Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics

Suzanne S. Lee
Georgia Institute of Technology

Per A. Mykland
Department of Statistics, University of Chicago

This article introduces a new nonparametric test to detect jump arrival times and realized jump sizes in asset prices up to the intra-day level. We demonstrate that the likelihood of misclassification of jumps becomes negligible when we use high-frequency returns. Using our test, we examine jump dynamics and their distributions in the U.S. equity markets. The results show that individual stock jumps are associated with prescheduled earnings announcements and other company-specific news events. Additionally, S&P 500 Index jumps are associated with general market news announcements. This suggests different pricing models for individual equity options versus index options. (*JEL* G12, G22, G14)

Financial markets sometimes generate significant discontinuities, so-called jumps, in financial variables. A number of recent empirical and theoretical studies proved the existence of jumps and their substantial impact on financial management, from portfolio and risk management to option and bond pricing and hedging (see Merton, 1976; Bakshi et al., 1997, 2000; Bates, 1996; Liu et al., 2003; Naik and Lee, 1990; Duffie et al., 2000, and Johannes, 2004). Despite advances in asset pricing models and their inference techniques, the studies have found that jumps are empirically difficult to identify, because only discrete data are available from continuous-time models, in which most of the aforementioned applications were studied. Our goal in this article is first to propose a new jump detection technique to resolve such identification problems. Further, we show that our technique provides a model-free tool for characterizing jump dynamics in individual equity and S&P 500 Index returns, which allows us to investigate different model structures for their option pricing.

There are primarily two motivations for this study. First, while researchers recognize the presence of jumps in order to better explain the excess kurtosis and skewness of return distributions and implied volatility smiles, we also note the fact that jumps do not come to markets regularly, but their arrivals tend to depend on market information. For instance, Piazzesi, (2003) shows that incorporating jumps related to market information improves bond pricing models. Given the difficulty of pinning down jump parameters, even with both time-series and cross-sectional data in parametric settings, we question how one can simply search observable information relevant to jumps that appear to have such strong impact on pricing securities. To learn about the stochastic features of irregular jump arrivals and their associated market information, it is critical to first develop a robust test to detect jumps. Once detected, one can examine what type of information is dynamically related to jumps to improve pricing models and better explain market phenomena. Our test is meant to be applicable to all kinds of financial time series, including equity returns and volatility, interest rates, and exchange rates, so long as high-frequency observations are available. Because jumps have bigger impacts on derivative securities than other securities (see Johannes, 2004, for discussion), our test would have greater implications in their financial management. We take a nonparametric approach for the results to be robust with respect to model specifications as well as to nonstationarity of price processes—a common feature of financial variables.

Another motivation for developing our test is the improvement of derivative hedging. The presence of jumps makes incomplete markets. The degree of market incompleteness depends on the size and intensity of jumps, which determines the magnitude of derivative hedging error (see Naik and Lee, 1990; and Bertsimas et al., 2001). As a by-product, our test yields both the direction and size of detected jumps, allowing the characterization of jump size distribution as well as stochastic jump intensity. These outcomes allow us to develop hedging strategies accordingly. Detecting arrival times is also important in order to rebalance hedging portfolios dynamically once a jump arrival is detected, as shown by Collin-Dufresne and Hugonnier (2001).

Our main methodological contribution is summarized as follows. We first explain the intuition on how we develop the test and its mathematical definition. Then, we provide its asymptotic distribution and detection criterion for practical use. Second, we prove that it can precisely distinguish actual jumps using high-frequency data and show that spurious detection of jumps becomes negligible. Hence, we show that stochastic jump estimates based on our test are accurate. Third, we compare our test to the existing nonparametric jump tests by Barndorff-Nielsen and Shephard (2006) and Jiang and Oomen (2005). Ours outperforms these others in terms of size and power. Monte Carlo simulation confirms the results as well.

After we discuss generally how our test can be applied in asset pricing models, we conduct an empirical study particularly for the U.S. equity markets. We collect from the Trade and Quote (TAQ) database high-frequency returns

of three individual stock and S&P 500 Index prices transacted on the New York Stock Exchange (NYSE) over the period of 1 September to 30 November 2005. We find that jumps do not occur regularly, so that stochastic jump intensity should be considered in equity markets. We observe more frequent jumps in the individual stock returns than in the index returns. Sizes of jumps in individual equities are greater than those in the index. We find that detected jumps are related to news releases from Factiva, a real-time financial news database. For the individual stocks, jumps are associated with company-specific news events like scheduled earnings announcements, as well as unscheduled news. For the index, jumps occur with general market news such as Federal Open Market Committee (FOMC) meetings and macroeconomic reports. It is noted that for individual equities, the majority of jumps occur with unscheduled news and their magnitudes are comparable to those that occur with earnings announcements. Therefore, we conclude that one should incorporate not only earnings announcements as in Dubinsky and Johannes, (2006), but also other company-specific news for individual equity option pricing. For the index options, general market announcements and events are to be taken into account. This evidence also provides support for the different model structures for individual equity and index options described in Bakshi et al., (2003).

The rest of the article is organized as follows. Section 1 sets up a theoretical framework to detect jumps and introduces the test. Section 2 discusses the precision of our test. Section 3 investigates the test's finite sample performance. Section 4 empirically examines jump dynamics in equity returns and their association with news releases. Finally, we conclude in Section 5.

1. A Theoretical Model for the Test and Its Asymptotic Theory

We employ a one-dimensional asset return process with a fixed complete probability space $(\Omega, \mathcal{F}_t, \mathcal{P})$, where $\{\mathcal{F}_t : t \in [0, T]\}$ is a right-continuous information filtration for market participants, and $\mathcal{P}$ is a data-generating measure. Let the continuously compounded return be written as $d \log S(t)$ for $t \geq 0$, where $S(t)$ is the asset price at t under $\mathcal{P}$. We are interested in finding jumps in the asset returns as follows. When there are no jumps in the market, $S(t)$ is represented as

$$d \log S(t) = \mu(t)dt + \sigma(t)dW(t) \tag{1}$$

where $W(t)$ is an $\mathcal{F}_t$-adapted standard Brownian motion. The drift $\mu(t)$ and spot volatility $\sigma(t)$ are $\mathcal{F}_t$-adapted processes, such that the underlying process is an Itô process that has continuous sample paths. When there are jumps, $S(t)$ is given by

$$d \log S(t) = \mu(t)dt + \sigma(t)dW(t) + Y(t)dJ(t), \tag{2}$$

where $J(t)$ is a counting process independent of $W(t)$. $Y(t)$ is the jump size, which is a predictable process (see Protter, 2004, for definition). Its mean $\mu_y(t)$

and standard deviation $\sigma_y(t)$ are also $\mathcal{F}_t$-predictable processes. We assume jump sizes $Y(t)$ are independent of each other and identically distributed. They are also independent of other random components $W(t)$ and $J(t)$. $J(t)$ may be a non-homogeneous Poisson-type jump process. Hence, for instance, scheduled (deterministic) events such as earnings announcements are allowed to affect jump intensity, in which case, the model can be the sum of a process with an intensity and a deterministic counting process.

Observation of $S(t)$, equivalently $\log S(t)$, occurs only at discrete times $0 = t_0 < t_1 < \cdots < t_n = T$. For simplicity, this article assumes that observation times are equally spaced: $\Delta t = t_i - t_{i-1}$. This simplified assumption can easily be generalized to non-equidistant cases by letting $\max_i(t_i - t_{i-1}) \to 0$. We also impose the following necessary assumption on price processes throughout this article:

Assumption 1. *For any $\epsilon > 0$,*

$$\textbf{A1.1} \quad \sup_i \sup_{t_i \leq u \leq t_{i+1}} |\mu(u) - \mu(t_i)| = O_p\left(\Delta t^{\frac{1}{2} - \epsilon}\right), \quad (3)$$

$$\textbf{A1.2} \quad \sup_i \sup_{t_i \leq u \leq t_{i+1}} |\sigma(u) - \sigma(t_i)| = O_p\left(\Delta t^{\frac{1}{2} - \epsilon}\right). \quad (4)$$

Following Pollard (2002), we use O_p notation throughout this article to mean that, for random vectors $\{X_n\}$ and non-negative random variable $\{d_n\}$, $X_n = O_p(d_n)$, if for each $\delta > 0$, there exists a finite constant M_δ such that $P(|X_n| > M_\delta d_n) < \delta$ eventually. One can interpret Assumption 1 as the drift and diffusion coefficients not changing dramatically over a short time interval. Furthermore, Assumption 1 allows drift and diffusion to depend on the process itself. This is reasonable in view of Lemma 2 in Mykland and Zhang (2006). It is satisfied for most Itô processes. The stochastic volatility model in Heston (1993) and its extended versions studied in Bakshi et al., (2005) are examples. This assumption also satisfies the stochastic volatility plus finite activity jump semi-martingale class in Barndorff-Nielsen and Shephard, (2004).

1.1 Intuition for and definition of the nonparametric jump test

In this subsection, we address the basic intuition behind our new detection technique and mathematically define the jump detection statistic $\mathcal{L}$. Our discussion here refers to a single test at time t_i. We do not make assumptions about whether there were or were not jumps before or after t_i. Generalization to a global test to determine whether a diffusion model (1) is rejected is straightforward by multiple tests (single tests over available times). A global test would be interesting in itself and achieving one is also part of our goal. However, as mentioned in the introduction, the main purpose of this article is to detect jumps over time in order to advance our knowledge of the stochastic mechanism of jump arrival dynamics. Multiple tests on time series allow us to extract such information.

The intuition that gives rise to our test is as follows. Imagine that asset prices evolve continuously over time. Suppose that a jump occurs in a market at some time, say t_i. We would expect the realized asset return at that time to be much greater than usual continuous innovations. What if the spot volatility at that time is also high? Even if there is no jump, if the volatility is high and if we can only observe prices in discrete times, the realized return we observe may be as high as the return that is actually due to a jump. To distinguish those two cases, it is natural to standardize the return by a measure that explains the local variation only from the continuous part of the process. We refer to this measure as *instantaneous volatility* in this article and denote it as $\sigma(t_i)$. This idea is incorporated into our test. In essence, we compare a realized return at any given time to a consistently estimated *instantaneous volatility* using corresponding local movements of returns. More specifically, the ratio of realized return to estimated volatility creates the test statistic for jumps.

