composite Index. This paper uses the former for time series analysis, but also considers the latter when assessing stock index options. The CRSP data used for parameter estimation consist of 21,519 daily cum-dividend

value-weighted returns over January 2, 1926 through December 29, 2006, based on stocks traded on the New York Stock Exchange and the American Stock Exchange. The out-of-sample tests use 1,008 additional daily observations over 2007–2010.

CRSP daily returns for each month were converted to daily log excess returns using Ibbotson Associates' (2006) data on monthly Treasury bill returns and the formula

$$y_t = \ln(1 + R_t) - \frac{n_t}{N} \ln(1 + i), \quad (33)$$

where R_t is the daily CRSP cum-dividend return; i is that month's return on Treasury bills of at least one month to maturity; N is the number of calendar days spanned by the monthly Treasury bill return; and n_t is the number of calendar days spanned by the business-day return R_t . The monthly interest rate data were downloaded from Ken French's website and extended backward through 1926 using data in Ibbotson Associates' *S&P Yearbook*.

For pricing options based on the S&P 500 index, I constructed a spliced CRSP/S&P 500 series: CRSP excess returns up through March 4, 1957, and S&P 500 excess returns thereafter. I prefer CRSP data to the early cum-dividend S&P data available within the Schwert (1990) database, for three reasons. First, the S&P Composite Index prior to March 4, 1957 is based upon 90 stocks, whereas the CRSP index is broader. Second, S&P data within the Schwert database begin in 1928, and it is important for volatility filtration around the 1929 stock market crash to have data over 1926 and 1927.¹⁰ Third, the S&P Composite Index is only reported to two decimal places, which creates significant rounding error issues for the low S&P values (around five) observed in the 1930s. I updated Schwert's cum-dividend S&P 500 returns through 2010 using his methodology. Ex-dividend daily S&P 500 returns from CRSP (from Yahoo for data in 2010) were augmented by an average daily dividend yield computed from a monthly S&P 500 dividend series from Bloomberg and the previous month's end-of-month index level. Cum-dividend returns were then converted into log excess returns using Eq. (33).

3.2. Parameter estimates

Table 3 describes and provides estimates of the time dummies from the most general time-changed CGMY

⁹ The Fast Fourier Transform approach in Carr, Geman, Madan, and Yor (2002) uses 16,384 evaluations.

¹⁰ The Schwert database also includes daily returns over 1885 to 1927, based on the (price-weighted) Dow Jones Industrial Average.

Table 3

Effective length of a business day relative to 1-day Wednesday returns, on a close-to-close basis.

Estimates are for the YY model, which is the most general version of the Carr, Geman, Madan, and Yor (2003) model of time-changed Lévy returns. Estimates for other models are almost identical. Data are daily value-weighted excess returns over 1926–2006, using cum-dividend returns from the Center for Research in Securities Prices. One 3-day weekday holiday on August 14–17, 1945 is included among the 3-day returns. The annualization factor is the sum over 1926–2006 of all daily time horizons divided by the 80.936-year span, excluding the 12- and 7-day market closings in 1933 and 2001. Standard errors are in parentheses.

Number of days	Description	Number of observations	Model 1 estimates	Model 2 estimates
1	Monday close → Tuesday close	3,831	1.02 (0.04)	1.03 (0.03)
1	Tuesday → Wednesday	4,037	1	1
1	Wednesday → Thursday	3,998	0.94 (0.03)	0.94 (0.03)
1	Thursday → Friday	3,924	0.93 (0.03)	0.92 (0.03)
1	Friday → Saturday (1926–1952)	1,141	0.43 (0.02)	0.44 (0.02)
2	Saturday → Monday (1926–1952)	1,120	1.05 (0.05)	1.07 (0.05)
2	Weekday holiday	341	1.25 (0.11)	1.26 (0.10)
2	Wednesday exchange holiday in 1968	22	0.73 (0.33)	0.81 (0.35)
3	Holiday weekend or holiday	2,755	1.10 (0.04)	1.10 (0.04)
4	Holiday weekend	343	1.58 (0.14)	1.56 (0.13)
5	Holiday weekend	6	1.31 (1.00)	1.25 (0.93)
	Total number of observations	21,518		
	Annualization factor: Wednesday → yearly		259.8 (5.6)	260.3 (5.5)

model, with Wednesday returns (Tuesday close to Wednesday close) arbitrarily selected as the benchmark day. Estimates from other Lévy models were typically within ± 0.01 of those in Table 3. Daily variance tended to be highest at the beginning of the week and decline thereafter, but day-of-the-week effects do not appear to be especially pronounced. The major exception is the Saturday returns generally available over 1926–1952.¹¹ Saturdays were effectively 43% as long as the typical Wednesday. Total weekend variance (Friday close to Monday close) was $(0.43 + 1.05) / 1.10 - 1 = 34.5\%$ higher when Saturday trading was available (over 1926–1952) than when it was not (over 1945–2006).¹² This is qualitatively similar to but less pronounced than the doubling of weekend variance found by Barclay, Litzenberger, and Warner (1990) in Japanese markets when Saturday half-day trading was feasible. Barclay, Litzenberger, and Warner lucidly discuss market microstructure explanations for the increase in variance.

Holidays generally did not have a strong impact on the effective length of a business day, with the exception of holiday weekends spanning four calendar days. Consistent with French and Roll (1986), two-day returns spanning the Wednesday exchange holidays in 1968 (Tuesday close to Thursday close) had a variance not statistically different from a typical one-day Wednesday return but substantially less than the $1 + 0.94 = 1.94$ two-day variance observed for returns from Tuesday close to Thursday close in other years. Overall, the common practice of ignoring day-of-the-week effects, weekends, and holidays when analyzing the time series properties of daily stock

market returns appears to be a reasonable approximation, provided the data exclude Saturday returns.

Table 4 contains estimates for various specifications listed in Table 1, while Fig. 1 presents associated normal probability plots for Model 2. (The plots for Model 1 are similar.) All specifications capture the leverage effect by a correlation ρ_{sv} with the diffusion shock to conditional variance. The specifications diverge in their modeling of the Lévy shocks dL_t orthogonal to the variance innovation. SV is the Heston (1993) stochastic volatility model, while SVJ1 and SVJ2 have an additional diffusion for small asset return shocks, plus finite-activity normally-distributed jumps to capture outliers. The other models examine the generalized time-changed CGMY model, along with specific parameter restrictions or relaxations indicated in Table 1.

Most specifications using either Model 1 or Model 2 have similar estimates for the parameters determining the conditional mean and stochastic variance evolution. The evidence for a variance-sensitive equity premium ($\mu_1 > 0$) is stronger for Model 2 specifications, but μ_1 is not typically significantly different from zero for either model. Latent permanent variance in Model 2 mean-reverts towards an estimated average level around $(0.172)^2$, with a half-life about 1.6 months. The SV and LS models are the outliers, with different estimates of the equity premia and variance process from other specifications. As discussed in Section 3.6, this reflects these two specifications' substantially different approach to variance filtration.

The various specifications diverge primarily in how they capture tail risk. The Merton-based SVJ1 and SVJ2 results in Panel B of Table 4 largely replicate the jump risk results in Bates (2006). The SVJ1 model has symmetric normally distributed jumps with a jump standard deviation of approximately 3.3% and time-varying jump intensities that average $\lambda_1 \theta = 3.4$ jumps per year. As shown in Fig. 1, this jump risk assessment fails to capture the substantial 1987 crash. By contrast, the SVJ2 model adds a second jump component that directly captures the 1987 outlier. The resulting increase

¹¹ Saturday morning trading (ten to noon) was standard before 1945. Over 1945–1951, it was increasingly eliminated in summer months, and was permanently eliminated on June 1, 1952.

¹² As the time dummy estimates are estimated jointly with the volatility and autocorrelation filtrations, the estimates of weekend variances with versus without Saturday trading control for divergences in volatility and autocorrelation levels in the two samples.

Table 4

Parameter estimates using daily CRSP value-weighted excess returns y_t over 1926–2006.

See Eqs. (14) and (25) for definitions of parameters affecting the conditional means and volatilities, and see Table 1 for the various specifications of Lévy shocks. All parameters in Panel A are in annualized units except for the half-life $HL = 12 \ln 2 / \beta$ of variance shocks, which is in months. Models with $f_{jump} < 1$ in Panel B combine Lévy jumps with an additional independent diffusion, with variance proportions $(f_{jump}, 1 - f_{jump})$, respectively. The fraction of total return variance attributable to jumps is $(1 - \rho_{sv}^2) f_{jump}$. Standard errors are in parentheses.

Panel A: Estimates of parameters affecting the conditional means and volatilities								
Models of Lévy shocks	ln L	Conditional mean			Stochastic volatility			
		μ_0	μ_1	$\sigma_\rho \sqrt{252}$	$\sqrt{\theta}$	σ	ρ_{sv}	HL (months)
Model 1: $y_{t+1} = \rho v_t + \eta_{t+1}$								
SV	74,940.85	0.013 (0.015)	2.16 (0.90)	0.029 (0.007)	0.153 (0.004)	0.452 (0.010)	−0.625 (0.018)	1.4 (0.1)
SVJ1	75,043.90	0.042 (0.015)	0.91 (0.91)	0.030 (0.006)	0.155 (0.005)	0.374 (0.011)	−0.641 (0.020)	1.9 (0.2)
SVJ2	75,048.49	0.042 (0.015)	0.87 (0.76)	0.030 (0.007)	0.155 (0.007)	0.371 (0.015)	−0.642 (0.018)	1.9 (0.2)
DEXP	75,047.33	0.043 (0.015)	0.87 (0.90)	0.031 (0.007)	0.155 (0.005)	0.368 (0.012)	−0.587 (0.020)	2.0 (0.2)
VG	75,049.09	0.043 (0.015)	0.92 (0.91)	0.030 (0.006)	0.155 (0.005)	0.366 (0.012)	−0.586 (0.020)	2.0 (0.2)
Y	75,049.63	0.042 (0.015)	0.90 (0.92)	0.030 (0.006)	0.156 (0.009)	0.351 (0.020)	−0.576 (0.032)	2.1 (0.2)
YY	75,052.56	0.041 (0.015)	0.87 (0.92)	0.030 (0.006)	0.158 (0.009)	0.360 (0.019)	−0.571 (0.031)	2.1 (0.2)
YY_D	75,052.81	0.042 (0.015)	0.93 (0.91)	0.030 (0.006)	0.156 (0.006)	0.355 (0.013)	−0.579 (0.021)	2.1 (0.2)
LS	75,007.86	0.018 (0.015)	1.50 (0.73)	0.031 (0.007)	0.171 (0.006)	0.431 (0.015)	−0.541 (0.020)	1.8 (0.2)
Model 2: $y_{t+1} = \rho v_t + (1 - \rho) \eta_{t+1}$								
SV	74,999.87	−0.014 (0.020)	3.04 (0.90)	0.043 (0.005)	0.170 (0.004)	0.562 (0.015)	−0.658 (0.017)	1.0 (0.1)
SVJ1	75,092.10	0.033 (0.020)	1.69 (1.04)	0.036 (0.005)	0.171 (0.004)	0.457 (0.015)	−0.674 (0.018)	1.4 (0.1)
SVJ2	75,096.68	0.037 (0.020)	1.25 (0.89)	0.036 (0.005)	0.172 (0.004)	0.456 (0.015)	−0.673 (0.018)	1.4 (0.1)
DEXP	75,094.20	0.034 (0.020)	1.44 (0.90)	0.036 (0.005)	0.171 (0.004)	0.452 (0.015)	−0.625 (0.018)	1.5 (0.1)
VG	75,094.70	0.034 (0.020)	1.42 (0.90)	0.037 (0.005)	0.171 (0.004)	0.447 (0.016)	−0.623 (0.018)	1.6 (0.1)
Y	75,093.68	0.036 (0.021)	1.35 (0.90)	0.036 (0.005)	0.172 (0.007)	0.432 (0.021)	−0.613 (0.027)	1.6 (0.1)
YY	75,097.20	0.033 (0.020)	1.44 (0.90)	0.036 (0.005)	0.172 (0.006)	0.437 (0.018)	−0.613 (0.022)	1.6 (0.1)
YY_D LS	75,097.49	0.035 (0.020)	1.36 (0.90)	0.036 (0.005)	0.172 (0.005)	0.436 (0.016)	−0.616 (0.020)	1.6 (0.1)
	75,045.48	0.053 (0.019)	1.50 (0.76)	0.031 (0.003)	0.174 (0.005)	0.436 (0.015)	−0.576 (0.019)	1.8 (0.2)
Panel B: Estimates of jump parameters								
Models of Lévy shocks	f_{jump}	CGMY parameters				Merton parameters		
		w_n	G	M	Y_n	Y_p	λ_i	$\bar{\gamma}_i$
Model 1: $y_{t+1} = \rho v_t + \eta_{t+1}$								
SVJ1	0.150 (0.017)					142.7 (22.7)	0.000 (0.002)	0.032 (0.002)
SVJ2	0.156 (0.054)					162.9 (30.9)	0.000 (0.002)	0.029 (0.002)
						0.5 (1.6)	−0.189 (0.094)	0.005 (0.189)
DEXP	0.253 (0.027)	0.49 (0.01)	66.1 (6.0)	45.4 (8.4)	−1			
VG	0.272 (0.030)	0.52 (0.07)	41.1 (5.4)	31.6 (9.1)	0			
Y	1	0.59 (0.06)	7.0 (4.6)	2.3 (7.3)	1.87 (0.03)			

YY	1	0.88 (0.03)	1.6 (4.5)	40.1 (31.3)	1.94 (0.01)	-0.24 (1.36)			
YY_D	0.436 (0.090)	0.72 (0.15)	6.9 (9.0)	49.2 (34.9)	1.71 (0.35)	-0.72 (1.57)			
LS	1	1	0.001		1.96 (0.01)				
Model 2: $y_{t+1} = \rho y_t + (1 - \rho) \epsilon_{t+1}$									
SVJ1	0.133 (0.015)						114.1 (19.2)	-0.001 (0.003)	0.034 (0.002)
SVJ2	0.140 (0.015)						126.4 (23.7)	0.000 (0.002)	0.031 (0.002)
							0.41 (0.04)	-0.198 (0.022)	0.010 (0.046)
DEXP	0.236 (0.026)	0.53 (0.06)	55.4 (5.4)	50.0 (4.8)					
VG	0.257 (0.030)	0.54 (0.07)	41.1 (4.9)	31.6 (10.7)					
Y	1	0.59 (0.05)	6.8 (4.2)	3.2 (8.2)	1.87 (0.03)				
YY	1	0.89 (0.03)	2.6 (4.1)	71.1 (57.8)	1.94 (0.01)	-1.96 (2.62)			
YY_D	0.380 (0.158)	0.90 (0.30)	8.1 (8.7)	51.8 (51.7)	1.62 (0.40)	-1.08 (2.36)			
LS	1	1	0.001		1.97 (0.01)				

in log likelihood is statistically significant under a likelihood ratio test, with a marginal significance level around 3% for Models 1 and 2.

The various CGMY models primarily diverge across the specification of the (Y_p, Y_n) parameters: whether they are set to specific levels and whether they diverge for positive versus negative jumps. The DEXP model with $Y_p = Y_n = -1$ is conceptually similar to the jump-diffusion model SVJ1, but it uses instead a finite-activity double exponential distribution for jumps. Despite the fatter-tailed specification, Fig. 1 indicates the DEXP model has difficulties comparable to SVJ1 in capturing the 1987 crash. The VG model has an infinite-activity variance process ($Y_p = Y_n = 0$) and has a slightly higher log likelihood. The VG normal probability plot is virtually identical to that of DEXP in Fig. 1, and is omitted to save space. Both models include a diffusion component, which captures 73–76% of the variance of the orthogonal Lévy shock dL_t .

Specifications Y, YY, and LS involve pure-jump processes for the orthogonal Lévy process L_t , without an additional diffusion component. Overall, higher values of Y fit the data better—especially the 1987 crash, which ceases to be an outlier under these specifications. Relaxing the restriction $Y_p = Y_n$ leads to a statistically significant improvement in fit, with the increase in log likelihood (YY versus Y) having *p*-values of 1.8% and 0.8% for Models 1 and 2, respectively. Point estimates of the jump parameters (w_n, G, Y_n) governing downward jump intensities diverge sharply from the parameters $(1 - w_n, M, Y_p)$ governing upward jump intensities when the $Y_p = Y_n$ restriction is relaxed, although standard errors are large. The dampening coefficient *G* is not significantly different from zero, implying one cannot reject the hypothesis that the downward jump intensity is from a stochastic-intensity version of the Carr and Wu (2003) log-stable process. By contrast, the upward intensity is estimated as a finite-activity jump process which, however, still overestimates the frequency of big positive outliers (YY plot in Fig. 1).