How do we estimate *instantaneous volatility*? A commonly used nonparametric estimator for variance in the literature is the *realized power (quadratic) variation*, defined as the sum of squared returns

$$\text{plim}_{n \to \infty} \sum_{i=2}^{n} (\log S(t_i) - \log S(t_{i-1}))^2. \tag{5}$$

Using high-frequency returns within some period just before our testing time may yield a variance estimate over that period. However, this well-known variance estimator is inconsistent in the presence of jumps in a return process. Alternatively, a slightly modified version called the *realized bipower variation*, defined as the sum of products of consecutive absolute returns

$$\text{plim}_{n \to \infty} \sum_{i=3}^{n} |\log S(t_i) - \log S(t_{i-1})| |\log S(t_{i-1}) - \log S(t_{i-2})|, \tag{6}$$

has been suggested and shown to be a consistent estimator for the integrated volatility, even when there are jumps in return processes (see Barndorff-Nielsen and Shephard, 2004; and Aït-Sahalia, 2004). Despite the intuition that jumps in a process may impact its volatility estimation, it remains consistent no matter how large or small jumps are mixed with the diffusive part of pricing models. Incorporating this estimator makes the detection procedure independent of the presence of jumps over time, especially those jumps used for volatility estimation.

Our test is based on this interesting insight. Even if a highly volatile market environment makes the detection of jumps more difficult, infrequent Poisson jumps will be detected by our procedure fairly accurately, as long as we use high-frequency observations.

Here, we describe the formulation of the statistic and provide its mathematical definition. Suppose we have a fixed time horizon T, and n is the number
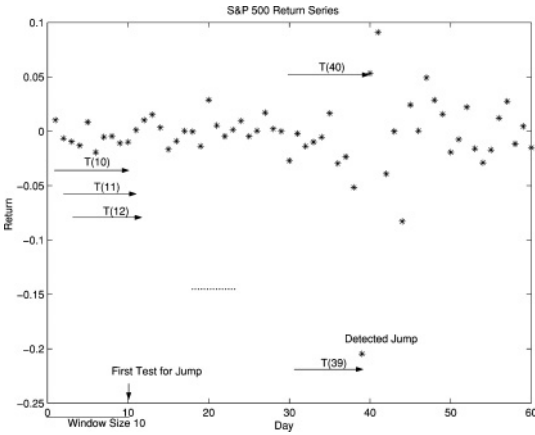
Figure 1
Intuition of jump detection test
This graph illustrates how our jump detection test distinguishes the jump arrivals. The example uses S&P 500 Index returns and a window size K is 10. The jump detection statistic $\mathcal{L}$ is formulated by taking the ratio of the last return in a window to the *instantaneous volatility*, estimated by bipower variation using the returns in the same window. By shifting the window over time, one can determine stochastic jump dynamics.

of observations in $[0, T]$. The distance between two successive observations is $\Delta t = \frac{T}{n}$. Consider a local movement of the process within a window size K. With realized returns in the window consisting of the previous $K - 1$ observations just before a testing time t_i, the *instantaneous volatility* is estimated based on the realized bipower variation, which was discussed earlier. We then take the ratio of this estimated volatility to the next realized return in order to determine whether there was a jump arrived at t_i and how large that jump size was. For example, if $\Delta t = 5$ minutes, $t_i = 10\text{-}05$ a.m., and $K = 10$, then we test for a jump by examining the relative magnitude of a realized return from 10:00 a.m. to 10:05 a.m. compared to instantaneous volatility estimated using 5-minute returns from 9:10 a.m. to 10:00 a.m. Figure 1 illustrates the construction of the test. Mathematical notation of the test statistics is as follows:

Definition 1. *The statistic $\mathcal{L}(i)$, which tests at time t_i whether there was a jump from t_{i-1} to t_i, is defined as*

$$\mathcal{L}(i) \equiv \frac{\log S(t_i)/S(t_{i-1})}{\widehat{\sigma(t_i)}}, \tag{7}$$

where

$$\widehat{\sigma(t_i)}^2 \equiv \frac{1}{K-2} \sum_{j=i-K+2}^{i-1} |\log S(t_j)/S(t_{j-1})||\log S(t_{j-1})/S(t_{j-2})|. \tag{8}$$

Notice that the realized bipower variation is used for the instantaneous volatility estimation in the denominator of the statistic, which makes our technique robust

6

to the presence of jumps in previous time periods. We choose the window size K in such a way that the effect of jumps on the volatility estimation disappears.

We state our results ignoring the drift. For our analysis using high-frequency data, the drift (of order dt) is mathematically negligible compared to the diffusion (of order $\sqrt{dt}$) and the jump component (of order 1). In fact, the drift estimates have higher standard errors, so that they cause the precision of variance estimates to decrease if included in variance estimation. Though we study a simplified version of the model without the drift term, i.e., $\mu = 0$, we also prove theoretically in Appendix A.1 that the main result continues to hold with the nonzero drift term. A modified statistic $\mathcal{L}_\mu(i)$ for the nonzero drift case is formally defined, and a corresponding theorem is presented therein as well.

1.2 Under the absence of jumps at testing time t_i

This subsection explains the asymptotic behavior of our jump detection statistics, $\mathcal{L}(i)$, when there is no jump at time t_i. We suppose our realized return from time t_{i-1} to t_i is from the diffusion part of the model (1) or (2). The asymptotic distribution of $\mathcal{L}$ is provided in Theorem 1.

Theorem 1. *Let $\mathcal{L}(i)$ be as in Definition 1 and the window size $K = O_p(\Delta t^\alpha)$, where $-1 < \alpha < -0.5$. Suppose the process follows (1) or (2) and Assumption 1 is satisfied. Let $\bar{A}_n$ be the set of $i \in \{1, 2, \ldots, n\}$ so that there is no jump in $(t_{i-1}, t_i]$. Then, as $\Delta t \to 0$,*

$$\sup_{i \in \bar{A}_n} |\mathcal{L}(i) - \widehat{\mathcal{L}}(i)| = O_p(\Delta t^{\frac{3}{2} - \delta + \alpha - \epsilon}), \tag{9}$$

where δ satisfies $0 < \delta < \frac{3}{2} + \alpha$ and

$$\widehat{\mathcal{L}}(i) = \frac{U_i}{c}. \tag{10}$$

Here $U_i = \frac{1}{\sqrt{\Delta t}}(W_{t_i} - W_{t_{i-1}})$, a standard normal variable and a constant $c = E|U_i| = \sqrt{2}/\sqrt{\pi} \approx 0.7979$. A related procedure in which the instantaneous mean is estimated is considered in Theorem 1.1 in Appendix A.1.

Proof of Theorem 1. *See Appendix A.2.*

Theorem 1 states that our detection statistic $\mathcal{L}(i)$ follows approximately the same distribution as $\widehat{\mathcal{L}}(i)$. And, we find that $\widehat{\mathcal{L}}(i)$ follows a normal distribution with mean 0 and variance $\frac{1}{c^2}$ because U_i is a standard normal random variable. When the absence of jumps in the whole price process is absolutely known a priori, the usage of realized quadratic variation for estimating *instantaneous volatility* would yield the same asymptotic distribution. As discussed when explaining our intuition, however, we do not require the absence of jumps in earlier or later time periods. Therefore, using quadratic variation does not

suffice in this case. As can be noted, $\mathcal{L}(i)$ is asymptotically independent and normally distributed over time; hence, one can easily find a joint asymptotic null distribution of the test statistics for various periods.

1.3 Under the presence of jumps at testing time t_i

We show in this subsection how the jump test reacts to the arrival of jumps and discuss the choice of window size. Suppose the realized return from time t_{i-1} to t_i is now from jump diffusion of Equation (2). Theorem 2 specifically shows that as the sampling interval Δt goes to 0, the test statistic becomes so large that we can detect the jump arrival at time t_i.

Theorem 2. *Let $\mathcal{L}(i)$ be as in Definition 1 and the window size $K = O_p(\Delta t^\alpha)$, where $-1 < \alpha < -0.5$. Suppose the process follows (2) and the Assumption 1 is satisfied. Suppose there is a jump at any time $\tau \in (t_{i-1}, t_i]$. Then,*

$$\mathcal{L}(i) \approx \frac{U_i}{c} + \frac{Y(\tau)}{c\sigma\sqrt{\Delta t}}, \tag{11}$$

where $Y(\tau)$ is actual jump size at actual jump time τ. Therefore, $\mathcal{L}(i) \to \infty$, as $\Delta t \to 0$. If there is no jump at any time $\tau \in (t_{i-1}, t_i]$, $\mathcal{L}(i)$ has the asymptotic behavior described in Theorem 1.

Proof of Theorem 2. *See Appendix A.3.*

The benefit of the bipower variation as an *instantaneous volatility* estimator in the denominator of the test statistic is that the presence of jumps at earlier times does not affect the consistency of estimation. Therefore, our test is robust to earlier jumps in detecting current jumps. This does not imply that we make no use of earlier jumps. One can learn about an earlier jump arrival by doing the same single test at that earlier time. Several single tests over time become a multiple test, which can provide us more information on jump dynamics.

In order to retain the benefit of bipower variation, the window size K must be large enough so that the effect of jumps on estimating *instantaneous volatility* disappears, but it must obviously be smaller than the number of observations n. The condition $K = O_p(\Delta t^\alpha)$ with $-1 < \alpha < -0.5$ satisfies this requirement. Therefore, the choice of sampling frequency Δt will determine the window size. In general, for *nobs* number of observations per day, $\Delta t = \frac{1}{252 \times nobs}$. The integers between $\sqrt{252 \times nobs}$ and $252 \times nobs$ are candidates for K. If daily data are used in the analysis, $\Delta t = \frac{1}{252}$, and $K = \Delta t^\alpha$, integers between 15.87 and 252 are within the required range. The unreported simulation study finds that if K is within the range, increasing K only elevates the computational burden without marginal contribution. We suggest the smallest integer K that satisfies the necessary condition as an optimal choice for K. Our specific recommendation of optimal window sizes for one-week, one-day, one-hour, 30-minute, 15-minute, and 5-minute data are 7, 16, 78, 110, 156, and 270, respectively.

1.4 Selection of rejection region

In this subsection, we address the rejection region for our proposed test. In effect, we demonstrate that the suggested rejection region allows us to distinguish jumps more precisely at higher frequencies of observation in Section 3.

As described in Theorems 1 and 2, our test statistics present completely different limiting behavior depending on the existence of jumps at the testing times. If there is no jump at the testing time, our test statistic follows approximately a normal distribution. If there is a jump, however, the test statistic becomes very large. To determine a reasonable rejection region, we raise a question of how large our test statistic can be when there is no jump. Hence, we first study the asymptotic distribution of maximums of our test statistics under the absence of jumps at any time in $(t_{i-1}, t_i]$. Such a distribution then guides us to choose the relevant threshold for the test to distinguish the presence of jumps at a testing time. Lemma 1 states the limiting distribution of the maximums as follows:

Lemma 1. *If the conditions for $\mathcal{L}(i)$, K, c, and $\bar{A}_n$ are as in Theorem 1, then as $\Delta t \to 0$,*

$$\frac{\max_{i \in \bar{A}_n} |\mathcal{L}(i)| - C_n}{S_n} \to \xi, \tag{12}$$

where ξ has a cumulative distribution function $P(\xi \leq x) = \exp(-e^{-x})$,

$$C_n = \frac{(2 \log n)^{1/2}}{c} - \frac{\log \pi + \log(\log n)}{2c(2 \log n)^{1/2}} \quad and \quad S_n = \frac{1}{c(2 \log n)^{1/2}}, \tag{13}$$

where n is the number of observations.

Proof of Lemma 1. *See Appendix A.4.*

In short, the main idea in selecting a rejection region is that if our observed test statistics are not even within the usual region of maximums, it is unlikely that the realized return is from the continuous part of the jump diffusion model. To apply this result for selecting a rejection region, for instance, we can set a significance level of 1%. Then, the threshold for $\frac{|\mathcal{L}(i)| - C_n}{S_n}$ is β^*, such that $P(\xi \leq \beta^*) = \exp(-e^{-\beta^*}) = 0.99$. Equivalently, $\beta^* = -\log(-\log(0.99)) = 4.6001$. Therefore, if $\frac{|\mathcal{L}(i)| - C_n}{S_n} > 4.6001$, then we reject the hypothesis of no jump at t_i.

2. Misclassifications

In this section, after defining four different types of misclassification that can occur in our study, we demonstrate that the probability of such a misclassification becomes negligible at higher frequencies of observation. For a single

testing time, say t_i, there can be two kinds of misclassification. The first is when there is a jump in interval $(t_{i-1}, t_i]$, but the test fails to reveal its existence. We call this a *failure to detect actual jump (FTD_i)* at t_i. The second kind of misclassification is when there is no jump in $(t_{i-1}, t_i]$, but the test wrongly concludes there is one. We call this a *spurious detection of jump (SD_i)* at t_i. It is usually the case that we do this test several times with time-series data. Global extension of these concepts is straightforward. If there are some jumps over the whole time horizon, $[0, T]$, but the test fails to detect any one of them, we call it a *global failure to detect actual jump (GFTD)*. If there are some returns that are not due to jumps, but the procedure wrongly declares any one of them as due to a jump, we call it a *global spurious detection of jump (GSD)*. We will use the following mathematical notations to explain the above situations. We let A_n be jump times among n observations and B_n be times at which the test declares the presence of a jump. We use J_i (J is for jumps) to denote the event that there is a jump in $(t_{i-1}, t_i]$. Note that $J_i = \{i \in A_n\}$. D_i (D is for declaring jumps) denotes the event that our test declares a jump in $(t_{i-1}, t_i]$. In this case, $D_i = \{i \in B_n\}$. Then, the following statements hold:

failure to detect actual jump at t_i (local property) $(FTD_i) = J_i \bigcap D_i^C$,

spurious detection of jump at t_i (local property) $(SD_i) = J_i^C \bigcap D_i$,

failure to detect actual jumps (global property) $(GFTD) = \bigcup_{i=1}^{n} \left(J_i \bigcap D_i^C \right),$

spurious detection of jumps (global property) $(GSD) = \bigcup_{i=1}^{n} \left(J_i^C \bigcap D_i \right).$

With these new notations, we now generalize the example at the end of Section 1, using a fixed significance level to any significance level α_n that approaches 0. Alternatively, β_n approaches ∞. In Theorem 3 and Corollary 1, we explicitly show that both the conditional and unconditional probabilities of global failure to detect actual jumps (*GFTD*) approach 0. They state specifically the convergence rates to show how fast these probabilities converge to 0.

Theorem 3. *Let β_n be the $(1 - \alpha_n)th$ percentile of the limiting distribution of ξ in Lemma 1, where α_n is the significance level of the test. Let N be the number of jumps in $[0, T]$. Then, the probability of global failure to detect actual jumps (GFTD) is*

$$P(GFTD|N) = \frac{2}{\sqrt{2\pi}} y_n N + o(y_n N), \qquad (14)$$

where $y_n = (\beta_n S_n + C_n)c\sigma\sqrt{\Delta t}$. *Therefore, as long as* $\beta_n \to \infty$ *slower than* $\sqrt{n \log n}$,

$$P(GFTD|N) \longrightarrow 0 \tag{15}$$

as $\Delta t \to 0$.

Proof of Theorem 3. *See Appendix A.5.*

If jumps arrive randomly, as in a Poisson process with constant intensity λ, one can expect the probability to be as in Corollary 1.

Corollary 1. *If $J(t)$ is a Poisson process with jump intensity λ, and the observation time horizon is from 0 to T, then*

$$E[P(GFTD)] = \frac{2}{\sqrt{2\pi}} y_n \lambda T + o(y_n \lambda T). \tag{16}$$

Theorem 4 also shows that the conditional probability of global spurious detection of jumps *(GSD)* approaches 0 quickly. The corresponding convergence rate is provided as the significance level α_n approaches 0 or, equivalently, the rejection threshold β_n approaches ∞.

Theorem 4. *Let β_n be as in Theorem 3. Again, let N be the number of jumps in $[0, T]$. Then, as $\Delta t \to 0$, the probability of spurious detection of jumps (GSD) is*

$$P(GSD|N) = \exp(-\beta_n) + o(\exp(-\beta_n)). \tag{17}$$

Therefore, as $\beta_n \to \infty$,

$$P(GSD|N) \longrightarrow 0, \tag{18}$$

as $\Delta t \to 0$.

Proof of Theorem 4. *See Appendix A.6.*

The immediate consequence of our findings in Theorems 3 and 4 is that we can classify actual and spurious jumps precisely so that we obtain an accurate stochastic jump intensity estimator based on our nonparametric test. If both probabilities in Theorems 3 and 4 become negligible, then the likelihood of global misclassification by either GFTD or GSD is also negligible, as stated in Theorem 5.

Theorem 5. *If $\widehat{\Lambda}(T)$ is the estimator of the number of jumps in $[0, T]$ using our test (which is formally a cumulative jump intensity estimator), and $\Lambda^{actual}(T) = N$ is the number of actually realized jumps in $[0, T]$, then the probability of global misclassification is*

$$P(\widehat{\Lambda} \neq \Lambda^{actual}|N) = P(GFTD \text{ or } GSD|N)$$
$$= \frac{2}{\sqrt{2\pi}} y_n N + \exp(-\beta_n) + o(\exp(-\beta_n)). \quad (19)$$

It can be minimized at $\beta_n^* = -\log(\frac{\sigma\sqrt{T}N}{\sqrt{2n\log n}})$. Moreover, the overall optimal convergence rate is $\frac{\sigma\sqrt{T}N}{\sqrt{2n\log n}}$.

Proof of Theorem 5. *See Appendix A.7.*

3. Monte Carlo Simulation

In this subsection, we examine the effectiveness of our test using a Monte Carlo simulation. Our asymptotic argument in previous sections requires that the sampling interval Δt converges to 0. This ideal requirement cannot be perfectly met in real applications. This subsection investigates the finite sample performance of the test. The main result from this simulation proves that as we increase the frequency of observation the precision of our test increases. For series generation, we used the Euler-Maruyama stochastic differential equation (SDE) discretization scheme (Kloeden and Platen, 1992), an explicit order 0.5 strong and order 1.0 weak scheme. We discard the burn-in period—the first part of the whole series—to avoid the starting value effect. Throughout, we use the notation $\Delta t = \frac{1}{252 \times nobs}$, with *nobs* as the number of observations per day.

3.1 Constant volatility
We first consider the simplest model in the class, with a fixed volatility. Table 1 presents the probability of spurious detection of jump (SD_i). We simulate two

Table 1
Probability of spurious detection $P(SD_i)$

freq	$\sigma = 0.3$	(SE)	$\sigma = 0.6$	(SE)	SV	(SE)
24-hour	1.3305e-03	(7.4050e-05)	1.3305e-03	(7.6239e-05)	3.9110e-03	(1.3319e-04)
12-hour	5.7380e-04	(3.4901e-05)	5.3222e-04	(3.3570e-05)	2.3306e-03	(7.7336e-05)
6-hour	2.0696e-04	(1.4460e-05)	2.1209e-04	(1.4790e-05)	1.3289e-03	(4.3670e-05)
2-hour	5.2879e-05	(4.3701e-06)	5.5911e-05	(4.2684e-06)	4.8131e-04	(1.6731e-05)
1-hour	2.1775e-05	(1.9032e-06)	2.5126e-05	(2.0353e-06)	2.7688e-04	(1.0062e-05)
30-minute	8.8436e-06	(8.3749e-07)	8.6768e-06	(8.3965e-07)	1.4467e-04	(7.2952e-06)
15-minute	3.4947e-06	(3.7430e-07)	4.1876e-06	(4.2736e-07)	8.9449e-05	(4.2818e-06)

The encompassing model is $d\log S(t) = \mu(t)dt + \sigma(t)dW(t)$. Constant volatility sets $\sigma(t) = \sigma$ at 30% and 60%. Stochastic volatility (SV) assumes the Affine model of Heston (1993), specified as $d\sigma^2(t) = (\theta_0 + \theta_1\sigma^2(t)) dt + \omega\sigma(t)dB(t)$, where $B(t)$ denotes a Brownian motion. The table shows means and standard errors (in parentheses) of probability of spurious detection $P(SD_i)$ at time t_i. The significance level α is 5%. *freq* denotes the frequency of observations.