Motivated by option pricing issues, Carr and Wu (2003) advocate using a log-stable distribution with purely downward jumps. An approximation to this model generated by setting $G = 0.001$ and $w_n = 1$ fits stock market excess returns very badly. The basic problem is that while the LS model does allow substantial positive excess returns, it severely underestimates the frequency of large positive values. This leads to a bad fit for the upper tail (LS plot in Fig. 1). However, the YY estimates indicate that the Carr and Wu specification can be a useful component of a model, provided the upward jump intensity function is modeled separately.

The Y and YY models generate at least one Y parameter in the infinite-activity, infinite-variation range [1, 2], and typically near the diffusion value of two. This suggests that the models may be trying to capture considerable near-zero activity. Adding an additional diffusion component to the time-changed YY Lévy specification to capture that activity separately (specification YY_D) alters the CGMY jump parameter estimates substantially but leads to no statistically significant improvement in fit.

Overall, Fig. 1 suggests the alternate fat-tailed specifications fit the data similarly over most of the data range ($|z| < 3$). The models SV, SVJ1, DEXP, VG, and LS appear less desirable, given their failure to capture the largest

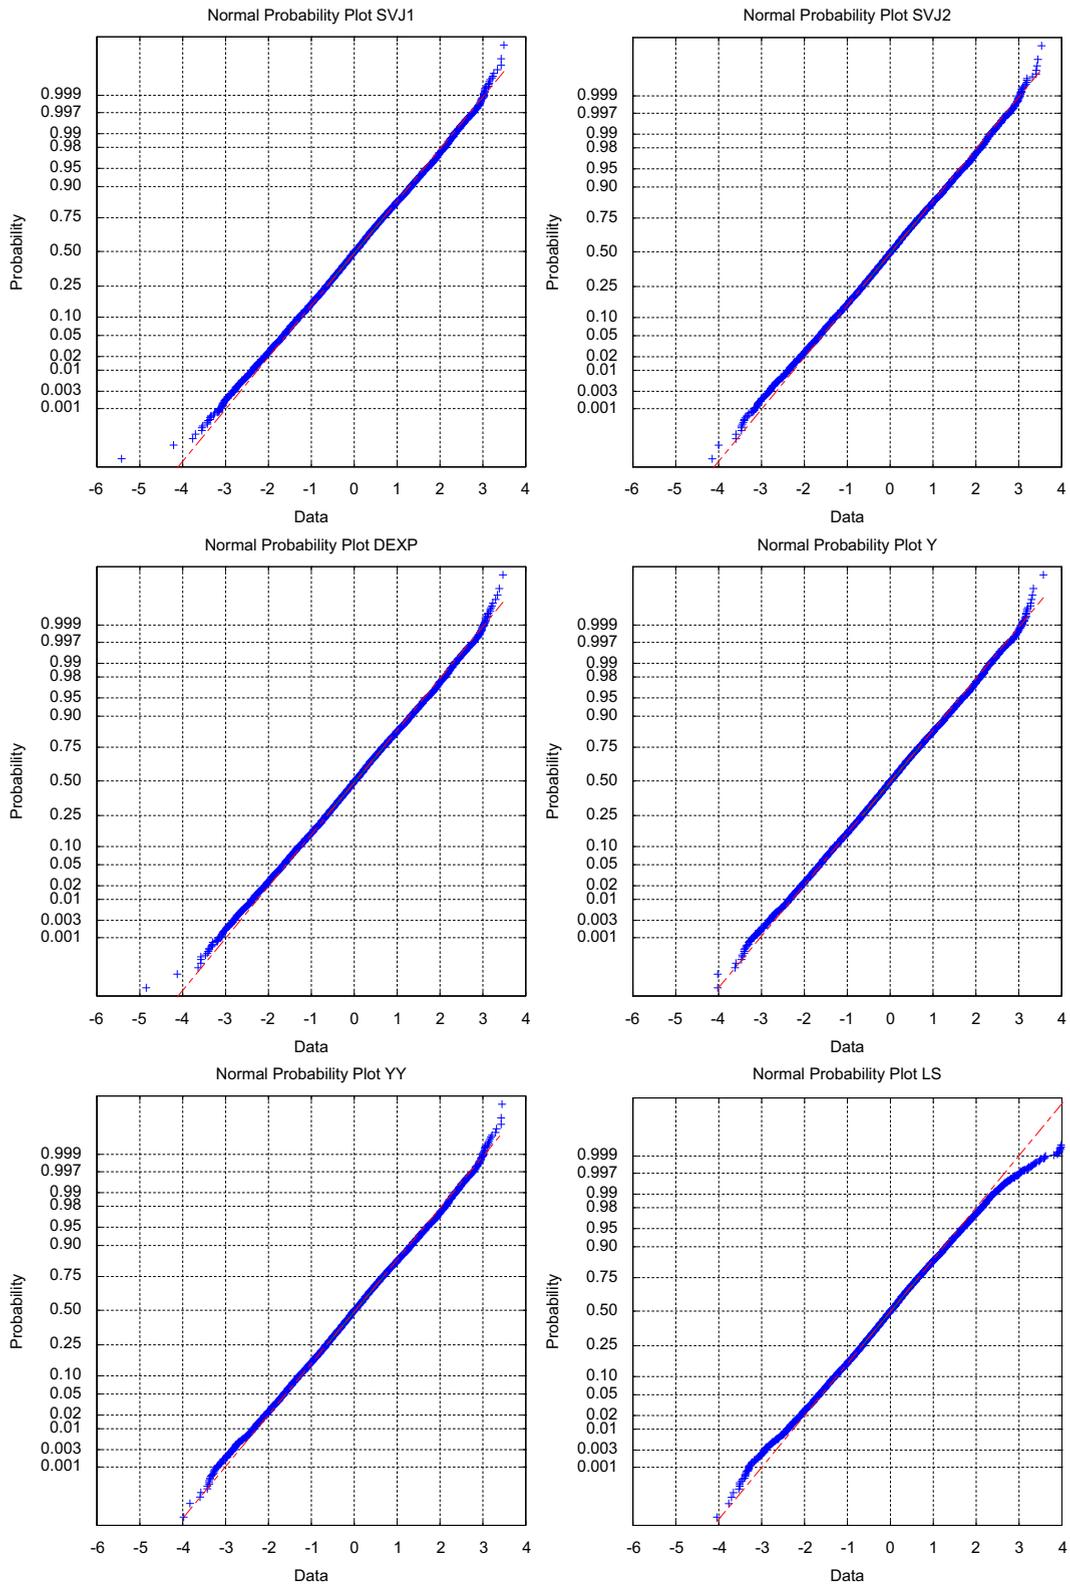


Fig. 1. Normal probability plots for the normalized returns $z_{t+1} = N^{-1}[CDF(y_{t+1}|Y_t)]$, for different Lévy specifications under Model 2. The diagonal lines are the theoretical quantiles conditional upon correct specification, whereas the data (+) are the empirical quantiles. See Eq. (A.12) for the computation of the conditional distribution function (CDF).

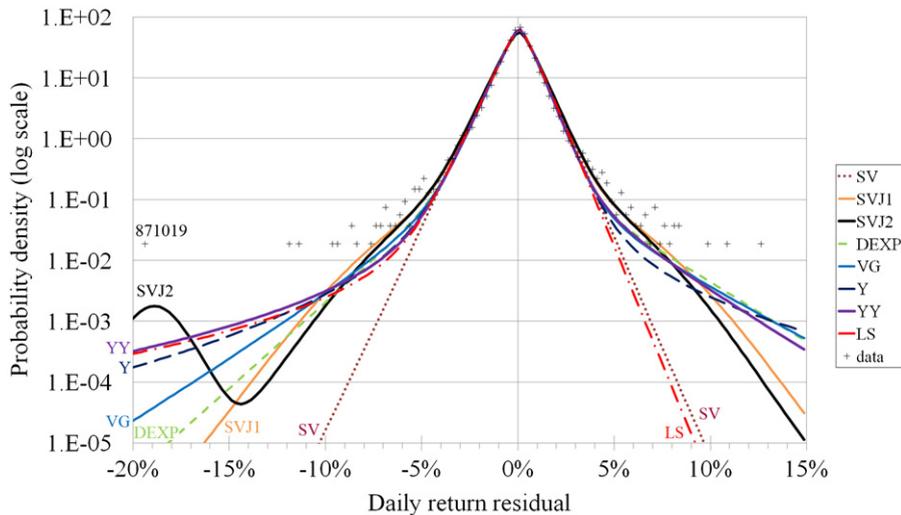


Fig. 2. Unconditional probability density functions of Model 1 return residuals, on a log scale. Empirical data frequencies (+) are from a histogram of 21,516 daily return residuals $\eta_{t+1} = y_{t+1} - \hat{\rho}_t y_t$ from the YY model, using a 0.25% cell width. Residuals for the two extended market closings in 1933 and 2001 were excluded.

outliers. The SVJ2, Y, and YY specifications appear to fit about the same. All models appear to have a small amount of specification error (deviations from linearity) in the $z \in [-3.5, -2]$ range and in the upper tail ($z > 3$).

3.3. Unconditional distributions

A further diagnostic of model specification is the models' ability or inability to match the unconditional distribution of returns; in particular, the tail properties of unconditional distributions. Mandelbrot (1963) and Mandelbrot and Hudson (2004) argue that empirical tails satisfy a power law: tail probabilities plotted against absolute returns approach a straight line when plotted on a log–log graph. This empirical regularity underlies Mandelbrot's advocacy of the stable Paretian distribution, which possesses this property and is nested within the CGMY model for $G=M=0$.

Mandelbrot's argument is premised upon identically and independently distributed returns, but the argument can be extended to time-changed Lévy processes. Lévy densities time-average; if the conditional intensity of moves of size x is $(1 - \rho_{sv}^2)k(x)V_t$, then the unconditional frequency of moves of size x is $(1 - \rho_{sv}^2)k(x)E(V_t)$. Because unconditional probability density functions asymptotically approach the unconditional Lévy densities for large $|x|$, while unconditional tail probabilities approach the corresponding integrals of the unconditional Lévy densities, examining unconditional distributions might still be useful under stochastic volatility.

Fig. 2 provides estimates of unconditional probability density functions of stock market excess return residuals $\eta_{t+1} = y_{t+1} - \rho_t y_t$ for various specifications under Model 1, as well as data-based estimates from a histogram of filtered residuals $\hat{\eta}_{t+1} = y_{t+1} - \hat{\rho}_t y_t$. Given the day-of-the-week effects reported in Table 3, the unconditional density functions are a mixture of horizon-dependent densities, with mixing weights set equal to the empirical frequencies. (The two shocks spanning the longer market closings in 1933 and 2001 are omitted.) The substantial

impact of the 1987 crash outlier upon parameter estimates is apparent. The SVJ2 estimates treat that observation as a unique outlier. The CGMY class of models progressively fatten the lower tail as greater flexibility is permitted for the lower tail parameter Y_n , with the lower tail approaching the Carr and Wu (2003) log-stable (LS) estimate for the YY specification. However, the LS model is unable to capture the frequency of large positive outliers and behaves similarly to the SV model in the upper tail. All models closely match the empirical unconditional density function within the $\pm 3\%$ range where most observations occur; and all models underestimate the unconditional frequency of moves of 3–7% in magnitude.

Fig. 3 provides similar estimates for unconditional lower and upper tail probabilities from the YY model, using the log–log graphs advocated by Mandelbrot. Residuals were selected from 19,663 days with time horizons within 11% of Wednesday's value, to make time horizons roughly identical. In addition, 1,000 sample paths of stock market excess return residuals over 1926–2006 were simulated via a Monte Carlo procedure that used YY parameter estimates, to provide confidence intervals.¹³ Unsurprisingly, the confidence intervals on extreme tail events are wide. The underestimation of moves of 3–7% in magnitude is again apparent, and is statistically significant. This rejection of the YY model does not appear attributable to misspecification of the Lévy density, which in Fig. 1 captures conditional densities quite well. Rather, the poor unconditional fits in Figs. 2 and 3 appear due to misspecification of volatility dynamics. Half of the 3–7% moves occurred over 1929–1935, a prolonged high-volatility period that simulated

¹³ Conditional intradaily variance shocks for data-based daily time horizons were simulated using the approach of Bates (2006, Appendix A.6). Lévy shocks η_{t+1} conditional upon intradaily average variance were generated via an inverse cumulative distribution function (CDF) methodology, with CDFs computed by Fourier inversion. The subset of days with roughly 1.0 days to maturity ($\pm 11\%$) were then selected.

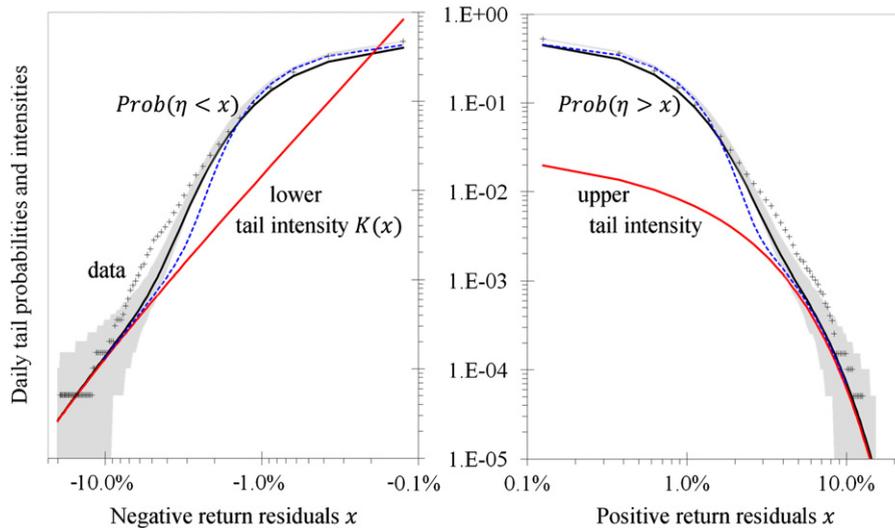


Fig. 3. Unconditional daily tail probabilities (solid black lines) and Lévy tail intensities of return residuals; log scales on both axes. The probability and intensity functions are based on YY Model 1 parameter estimates. Data-based estimates of tail probabilities (+) are the empirical quantiles of excess return residuals $\eta_{t+1} = y_{t+1} - \hat{\rho}_{t|t} y_t$ for 19,663 business days with estimated time horizons within $\pm 11\%$ of the typical Wednesday. Grey areas are 95% confidence intervals from 1,000 simulated sample paths of the stochastic-intensity YY process. The dashed lines give the tail probabilities from a comparable constant-intensity Lévy process. Tail probabilities converge to tail intensities for large $|x|$, and converge faster when volatility is not stochastic.

volatility realizations from the one-factor variance process of Eq. (14) generally do not match.

Fig. 3 also shows the unconditional tail intensities, and shows that the estimated tail probabilities converge to the estimated tail intensities. The lower tail intensity is

$$K(x) \equiv \int_{-\infty}^x k(y) dy = C_n G^{Y_n} \Gamma(-Y_n, G|x|), \quad (34)$$

where $C_n = w_n(1 - \rho_{sv}^2)\theta\tau / [\Gamma(2 - Y_n)G^{Y_n - 2}]$ and $\Gamma(a, z)$ is the incomplete gamma function. Furthermore, given G estimates near zero, $K(x) \approx C_n x^{-Y_n} / Y_n$ is roughly a power function in x , implying the near linearity (when plotted on a log-log scale) evident in Fig. 3.

However, the graph indicates that the convergence of tail probabilities to the tail intensity $K(x)$ occurs only for observations in excess of 5% in magnitude, or roughly 5 standard deviations. As this is outside the range of almost all data, it does not appear that graphing empirical tail probabilities on log-log scales provides a useful diagnostic of model specification and tail properties for daily data. This is partly due to stochastic volatility, which significantly slows the asymptotic convergence of unconditional tail probabilities to $K(x)$ for large $|x|$. Absent stochastic volatility ($\sigma = 0$), the tail probabilities of an i.i.d. YY Lévy process converge to $K(x)$ for daily observations roughly in excess of 3% in magnitude (3 standard deviations), as illustrated by the dashed lines in Fig. 3.