Table 2
Probability of detecting actual jump $[1 - P(FTD_i)]$

Jump Size	3σ	2σ	1σ	0.5σ	0.25σ	0.1σ
			Constant volatility σ at 30%			
freq = 24-hour	0.9920	0.9880	0.9810	0.9270	0.4690	0.0260
	(0.0028)	(0.0034)	(0.0043)	(0.0082)	(0.0158)	(0.0050)
freq = 6-hour	0.9860	0.9780	0.9820	0.9700	0.9050	0.1520
	(0.0037)	(0.0046)	(0.0042)	(0.0054)	(0.0093)	(0.0114)
freq = 1-hour	0.9950	0.9860	0.9890	0.9890	0.9770	0.8880
	(0.0022)	(0.0037)	(0.0033)	(0.0033)	(0.0047)	(0.0100)
freq = 15-minute	0.9980	0.9970	0.9960	0.9920	0.9970	0.9820
	(0.0014)	(0.0017)	(0.0020)	(0.0028)	(0.0017)	(0.0042)

Jump Size	$3\widetilde{\sigma(t)}$	$2\widetilde{\sigma(t)}$	$1\widetilde{\sigma(t)}$	$0.5\widetilde{\sigma(t)}$	$0.25\widetilde{\sigma(t)}$	$0.1\widetilde{\sigma(t)}$
			Stochastic volatility (SV)			
freq = 24-hour	0.9470	0.9330	0.8540	0.5720	0.2500	0.0320
	(0.0071)	(0.0079)	(0.0112)	(0.0157)	(0.0137)	(0.0056)
freq = 6-hour	0.9770	0.9690	0.9410	0.8480	0.5320	0.1400
	(0.0047)	(0.0055)	(0.0075)	(0.0114)	(0.0158)	(0.0110)
freq = 1-hour	0.9870	0.9860	0.9830	0.9610	0.8770	0.5260
	(0.0036)	(0.0037)	(0.0041)	(0.0061)	(0.0104)	(0.0158)
freq = 15-minute	0.9970	0.9990	0.9980	0.9920	0.9610	0.8100
	(0.0017)	(0.0010)	(0.0014)	(0.0028)	(0.0061)	(0.0130)

The encompassing model is $d\log S(t) = \mu(t)dt + \sigma(t)dW(t) + Y(t)dJ(t)$. Constant volatility sets $\sigma(t) = \sigma$ at 30%. Stochastic volatility (SV) assumes the Affine model of Heston (1993), specified as $d\sigma^2(t) = \left(\theta_0 + \theta_1\sigma^2(t)\right)dt + \omega\sigma(t)dB(t)$, where $B(t)$ denotes a Brownian motion. The table contains means and standard errors (in parentheses) of probability of detecting actual jumps $[1 - P(FTD_i)]$ at time t_i. The significance level α is 5%. The jump sizes are set in comparison with volatility level: 3σ means jump sizes are set at three times of volatility level. For stochastic volatility, the jump size depends on the mean of volatility $\widetilde{\sigma(t)} = E[\sigma(t)]$. *freq* denotes the frequency of observations.

constant volatility diffusion processes with fixed spot volatilities at realistic annualized values of 30% and 60%, respectively. One thousand series of returns over one year are simulated at several different frequencies from daily to 15-minute returns. The significance level for this study is 5%. Table 1 shows that increasing the frequency of observations reduces the probability of spurious detection of jump (SD_i).

Table 2 shows the probability of success in detecting an actual jump, that is, 1 minus the probability of failure to detect actual jump (FTD_i). One thousand simulated tests at different frequencies, from daily to 15-minute, are performed. Arrivals of jumps with six different magnitudes are assumed at 300–10% of the given volatility level of 30%. We choose different jump sizes to present that it is harder to detect smaller-sized jumps at low frequency. However, we show that as we increase the frequency, we obtain very high detection power (above 98%), even for very small-sized jumps. For instance, we observe from Table 2 that when the magnitudes of relevant jumps are equal to 10% of volatility, econometricians are less likely to distinguish between price changes that result from diffusion and those from the actual jump in a low-frequency environment, such as with daily observations. This study shows that only 2% of jumps can be detected using daily returns. In contrast, at observation frequencies of every

30 minutes or higher, we can distinguish the difference more than 95% of the time.

3.2 Stochastic volatility

We examine how the test performs differently for stochastic volatility. We first consider a one-factor Affine model and compare with the constant volatility case, studied in Subsection 3.1. The model is as in Heston (1993), which is specified as the square root processes (used by Cox, Ingersoll, and Ross, 1985)

$$d\sigma^2(t) = \left(\theta_0 + \theta_1\sigma^2(t)\right)dt + \omega\sigma(t)dB(t), \tag{20}$$

where $B(t)$ denotes a Brownian motion. For the simulation, we used the parameter estimates from equity markets found in the empirical study by Bakshi et al., (2005), Table 2, to mimic the real equity markets. We assume no correlation between the Brownian motion in volatility and the random terms in the return process, which leaves us with no leverage effect in this simulation. The jump size and the Poisson jump counting process are set to be the same as in the case of constant volatility. We include the result when the volatility is stochastic in Tables 1 and 2 in order to allow direct comparison with the constant volatility case. It confirms our intuition that if the volatility changes over time, it would be more difficult to disentangle jumps. At every frequency, the corresponding spurious detection probability for stochastic volatility is greater than that for fixed volatility. We find the similar result from the comparison in Table 2: the probability of successful detection of a jump at some given time decreases under stochastic volatility. For another robustness check, we also study the case where the stochastic volatility is driven by a general model that accommodates stochastic elasticity of variance and a nonlinear drift (e.g., Aït-Sahalia, 1996; Aït-Sahalia, 2004; and Bakshi et al., 2005) as

$$\begin{aligned} d\sigma^2(t) = & \left(\theta_0 + \theta_1\sigma^2(t) + \theta_2\sigma^4(t) + \theta_3\sigma^{-2}\right)dt \\ & + \sqrt{\omega_0 + \omega_1\sigma^2(t) + \omega_2\sigma^{2\omega_3}}\,dB(t). \end{aligned} \tag{21}$$

Again, using the estimates for all the nested models in Bakshi et al., (2005), we obtain results similar to those under the Affine models presented earlier. The only impact of nonlinearity of drift is in the rate of convergence. In other words, at higher frequencies, such as 15-minute or higher, we find that the success rates of detecting actual jumps become similar to those for the Affine case, though at lower frequencies, such as daily, we have lower success rates. We also calculate the sample skewness and kurtosis that surrogate tail asymmetry and tail size and find that both decrease (in absolute terms for skewness) to normal cases as we increase frequency. Hence, our asymptotic result is not affected, although small-sample distributions of our test will be affected at lower frequencies such that the precision of our test at lower frequency decreases, as shown in the simulation results of Tables 1 and 2. To capture the effect of skewness and kurtosis at

Table 3
Global misclassification

Jump size	3σ	2σ	1σ	0.5σ	0.25σ
	Probability of global misclassification				
freq = 24-hour	0.2140	0.2120	0.1640	0.3540	0.9860
	(0.0184)	(0.0183)	(0.0166)	(0.0214)	(0.0053)
freq = 12-hour	0.1252	0.1312	0.1531	0.1173	0.7714
	(0.0148)	(0.0151)	(0.0161)	(0.0144)	(0.0188)
freq = 6-hour	0.0477	0.0407	0.0407	0.0506	0.0755
	(0.0095)	(0.0088)	(0.0088)	(0.0098)	(0.0118)
freq = 2-hour	0.0036	0.0046	0.0023	0.0063	0.0050
	(0.0027)	(0.0030)	(0.0021)	(0.0035)	(0.0031)
freq = 1-hour	0.0012	0.0012	0.0005	0.0007	0.0008
	(0.0015)	(0.0015)	(0.0010)	(0.0011)	(0.0013)
freq = 30-minute	0.0001	0.0001	0.0006	0.0004	0.0002
	(0.0003)	(0.0004)	(0.0011)	(0.0009)	(0.0006)
freq = 15-minute	0.0001	0.0001	0.0001	0.0002	0.0000
	(0.0001)	(0.0002)	(0.0005)	(0.0006)	(0.0001)

The table contains means and standard errors (in parentheses) of probability of global misclassification of jumps, discussed in Theorem 5. $P\left(\widehat{\Lambda(T)} \neq \Lambda^{\text{actual}}(T)\right)$ is by either global spurious detection of jumps (*GSD*) or global failure to detect actual jumps (*GFTD*), where $\widehat{\Lambda(T)}$ is the number of jumps in $[0, T]$ using our test, and $\Lambda^{\text{actual}}(T)$ is the number of actually realized jumps. The model under consideration is a jump diffusion process with Affine stochastic volatility and jumps of various sizes. The significance level α is 5%. *freq* denotes the frequency of observations.

lower frequencies, one may apply the Edgeworth expansion method. For an application of Edgeworth expansion in the context of realized volatility, see Zhang, Mykland, and Aït-Sahalia (2005a). For general discussion of Edgeworth expansion for dependent data, see Goetze and Hipp (1983) and Mykland (1993). In summary, the results show that stochastic volatility reduces the precision of jump detection. However, this study does confirm that if we increase the frequency of observation, we obtain improved power of jump detection.