No power law properties are observed for upper tail probabilities, given substantial estimated exponential dampening. The failure of both lower and upper unconditional tail probabilities to capture the frequency of moves of 3–7% in magnitude is again apparent and statistically significant.

3.4. Subsample estimates

Table 5 provides estimates for data subsamples, as a test of the stability of the time series process. The mean, stochastic volatility, and jump parameters were allowed to differ before and after March 5, 1957.¹⁴ The time dummies (similar to those in Table 3) that capture day-of-the-week effects were kept common across subsamples; but some of those dummies also capture subsample-specific phenomena (Saturday trading before 1953; exchange holidays in 1968). The estimation and filtration over the two subsamples nest the full-sample estimates of Table 4, so that standard likelihood ratio tests can be used to test whether the divergences in subsample parameter estimates are statistically significant.

Parameter estimates diverge strongly across subsamples, with p -values less than 10^{-16} , but in different fashions for the SVJ1 and YY models. For the SVJ1 model, the major divergence was clearly in the estimated volatility process. The 1926–1957 period includes the highly volatile 1930s, yielding an overall average variance of $(0.202)^2$ over 1926–1957 versus $(0.149)^2$ over 1957–2006. The volatility dynamics also diverge, with volatility more volatile and with faster mean reversion over 1926–1957 than over 1957–2006. Jump risk estimates diverge as well, with more frequent but smaller jumps in the first half than in the second half. Progressively relaxing full-sample constraints on parameter categories (mean, σ_ρ , stochastic volatility parameters, jump parameters) indicates that

¹⁴ The data split was chosen so that the second subsample's estimates could be compared with estimates from S&P 500 returns, as well as with other studies that use data starting in the 1950s, such as Andersen, Benzoni, and Lund (2002), Chernov, Gallant, Ghysels, and Tauchen (2003), and Bates (2006).

Table 5

Subsample parameter estimates for Model 2.

Data are daily CRSP value-weighted excess returns over 1926–2006. Split-sample estimates involve different parameter values before/after March 5, 1957, apart from time dummies. See Eq. (14) and the notes to Table 4 for parameter definitions. Standard errors are in parentheses.

Panel A: Estimates of parameters affecting the conditional means and volatilities										
Model	Period	ln L	Conditional mean			Stochastic volatility				
			μ_0	μ_1	$\sigma_\rho\sqrt{252}$	$\sqrt{\theta}$	σ	ρ_{sv}	HL (months)	
SVJ1	1926–2006	75,092.10	0.033 (0.020)	1.69 (1.04)	0.036 (0.005)	0.171 (0.004)	0.457 (0.015)	–0.674 (0.018)	1.4 (0.1)	
SVJ1	1926–1957		0.051 (0.034)	1.35 (1.38)	0.050 (0.009)	0.202 (0.007)	0.678 (0.035)	–0.661 (0.027)	0.9 (0.1)	
YY	1957–2006	75,183.99	0.003 (0.027)	2.90 (1.61)	0.027 (0.005)	0.149 (0.005)	0.314 (0.015)	–0.725 (0.023)	1.7 (0.2)	
YY	1926–2006		0.033 (0.020)	1.44 (0.90)	0.036 (0.005)	0.172 (0.006)	0.437 (0.018)	–0.613 (0.022)	1.6 (0.1)	
YY	1926–1957	75,196.14	0.056 (0.034)	1.03 (1.15)	0.051 (0.009)	0.201 (0.008)	0.657 (0.033)	–0.585 (0.026)	1.2 (0.2)	
YY	1957–2006		0.012 (0.027)	0.35 (0.67)	0.025 (0.005)	0.365 (0.320)	0.404 (0.026)	–0.281 (0.247)	1.7 (0.2)	
Panel B: Estimates of jump parameters										
Model	Period	f_{jump}	CGMY parameters					Merton parameters		
			w_n	G	M	Y_n	Y_p	λ_1	$\bar{\gamma}_1$	δ_1
SVJ1	1926–2006	0.133 (0.015)						114.1 (19.2)	–0.001 (0.003)	0.034 (0.002)
SVJ1	1926–1957	0.167 (0.015)						216.8 (54.9)	0.000 (0.003)	0.028 (0.003)
SVJ1	1957–2006	0.093 (0.020)						49.5 (12.0)	–0.003 (0.007)	0.043 (0.004)
YY	1926–2006	1	0.89 (0.03)	2.6 (4.1)	71.1 (57.8)	1.94 (0.01)	–1.96 (2.6)			
YY	1926–1957	1	0.86 (0.04)	20.7 (7.9)	97.8 (107.2)	1.82 (0.05)	–3.1 (4.5)			
YY	1957–2006	1	0.92 (0.15)	0.0 (0.0)	6.0 (13.4)	1.54 (0.41)	1.94 (0.03)			

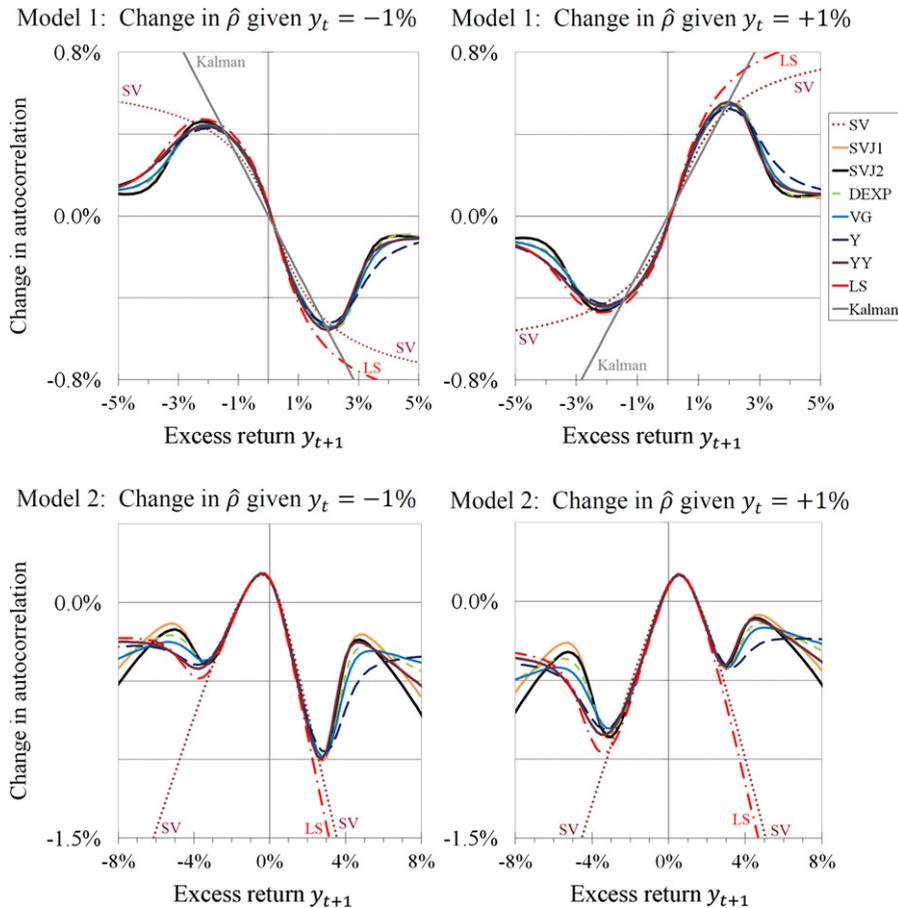


Fig. 4. Revision in autocorrelation estimate $\hat{\rho}_{t+1|t+1} - \hat{\rho}_{t|t}$ conditional upon observing excess return y_{t+1} and conditional upon $y_t = \pm 1\%$.

between 71% and 86% of the subsample improvement in log likelihood comes from using subsample stochastic volatility parameters. Between 8% and 22% of the change in log likelihood comes from using subsample jump parameters, depending on whether stochastic volatility or jump parameters are relaxed first.

The 1957–2006 subsample estimates for the YY model are even more heavily affected by the 1987 crash than are the full-sample estimates. The parameter G approaches its lower bound of zero, implying that the lower tail density is approaching the time-changed version of the infinite-variance log-stable distribution. Correspondingly, the subsample unconditional variance estimate $\hat{\theta} = (0.365)^2$ becomes substantially meaningless and cannot be compared with estimates from other models or other periods. By contrast, the estimates over 1926–1957 are strictly finite-variance. Given strong interactions between stochastic volatility and jump parameters, it is not clear which is more responsible for the strong rejections of parameter stability across subsamples.

3.5. Autocorrelation filtration

Because the prior distribution of ρ_t given past data \mathbf{Y}_t is assumed $N(\hat{\rho}_{t|t}, W_{t|t})$, it can be shown that the autocorrelation filtration algorithm of Eq. (31) for Model 1 updates

conditional moments via the robust Kalman filtration approach of Masreliez (1975):

$$\hat{\rho}_{t+1|t+1} = \hat{\rho}_{t|t} - y_t W_{t|t} \frac{\partial \ln p(y_{t+1} | \mathbf{Y}_t)}{\partial y_{t+1}} \quad (35)$$

$$W_{t+1|t+1} = W_{t|t} + \sigma_\rho^2 + (y_t W_{t|t})^2 \frac{\partial^2 \ln p(y_{t+1} | \mathbf{Y}_t)}{\partial y_{t+1}^2}. \quad (36)$$

If y_{t+1} were conditionally normal, the log density would be quadratic in y_{t+1} , and Eq. (35) would be the linear updating of standard Kalman filtration. More generally, the conditionally fat-tailed properties of y_{t+1} are explicitly recognized in the filtration.¹⁵ The partial derivatives of log densities can be computed numerically by Fourier inversion.

Fig. 4 illustrates the autocorrelation filtrations estimated under various models. For Model 1, the autocorrelation revision is fairly similar to standard Kalman filtration for observations within a $\pm 2\%$ range, which captures most observations given a unconditional daily standard deviation around 1%. However, the optimal filtration for fat-tailed distributions is to downweight

¹⁵ See Schick and Mitter (1994) for a literature review of robust Kalman filtration.

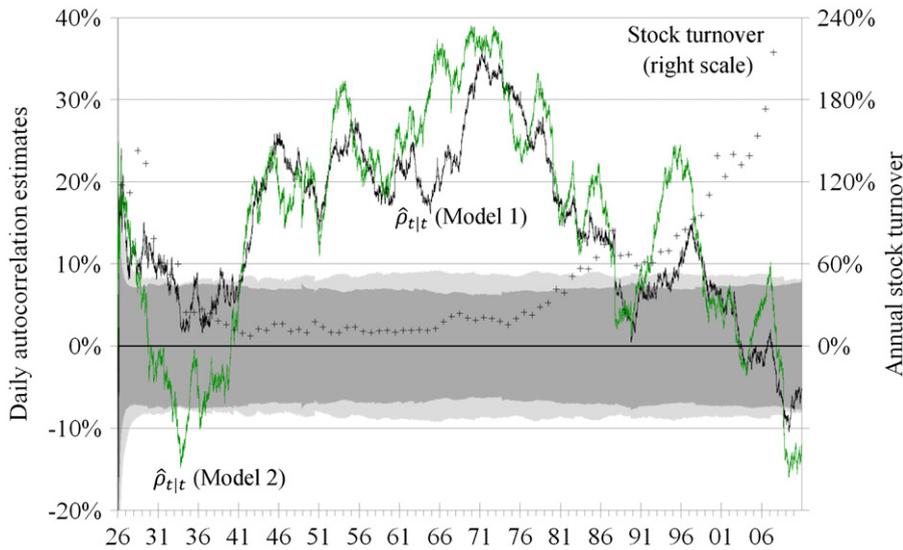


Fig. 5. Daily filtered autocorrelation estimates over 1926–2010 from the YY model, and stocks’ annual turnover (+) from French (2008). The light and dark grey areas are 95% confidence intervals for the zero autocorrelation hypothesis for Models 1 and 2, respectively.

the information from returns larger than 2% in magnitude. The exceptions are the stochastic volatility (SV) and Carr and Wu log-stable (LS) specifications. Those specifications do not particularly downweight outliers occurring in non-fat tails: in both tails for SV, in the upper tail for LS.

The autocorrelation filtration under Model 2 is different. Because $y_{t+1} = \rho_t + (1 - \rho_t)\eta_{t+1}$ in that model, large observations of y_{t+1} are attributable either to large values of $1 - \rho_t$ (small values of ρ_t), or to large values of the Lévy shocks captured by η_{t+1} . The resulting filtration illustrated in the lower panels of Fig. 4 is consequently sensitive to medium-size movements in a fashion substantially different from the Model 1 specifications.

Fig. 5 presents filtered estimates of the daily autocorrelation from the YY model. The most striking result is the extraordinarily pronounced increase in autocorrelation estimates from 1941 to 1972, reaching a peak around 35% in May 1970 for the Model 1 estimates. Estimates from both models give similar results, as do crude sample autocorrelation estimates using a one- or two-year moving window.¹⁶ Autocorrelation estimates fell steadily after 1971, and became insignificantly different from zero over 1999–2007. Filtered autocorrelation estimates appear inversely related to measures of annual stock turnover computed by French (2008), attaining values closer to zero in the higher-turnover periods before 1933 and after 1982. This inverse relationship is consistent with the standard stale-price explanation of autocorrelation in stock index returns.

More puzzling are the negative autocorrelation estimates for stock market excess returns in 2008–2010. Autocorrelation estimates from Model 2 can be affected by the overall level of stock market volatility; e.g., in the 1930s. However, negative $\hat{\rho}_{t|t}$ estimates from Model 1 are

unprecedented over 1926–2006, and are statistically significant in late 2008 and early 2009.¹⁷ Fig. 5 also indicates that estimates of daily autocorrelation are essentially nonstationary, indicating that fitting ARMA processes with time-invariant parameters to stock market excess returns is fundamentally pointless.

3.6. Volatility filtration

When returns follow an autocorrelated process with i.i.d. shocks of the form

$$y_{t+1} = \rho y_t + (1 - \rho)u_{t+1}, \quad u_{t+1} \sim (\bar{u}, \sigma_u^2), \quad (37)$$

variance can be measured in various ways.

1. Conditional or residual variance:

$$\text{Var}[y_{t+1}|y_t] = (1 - \rho)^{2m} \sigma_u^2. \quad (38)$$

2. Unconditional variance of returns:

$$\text{Var}[y_{t+1}] = \text{Var}\left[\frac{(1 - \rho)^m u_{t+1}}{1 - \rho L}\right] = \frac{(1 - \rho)^{2m}}{1 - \rho^2} \sigma_u^2, \quad (39)$$

where L is the lag operator.

3. The contribution of u_{t+1} to permanent variance:

$$\begin{aligned} & \text{Var}\left[\sum_{\tau=1}^{\infty} y_{t+\tau}|y_t\right] - \text{Var}\left[\sum_{\tau=1}^{\infty} y_{t+\tau}|y_t, u_{t+1}\right] \\ &= \text{Var}\left[\frac{(1 - \rho)^m u_{t+1}}{1 - \rho}\right] = (1 - \rho)^{2(m-1)} \sigma_u^2. \end{aligned} \quad (40)$$

These measures of variance are also approximately relevant in the above models with stochastic conditional volatility and slow-moving autocorrelations. The θ values

¹⁶ See LeBaron (1992, Fig. 1) for annual estimates of the daily autocorrelation of S&P composite index returns over 1928–1990.

¹⁷ A market with negatively autocorrelated returns is sometimes called a choppy market.

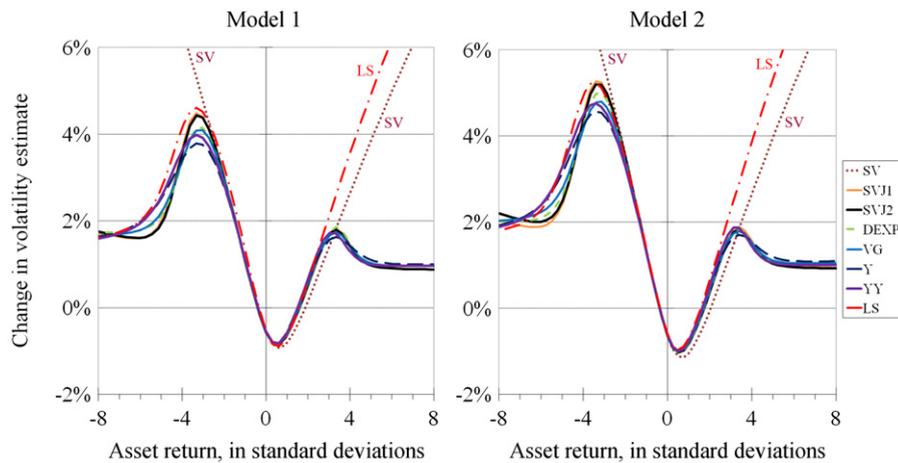


Fig. 6. News impact curves for various models. The graphs show the revision in estimated annualized standard deviation $(E_{t+1} - E_t)\sqrt{V_{t+1}}$ conditional upon observing a standardized excess return $y_{t+1}/\sqrt{\hat{V}_{t|t}/252}$.

in Panel A of Table 4 are estimates of the average level of residual variance for Model 1 ($m=0$) but are estimates of shocks' contribution to permanent variance for Model 2 ($m=1$). Furthermore, the unconditional variance of returns is an increasing function of ρ under Model 1 but a decreasing function under Model 2. The latter is more consistent with the inverse relationship between annual estimates of daily autocorrelation and volatility over 1928–1990 reported in LeBaron (1992).

The left panel of Fig. 6 illustrates how the estimated conditional volatility $E_{t+1}\sqrt{V_{t+1}}$ is updated for the various specifications under Model 1. The conditional volatility revisions use median parameter values $(\kappa_t, \nu_t) = (0.00294, 5.85)$ for the prior gamma distribution of V_t , implying a conditional mean $\kappa_t \nu_t = (0.131)^2$ that is close to the $(0.129)^2$ median value observed for $\hat{V}_{t|t}$ estimates from the YY model.¹⁸ For comparability with GARCH analyses such as Hentschel (1995), Fig. 6 shows the news impact curve, or revision in conditional volatility estimates upon observing a given excess return, using the expected volatility formula in Bates (2006, p. 932).

All news impact curves are tilted, with negative returns having a larger impact on volatility revisions than positive returns. This reflects the leverage effect, or estimated negative correlation between asset returns and volatility shocks. All specifications process the information in small-magnitude asset returns similarly. Furthermore, almost all specifications truncate the information from returns larger than three standard deviations. This was also found in Bates (2006, Fig. 1) for the SVJ1 model, indicating such truncation appears to be generally optimal for arbitrary fat-tailed Lévy processes. The SV and LS exceptions support this rule. The LS model has a fat lower tail but not an especially fat upper tail, and it truncates the volatility impact of large negative returns but not of large positive returns. The fact that volatility revisions are not monotonic in the magnitude of

asset returns is perhaps the greatest divergence of these models from GARCH models, which almost invariably specify a monotonic relationship.¹⁹ However, because moves in excess of ± 3 standard deviations are rare, all specifications generate similar volatility estimates most of the time. The volatility filtrations for Model 2 shown in the right panel of Fig. 6 using the YY model's median parameters $(\kappa_t, \nu_t) = (0.00385, 6.01)$ are qualitatively similar to those for Model 1.

Volatility filtration does appear sensitive to the data interval used in estimation, via the underlying parameter estimates. For instance, the subsample SVJ1 estimates in Table 5 yield filtered annualized volatilities that are 1.86% higher on average over 1926–1957 than the full-sample estimates, and 1.29% lower over 1957–2006. A major contributing factor is the estimate of average variance θ in Table 5, which is higher in the first than in the second subsample. A similar influence of average variance estimates upon conditional variance estimates is observed in GARCH models.²⁰

3.7. Multifactor variance models

The analysis above notes two areas where the one-factor square-root process for conditional variance appears inconsistent with the data: the subsample instability of the average variance estimate $\hat{\theta}$, and the substantial number of 3–7% stock market movements (over 1929–1935) shown in Fig. 3 that a one-factor model is unlikely to generate. More generally, graphs of volatility estimates from the various models show more persistent

¹⁹ An exception is Maheu and McCurdy (2004), who put a jump filter sensitive to outliers into a GARCH model. They find that the sensitivity of variance updating to the latest squared return should be reduced for outliers, for both stock and stock index returns.

²⁰ See, e.g., Andersen, Bollerslev, Christoffersen, and Diebold (2007, p. 521), who note that GARCH(1,1) models diverge from the RiskMetrics approach in taking into account reversion of conditional variance toward its mean. This implies that GARCH conditional variance estimates are affected by sample-specific estimates of the average variance.

¹⁸ As $\hat{V}_{t|t}$ estimates have substantial positive skewness, the median is substantially below the mean estimate of $(0.158)^2$ reported in Panel A of Table 4.

swings than are consistent with the 1½- to 2-month estimated half-life of variance shocks, as is shown below. It consequently appears worthwhile to explore multi-factor models of variance evolution.

Multiple option pricing papers explore whether multi-factor models can better capture various features of observed option prices. Taylor and Xu (1994) consider whether a two-factor model could better capture the evolving term structure of implicit variances for foreign currency options, while Carr and Wu (2007) examine two- and three-factor models of the skewness implicit in currency options. Multifactor models of stock index options include Bates (2000), Huang and Wu (2004), and Santa-Clara and Yan (2010). However, few papers estimate multifactor stochastic volatility models of stock market returns using only time series data, with the notable exception of Chernov, Gallant, Ghysels, and Tauchen (2003).

Two-factor models can be specified in various ways, depending upon which features of conditional distributions one wishes to capture. Additive approaches typically specify independently evolving determinants of the various sources of volatility: diffusion versus jump risk, for instance (Huang and Wu, 2004; Santa-Clara and Yan, 2010) or left tail versus right tail (Carr and Wu, 2007). Bates (2000) effectively estimates separate diffusion and jump-intensity loadings on two underlying and independently evolving square-root factors. Concatenated models typically have one factor being a slow-moving central tendency toward which spot variance evolves. While examples from stock index options are somewhat lacking, Egloff, Leippold, and Wu (2010) have a concatenated model of S&P 500 variance swaps. Concatenated models of currency options include Taylor and Xu (1994) and the last model in Carr and Wu (2007), while Dai and Singleton (2000) discuss additive versus concatenated models of the term structure of interest rates.

The above additive approaches model instantaneous shocks as the sum of two independent shocks that are driven by separate underlying variance factors. The resulting conditional cumulant generating function (CGF) of instantaneous shocks $ds_t = ds_{1t} + ds_{2t}$ used in Eq. (15) takes the form

$$\ln E[e^{\Phi(ds_{1t} + ds_{2t})} | V_{1t}, V_{2t}] = [\Phi\mu_0 + h_1(\Phi)V_{1t} + h_2(\Phi)V_{2t}]dt \quad (41)$$

for

$$h_i(\Phi) \equiv \mu_i\Phi + \frac{1}{2}\rho_{i,sv}^2(\Phi^2 - \Phi) + (1 - \rho_{i,sv}^2)g_{dL_i}(\Phi), \quad i = 1, 2, \quad (42)$$

with associated instantaneous conditional variance $V_t = V_{1t} + V_{2t}$. By contrast, discrete-time cumulant generating functions for concatenated models can be computed directly using the conditional CGF specification $[\Phi\mu_0 + h(\Phi)V_t]dt$ in Eq. (15), but with a more complicated concatenated two-factor model describing the evolution of V_t .

Two-factor additive and concatenated models both specify ARMA(2,1) dynamics for total variance V_t and, consequently, have similar implications for spot volatility forecasts and for expected quadratic variation when those ARMA specifications are similar. The option pricing implications for the term structures of the model-free implicit

volatilities of Britten-Jones and Neuberger (2000) are therefore also virtually identical, as are (approximately) the implications for the term structures of at-the-money implicit volatilities. Indeed, additive and concatenated approaches have substantially identical pricing implications for all options in the special case of identical $h_1(\Phi) = h_2(\Phi) = h(\Phi)$ and identical ARMA specifications. Additive models such as the papers cited above with $h_1(\Phi) \neq h_2(\Phi)$ introduce time-varying higher-moment phenomena such as stochastic skewness, for daily market returns and for the longer maturities relevant for option pricing. Such models when fitted to option prices can consequently capture evolution over time in moneyness phenomena, such as the tilt and curvature of options' volatility smirk. Concatenated models, by contrast, primarily have implications for the shape and evolution of the term structures of implicit volatilities.²¹

While two-factor additive models have been estimated using options data, I find it somewhat difficult to estimate them using only time series data.²² Consequently, I focus here upon the dynamics of total variance V_t , and estimate the following concatenated two-factor model for the overall conditional variance of intradaily shocks:

$$\begin{aligned} ds_t &= (\mu_0 + \mu_1 V_t + \mu_2 \theta_{t^*})dt + (\rho_{sv} \sqrt{V_t} dW_{1t} - \frac{1}{2} \rho_{sv}^2 V_t dt) + dL_t, \\ dV_t &= \beta(\theta_{t^*} - V_t)dt + \sigma \sqrt{V_t} dW_{1t}, \\ d\theta_t &= \beta_2(\bar{\theta} - \theta_t)dt + \sigma_2 \sqrt{\theta_t} dW_{2t}, \end{aligned} \quad (43)$$

where θ_{t^*} is the central tendency each day toward which intradaily spot variance reverts, dL_t is again a compensated Lévy increment with instantaneous variance $(1 - \rho_{sv}^2)V_t dt$, and W_{1t} and W_{2t} are independent Wiener processes. The central tendency θ_{t^*} is assumed constant intradaily for analytic tractability but resets at the end of each day to its underlying value θ_t . Given that I estimate a slow-moving process for θ_t , the specified process approximates a fully continuous-time model of the central tendency.²³ Alternately, the process for θ_{t^*} can be

²¹ Principal components analyses of index options' implicit volatility (IV) surfaces such as Skiadopoulos, Hodges, and Clewlow (1999) typically identify the first principal component as roughly a level effect for IVs of all strike prices and maturities. The secondary and tertiary principal components reflect shifts across strike prices and across maturities, but results are sensitive to the data used. The referee notes that analyses of exchange-traded options with a wide range of strike prices give heavier weight to shifts in the tilt of volatility smirks, whereas analyses of over-the-counter data with fewer strike prices but more maturities give greater weight to shifts in IV term structures.

²² In a two-factor additive variance model, one variance factor is estimated as highly volatile but with rapid mean reversion, while the other has less volatility and slower mean reversion. However, maximum likelihood estimation does not appear well behaved unless the Feller condition that ensures strictly positive variance is imposed on the first variance process. This constraint influences all parameter and latent variable estimates, making the results difficult to interpret.

²³ As discussed in Carr and Wu (2007, pp. 243–244), the generalized Fourier transform for a continuous-time central-tendency model must be evaluated numerically, which slows the AML estimation methodology considerably. One such numerical procedure involves subdividing individual days and concatenating the intradaily cumulant generating function, as described in Appendix B. Using 64 intradaily subintervals, conditional log densities for the YY model diverge on average by 10^{-7} from those computed using the process in Eq. (43), while the log likelihood estimate for 21,519 observations changes by only 0.025.

Table 6

Parameter estimates for the two-factor variance process.

Data are daily CRSP value-weighted excess returns over 1926–2006. See Eq. (43) and the notes to Table 4 for parameter definitions. All parameters are in annualized units except the half-lives $HL = 12 \ln 2/\beta$ and $HL_2 = 12 \ln 2/\beta_2$ of spot variance and central tendency shocks, respectively, which are in months. The fraction of total return variance attributable to jumps is $(1 - \rho_{sv}^2)/f_{jump}$. Standard errors are in parentheses.

	Models of Lévy shocks						
	SV	SVJ1	SVJ2	DEXP	VG	Y	YY
Conditional mean							
μ_0	0.006 (0.018)	0.047 (0.019)	0.038 (0.019)	0.051 (0.019)	0.052 (0.019)	0.047 (0.019)	0.041 (0.019)
μ_1	7.3 (1.0)	4.4 (0.9)	4.4 (1.0)	4.0 (1.0)	4.1 (0.9)	4.3 (1.0)	5.1 (0.9)
μ_2	-4.2 (0.5)	-3.0 (0.4)	-2.9 (0.4)	-2.9 (0.4)	-2.9 (0.4)	-2.9 (0.4)	-3.3 (0.5)
$\sigma_\rho \sqrt{252}$	0.074 (0.007)	0.057 (0.007)	0.056 (0.007)	0.055 (0.007)	0.057 (0.007)	0.057 (0.007)	0.057 (0.007)
V_t process							
HL	0.212 (0.015)	0.245 (0.020)	0.249 (0.020)	0.240 (0.020)	0.241 (0.020)	0.231 (0.019)	0.219 (0.018)
σ	0.766 (0.028)	0.623 (0.029)	0.622 (0.031)	0.620 (0.029)	0.617 (0.029)	0.622 (0.047)	0.637 (0.032)
ρ_{sv}	-0.684 (0.017)	-0.746 (0.019)	-0.749 (0.019)	-0.701 (0.019)	-0.703 (0.020)	-0.707 (0.042)	-0.720 (0.022)
θ_t process							
$\sqrt{\bar{\theta}}$	0.154 (0.006)	0.157 (0.006)	0.158 (0.007)	0.157 (0.007)	0.157 (0.007)	0.160 (0.010)	0.157 (0.007)
HL_2	10.3 (1.8)	11.0 (2.1)	11.0 (2.1)	11.5 (2.2)	11.5 (2.3)	10.7 (1.3)	10.7 (1.8)
σ_2	0.191 (0.014)	0.211 (0.017)	0.210 (0.018)	0.208 (0.017)	0.210 (0.017)	0.221 (0.012)	0.211 (0.015)
CGMY parameters							
w_n				0.43 (0.06)	0.43 (0.06)	0.40 (0.08)	0.13 (0.02)
G				46.8 (5.5)	29.7 (5.2)	1.4 (4.8)	37.9 (14.3)
M				62.2 (10.5)	41.9 (9.3)	6.5 (7.8)	4.9 (3.6)
Y_n				-1	0	1.84 (0.04)	-0.38 (0.98)
Y_p				-1	0	1.84 (0.04)	1.90 (0.02)
Merton parameters							
λ_1, λ_2		121.0 (20.6)	128.3 (24.8), 0.69 (1.20)				
$\bar{\gamma}_1, \bar{\gamma}_2$		0.003 (0.003)	0.003 (0.002), -0.178 (0.020)				
δ_1, δ_2		0.033 (0.002)	0.030 (0.002), 0.011 (0.016)				
f_{jump}	0	0.130 (0.015)	0.139 (0.033)	0.281 (0.032)	0.319 (0.035)	1	1
$\ln L$	75,166.90	75,242.88	75,247.50	75,245.88	75,246.17	75,242.29	75,245.21

interpreted as a model of (stationary) parameter drift in the unconditional mean θ of a one-factor variance process, similar in spirit to the models of parameter drift in the daily autocorrelation parameter in Eqs. (23) and (24).

The multifactor generalized Fourier transform for future $(s_{t+1} - s_t, V_{t+1}, \theta_{t+1})$ conditional upon current (V_t, θ_t) is exponentially affine in those state variables and is given in Eq. (B.1) in Appendix B. Modifying this transform for autocorrelated returns under Models 1 and 2 is a straightforward extension of Eqs. (26) and (27), while recursive filtration of (V_t, θ_t) is also a straightforward extension of the approach used for one-factor variance processes. I report below only results for Model 2, using a filtration procedure outlined at the end of Appendix A.

Table 6 contains parameter estimates for the various Lévy specifications. The major change relative to the one-factor variance estimates in Table 4 is, unsurprisingly, in the variance dynamics—which are ARMA(2,1), as opposed to the AR(1) dynamics of the one-factor variance model in Eq. (14). Spot variance shocks are highly correlated with market returns ($\rho_{sv} \approx -0.7$), but V_t mean-reverts rapidly to its central tendency θ_t , with an estimated half-life of roughly one week. The central tendency, by contrast, is less volatile and more persistent, with an estimated half-life of almost a year. Estimates of the volatility process are virtually identical for most models. The exceptions are

again the SV and LS models, which have more volatile spot variance processes given greater sensitivity to return outliers in the upper tail (LS) or both tails (SV).

The conditional mean of stock market returns in Table 6 appears sensitive to the variance divergence $V_t - \theta_t$, not to the level of spot variance V_t observed in Panel A of Table 4. This is not necessarily evidence of return predictability, however. The conditional mean parameters (μ_0, μ_1, μ_2) influence the filtration algorithm for estimating the latent state variables (ρ_t, V_t, θ_t), as well as directly affecting conditional densities. Jump parameter estimates in Table 6 are broadly similar to those in Panel B of Table 4, except for the YY model.