3.3 Global misclassification

In this subsection, we examine the global likelihood of misclassification by either global spurious detection of jumps (*GSD*) or global failure to detect actual jumps (*GFTD*), as studied in Theorem 5. This tells us how accurately we can locate actual jump arrival times. We simulate 500 different series of one-year observations at different frequencies from daily to 15-minute returns. We consider five different jump sizes from 300% to 25% of volatility level. As discussed in Theorem 5, we find that the likelihood of misclassification becomes negligible at higher frequencies. Table 3 presents the probability that the number of jumps counted by our test is not equal to the actual number of jumps; this shows that global misclassification of jumps is negligible.

3.4 Comparison with other jump tests

Comparable tests to ours are those introduced by Barndorff-Nielsen and Shephard (2006; hereafter BNS) and Jiang and Oomen (2005; hereafter JO), both of which are also nonparametric approaches, making test results robust to

model specification. This subsection discusses the differences between these tests (BNS and JO) and ours (LM). Furthermore, we argue, by simulation, that our test outperforms these two tests. Our study measures the performance of tests by the probability of global success in detecting actual jumps and global spurious detection of jumps in one day, which can be alternatively taken as power and size for each test.

The main differences between the two others and ours are as follows. BNS takes the difference (or ratio) between the realized quadratic variation and bipower variation during a time interval to indicate the presence of jumps in that interval. JO's swap variance test solves the problem with a similar approach to BNS's. The only difference between BNS and JO is that instead of using bipower variation, JO uses cumulative delta-hedged gain or loss of a variance swap replicating strategy. They both provide asymptotic null distributions for their jump test statistics. One of the features of their tests, as described in BNS, is that they cannot distinguish two jumps a day with low variance and one jump with high variance in terms of detection rates. The reason for this problem is that these tests depend on integrated quantities. JO's swap variance test shares this feature because it also employs integrated quantities. For instance, suppose there are two jumps in a day and an analyst chooses one day as the interval for their test. Even if the presence of jumps in that day can be recognized, the other two tests cannot determine how many jumps occurred, whether the jump(s) was (were) negative or positive, at what time of day the jump(s) occurred, and how large each jump was. These issues can, however, be resolved by our test.

Furthermore, our test also outperforms the BNS and OJ tests. We compare the performance of the three tests in terms of probability of global success in detecting actual jumps within a given interval and probability of global spurious detection of jumps in that interval. Our test does not use the conventional terms of size and power, but introduces the misclassification of jumps in Section 2, because the detection criterion for our test is based on the distribution of *maximums* of null distribution, which is different from the usual hypothesis tests, and we do not select one model over another when we do a single test for jump detection. In essence, the probability of global success in detecting actual jumps is the *power* of the test, and the probability of global spurious detection of jumps is the *size* of the test.

We now report simulation results comparing the power and size of our new test to BNS's linear test and OJ's difference test. For the probability of global success to detect actual jumps (power), the simulation introduces one jump within one day to a diffusion process with a constant volatility. We consider 3000 simulated series of a process in one day, with each jump arriving randomly as a Poisson process. The number of jumps was set to be one per day. If there is a given number of jumps for some given time (one day in this study), then the Poisson jump arrival time is uniformly distributed (see Ross, 1995). Hence, we randomly select the arrival times from a uniform distribution. σ is set at 30%, as in the previous simulation. The standard deviation of jump size distribution is

Table 4
Comparison with other jump tests

Probability of global success to detect actual jumps $[1 - P(GFTD)]$					
100%	Our test (LM)		Linear test (BNS)		Difference test (OJ)	
freq	$1 - P(GFTD)$	(SE)	$1 - P(GFTD)$	(SE)	$1 - P(GFTD)$	(SE)
2-hour	0.8260	(0.0120)	0.8110	(0.0124)	0.7690	(0.0133)
1-hour	0.8720	(0.0106)	0.8460	(0.0114)	0.8210	(0.0121)
30-minute	0.9190	(0.0086)	0.8990	(0.0095)	0.8820	(0.0102)
15-minute	0.9410	(0.0075)	0.9260	(0.0083)	0.9200	(0.0086)
50%	Our test (LM)		Linear test (BNS)		Difference test (OJ)	
freq	$1 - P(GFTD)$	(SE)	$1 - P(GFTD)$	(SE)	$1 - P(GFTD)$	(SE)
2-hour	0.7480	(0.0137)	0.7110	(0.0143)	0.6470	(0.0151)
1-hour	0.8410	(0.0116)	0.8030	(0.0126)	0.7550	(0.0136)
30-minute	0.9070	(0.0092)	0.8590	(0.0110)	0.8390	(0.0116)
15-minute	0.9140	(0.0089)	0.8840	(0.0101)	0.8630	(0.0109)
25%	Our test (LM)		Linear test (BNS)		Difference test (OJ)	
freq	$1 - P(GFTD)$	(SE)	$1 - P(GFTD)$	(SE)	$1 - P(GFTD)$	(SE)
2-hour	0.6140	(0.0154)	0.5330	(0.0158)	0.4400	(0.0157)
1-hour	0.7520	(0.0137)	0.6380	(0.0152)	0.5540	(0.0157)
30-minute	0.8060	(0.0125)	0.7220	(0.0142)	0.6560	(0.0150)
15-minute	0.8690	(0.0107)	0.7960	(0.0127)	0.7470	(0.0138)

Probability of global spurious detection of jumps [P(GSD)]					
	Our test (LM)		Linear test (BNS)		Difference test (OJ)	
freq	$P(GSD)$	(SE)	$P(GSD)$	(SE)	$P(GSD)$	(SE)
2-hour	0.0000	(0.0000)	0.2426	(0.0061)	0.0484	(0.0030)
1-hour	0.0000	(0.0000)	0.1686	(0.0053)	0.0076	(0.0012)
30-minute	0.0000	(0.0000)	0.1190	(0.0046)	0.0010	(0.0004)
15-minute	0.0000	(0.0000)	0.0840	(0.0039)	0.0000	(0.0000)

The table shows the probability of global success in detecting actual jumps $[1 - P(GFTD)]$ (equivalent to power) and the probability of global spurious detection of jumps $P(GSD)$ (equivalent to size) within one day by our test (LM), linear test of Barndorff-Nielsen and Shephard (2006) (BNS), and difference test of Jiang and Oomen (2005) (JO). The null model is a diffusion process with constant volatility at 30% and the alternative of a jump diffusion process containing one jump per day with constant volatility at 30%. The time interval for integration of linear test (BNS) and difference test (OJ) is one day. *freq* denotes the frequency of observations. The jump sizes are 100%, 50%, and 25% of the volatility level.

chosen at 100%, 50%, 25%, and 10% of σ. Our test (LM) outperforms BNS and JO at all jump sizes. For the probability of global spurious detection of jumps (size), we apply the same analysis without introducing jumps in the one-day interval. Table 4 shows the superior performance of our test. In conclusion, our test performs better than BNS and JO for all jump sizes, numbers of jumps, and frequencies.

4. Implications of Identifying Jump Dynamics

Precise identification of jump arrival dynamics, including timing and jump size distribution, using our new test creates important implications for, among others, financial derivative security pricing. Because the test is applicable to all kinds of financial variables, such as stock returns and volatility, interest rates, and exchange rates, as long as high-frequency observations are available, we first discuss general implications for pricing equity options, index options, currency options, interest rate derivatives, volatility derivatives, and bond pricing.

Then we apply our test to real individual equity and S&P 500 Index returns and focus our discussion specifically on equity option markets based on our empirical findings.

4.1 Application to general option and bond pricing models

Much of the option and bond pricing literature estimates jumps by various parametric inference methods such as the Calibration method, (Implied State) Generalized Method of Moments (GMM), (Simulated) Maximum Likelihood Estimation (MLE), Efficient Method of Moment (EMM), or Bayesian approach (see Bakshi et al., 1997; Bates, 2000; Pan 2002; Schaumburg 2001; Chernov et al., 2003; Eraker, Johannes, and Polson, 2003; Aït-Sahalia, 2004; Piazzesi, 2003; and Aït-Sahalia and Jacod, 2005). As is well known, parametric approaches run the risk of incorrect specification for functionals in their chosen models. This is not the case with our nonparametric test. Many of the above-mentioned studies realize that jump parameters are difficult to pin down even with both time-series and cross-sectional data. Biased estimates can lead to increased security pricing errors. They also conclude that the models can be improved by introducing the stochastic jump intensity that captures the time-varying nature of the financial markets. Some studies allow jump arrival rates to depend on variables such as latent volatility, latent jump size, or market information in an Affine or non-Affine fashion (see Chernov et al., 2003; Piazzesi, 2003; and Dubinsky and Johannes, 2006, for instance). Though incorporating state variables is intuitively attractive, these existing (computationally intensive) parametric procedures become quite complicated to implement due to the discontinuity introduced by jumps and the increased number of parameters. Moreover, they often end up finding counterintuitively insignificant outcomes.

If our test is applied initially to determine the jump intensity and jump size distribution, and to search for observable variables that are significantly associated with jump arrivals, the inference for the whole model can become much simpler. This benefit is discussed in Tauchen and Zhou (2005), who applied Barndorff-Nielsen and Shephard (2006) to empirically prove that jump volatilities explain credit spread and credit default swap (CDS) spread better than other risk factors in the markets. Given the market information determined by the jump test, which can predict or explain jump arrivals, such as earnings announcements and other company-specific news releases for equities discussed in Subsection 4.2, equity option pricing models can be improved with reduced uncertainty involved with jumps.

4.2 Empirical analysis for U.S. equity option markets

In this subsection, we apply our new test on three major U.S. individual equity and S&P 500 Index returns to determine their jump dynamics, which may suggest different option pricing models as well as provide supporting evidence for flatter implied volatility smiles of individual equity options compared to index options (see Bakshi et al., 2003).