Fig. 7 presents the filtered estimates of annualized volatilities from the best-fitting SVJ2 model over 1926–2010, as well as the associated annualized conditional volatilities.²⁴ The graphs from other fat-tailed models are

²⁴ Annualized volatility refers to the choice of units. V_t is variance per year, while the daily volatility estimate for a Wednesday return with an estimated length of 1/260 years (from Table 3) is approximately $E_t \sqrt{V_t/260}$. Because variance mean-reverts, it is not appropriate to interpret spot volatility estimates in Fig. 7 as showing the volatility estimate for a one-year investment. However, the estimated central tendency $E_t \sqrt{\theta_t}$ is roughly the predicted volatility for a one-year investment.

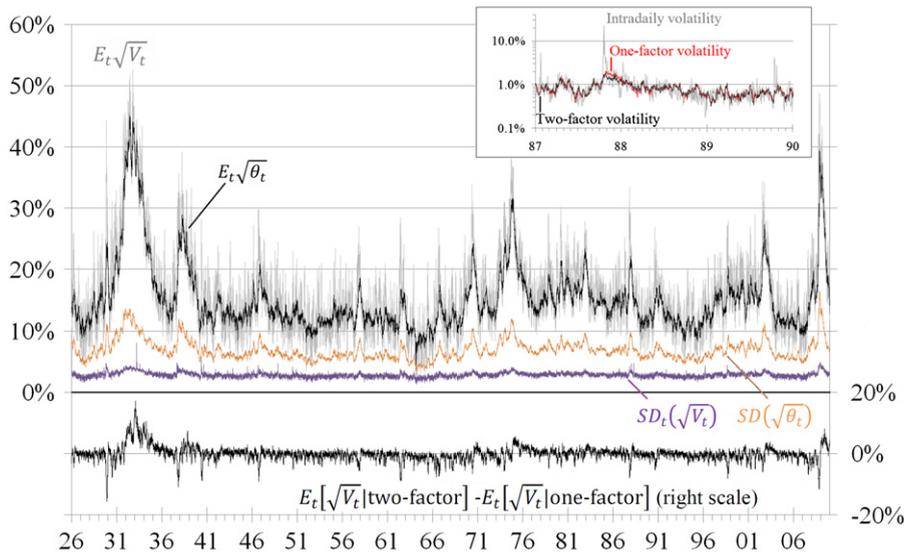


Fig. 7. Filtered annualized spot volatility and central tendency estimates over 1926–2010, associated conditional standard deviations of volatility estimates, and comparisons with one-factor estimates and with intradaily realized volatility. V_t is the annualized instantaneous (spot) variance, while θ_t is the central tendency toward which it reverts. Filtered estimates and standard deviations are from the SVJ2 two-factor variance Model 2, based upon CRSP value-weighted excess returns up to the latest date. The lower panel shows the difference between spot volatility estimates from two- and one-factor models. The inset shows daily realized volatilities of 15-minute S&P 500 futures returns over 1987–1989 on a log scale, and filtered forecasts of those realized volatilities from the one- and two-factor volatility models.

similar, apart from the above-noted sensitivity of SV and LS volatility estimates to return outliers.²⁵ The graph highlights the longer-term evolution of underlying volatility $\sqrt{\theta_t}$; in particular, the persistently turbulent market conditions of the 1930s, unmatched by any period in the post-1945 era. Spot volatility estimates $E_t\sqrt{V_t}$ oscillate around the $\sqrt{\theta_t}$ estimates, at a frequency too high to be discernable in this 85-year graph. The conditional standard deviations of spot volatilities are on average about 2.9%, indicating a 95% confidence interval of roughly $\pm 5.7\%$ in the annualized volatility estimates. Underlying volatility $\sqrt{\theta_t}$ is much less precisely estimated than spot volatility, with an average conditional standard deviation of 6.9%.

The lower graph in Fig. 7 shows that spot volatility estimates are generally similar for the one- and two-factor variance specifications, apart from substantial divergences in the 1930s. Spot volatility estimates from both the one- and two-factor specifications averaged 15.9% over the full 1926–2006 period. Two-factor filtration does, however, yield more accurate volatility estimation: a 2.9% conditional standard deviation of spot volatility on average, as opposed to 3.2% for the one-factor filtration.

The inset to Fig. 7 compares adjusted filtered volatility estimates $0.8655E_t\sqrt{V_t}$ over 1987–1989 with subsequent realized volatilities computed daily from intradaily 15-minute log-differenced S&P 500 futures prices. (Open-to-close futures returns were 86.55% as volatile as close-to-close futures returns over 1982–2001.) The inset shows that the one- and two-factor variance estimates using daily data

generally track and predict realized intradaily volatility closely—except for the occasional volatility spikes, which are how major daily outliers manifest themselves in intradaily data. The cluster of high intradaily volatility values over October 19–28, 1987 is inconsistent with the models' diffusive-volatility assumption and is the major time series evidence supporting an alternate volatility-jump specification.²⁶

4. Option pricing implications

Do these alternative models imply different option prices? Exploring this issue requires identifying the appropriate pricing of equity, jump, and stochastic volatility risks. Furthermore, the presence of substantial and stochastic autocorrelation raises issues not previously considered when pricing options. In particular, the observed stock index level underlying option prices can be stale, while the relevant volatility measure over the option's lifetime is also affected. The variance of the sum of future stock market returns is not the sum of the variances when returns are autocorrelated.

4.1. Equilibrium stock index and futures processes

To examine these issues, I focus upon Model 2, with its interpretation in Eqs. (24)–(25) of ds_t as the permanent shock to the log stock market level. Furthermore, I address

²⁵ See Bates (2006, Fig. 4) for the divergence between SV and SVJ1 volatility estimates from a one-factor variance model.

²⁶ Johannes, Polson, and Stroud (2009, p. 2788) find that the superior fit of a volatility-jump model is entirely attributable to the increase in log likelihood on days shortly after the 1987 crash.

the potential impact of autocorrelations upon option prices by examining prices of options on S&P 500 futures. I assume that futures prices respond instantaneously and fully to the arrival of news, whereas lack of trading in the underlying stocks can delay the incorporation of that news into the reported S&P 500 stock index levels. I also assume that index arbitrageurs effectively eliminate any stale prices in the component stocks on days when futures contracts expire, so that stale stock prices do not affect the cash settlement of stock index futures. MacKinlay and Ramaswamy (1988) provide evidence supportive of both assumptions.

These assumptions have the following implications under Model 2 [Eqs. (24) and (25)].

1. The observed futures price F_t underlying options on S&P 500 futures is not stale.
2. Log futures price innovations equal the permanent innovations ds_t :

$$d \ln F_t = ds_t. \tag{44}$$

Consequently, European options on stock index futures can be priced directly using risk-neutral versions of Eqs. (14) or (43)—which are affine, simplifying option evaluation considerably. Furthermore, option prices do not depend upon ρ_t , except indirectly through the impact of autocorrelation filtration upon the filtration of the latent variance state variables V_t and θ_t .

A risk-neutral version of the s_t process for use in pricing options can be derived using the myopic power utility pricing kernel of Bates (2006) to price the various risks:

$$d \ln M_t = \mu_m^* dt - R ds_t. \tag{45}$$

This pricing kernel constrains both the equity premium estimated under the objective time series model and the transformation of the s_t process into the risk-neutral process appropriate for pricing options. In particular, the instantaneous equity premium for one-factor variance models is

$$(\mu_0 + \mu_1 V_t) dt = -E_t(e^{ds_t} - 1)(e^{-R ds_t} - 1), \tag{46}$$

which implies

$$\begin{aligned} \mu_0 = 0, \mu_1 V_t = R \left[\rho_{sv}^2 + (1 - \rho_{sv}^2)(1 - f_{jump}) \right] V_t \\ - \int_{-\infty}^{\infty} (e^x - 1)(e^{-Rx} - 1)k(x)dx \approx RV_t. \end{aligned} \tag{47}$$

The approximation in Eq. (47) follows from first-order Taylor expansions of the exponential terms and from the fact that jumps account for a fraction $f_{jump}(1 - \rho_{sv}^2)$ of overall variance V_t . For two-factor variance models, the assumption of nonsystematic θ_t risk implies $\mu_2 = 0$.

The equity premium in Eq. (47) is well defined for the SVJ1 and SVJ2 models. For the CGMY models, the restriction $G \geq R$ is required for a finite equity premium; the intensity of downward jumps must fall off faster than an investor's risk aversion to such jumps. The log-stable

process is inconsistent with a finite equity premium for $R > 0$.²⁷

The key risk aversion parameter R used for change of probability measure was estimated by imposing the equity premium restrictions in Eq. (47) and re-estimating all times series models. The additional parameter restriction $G \geq R$ was imposed upon all CGMY models and was binding for the one-factor variance YY model.²⁸ Parameter estimates reported in Panel A of Table 7 for the one-factor variance models changed little relative to those in Table 4, while risk aversion was estimated at roughly 2.5 for all models. The restriction of a purely variance-sensitive equity premium ($\mu_0 = 0$) was not rejected for any one-factor variance model.

By contrast, the equity premium restriction $\mu_0 = \mu_2 = 0$ was strongly rejected for all two-factor variance models. These rejections were examined further by estimating and restricting two separate sets of (μ_0, μ_1, μ_2) parameters: one set for generating filtered conditional distributions of the latent state variables (ρ_t, V_t, θ_t) , the other for evaluating the conditional means and densities of excess returns y_{t+1} given those conditional distributions. The reduction in log likelihood from imposing the equity premium restriction is entirely attributable to poorer filtration of state variables, not to reduced predictability of future stock market returns.²⁹

4.2. Option pricing

The pricing kernel also determines the risk-neutral cumulant generating function (CGF) for pricing European options on stock index futures over a multiday time gap $\tau_{t,T}$. For one-factor variance processes, the risk-neutral CGF conditional upon knowing the objective variance V_t is

$$\begin{aligned} C_{t,T}^*(\Phi) + D_{t,T}^*(\Phi)V_t = \ln \left\{ \frac{E[e^{\ln(M_T/M_t)} e^{\Phi(s_T - s_t)} | V_t]}{E[e^{\ln(M_T/M_t)} | V_t]} \right\} \\ = C(\tau_{t,T}; \Phi - R, 0) - C(\tau_{t,T}; -R, 0) + [D(\tau_{t,T}; \Phi - R, 0) \\ - D(\tau_{t,T}; -R, 0)]V_t, \end{aligned} \tag{48}$$

where $C(\cdot)$ and $D(\cdot)$ are defined in Eqs. (18)–(22) and $\tau_{t,T}$ is the sum of the relevant future daily intervals. Two-factor variance processes have a similar risk-neutral CGF:

$$\begin{aligned} C_{t,T}^*(\Phi) + D_{t,T}^*(\Phi)V_t + DD_{t,T}^*(\Phi)\theta_t \\ = CC_{t,T}(\Phi - R, 0, 0) - CC_{t,T}(-R, 0, 0) + [D_{t,T}(\Phi - R, 0) \\ - D_{t,T}(-R, 0)]V_t + [DD_{t,T}(\Phi - R, 0, 0) - DD_{t,T}(-R, 0, 0)]\theta_t, \end{aligned} \tag{49}$$

²⁷ Carr and Wu (2003) specify a log-stable process for the risk-neutral process underlying option prices. This can be generated from a CGMY process for the actual process with $w_n = 0$ and $G = R$.

²⁸ Wu (2006) proposes an alternate pricing kernel with negative risk aversion for downside risk, thereby automatically imposing $G \geq R$.

²⁹ For instance, the log likelihood of the unconstrained SVJ1 model on spliced CRSP/S&P 500 data was 74,259, which rose slightly to 74,266 when two separate sets of conditional mean parameters were estimated. Imposing the equity premium constraints only on the second set of conditional mean parameters reduced log likelihood to 74,258. Imposing the constraint also on the conditional mean parameters that affect filtration reduced log likelihood substantially, to 74,224.

Table 7

Parameter estimates on spliced CRSP/S&P 500 data over 1926–2006, with constrained equity premium.

Data are daily CRSP value-weighted excess returns before March 5, 1957, and S&P 500 excess returns thereafter. See the notes to Tables 4 and 6 for parameter definitions of the one- and two-factor variance models, respectively.

Estimates of the one-factor variance models in Panel A impose the conditional mean restriction $\mu_0=0$; μ_1 is roughly equal to the risk aversion parameter R . Standard errors of parameter estimates could not be computed for the YY model in Panel A because the parameter constraint $G \leq R$ is binding. Estimates of the two-factor variance models in Panel B impose the restrictions $(\mu_0, \mu_2)=0$, a restriction that was rejected at significance levels of less than 10^{-16} for all models.

Panel A: Parameter estimates for one-factor variance models							
	Models of Lévy shocks						
	SV	SVJ1	SVJ2	DEXP	VG	Y	YY
Conditional mean							
R	2.49 (0.62)	2.44 (0.61)	2.43 (0.57)	2.44 (0.61)	2.50 (0.61)	2.42 (0.61)	2.38
$\sigma_\rho \sqrt{252}$	0.043 (0.005)	0.038 (0.005)	0.037 (0.005)	0.037 (0.005)	0.037 (0.005)	0.037 (0.005)	0.037
$H_0: \mu_0=0$ (p -value)	0.383	0.265	0.507	0.247	0.206	0.185	0.212
V_t process							
$\sqrt{\bar{\theta}}$	0.172 (0.004)	0.173 (0.004)	0.174 (0.004)	0.174 (0.004)	0.174 (0.004)	0.174 (0.006)	0.175
HL	1.2 (0.1)	1.4 (0.1)	1.4 (0.1)	1.5 (0.1)	1.5 (0.1)	1.6 (0.1)	1.6
σ	0.534 (0.014)	0.448 (0.015)	0.449 (0.015)	0.444 (0.015)	0.441 (0.015)	0.427 (0.018)	0.429
ρ_{sv}	-0.649 (0.016)	-0.678 (0.017)	-0.679 (0.017)	-0.632 (0.017)	-0.631 (0.017)	-0.623 (0.022)	-0.621
CGMY parameters							
w_n				0.55 (0.06)	0.53 (0.07)	0.58 (0.05)	0.90
G				51.6 (5.0)	35.9 (4.6)	5.4 (4.0)	2.4
M				53.9 (12.7)	34.4 (11.1)	4.5 (8.7)	64.3
Y_n				-1	0	1.87 (0.03)	1.94
Y_p				-1	0	1.87 (0.03)	-1.29
Merton parameters							
λ_1, λ_2		108.7 (18.2)	122.7 (23.1), 0.43 (0.38)				
\bar{v}_1, \bar{v}_2		-0.001 (0.003)	0.000 (0.002), -0.219 (0.027)				
δ_1, δ_2		0.034 (0.002)	0.031 (0.002), 0.003 (0.150)				
f_{jump}	0	0.126 (0.015)	0.138 (0.021)	0.228 (0.026)	0.247 (0.028)	1	1
$\ln L$	74,028.53	74,119.26	74,125.33	74,121.73	74,122.51	74,122.19	74,124.33
Panel B: Parameter estimates for two-factor variance process							
	Models of Lévy shocks						
	SV	SVJ1	SVJ2	DEXP	VG	Y	YY
Conditional mean							
R	2.53 (0.63)	2.46 (0.62)	2.45 (0.65)	2.45 (0.62)	2.44 (0.62)	2.38 (0.64)	2.43 (0.62)
$\sigma_\rho \sqrt{252}$	0.060 (0.006)	0.051 (0.006)	0.051 (0.006)	0.051 (0.006)	0.050 (0.006)	0.050 (0.006)	0.051 (0.006)
V_t process							
HL	0.330 (0.020)	0.369 (0.025)	0.365 (0.025)	0.358 (0.024)	0.357 (0.025)	0.354 (0.025)	0.342 (0.024)
σ	0.662 (0.021)	0.557 (0.023)	0.560 (0.025)	0.558 (0.023)	0.556 (0.024)	0.553 (0.028)	0.560 (0.026)
ρ_{sv}	-0.664 (0.017)	-0.738 (0.019)	-0.743 (0.019)	-0.693 (0.019)	-0.695 (0.019)	-0.699 (0.026)	-0.706 (0.023)
θ_t process							
$\sqrt{\bar{\theta}}$	0.164 (0.008)	0.165 (0.007)	0.166 (0.008)	0.165 (0.007)	0.165 (0.007)	0.166 (0.008)	0.166 (0.008)
HL_2	23.8 (7.1)	19.3 (5.0)	19.2 (4.9)	18.9 (4.8)	18.7 (4.8)	18.4 (4.6)	18.2 (4.6)
σ_2	0.154 (0.019)	0.177 (0.022)	0.177 (0.022)	0.179 (0.022)	0.180 (0.022)	0.185 (0.023)	0.183 (0.022)
CGMY parameters							
w_n				0.44 (0.05)	0.44 (0.05)	0.39 (0.05)	0.16 (0.03)
G				47.5 (5.0)	30.4 (4.7)	2.7 (4.3)	24.9 (10.7)
M				69.8 (10.5)	47.8 (9.2)	10.2 (7.7)	7.1 (7.3)
Y_n				-1	0	1.82 (0.04)	0.46 (0.74)
Y_p				-1	0	1.82 (0.04)	1.89 (0.03)
Merton parameters							
λ_1, λ_2		143.0 (24.0)	173.3 (33.2), 0.50 (0.82)				
\bar{v}_1, \bar{v}_2		0.002 (0.002)	0.003 (0.002), -0.190 (0.020)				
δ_1, δ_2		0.030 (0.005)	0.026 (0.002), 0.023 (0.027)				
f_{jump}	0	0.134 (0.015)	0.140 (0.028)	0.288 (0.031)	0.319 (0.034)	1	1
$\ln L$	74,140.11	74,224.13	74,231.22	74,228.56	74,229.37	74,227.79	74,229.69

where $CC_{t,T}(\cdot)$, $D_{t,T}(\cdot)$, and $DD_{t,T}(\cdot)$ are computed via a multi-day recursive procedure given in Appendix B, using the estimated daily intervals over the lifetime of the options.