We use ultra-high-frequency data from transactions on the NYSE that are collected from the TAQ database. We calculate stock returns by taking differences of log transaction prices and multiplying all returns by 100 to present them as percentages. The time span is three months, from 1 September to 30 November 2005, which represent the most recent data available, and which have never been investigated in the literature. We choose three equities (Wal-Mart (WMT), IBM (IBM), and General Electric (GE)) and the S&P 500 Index to compare their different jump dynamics. Because the simulation study in Section 3 shows that a 15-minute frequency is high enough for our test to achieve a sufficient power to detect actual jumps, we choose 15-minute returns to generate the empirical evidence, and we present the results in Table 5. With this choice of frequency, results are not greatly affected by the presence of market microstructure noise. Therefore, the evidence presented here is robust to the effect of the noise. The significance level for all analyses is 5%. The outcomes of the tests are: each jump arrival date and time, jump size, and mean and variance of the detected jump size distribution. We do not assume that there is no more than one jump per day, but we do assume that when a jump occurs, the jump size dominates the return. After detecting the realized jumps, we search for real-time financial news and information releases around jump arrival times using the Factiva database in order to examine their association with jump arrivals. Sources used in the Factiva search include the *Wall Street Journal*, the *Financial Times*, and the Dow Jones and Reuters newswires.

Overall results indicate that jumps do not occur regularly, so that existing option pricing models with constant jump intensity are not appropriate. Detailed results for individual equities are explained as follows. First, we find that most jumps in equity prices under consideration arrive around the time of a market's opening. Except in one or two cases, jumps were always associated with news events. During the three months, there was one scheduled corporate news event for each company—the third-quarter earnings announcement. On the earnings announcement date (EAD), there was always a jump in the corresponding series. This evidence coincides with the individual equity option pricing model in Dubinsky and Johannes, (2006), who incorporate deterministic jump events conditional on scheduled EADs. However, as in Table 5, in addition to scheduled announcements, we find that a majority of stock price jumps are associated with other unscheduled company-specific news events. The implication of this finding is that scheduled earnings announcements as a deterministic jump predictor in earlier studies are not sufficient to model stochastic jump arrivals in individual equity prices. Other company-specific news releases, in addition to EADs, should be incorporated in accordance with the evidence in Table 5. Similarly, jumps in the S&P 500 Index are associated with market-wide news such as FOMC meetings and macroeconomic reports. This empirical finding is consistent with the notion that significant financial market jumps are related to responses to macroeconomic news announcements (see Andersen et al., 2003).

Table 5
Dynamics of jump arrival in individual equities and the S&P 500 index

Wal Mart (WMT)				
Date	Time	Size (%)	News	S
Sep 26	9:30 a.m.	1.29	Law Suit[1]	N
Oct 06	9:31 a.m.	1.12	Sales Up[3]	Y
Oct 10	9:30 a.m.	1.53	New Launch[5]	N
Oct 14	9:30 a.m.	0.93	Law Suit[7]	N
Oct 19	10:45 a.m.	0.84	Investment[9]	N
Oct 31	9:30 a.m.	1.33	Sales Up[11]	N
Nov 15	9:48 a.m.	−1.25	EAD[13]	N
Nov 18	9:30 a.m.	1.11	Announce[15]	N

$\mu_y^{WMT} = 0.8625$ $\sigma_y^{WMT} = 0.8817$

IBM (IBM)				
Date	Time	Size (%)	News	S
Sep 13	9:45 a.m.	−0.76	Announce[2]	N
Sep 19	9:30 a.m.	−0.74	Deal News[4]	N
Sep 21	9:30 a.m.	−0.96	Option Market[6]	N
Sep 27	9:45 a.m.	1.03	Good News[8]	N
Oct 07	9:45 a.m.	1.04	Good News[10]	N
Oct 10	9:30 a.m.	0.98	New Product[12]	N
Oct 11	9:30 a.m.	1.24	Expansion[14]	Y
Oct 18	9:30 a.m.	1.99	EAD[16]	Y
Oct 19	9:30 a.m.	−1.29	Rival EAD[17]	Y
Nov 18	9:30 a.m.	1.30	Announcement[18]	N

$\mu_y^{IBM} = 0.3830$ $\sigma_y^{IBM} = 1.1801$

General Electric (GE)				
Date	Time	Size (%)	News	S
Sep 15	9:45 a.m.	0.70	Good News[19]	N
Sep 21	9:30 a.m.	−0.95	Bad News[21]	N
Oct 06	9:31 a.m.	2.03	Announce[23]	N
Oct 07	9:30 a.m.	0.92	Good News[24]	Y
Oct 14	9:30 a.m.	1.11	EAD[25]	N
Nov 18	9:30 a.m.	2.20	Announce[26]	N

$\mu_y^{GE} = 1.0017$ $\sigma_y^{GE} = 1.1324$

S&P 500				
Date	Time	Size (%)	News	S
Sep 20	3:00 p.m.	−0.45	FOMC[20]	Y
Nov 29	9:30 a.m.	0.96	Good News[22]	Y

$\mu_y^{S\&P} = 0.28$ $\sigma_y^{S\&P} = 1.6260$

The above table contains results of our jump test: detected jump dates, jump times, and observed jump sizes with associated news events of three U.S. individual equities and the S&P 500 Index. They are based on transaction prices from the New York Stock Exchange (NYSE) during 3 months from 1 September to 30 November 2005. μ_y^{tic} and σ_y^{tic} are the mean and standard deviation of realized jump sizes of a company *tic* during one year from 1 July 2005 to 30 June 2006. The *p*-values of all the detected jumps reported here are less than 5%. S denotes whether the news was scheduled (Y) or not (N). FOMC denotes the Federal Open Market Committee meeting date and EAD denotes the earnings announcement dates of corresponding companies. The specific news-retrieval sources used in the Factiva search include the *Wall Street Journal*, the *Financial Times*, and the Dow Jones and Reuters newswires. Related news reported are as follows. 1. A potential Tommy Hilfiger suitor, 2. Global partnership pact with Nu Horizons Electronics Corp. 3. Sales meet its expectation, 4. Negative news of Siebel deal with Oracle, its rival in database software, 5. Metro 7, a new clothing brand launch, 6. Unusual option activity, 7. Uncollected sales taxes, 8. Delivery of most secure version of Linux available, 9. $25 million investment to establish a private-equity fund, 10. "Healthy" hardware sales booking of $12 billion or more, 11. Sales up above its prior forecast, 12. New virtualization technologies from software to systems, 13. Smallest profit gain in more than 4 years, 14. Speech solution offering, 15. Partnership with Fossil to produce fashion watch line, 16. Earnings fall on items, 17. EMC (its key rival) earnings net up 94%, 18. Global movement management initiative, 19. Financial conglomerate backed its earnings forecast unexpectedly, 20. Raise its target for the federal fund rate by 25 basis points, 21. Defense contracts, 22. Economic report showing a rebound in durable goods orders and consumer confidence improvement, 23. Syndicate with Morgan Stanley, 24. Solid growth in equipment and services, 25. Earnings rise 15%, 26. Final phase of insurance portfolio transformation.

Second, we compare the results for individual equities and the S&P 500 Index in Table 5. It shows that the frequency of jump occurrences in the index is much less than in the three individual equities. Another difference between jumps in individual equities and in the S&P 500 Index is the jump size distribution. The mean of jumps in the index (0.2036) is smaller than those in the three stocks (0.3014, 0.3487, and 0.4278). The variance of jumps in the index (0.5521) is also smaller than that in the stocks (1.2689, 0.9745, and 1.2050). We used one-year observations from 1 July 2005 to 30 June 2006 to obtain a reliable mean and variance, avoiding problems with small samples. In summary, individual stocks have more frequent jumps of greater size than the S&P 500 Index. According to definition of the index, the index should in theory jump when the individual component stock jumps. We refer, however, to detected jumps as those of empirically and economically significant magnitudes. The presence of jumps in both series implies that both individual stocks and the index have fat-tailed return distributions in physical measures. More frequent and larger-sized jumps in individual stock returns are likely to make the tails fatter.

Finally, this empirical outcome supports the explanation of Bakshi et al., (2003) for the implied volatility smiles of equity options. They economically prove that the negative risk-neutral skewness is feasible as long as the physical return distributions are fat-tailed and investors are risk averse. The evidence of different structure of fat-tailed return distributions supports their economic reasoning for the implied volatility curve, coupled with different levels of risk aversion of investors participating in the two different option markets.

5. Concluding Remarks

Realizing the difficulty in estimating jump parameters and the importance of incorporating market information into option pricing models, we introduce a new jump test to characterize the dynamic jump mechanism. Monte Carlo simulation proves that our test can precisely disentangle jump arrivals using high-frequency observations, allowing us to investigate the stochastic nature of jump dynamics and to search for observable information relevant to jump arrivals in any kind of financial time series. We perform an empirical study using our test and provide evidence of strong association between jumps and news events in the U.S. equity markets. Individual stock jumps occur with earnings announcements and other company-specific news releases, whereas the S&P 500 Index jumps with more overall market news such as FOMC meetings. These results suggest the need to consider company-specific news releases for individual equity option pricing and market-wide news releases for the index option pricing. We also find more frequent jumps with greater mean and variance in individual equities than the index returns, which supports the economic explanation by Bakshi et al., (2003) on negative risk-neutral skewness.

Though our test is designed to detect Poisson-type rare jumps in jump diffusion frameworks, it is still applicable to discrete observations from suitably modified pure jump models such as infinite-activity Lévy processes as in Carr and Wu (2004) and Wu (2006). Distinguishing very small magnitude jumps from such models is beyond the scope of the current article and is one possible direction for future research. Another extension of the test is to take into account the presence of market microstructure noise in ultra-high-frequency data when testing as in Andersen et al. (2006), which we are currently developing. Finally, since the test is designed to use high-frequency observations, its application on market microstructure problems would also be interesting.

Appendix A

A.1 The Nonzero Drift
The main conclusion of Theorem 1 is not altered for the case where we estimate the nonzero drift. A modified version of Definition 1 for this case is as follows:

Definition 1.1. *The statistic $\mathcal{L}_\mu(i)$, which tests at time t_i whether there was a jump from t_{i-1} to t_i, is defined as*

$$\mathcal{L}_\mu(t_i) \equiv \frac{\log S(t_i)/S(t_{i-1}) - \widehat{m}_i}{\widehat{\sigma(t_i)}}, \tag{A1}$$

where

$$\widehat{\sigma(t_i)} \text{ is as in Equation (8) and } \widehat{m}_i = \frac{1}{K-1} \sum_{j=i-K+1}^{i-1} (\log S(t_j)/S(t_{j-1})).$$

This demeans the return at time t_i using the average return in the window. The result similar to Theorem 1 holds for the case of nonzero drift in Theorem 1.1.