Following Bates (2006), European call prices on an S&P 500 futures contract with exercise price X are evaluated at their filtered values:

$$\hat{c}(F_t, X; t, T | \mathbf{Y}_t) = e^{-r_t T} F_t - e^{-r_t T} X \times \left\{ \frac{1}{2} + \frac{1}{\pi} \operatorname{Re} \left[\int_0^\infty \frac{e^{C_{t,T}^*(i\Phi) + g_{t,T}[D_{t,T}^*(i\Phi)] + g_{t,T}^0[DD_{t,T}^*(i\Phi)] - i\Phi \ln(X/F_t)}}{i\Phi(1-i\Phi)} d\Phi \right] \right\}, \quad (50)$$

where $\operatorname{Re}[c]$ is the real component of complex-valued c ; $T_t = n_t/365$ is the maturity associated with the continuously compounded Treasury bill yield r_t , given n_t calendar days until option maturity; and $\{g_{t,T}(\psi), g_{t,T}^0(\psi_2)\}$ are the filtered gamma cumulant generating functions that summarize what is currently known about (V_t, θ_t) given past data \mathbf{Y}_t . (The $g_{t,T}^0$ term is omitted for one-factor variance models.) The associated annualized implicit volatilities IV_t from option prices with N_t business days until maturity are then computed using the French (1984) business-day approach of a per period volatility of $IV_t \sqrt{N_t/252}$, given the evidence in Table 3 that business days are roughly comparable in length.

Implicit volatilities from the one- and two-factor models are graphed in Fig. 8 and are compared with observed IVs from settlement prices for American options on S&P 500 futures with nonzero trading volume on January 3, 2007. Fig. 8 also shows 95% confidence intervals for the SVJ2 estimates, computed in three ways: based on parameter uncertainty alone, based on parameter and V_t state uncertainty, and (for the two-factor model) based on parameter and (V_t, θ_t) state uncertainty.

Panels A and B of Fig. 8 show the estimated and observed volatility smirks from the one- and two-factor variance models, respectively, for January 2007 options with 16 days (11 business days) to maturity. The most striking result is that all Lévy specifications – including the SV model – generate virtually identical option prices and IVs over a range of ± 2 standard deviations, a range that contains the most actively traded options. The estimated at-the-money IVs were virtually identical across all specifications on January 3, 2007, reflecting no recent major return outliers that would induce different volatility estimates from different specifications. Some divergences emerge between one- and two-factor estimates, but both show a substantial volatility smirk. Given the similarity to SV estimates, the tilt of the estimated volatility smirk for near-the-money options appears to be driven primarily by the correlation between shocks to variance and stock market returns. Only for deep

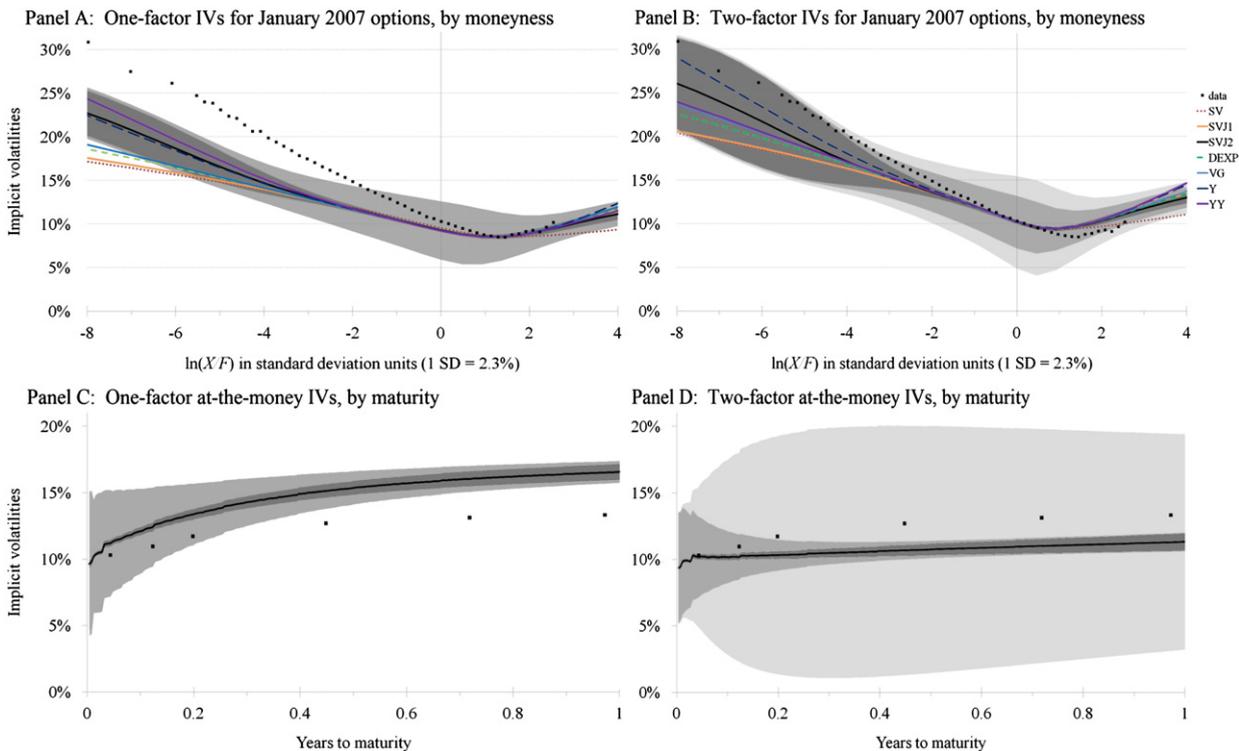


Fig. 8. Estimated and observed implicit volatilities (IVs) for options on S&P 500 futures on January 3, 2007, by moneyness and maturity. The IV data (■) are from options' settlement prices. Estimated IVs (lines) are from various models' filtered estimates of end-of-day option prices, based upon daily spliced CRSP/S&P 500 excess returns from 1926 to January 3, 2007. The moneyness graphs (Panels A and B) are for January 2007 options with 16 days (11 business days) to maturity. At-the-money IVs by maturity (Panels C and D) are shown for the SVJ2 model only; estimates from other models are nearly identical. 95% confidence intervals for the SVJ2 estimates are shown in dark grey for parameter uncertainty, medium grey for combined parameter and V_t uncertainty, and light grey for combined parameter and (V_t, θ_t) uncertainty.

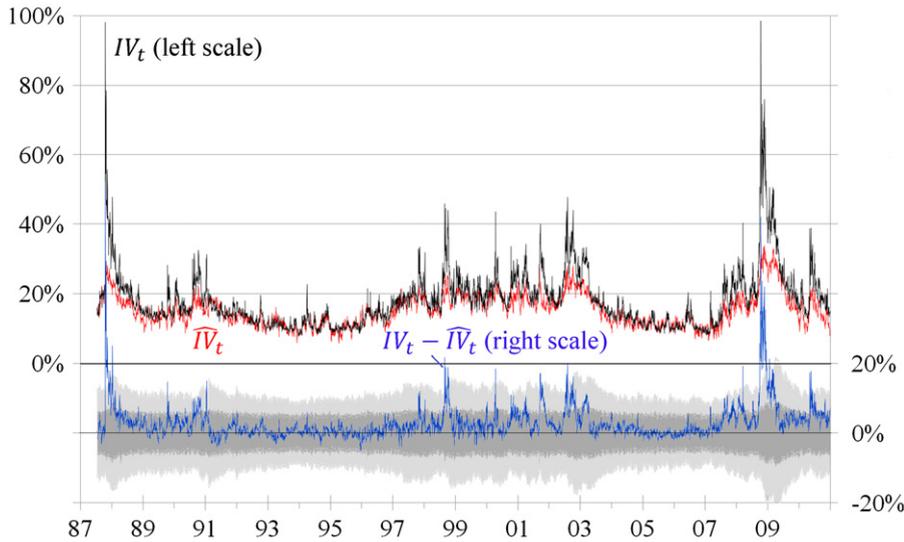


Fig. 9. Estimated and observed at-the-money implicit volatilities (IVs), 1987–2010. Observed IV_t is from settlement prices on short-term options on S&P 500 futures with at least four business days to maturity. The estimated \widehat{IV}_t is from the two-factor SVJ2 model's filtered estimates of option prices, using spliced CRSP/S&P 500 daily excess returns from 1926 to date t . The lower panel shows the difference between the two. The dark and light grey areas identify the 95% confidence intervals for the volatility gap $IV_t - \widehat{IV}_t$ for one- and five-week options respectively, given parameter and state uncertainty.

out-of-the-money put options do the divergences in estimated tail properties generate substantially different IV patterns across different Lévy specifications.³⁰ The near-the-money IV estimates have substantial state uncertainty, given the difficulty of identifying current variance V_t (and central tendency θ_t for the two-factor model) from daily returns. By contrast, the confidence intervals for deep out-of-the-money put options' IV estimates are driven primarily by parameter uncertainty, given the difficulty of identifying the precise distribution of extreme events even with 81 years of daily data.

Panels C and D of Fig. 8 show at-the-money IVs by maturity for the SVJ2 model. (Estimates from other Lévy specifications are nearly identical.) The two graphs show that the one- and two-factor models diverge strongly by maturity. In the one-factor model, risk-neutral spot variance V_t^* is projected to mean-revert within months to its unconditional level of $(0.187)^2$, which is tightly estimated given 81 years of data. Consequently, state uncertainty matters for short-maturity IVs, but not for longer maturities. In the two-factor concatenated model, by contrast, V_t^* is projected to revert within weeks to a slowly moving risk-neutral central tendency θ_t^* that is imprecisely estimated. The uncertainty regarding where V_t^* is headed generates wide confidence intervals for filtered at-the-money IV estimates for options of one month to maturity or longer.

Fig. 9 chronicles estimated and observed at-the-money IVs over July 1987 through December 2010 for the short-term options with maturities of six days (four business

days) or more, beginning shortly after the introduction of serial (monthly) options in July 1987. The options' maturities begin at 22–29 business days (31–39 calendar days) and typically shrink to five business days, before being replaced by a longer maturity.³¹ Observed and estimated IVs are generally close during periods of low volatility but diverge substantially during high-volatility periods—especially after the 1987 stock market crash and at the height of the financial crisis in 2008. The divergences are sometimes large enough to be statistically significant for options with only a week (five business days) to maturity. By contrast, it is rare that the divergences in IVs for five-week (25 business day) options are statistically significant, given larger standard errors attributable to state uncertainty about the central tendency.

Table 8 compares the estimated term structures of at-the-money IVs over 1988 through 2006 with observed IVs, using values linearly interpolated or extrapolated to standardized monthly maturities. The table also reports properties of some risk-neutral volatility measures that influence the estimated term structures: the filtered spot volatility $E_t \sqrt{V_t^*}$, the filtered central tendency volatilities $E_t \sqrt{\theta_t^*}$, and the unconditional volatility $\sqrt{E(V_t^*)}$. These volatilities are upwardly scaled versions of the corresponding objective volatilities and affect short-, medium-, and long-term IVs, respectively. For the SVJ2 model used in Table 8, the scaling factors are

$$\frac{V_t^*}{V_t} = \frac{1 + \sum_{i=1}^2 \lambda_i e^{-R\bar{\gamma}_i + \frac{1}{2}R^2 \delta_i^2} [(\bar{\gamma}_i - R\delta_i^2)^2 + \delta_i^2]}{1 + \sum_{i=1}^2 \lambda_i (\bar{\gamma}_i^2 + \delta_i^2)} \equiv R^*$$

³⁰ These results differ from Bates (2000, Table 3), who reports substantial divergences between the SV and SVJ models based substantially upon implicit parameter estimation. It would appear that those estimates are strongly influenced by the prices of deep out-of-the-money options.

³¹ Good Friday occurred on the third Friday of the month in 1992, 2000, 2003, and 2008, shortening corresponding serial options' maturities by one day.

Table 8

Statistical properties of risk-neutral volatilities and implicit volatilities (IVs) over 1988–2006, in percent.

Observed IVs are from options on S&P 500 futures, while filtered volatilities and \widehat{IV} s are from the one- and two-factor SVJ2 models estimated on spliced CRSP/S&P 500 daily excess returns over the specified intervals. Observed IVs and model-specific fitted \widehat{IV} s for fixed maturities were linearly interpolated from values at observed maturities. Out of 4,789 observations, 330 (36) of the 21-day (126-day) IVs were linearly extrapolated.

Asterisks (*) for the average divergences between observed and fitted IVs indicate statistical significance at the 1% level, based upon Newey-West standard errors computed using 1,008 daily lags. The standard errors range from 0.6% to 1.0% for the one-factor estimates and are 0.6% for the two-factor estimates. The R^2 for each fitted \widehat{IV}_t series was computed by $R^2 = 1 - \text{MSE}(IV_t - \widehat{IV}_t) / \text{Var}(IV_t)$.

Statistic	Estimation interval	Filtered volatilities			Implicit volatilities by maturity (number of business days)					
		$E_t \sqrt{V_t^*}$	$E_t \sqrt{\theta_t^*}$	$\sqrt{E[V_t^*]}$	21	42	63	84	105	126
Average values										
Options' IVs	1926–2006				16.6	17.1	17.4	17.5	17.5	17.6
One-factor \widehat{IV} s	1926–2006	15.4		18.7	15.5	16.0	16.5	16.6	16.7	16.8
One-factor \widehat{IV} s	1957–2006	14.2		16.3	14.1	14.4	14.7	14.7	14.8	14.9
Two-factor \widehat{IV} s	1926–2006	15.0	14.5	17.0	14.7	14.7	14.7	14.7	14.7	14.7
Average $(IV_t - \widehat{IV}_t)$										
One-factor	1926–2006				1.1	1.0	0.9	0.9	0.8	0.8
One-factor	1957–2006				2.5*	2.7*	2.7*	2.7*	2.7*	2.7*
Two-factor	1926–2006				1.9*	2.3*	2.7*	2.8*	2.8*	2.9*
Standard deviations										
Options' IVs	1926–2006				5.8	5.5	5.1	5.1	5.0	4.8
One-factor \widehat{IV} s	1926–2006	5.1			3.9	3.3	2.4	2.4	2.3	1.9
One-factor \widehat{IV} s	1957–2006	4.8			3.9	3.5	2.8	2.7	2.6	2.3
Two-factor \widehat{IV} s	1926–2006	4.8	3.0		3.7	3.3	3.0	3.0	2.9	2.8
RMSE $(IV_t - \widehat{IV}_t)$										
One-factor	1926–2006				3.2	3.3	3.5	3.5	3.5	3.6
One-factor	1957–2006				3.0	3.0	3.1	3.2	3.1	3.2
Two-factor	1926–2006				2.9	2.9	2.8	2.8	2.7	2.7
R^2										
One-factor	1926–2006				70	64	53	53	51	44
One-factor	1957–2006				73	70	63	61	62	56
Two-factor	1926–2006				75	72	70	70	71	68

$$\frac{\theta_t^*}{\bar{\theta}_t} = R^* \frac{\beta}{\beta + R\rho_{sv}\sigma}. \quad (51)$$

Implicit volatilities from the one-factor SVJ2 model estimated from the full 1926–2006 data set appear to have been substantially affected by the high-volatility conditions in the 1930s. IVs are higher at all maturities than are IVs based on SVJ2 parameter estimates over 1957–2006, while convergence toward the higher unconditional volatility at longer maturities imparts a more pronounced upward slope on average to the term structures. The two-factor variance model by contrast has flat term structures on average, because estimates of risk-neutral medium-term volatility $\sqrt{\theta_t^*}$ are on average close to the spot volatility estimates.

While it can be difficult to identify ex ante whether an estimated (filtered) IV gap on any given day is statistically significant, the average IV gaps over 1988 through 2006 are statistically different from zero ex post at all maturities for two of the estimated models: the one-factor SVJ2 model estimated on stock market returns over 1957–2006 (2.5–2.7% gaps), and the two-factor SVJ2 model estimated on returns over 1926–2006 (1.9–2.7% gaps).³² By contrast,

the smaller 0.8–1.1% gaps between observed and estimated IVs at various maturities for the full-sample one-factor variance model are not statistically significant. Despite substantial bias, the two-factor model tracks the term structure of at-the-money IVs more closely than do the two estimates of the one-factor model, with lower root mean squared errors and higher R^2 s at all maturities and especially at longer maturities.

Overall, the above results appear broadly compatible with previous studies that have compared IVs with subsequently realized volatility over the options' lifetime. Observed IVs from options on index futures do appear higher on average over the post-1987 period than is justified by risk-adjusted valuations based upon time series analysis. Furthermore, time series plots indicate the gap is especially pronounced during periods of high volatility.

4.3. The 2007–2008 financial crisis

The impact of the financial crisis upon index options can also be assessed by comparing model-specific estimates of the Volatility Index (VIX) with historical VIX data available at <http://www.cboe.com/micro/vix/historical.aspx>, on the website of the Chicago Board Options Exchange. The estimated VIX is computed by evaluating the log contract synthesized by the VIX portfolio of

³² The significance tests use Newey-West standard errors with 1,008 daily lags (four years) to adjust for substantial positive autocorrelations in IV gaps present even at yearly intervals.

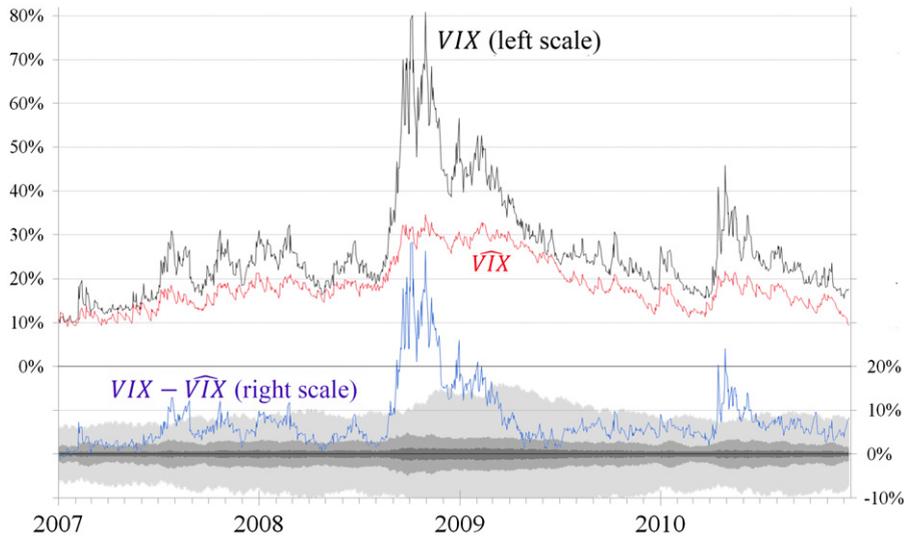


Fig. 10. Estimated and observed VIX, 2007–10. VIX is the volatility index computed from S&P 500 options by the Chicago Board Options Exchange, converted to a percentage. \widehat{VIX} is the filtered out-of-sample estimate based on two-factor SVJ2 parameter estimates over 1926–2006 and on subsequent S&P 500 excess returns. The lower panel shows the difference between the two. 95% confidence intervals are shown for the difference between observed and estimated VIX based upon parameter uncertainty (dark grey), combined parameter and V_t uncertainty (medium grey), and combined parameter and (V_t, θ_t) uncertainty (light grey).

options³³ :

$$\begin{aligned} \widehat{VIX}_t^2 &= -\frac{2}{\tau_{t,T}} E\{E_t^*[\ln(F_T/F_t)|V_t, \theta_t] | \mathbf{Y}_t\} \\ &= -2 \frac{252}{21} \frac{\partial [C_{t,T}^*(\Phi) + D_{t,T}^*(\Phi) \hat{V}_{t|t} + DD_{t,T}^*(\Phi) \hat{\theta}_{t|t}]}{\partial \Phi} \Big|_{\Phi=0}, \end{aligned} \quad (52)$$

where $C_{t,T}^*$, $D_{t,T}^*$, and $DD_{t,T}^*$ are computed from Eqs. (48) and (B.6) using a maturity of 21 business days. The state variable estimates $(\hat{V}_{t|t}, \hat{\theta}_{t|t})$ over 2007–2010 are out-of-sample filtered estimates based upon the SVJ2 parameter estimates in Table 7, and upon daily excess returns up to date t .

Observed and estimated end-of-day VIX values are shown in Fig. 10, as well as the gap between the two. While the VIX has been called a fear index or fear gauge, this gap identifies when VIX levels derived from traded option prices appear excessive relative to the underlying volatility of the stock market. Over 2007, for instance, the estimated VIX gradually increases from 10% to 20% but does not exhibit the dramatic swings observed in the actual VIX. Those swings appear related to major events discussed in the Brunnermeier (2009) account of the financial crisis, beginning with the increase in subprime mortgage defaults in February 2007. The VIX gap increased further in late July following the first recognitions of substantial bank exposure, and again in the first ten days of August following a volatile week that generated large losses for quantitative hedge funds. The VIX gap fluctuated considerably over the remainder of 2007 and in

the first half of 2008, in parallel with the Eurodollar-Treasury interest rate spread that Brunnermeier (2009, Fig. 3) uses to gauge the evolving liquidity crisis.

The VIX gap increased to 12% shortly after Lehman Brothers' bankruptcy on September 15, 2008 and it temporarily exceeded 21% following the initial rejection of the Troubled Asset Relief Program bill by the House of Representatives on September 29. October's turbulent and falling stock market led to further increases in observed and estimated VIX, with the difference between the two peaking at 48% on October 24. The divergence subsided in following months, especially after the stock market bottomed out in March 2009, before flaring up again with the European sovereign debt crisis in the spring of 2010. The VIX gap stabilized around 4–9% for the remainder of 2010, with observed and estimated VIX moving roughly in parallel. Because of central tendency uncertainty regarding where daily volatility is headed, only the largest VIX gaps can be identified as statistically significant.

5. Summary and conclusions

This paper provides estimates of the time-changed Carr, Geman, Madan, and Yor (2003) CGMY Lévy process based on stock market excess returns, and compares them to the time-changed finite-activity jump-diffusions previously examined by Bates (2006). I draw the following three conclusions.

First, it is important to recognize the fat-tailed properties of returns when filtering latent variables. Failure to do so makes latent variable estimates excessively sensitive to daily outliers larger than three standard deviations and affects parameter estimates—especially the parameters of the volatility process. However, such major outliers are relatively rare. Conditional volatility estimates from the

³³ Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) explain how the log contract approximately prices realized variance over the lifetime of the underlying options.

less fat-tailed distributions [the Heston (1993) stochastic volatility model; the Carr and Wu (2003) log-stable model] diverge substantially from those of other distributions only in the weeks following large outliers.

Second, it is not particularly important which fat-tailed distribution one uses. Estimates of the volatility process parameters and realizations are virtually unchanged across most specifications, while the option pricing implications are virtually identical for all but the deepest out-of-the-money options.

Third, conditional upon no recent outliers, even the Heston stochastic volatility model fits option prices similarly to the jump models for all but deep out-of-the-money options. For these stochastic volatility or stochastic intensity models, the estimated tilt of the volatility smirk for near-the-money options (± 2 standard deviations) appears primarily driven by the leverage effect.

I also present evidence of some structural shifts over time in the data generating process. Most striking is the apparently nonstationary evolution of the first-order autocorrelation of daily stock market returns, which rose from near-zero in the 1930s to around 35% in 1971 before drifting down again to near-zero values at the end of the 20th century, and even negative in the 21st. The high autocorrelation estimates in the 1960s and 1970s are clearly attributable to a stale-price problem from low stock turnover and are of substantial importance when assessing historical stock market volatility. The paper develops methods of dealing with time-varying autocorrelation, by treating it as an additional latent state variable to be filtered from observed data. Furthermore, the paper develops a nonaffine model (Model 2) of evolving autocorrelation that can nevertheless be easily estimated on time series data. The model generates an affine risk-neutral process for pricing index options and is consistent with the inverse relationship between autocorrelation and volatility found by LeBaron (1992).

Finally, the paper also shows longer-term swings in volatility, which are modeled using a two-factor concatenated volatility model. Estimating a latent variable (the central tendency) underlying another latent variable (spot variance) underlying daily stock market returns is perforce imprecise. Nevertheless, the two-factor model usefully highlights the misleading precision of multi-period forecasts from one-factor variance models. One-factor models erroneously predict tight confidence intervals for implicit volatility at longer maturities, given hypothesized volatility mean reversion to an identifiable mean. The two-factor model estimates spot volatilities and term structures of implicit volatilities more accurately than the one-factor model—with, however, substantial gaps remaining on average between observed and estimated at-the-money implicit volatilities.

Alternate data sources could yield more accurate assessments of spot variance and of its central tendency: intradaily realized variances, for instance, or the high-low range data examined by Alizadeh, Brandt, and Diebold (2002). This paper has focused upon daily returns because of its focus on daily crash risk, but the AML methodology can equally be applied to estimating conditional volatilities from those alternative data. Realized variances

are noisy signals of latent conditional variance when intradaily jumps are present, indicating the need for filtration methodologies such as AML. Such applications are potential topics for future research.

Appendix A. Filtration under Model 2

A.1. One-factor variance process

From Eq. (27), the cumulant generating function (CGF) for future $(y_{t+1}, \rho_{t+1}, V_{t+1})$ conditional upon current values under Model 2 is

$$\ln F(\Phi, \xi, \psi | y_t, \rho_t, V_t) = \mu_0 \tau (1 - \rho_t) \Phi + \theta C(\tau; (1 - \rho_t) \Phi, \psi) + \frac{1}{2} \sigma_\rho^2 \xi^2 + (\xi + \Phi y_t) \rho_t + D(\tau; (1 - \rho_t) \Phi, \psi) V_t. \tag{A.1}$$

The filtered CGF conditional upon only observing past data \mathbf{Y}_t can be computed by integrating this over the independent conditional distributions of the latent variables (ρ_t, V_t) :

$$F(\Phi, \xi, \psi | \mathbf{Y}_t) = \iint e^{\ln F(\cdot | y_t, \rho_t, V_t)} p(V_t | \mathbf{Y}_t) p(\rho_t | \mathbf{Y}_t) dV_t d\rho_t = \int e^{\mu_0 \tau (1 - \rho_t) \Phi + \theta C(\tau; (1 - \rho_t) \Phi, \psi) + \frac{1}{2} \sigma_\rho^2 \xi^2 + (\xi + \Phi y_t) \rho_t + g_{t|t}[D(\tau; (1 - \rho_t) \Phi, \psi)]} \times p(\rho_t | \mathbf{Y}_t) d\rho_t \tag{A.2}$$

where $g_{t|t}(\psi) \equiv -v_t \ln(1 - \kappa_t \psi)$ is the gamma conditional CGF for latent V_t . Under the change of variables $(z, x) \equiv [(1 - \rho_t) \Phi, (1 - \rho_t)]$, and under the assumption that the scaling term $x = (1 - \rho_t) > 0$, the Fourier inversion used in evaluating $p(y_{t+1} | \mathbf{Y}_t)$ from Eq. (A.2) becomes

$$p(y_{t+1} | \mathbf{Y}_t) = \frac{1}{\pi} \operatorname{Re} \left[\int_{\Phi=0}^{\infty} F(i\Phi, 0, 0 | \mathbf{Y}_t) e^{-i\Phi y_{t+1}} d\Phi \right] = \frac{1}{\pi} \operatorname{Re} \left[\int_{z=0}^{\infty} e^{\mu_0 \tau iz + \theta C(\tau; iz, 0) + g_{t|t}[D(\tau; iz, 0)] - iy_t z} \times \left(\int_{x=0}^{\infty} \frac{1}{x} e^{-iz(y_{t+1} - y_t)/x} p(x | \mathbf{Y}_t) dx \right) dz \right] \equiv \frac{1}{\pi} \operatorname{Re} \left[\int_{z=0}^{\infty} F^*(iz, 0 | \mathbf{Y}_t) \left(\int_{x=0}^{\infty} \frac{1}{x} e^{(-iz(y_{t+1} - y_t)/x)} p(x | \mathbf{Y}_t) dx \right) \right] \tag{A.3}$$

where $\operatorname{Re}[c]$ denotes the real component of complex-valued c , the $1/x$ term in the integrand reflects the Jacobean from the change of variables, and

$$F^*(z, \psi | \mathbf{Y}_t) \equiv \exp\{(\mu_0 \tau_t)z + \theta C(\tau; z, \psi) + g_{t|t}[D(\tau; z, \psi)] - y_t z\} \tag{A.4}$$

is the joint moment generating function of $(s_{t+1} - s_t - y_t, V_{t+1})$ conditional upon past data \mathbf{Y}_t .

It is convenient to use the two-parameter inverse Gaussian distribution to approximate $p(x | \mathbf{Y}_t)$:

$$p(x | \mathbf{Y}_t) = \sqrt{\frac{\lambda_t}{2\pi x^3}} \exp \left[-\frac{\lambda_t (x - \bar{x}_t)^2}{2\bar{x}_t^2 x} \right], \quad x > 0 \tag{A.5}$$

where $\bar{x}_t = E(x | \mathbf{Y}_t)$ and $\lambda_t = \bar{x}_t^3 / \operatorname{Var}(x | \mathbf{Y}_t)$ summarize what is known about x (and about ρ_t) at time t . For this distribution, the inner integration in Eq. (A.3) can be replaced by

$$M_{-1}(a) \equiv \int_{x=0}^{\infty} x^{-1} e^{ax} p(x | \mathbf{Y}_t) dx$$

$$= \sqrt{\frac{\lambda_t}{(\lambda_t - 2a)^3}} \left[1 + \frac{\sqrt{\lambda_t(\lambda_t - 2a)}}{\bar{x}_t} \right] \exp \left[\frac{\lambda_t - \sqrt{\lambda_t(\lambda_t - 2a)}}{\bar{x}_t} \right] \quad (\text{A.6})$$

for $a = -iz(y_{t+1} - y_t) \equiv -iz\Delta y$. Consequently, evaluating Eq. (A.3) involves only univariate numerical integration.³⁴

Similar univariate integrations are used for filtering V_{t+1} and ρ_{t+1} conditional upon observing y_{t+1} . The noncentral posterior moments of V_{t+1} are given by

$$E(V_{t+1}^m | \mathbf{Y}_{t+1}) = \frac{1}{\pi p(y_{t+1} | \mathbf{Y}_t)} \text{Re} \left[\int_{z=0}^{\infty} \frac{\partial^m F^*(iz, \psi | \mathbf{Y}_t)}{\partial \psi^m} \Big|_{\psi=0} M_{-1}(-iz\Delta y) dz \right], \quad (\text{A.7})$$

where the derivatives of $F^*(\cdot)$ with respect to ψ inside the integrand can be easily evaluated from Eq. (A.4) given the specifications for $C(\cdot)$ and $D(\cdot)$ in Eqs. (18)–(19). The posterior moments of ρ_{t+1} can be computed by taking partials of Eq. (A.2) with respect to ζ and then again using the change of variables to reduce the Fourier inversion to a univariate integration. The resulting posterior mean and variance of ρ_{t+1} are

$$\hat{\rho}_{t+1|t+1} = 1 - \frac{1}{\pi p(y_{t+1} | \mathbf{Y}_t)} \text{Re} \left[\int_{z=0}^{\infty} F^*(iz, 0 | \mathbf{Y}_t) M_0(-iz\Delta y) dz \right] \quad (\text{A.8})$$

$$W_{t+1|t+1} = \sigma_\rho^2 + \frac{1}{\pi p(y_{t+1} | \mathbf{Y}_t)} \text{Re} \left[\int_{z=0}^{\infty} F^*(iz, 0 | \mathbf{Y}_t) M_1(-iz\Delta y) dz \right], \quad (\text{A.9})$$

where

$$M_0(a) \equiv \int_{x=0}^{\infty} e^{a/x} p(x | \mathbf{Y}_t) dx = \sqrt{\frac{\lambda_t}{\lambda_t - 2a}} \exp \left[\frac{\lambda_t - \sqrt{\lambda_t(\lambda_t - 2a)}}{\bar{x}_t} \right] \quad (\text{A.10})$$

and

$$M_1(a) \equiv \int_{x=0}^{\infty} x e^{a/x} p(x | \mathbf{Y}_t) dx = \bar{x}_t \exp \left[\frac{\lambda_t - \sqrt{\lambda_t(\lambda_t - 2a)}}{\bar{x}_t} \right]. \quad (\text{A.11})$$