Theorem 1.1. *Let $\mathcal{L}_\mu(i)$ be as in Definition 1.1 and the window size $K = O_p(\Delta t^\alpha)$, where $-1 < \alpha < -0.5$. Suppose the process follows (1) or (2) and Assumption 1 is satisfied. Let $\bar{A}_n$ be the set of $i \in \{1, 2, \ldots, n\}$ so that there is no jump in $(t_{i-1}, t_i]$. Then, as $\Delta t \to 0$,*

$$\sup_{i \in \bar{A}_n} |\mathcal{L}_\mu(i) - \widehat{\mathcal{L}}_\mu(i)| = O_p(\Delta t^{\frac{3}{2} - \delta + \alpha - \epsilon}), \tag{A2}$$

where

$$\widehat{\mathcal{L}}_\mu(i) = \frac{U_i - \bar{U}_{i-1}}{c}.$$

Here, $\bar{U}_{i-1} = \frac{1}{K-1} \sum_{j=i-K+1}^{i-1} U_j$ and the rest of notations are the same as in Theorem 1.

Proof of Theorem 1.1. *See Appendix A.2.*

The error rate is the same as in the zero drift case, because the error due to the drift term is dominated by the error due to the diffusion part. $\widehat{\mathcal{L}}_\mu(i)$ asymptotically follows a normal distribution with its mean 0 and variance $\frac{K}{c^2(K-1)} \to \frac{1}{c^2}$, since $K \to \infty$. Our choice of window size K makes the effect of jumps on $\hat{m}_i$ vanish because of the property of the Poisson process for rare jumps: there can be no more than a single jump in an infinitesimal time interval. Because there can be only a finite number, say F, of jumps in the window, the statistic for the nonzero drift case becomes

$$\mathcal{L}_\mu(i) \approx \frac{U_i - \bar{U}_{i-1}}{c} - \frac{F \times Y(\tau)}{c\sigma \times (K-1)\sqrt{\Delta t}} I_{\tau \in (t_{i-K}, t_{i-1}]}. \tag{A3}$$

The second term will disappear because of the condition $K\sqrt{\Delta t} \to \infty$. This proves that jumps in the window have asymptotically negligible effect on testing jumps with our choice of K.

A.2 Proof of Theorems 1 and 1.1 in Appendix A.1
For $t_{i-K} < t < t_i$,

$$\log S(t) - \log S(t_{i-K}) = \int_{t_{i-K}}^t \mu(u)du + \int_{t_{i-K}}^t \sigma(u)dW(u). \tag{A4}$$

Given the imposition of Assumption 1, we have (from A1.1)

$$\int_{t_{i-1}}^{t_i} \mu(u)du - \mu(t_{i-1})\Delta t = O_p(\Delta t^{\frac{3}{2}-\epsilon}) \tag{A5}$$

and

$$\int_{t_{i-K}}^{t_i} \mu(u)du - \mu(t_{i-K})K\Delta t = O_p(\Delta t^{\frac{3}{2}+\alpha-\epsilon}) \text{ uniformly in all } i. \tag{A6}$$

This implies

$$\sup_{i,t \leq t_i} \left| \int_{t_{i-K}}^t \{\mu(u) - \mu(t_{i-K})\}du \right| = O_p(\Delta t^{\frac{3}{2}+\alpha-\epsilon}). \tag{A7}$$

Similarly to *Lemma* 1 in Mykland and Zhang (2006), under the condition A1.2, we can apply Burkholder's Inequality (Protter, 2004) to get

$$\sup_{i,t \leq t_i} \left| \int_{t_{i-K}}^t \{\sigma(u) - \sigma(t_{i-K})\}dW(u) \right| = O_p(\Delta t^{\frac{3}{2}-\delta+\alpha-\epsilon}), \tag{A8}$$

where δ can be any number in $0 < \delta < \frac{3}{2} + \alpha$. This result is also uniform in i for $K = O_p(\Delta t^\alpha)$ as specified. Therefore, over the window, for $t \in [t_{i-K}, t_i]$, $d \log S(t)$ can be approximated by $d \log S^i(t)$, such that

$$d \log S^i(t) = \mu(t_{i-K})dt + \sigma(t_{i-K})dW(t), \tag{A9}$$

because

$$
\begin{aligned}
&|(\log S(t) - \log S(t_{i-K})) - (\log S^i(t) - \log S^i(t_{i-K}))| \\
&= \left| \int_{t_{i-K}}^{t} (\mu(u) - \mu(t_{i-K}))du + \int_{t_{i-K}}^{t} (\sigma(u) - \sigma(t_{i-K}))dW(u) \right| \\
&= O_p(\Delta t^{\frac{3}{2} - \delta + \alpha - \epsilon}).
\end{aligned}
\tag{A10}
$$

For all i, j and $t_j \in [t_{i-K}, t_i]$, the numerator is

$$
\begin{aligned}
&\log S(t_j) - \log S(t_{j-1}) - \widehat{m}_i \\
&= \log S^i(t_j) - \log S^i(t_{j-1}) - \frac{1}{K-1} \sum_{l=i-K+1}^{i-1} \left(\log S^i(t_l) - \log S^i(t_{l-1}) \right) \\
&\quad + O_p(\Delta t^{\frac{3}{2} - \delta + \alpha - \epsilon}) \\
&= \sigma(t_{i-K})W_{\Delta t} - \frac{1}{K-1} \sum_{l=i-K+1}^{i-1} \sigma(t_{i-K})W_{\Delta t} + O_p(\Delta t^{\frac{3}{2} - \delta + \alpha - \epsilon}) \\
&= \sigma(t_{i-K})\sqrt{\Delta t}(U_j - \bar{U}_{i-1}) + O_p(\Delta t^{\frac{3}{2} - \delta + \alpha - \epsilon}),
\end{aligned}
\tag{A11}
$$

where $U_j = \frac{1}{\sqrt{\Delta t}}(W_{t_j} - W_{t_{j-1}}) \sim iid$ Normal $(0, 1)$ and $\bar{U}_{i-1} = \frac{1}{K-1} \sum_{j=i-K+1}^{i-1} U_j$.

For the denominator, we use the instantaneous volatility estimator based on the realized bipower variation (see Barndorff-Nielsen and Shephard, 2004). According to *Proposition* 2 in Barndorff-Nielsen and Shephard, (2004), the impact of the drift term is negligible; hence, it does not affect the asymptotic limit behavior. Then, we prove the following approximation of the scaled volatility estimator:

$$
\begin{aligned}
&\text{plim}_{\Delta t \to 0} c^2 \hat{\sigma}^2(t) \\
&= \text{plim}_{\Delta t \to 0} \frac{1}{(K-2)\Delta t} \sum_{j} |\log S(t_j)/S(t_{j-1})| |\log S(t_{j-1})/S(t_{j-2})| \\
&= c^2 \sigma^2(t).
\end{aligned}
\tag{A12}
$$

It is due to

$$\frac{1}{(K-2)\Delta t} \sum_{j=i-K+3}^{i-1} |\log S(t_j)/S(t_{j-1})||\log S(t_{j-1})/S(t_{j-2})|$$

$$= \frac{1}{(K-2)\Delta t} \sum_{j=i-K+3}^{i-1} \left|\log S^i(t_j)/S^i(t_{j-1}) + O_p(\Delta t^{\frac{3}{2}-\delta+\alpha-\epsilon})\right|\right|$$

$$\times \log S^i(t_{j-1})/S^i(t_{j-2}) + O_p(\Delta t^{\frac{3}{2}-\delta+\alpha-\epsilon})\right|$$

$$= \frac{1}{(K-2)\Delta t} \sum_{j=i-K+3}^{i-1} |\log S^i(t_j)/S^i(t_{j-1})||\log S^i(t_{j-1})/S^i(t_{j-2})|$$

$$+ O_p(\Delta t^{\frac{3}{2}-\delta+\alpha-\epsilon})$$

$$= \frac{1}{K-2} \sum_{j=i-K+3}^{i-1} \sigma^2(t_{i-K})|\sqrt{\Delta t}U_j||\sqrt{\Delta t}U_{j-1}| + O_p(\Delta t^{\frac{3}{2}-\delta+\alpha-\epsilon})$$

$$= \sigma^2(t_{i-K})c^2 + O_p(\Delta t^{\frac{3}{2}-\delta+\alpha-\epsilon}), \tag{A13}$$

where U_i's are iid *Normal*(0, 1) and $c = E(|U_i|) \approx 0.7979$. Then,

$$\mathcal{L}(i) = \frac{(U_i - \bar{U}_{i-1})}{c} + O_p(\Delta^{\frac{3}{2}-\delta+\alpha-\epsilon}). \tag{A14}$$

This proves Theorems 1 and 1.1.

Alternatively, for the nonzero drift case, we can use Girsanov's theorem to suppose $\mu(t) = 0$, as in Zhang et al. (2005b).

A.3 Proof of Theorem 2
When there is a possibility of rare Poisson jumps in the window, the scaled bipower variation $c^2\hat{\sigma}^2(t)$ can be decomposed into two parts: one with jump terms and other without jump terms, as follows:

$$c^2\hat{\sigma}^2(t)$$

$$= \text{terms without jumps} + \frac{1}{(K-2)\Delta t} \sum_{\text{terms with jumps}} \sigma(t_j)|\Delta W_{t_j}||Jump|$$

$$= \text{terms without jumps} + \frac{1}{(K-2)\Delta t} O_p(\sqrt{\Delta t}) \sum_{\text{terms with jumps}} \sigma(t_j)|Jump|$$

$$= \text{terms without jumps} + \frac{1}{(K-2)\Delta t} O_p(\sqrt{\Delta t}). \tag{A15}$$

The order of the second term is due to the property of Poisson jump processes that allows a finite number of jumps over the window. Since $\sigma(t_j)|Jump| = O_p(1)$, $\sum_{\text{terms with jumps}} \sigma(t_j)|Jump| = O_p(1)$. For the effect of jump terms to

become negligible, the second term must be $o_p(1)$ as Δt goes to 0. The window size K that satisfies $K\sqrt{\Delta t} \to \infty$ and $K\Delta t \to 0$ as Δt goes to 0 will work. If we assume the window size to be $K = \Delta t^\alpha$, then the necessary condition for α is $-1 < \alpha < -0.5$. Accordingly,

$$\lim_{\Delta t \to 0} \hat{\sigma}^2|_{\text{alternative}} = \lim_{\Delta t \to 0} \hat{\sigma}^2|_{\text{null}} = \sigma^2(t). \tag{A16}$$

Then, putting the approximation for return above in the statistic yields

$$\mathcal{L}(i) \approx \frac{\sigma_{t_{i-K}}\sqrt{\Delta t}U_i + Y(\tau)I_{\tau \in (t_{i-1}, t_i)}}{c\sigma(t_{i-K})\sqrt{\Delta t}} = \frac{U_i}{c} + \frac{Y(\tau)}{c\sigma\sqrt{\Delta t}}I_{\tau \in (t_{i-1}, t_i)}. \tag{A17}$$

A.4 Proof of Lemma 1
Proof of Lemma 1 follows from Aldous (1989) and the proof in Galambos (1978).

A.5 Proof of Theorem 3
Let N be the number of jumps from time $t = 0$ to $t = T$. We claim there is a jump if $|\mathcal{L}(i)| > \beta_n S_n + C_n$. Fix a set of jump times as $A_n = \{i : \text{there is a jump in } (t_{i-1}, t_i]\}$. Then,

$P(\text{We correctly classify all N jumps}|N)$
$$= P(\text{For all } i \in A_n, |\mathcal{L}(i)| > \beta_n S_n + C_n)$$
$$\approx \prod_{i \in A_n} P(|\mathcal{L}(i)| > \beta_n S_n + C_n) \approx \prod_{i \in A_n} P(|Y(t_i)| > (\beta_n S_n + C_n)c\sigma\sqrt{\Delta t})$$
$$= \prod_{i \in A_n} \left(1 - F_{|Y|}(y_n)\right) \sim \left(1 - \frac{2}{\sqrt{2\pi}}y_n + o(y_n^2)\right)^N$$
$$= 1 - \frac{2}{\sqrt{2\pi}}y_n N + o(y_n^2 N). \tag{A18}$$

A.6 Proof of Theorem 4
Let $A_n^C = \{1, \ldots, n\} - A_n$ be a set of non-jump times. Then,

$P(\text{We incorrectly reject any non-jumps}|N)$
$$= P(\text{for some } i \in A_n^C, |\mathcal{L}(i)| > \beta_n S_n + C_n|N)$$
$$= P\left(\max_{i \in A_n^C} |\mathcal{L}(i)| > \beta_n S_n + C_n|N\right)$$
$$\approx P\left(\max_{i \in A_n^C} |\widehat{\mathcal{L}}(i)| > \beta_n S_n + C_n\right) = 1 - F_\xi(\beta_n) = \exp(-\beta_n) + o(\exp(-\beta_n)).$$
$$\tag{A19}$$

By L'Hopital's rule, we obtain the last step because

$$\lim_{\beta_n \to \infty} \frac{1 - F_\xi(\beta_n)}{\exp(-\beta_n)} = 1. \tag{A20}$$

A.7 Proof of Theorem 5

Prob(*GMJ or GMNJ*) = Prob(*GMJ*) + Prob(*GMNJ*) and the results follow from Theorems 3 and 4. Minimum probability can be achieved at β_n^*, which can be obtained by taking the first derivative of probability with respect to β_n and setting it equal to 0, as

$$\frac{\partial P}{\partial \beta_n} = \frac{2}{\sqrt{2\pi}} S_n \sigma c \sqrt{\Delta t} N - \exp(-\beta_n) = 0. \tag{A21}$$

References

Aït-Sahalia, Y. 1996. Testing Continuous-Time Models of the Spot Interest Rate. *Review of Financial Studies* 9:385–426.

Aït-Sahalia, Y. 2004. Disentangling Diffusion from Jumps. *Journal of Financial Economics* 74:487–528.

Aït-Sahalia, Y., and J. Jacod. 2005. Fisher's Information for Discretely Sampled Levy Processes. Working Paper, Department of Economics, Princeton University.

Aldous, D. 1989. *Probability Approximations via the Poisson Clumping Heuristic*. New York: Springer-Verlag.

Andersen, T., T. Bollerslev, F. Diebold, and C. Vega. 2003. Micro Effects Macro Announcements: Real-Time Price Discovery in Foreign Exchange. *American Economic Review* 93:38–62.

Andersen, T., T. Bollerslev, and D. Dobrev. 2006. No-Arbitrage Semi-Martingale Restrictions for Continuous-Time Volatility Models Subject to Leverage Effects, Jumps, and i.i.d. Noise: Theory and Testable Distributional Implications. *Journal of Econometrics*, 138:125–80.

Bakshi, G., C. Cao, and Z. Chen. 1997. Empirical Performance of Alternative Option Pricing Models. *Journal of Finance* 52:2003–49.

Bakshi, G., C. Cao, and Z. Chen. 2000. Pricing and Hedging Long-Term Options. *Journal of Econometrics* 94:277–318.

Bakshi, G., N. Kapadia, and D. Madan. 2003. Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. *Review of Financial Studies* 16:101–43.

Bakshi, G., N. Yu, and H. Ou-Yang. 2005. Estimation of Continuous-Time Models with an Application to Equity Volatility. *Journal of Financial Economics*, 82:227–49.

Barndorff-Nielsen, O. E., and N. Shephard. 2004. Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics* 2:1–48.

Barndorff-Nielsen, O. E., and N. Shephard. 2006. Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation. *Journal of Financial Econometrics* 4:1–30.

Bates, D. S. 1996. Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsch Mark Options. *Rievew of Financial Studies* 9:69–107.

Bates, D. S. 2000. Post-'87 Crash Fears in the S&P 500 Futures Option Market. *Journal of Econometrics* 94:181–238.

Bertsimas, D., L. Kogan, and A. Lo. 2001. Hedging Derivative Securities and Incomplete Markets: An Epsilon-Arbitrage Approach. *Operations Research* 49:372–97.

Carr, P., and L. Wu. 2004. Time-Changed Lévy Processes and Option Pricing. *Journal of Financial Economics* 71:113–41.

Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen. 2003. Alternative Models for Stock Price Dynamics. *Journal of Econometrics* 116:225–57.

Collin-Dufresne, P., and J. Hugonnier. 2001. Event Risk, Contingent Claims and the Temporal Resolution of Uncertainty. Working Paper, Tepper School of Business, Carnegie Mellon University.

Cox, J. T., J. E. Ingersoll, and S. A. Ross. 1985. A Theory of the Term Structure of Interest Rates. *Econometrica* 53:385–407.

Dubinsky, A., and M. Johannes. 2006. Earnings Announcements and Option Prices. Working Paper, Graduate School of Business, Columbia University.

Duffie, D., J. Pan, and K. Singleton. 2000. Transform Analysis and Asset Pricing for Affine Jump Diffusions. *Econometrica* 68:1343–76.

Eraker, B., M. Johannes, and N. Polson. 2003. The Impact of Jumps in Volatility and Returns. *Journal of Finance* 53:1269–300.

Galambos, J. 1978. *The Asymptotic Theory of Extreme Order Statistics*. New York: John Wiley & Sons.

Goetze, F., and C. Hipp. 1983. Asymptotic Expansions for Sums of Weakly Dependent Random Vectors. *Z. Wahrsch. Verw. Gebiete* 64:211–39.

Heston, S. 1993. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bonds and Currency Options. *Review of Financial Studies* 6:327–43.

Jiang, G. J., and R. Oomen. 2005. A New Test for Jumps in Asset Prices. Working Paper, Eller College of Management, University of Arizona.

Johannes, M. 2004. The Statistical and Economic Role of Jumps in Interest Rates. *Journal of Finance* 59:227–60.

Kloeden, P. E., and E. Platen. 1992. *Numerical Solution of Stochastic Differential Equations*. New York: Springer-Verlag.

Liu, J., F. Longstaff, and J. Pan. 2003. Dynamic Asset Allocation with Event Risk. *Journal of Finance* 58:231–59.

Merton, R. C. 1976. Option Pricing When Underlying Stock Returns Are Discontinuous. *Journal of Financial Economics* 3:125–44.

Mykland, P. 1993. Asymptotic Expansions for Martingales. *Annals of Probability* 21:800–18.

Mykland, P. A., and L. Zhang. 2006. ANOVA for Diffusions and Ito processes. *Ann. Statist.*, 34:1931–63.

Naik, V., and M. Lee. 1990. General Equilibrium Pricing of Options on the Market Portfolio with Discontinuous Returns. *Review of Financial Studies* 3:493–521.

Pan, J. 2002. The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics* 63:3–50.

Piazzesi, M. 2003. Bond Yields and the Federal Reserve. *Journal of Political Economy* 113:311–44.

Pollard, D. 2002. *A User's Guide of Measure Theoretic Probability*. Cambridge: Cambridge University Press.

Protter, P. 2004. *Stochastic Integration and Differential Equations: A New Approach*. New York: Springer.

Ross, S. M. 1995. *Stochastic Processes*. New York: John Wiley & Sons.

Schaumburg, E. 2001. Maximum Likelihood Estimation of Jump Processes with Applications to Finance. Ph.D. Thesis, Department of Economics, Princeton University.

Tauchen, G., and H. Zhou. 2005. Identifying Realized Jumps on Financial Markets. Working Paper, Department of Economics, Duke University.

Wu, L. 2006. Dampened Power Law: Reconciling the Tail Behavior of Financial Asset Returns. *Journal of Business* 79:1445–74.

Zhang, L., P. A. Mykland, and Y. Aït-Sahalia. 2005a. Edgeworth Expansions for Realized Volatility and Related Estimators. Working Paper, University of Illinois at Chicago.

Zhang, L., P. A. Mykland, and Y. Aït-Sahalia. 2005b. A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data. *Journal of American Statistical Association* 472:1394–411.