Finally, the conditional distribution function (CDF) used in the normal probability plots of Fig. 1 takes the form

$$\text{CDF}(y_{t+1} | \mathbf{Y}_t) = \frac{1}{2} - \frac{1}{\pi} \text{Re} \left[\int_{z=0}^{\infty} \frac{F^*(iz, 0 | \mathbf{Y}_t)}{iz} M_0(-iz\Delta y) dz \right]. \quad (\text{A.12})$$

A.2. Two-factor variance processes

Filtration under the two-factor variance process of Eq. (43) is similar. The joint moment generating function $F^*(z, \psi | \mathbf{Y}_t)$ of Eq. (A.4) gets replaced in Eqs. (A.3), (A.7),

(A.8), (A.9), and (A.12) by

$$F^*(z, \psi, \psi_2 | \mathbf{Y}_t) = \exp \left\{ CC(\tau; z, \psi, \psi_2) + g_{t|t} [D(\tau, z, \psi)] \right\} \times \exp \left\{ g_{t|t}^\theta [DD(\tau; z, \psi, \psi_2)] - y_t z \right\} \quad (\text{A.13})$$

evaluated at $\psi_2 = 0$. $CC(\cdot)$ and $DD(\cdot)$ are defined below in Eqs. (B.2) through (B.4), while $g_{t|t}^\theta(\psi_2) = -v_{t|t}^\theta \ln(1 - \kappa_{t|t}^\theta \psi_2)$ is the gamma conditional CGF that summarizes what is known at time t about latent θ_t . Filtration of θ_t proceeds similarly to that of V_t in Eq. (A.7), with posterior moments

$$E(\theta_{t+1}^m | \mathbf{Y}_{t+1}) = \frac{1}{\pi p(y_{t+1} | \mathbf{Y}_t)} \times \text{Re} \left[\int_{z=0}^{\infty} \frac{\partial^m F^*(iz, 0, \psi_2 | \mathbf{Y}_t)}{\partial \psi_2^m} \Big|_{\psi_2=0} M_{-1}(-iz\Delta y) dz \right]. \quad (\text{A.14})$$

Appendix B. Cumulant generating functions under a two-factor variance process

The approximate concatenated two-factor variance model holds the central tendency θ_t constant intraday and assumes that its daily evolution is noncentral gamma and is independent of other state variables. Consequently, the one-day conditional cumulant generating function (CCGF) for $s_{t+1} - s_t$ and future state variables (V_{t+1}, θ_{t+1}) conditional upon (V_t, θ_t) and a one-day interval of length τ involves replacing θ in Eq. (17) by θ_t and adding the noncentral gamma CCGF that describes the evolution of θ_{t+1} :

$$\begin{aligned} \ln F(\Phi, \psi, \psi_2 | V_t, \theta_t) &\equiv \ln E[e^{\Phi(s_{t+1} - s_t) + \psi V_{t+1} + \psi_2 \theta_{t+1}} | V_t, \theta_t] \\ &= \mu_0 \tau \Phi + \theta_t C(\tau; \Phi, \psi) + D(\tau; \Phi, \psi) V_t \\ &\quad - \frac{2\beta_2 \bar{\theta}}{\sigma_2^2} \ln [1 - K_2(\tau) \psi_2] + \frac{e^{-\beta_2 \tau} \psi_2}{1 - K_2(\tau) \psi_2} \theta_t \\ &\equiv CC(\tau; \Phi, \psi_2) + D(\tau; \Phi, \psi) + DD(\tau; \Phi, \psi, \psi_2) \theta_t, \end{aligned} \quad (\text{B.1})$$

where $C(\cdot)$ and $D(\cdot)$ are defined in Eqs. (18)–(22),

$$CC(\tau; \Phi, \psi_2) = \mu_0 \tau \Phi - \frac{2\beta_2 \bar{\theta}}{\sigma_2^2} \ln [1 - K_2(\tau) \psi_2], \quad (\text{B.2})$$

$$DD(\tau; \Phi, \psi, \psi_2) = C(\tau; \Phi, \psi) + \frac{e^{-\beta_2 \tau} \psi_2}{1 - K_2(\tau) \psi_2}, \quad (\text{B.3})$$

$$K_2(\tau) = \frac{\sigma_2^2}{2\beta_2} (1 - e^{-\beta_2 \tau}). \quad (\text{B.4})$$

By iterated expectations, the multiperiod CCGF is also affine and satisfies

$$\begin{aligned} CC_{t,T} + D_{t,T} V_t + DD_{t,T} \theta_t &\equiv \ln E[e^{\Phi(s_T - s_t) + \psi V_T + \psi_2 \theta_T} | V_t, \theta_t] \\ &= \ln E[e^{\Phi(s_{t+1} - s_t) + CC_{t+1,T} + D_{t+1,T} V_{t+1} + DD_{t+1,T} \theta_{t+1}} | V_t, \theta_t]. \end{aligned} \quad (\text{B.5})$$

This implies the coefficients satisfy the backward recursion

$$\begin{aligned} CC_{t,T} &= CC_{t+1,T} + CC(\tau; \Phi, DD_{t+1,T}), \\ D_{t,T} &= D(\tau, \Phi, D_{t+1,T}), \\ DD_{t,T} &= DD(\tau; \Phi, D_{t+1,T}, DD_{t+1,T}) \end{aligned} \quad (\text{B.6})$$

³⁴ Alternate distributions could be used for $p(x | \mathbf{Y}_t)$; e.g., a beta distribution over the range $[0, 2]$. That would constrain $|\rho_t| < 1$ and results in an $M_{-1}(a)$ term that involves the confluent hypergeometric U -function. However, I could not find a method for evaluating that function that is fast, accurate, and robust to all parameter values.

subject to the terminal condition $(CC_{T,T}, D_{T,T}, DD_{T,T}) = (0, \psi, \psi_2)$. $(CC_{t,T}, D_{t,T}, DD_{t,T})$ are functions of (Φ, ψ, ψ_2) , a dependency omitted in Eqs. (B.5) and (B.6) to simplify the notation.

The multiperiod CCGF is useful in two contexts. First, subdividing individual days into increasingly fine subintervals yields a numerical algorithm for computing the joint CCGF for continuously evolving θ_t , which can then be compared with the daily discrete-time specification. Second, risk-neutral versions of Eqs. (B.5) and (B.6) are used for pricing options with multiple days to maturity, as described in Eq. (49).

The unconditional cumulant generating function of (V_t, θ_t) satisfies a recursion similar to Eq. (B.5). Ignoring daily seasonals in τ yields an approximate fixed-point characterization

$$g(\psi, \psi_2) \equiv \ln E \left[e^{\psi V_t + \psi_2 \theta_t} \right] = \ln E \left[E \left(e^{\psi V_{t+1} + \psi_2 \theta_{t+1}} | V_t, \theta_t \right) \right] \\ = CC(\tau; 0, \psi_2) + g[D(\tau; 0, \psi), DD(\tau; 0, \psi, \psi_2)]. \quad (B.7)$$

Derivatives of $g(\cdot)$ in Eq. (B.7) evaluated at $\psi = \psi_2 = 0$ implicitly give the unconditional moments of (V_t, θ_t) . The unconditional means are $EV_t = E\theta_t = \bar{\theta}$, while the unconditional variances and covariance are solutions to the recursive system of equations

$$\begin{pmatrix} 1 - e^{-2\beta\tau} & -2e^{-\beta\tau}(1 - e^{-\beta\tau}) & (1 - e^{-\beta\tau})^2 \\ 0 & 1 - e^{-(\beta+\beta_2)\tau} & -(1 - e^{-\beta\tau})e^{-\beta_2\tau} \\ 0 & 0 & 1 \end{pmatrix} \\ \times \begin{pmatrix} \text{Var}(V_t) \\ \text{Cov}(V_t, \theta_t) \\ \text{Var}(\theta_t) \end{pmatrix} = \begin{pmatrix} \frac{\sigma^2 \bar{\theta}}{2\beta} (1 - e^{-2\beta\tau}) \\ 0 \\ \frac{\sigma^2 \bar{\theta}}{2\beta_2} \end{pmatrix}. \quad (B.8)$$

The resulting unconditional $\text{Var}(V_t)$ is insensitive to the precise value of τ and was computed using $\tau = 1/252$. The unconditional means and variances are used to parameterize independent gamma distributions describing what is known about the initial latent variables (V_1, θ_1) when starting the recursive AML filtration procedure.

References

- Alizadeh, S., Brandt, M.W., Diebold, F.X., 2002. Range-based estimation of stochastic volatility models. *Journal of Finance* 57, 1047–1091.
- Andersen, T.G., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time equity return models. *Journal of Finance* 57, 1239–1284.
- Andersen, T.G., Bollerslev, T., Christoffersen, P.F., Diebold, F.X., 2007. Practical volatility and correlation modeling for financial market risk management. In: Carey, M., Stulz, R. (Eds.), *The Risks of Financial Institutions*, University of Chicago Press, Chicago, IL, pp. 513–544.
- Bakshi, G., Madan, D.B., 2000. Spanning and derivative-security valuation. *Journal of Financial Economics* 55, 205–238.
- Barclay, M.J., Litzenberger, R.H., Warner, J.B., 1990. Private information, trading volume, and stock-return variances. *Review of Financial Studies* 3, 233–254.
- Bates, D.S., 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94, 181–238.
- Bates, D.S., 2006. Maximum likelihood estimation of latent affine processes. *Review of Financial Studies* 19, 909–965.
- Bertoin, J., 1996. In: *Lévy Processes* Cambridge University Press, Cambridge, England.
- Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. *Journal of Finance* 55, 839–866.
- Broadie, M.N., Chernov, M., Johannes, M., 2009. Understanding index option returns. *Review of Financial Studies* 22, 4492–4529.
- Brunnermeier, M.K., 2009. Deciphering the liquidity and credit crunch 2007–2008. *Journal of Economic Perspectives* 23, 77–100.
- Carr, P., Geman, H., Madan, D.B., Yor, M., 2002. The fine structure of asset returns: an empirical investigation. *Journal of Business* 75, 305–332.
- Carr, P., Geman, H., Madan, D.B., Yor, M., 2003. Stochastic volatility for Lévy processes. *Mathematical Finance* 13, 345–382.
- Carr, P., Wu, L., 2003. The finite moment log stable process and option pricing. *Journal of Finance* 58, 753–777.
- Carr, P., Wu, L., 2004. Time-changed Lévy processes and option pricing. *Journal of Financial Economics* 71, 113–141.
- Carr, P., Wu, L., 2007. Stochastic skew in currency options. *Journal of Financial Economics* 86, 213–247.
- Chernov, M., Gallant, A.R., Ghysels, E., Tauchen, G., 2003. Alternative models for stock price dynamics. *Journal of Econometrics* 116, 225–257.
- Clark, P.K., 1973. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* 41, 135–155.
- Cox, J.C., Ross, S.A., Rubinstein, R., 1979. Option pricing: a simplified approach. *Journal of Financial Economics* 7, 229–263.
- Dai, Q., Singleton, K.J., 2000. Specification analysis of affine term structure models. *Journal of Finance* 55, 1943–1978.
- Dimson, E., 1979. Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics* 7, 197–226.
- Duffie, D., Pan, J., Singleton, K.J., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Eberlein, E., Keller, U., Prause, K., 1998. New insights into smile, mispricing, and value at risk: the hyperbolic model. *Journal of Business* 71, 371–405.
- Egloff, D., Leippold, M., Wu, L., 2010. The term structure of variance swap rates and optimal variance swap investments. *Journal of Financial and Quantitative Analysis* 45, 1279–1310.
- Eraker, B., Johannes, M., Polson, N.G., 2003. The impact of jumps in volatility and returns. *Journal of Finance* 58, 1269–1300.
- French, D.W., 1984. The weekend effect on the distribution of stock prices: implications for option pricing. *Journal of Financial Economics* 13, 547–559.
- French, K.R., 2008. The cost of active investing. *Journal of Finance* 63, 1537–1573.
- French, K.R., Roll, R., 1986. Stock return variances: the arrival of information and the reaction of traders. *Journal of Financial Economics* 17, 5–26.
- Gallant, A.R., Rossi, P.E., Tauchen, G., 1992. Stock prices and volume. *Review of Financial Studies* 5, 199–242.
- Hentschel, L., 1995. All in the family: nesting symmetric and asymmetric GARCH models. *Journal of Financial Economics* 39, 71–104.
- Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–344.
- Huang, J., Wu, L., 2004. Specification analysis of option pricing models based on time-changed Lévy processes. *Journal of Finance* 59, 1405–1439.
- Ibbotson Associates, 2006. *S&P 500 Yearbook*. R. G. Ibbotson Associates, Chicago, IL.
- Jiang, G.J., Tian, Y.S., 2005. Model-free implied volatility and its information content. *Review of Financial Studies* 18, 1305–1342.
- Johannes, M., Polson, N.G., Stroud, J., 2009. Optimal filtering of jump-diffusions: extracting latent states from asset prices. *Review of Financial Studies* 22, 2759–2799.
- Jukuivouille, E., 1995. Measuring true stock index value in the presence of infrequent trading. *Journal of Financial and Quantitative Analysis* 30, 455–464.
- Kou, S., 2002. A jump diffusion model for option pricing. *Management Science* 48, 1086–1101.
- LeBaron, B.D., 1992. Some relations between volatility and serial correlations in stock returns. *Journal of Business* 65, 199–219.
- Li, H., Wells, M.T., Yu, C.L., 2008. A Bayesian analysis of return dynamics with stochastic volatility and Lévy jumps. *Review of Financial Studies* 21, 2345–2378.
- Lo, A.W., MacKinlay, A.C., 1988. Stock market prices do not follow random walks: evidence from a new specification test. *Review of Financial Studies* 1, 41–66.
- MacKinlay, A.C., Ramaswamy, K., 1988. Index-futures arbitrage and the behavior of stock index futures prices. *Review of Financial Studies* 1, 137–158.
- Madan, D.B., Seneta, E., 1990. The variance gamma (v.g.) model for share market returns. *Journal of Business* 63, 511–525.
- Maheu, J.M., McCurdy, T.H., 2004. News arrival, jump dynamics and volatility components for individual stock returns. *Journal of Finance* 59, 755–793.
- Mandelbrot, B.B., 1963. The variation of certain speculative prices. *Journal of Business* 36, 394–419.
- Mandelbrot, B.B., Hudson, R.L., 2004. *The (mis) Behavior of Markets: A Fractal View of Risk, Ruin, and Reward* Basic Books, New York.

- Masreliez, C.J., 1975. Approximate non-Gaussian filtering with linear state and observation relations. *IEEE Transactions on Automatic Control* 20, 107–110.
- Merton, R.C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- Santa-Clara, P., Yan, S., 2010. Crashes, volatility, and the equity premium: lessons from S&P 500 options. *Review of Economics and Statistics* 92, 435–451.
- Schick, I.C., Mitter, S.K., 1994. Robust recursive estimation in the presence of heavy-tailed observation noise. *Annals of Statistics* 22, 1045–1080.
- Schwert, G.W., 1990. Indexes of U.S. stock prices from 1802 to 1987. *Journal of Business* 63, 399–426.
- Skiadopoulos, G., Hodges, S., Clewlow, L., 1999. The dynamics of the S&P 500 implied volatility surface. *Review of Derivatives Research* 3, 263–282.
- Taylor, S.J., Xu, X., 1994. The term structure of volatility implied by foreign exchange options. *Journal of Financial and Quantitative Analysis* 29, 57–74.
- Wu, L., 2006. Dampened power law: reconciling the tail behavior of financial security returns. *Journal of Business* 79, 1445–1473.