CONTINUOUS-TIME MODELS, REALIZED VOLATILITIES, AND TESTABLE DISTRIBUTIONAL IMPLICATIONS FOR DAILY STOCK RETURNS

TORBEN G. ANDERSEN,[a] TIM BOLLERSLEV,[b]* PER FREDERIKSEN[c] AND
MORTEN ØRREGAARD NIELSEN[d]

[a] *Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL, USA, NBER and CREATES*
[b] *Department of Economics, Duke University, Durham, NC, USA, and CREATES*
[c] *Equity Trading and Derivatives, Nordea Markets, Copenhagen, Denmark*
[d] *Department of Economics, Queen's University, Kingston, Ontario, Canada, and CREATES*

SUMMARY

We provide an empirical framework for assessing the distributional properties of daily speculative returns within the context of the continuous-time jump diffusion models traditionally used in asset pricing finance. Our approach builds directly on recently developed realized variation measures and non-parametric jump detection statistics constructed from high-frequency intra-day data. A sequence of simple-to-implement moment-based tests involving various transformations of the daily returns speak directly to the importance of different distributional features, and may serve as useful diagnostic tools in the specification of empirically more realistic continuous-time asset pricing models. On applying the tests to the 30 individual stocks in the Dow Jones Industrial Average index, we find that it is important to allow for both time-varying diffusive volatility, jumps, and leverage effects to satisfactorily describe the daily stock price dynamics. Copyright © 2009 John Wiley & Sons, Ltd.

Supporting information may be found in the online version of this article.

1. INTRODUCTION

The distributional properties of speculative prices, and stock returns in particular, rank among the most studied empirical phenomena in all of economics. We add to this burgeoning literature by showing how high-frequency intra-day data and realized variation measures may be used in the construction of simple-to-implement tests for the importance of jumps and so-called leverage effects. Our empirical results for the 30 individual stocks in the Dow Jones Industrial Average (DJIA) index support the notion that daily stock prices may be viewed as discretely sampled observations from an arbitrage-free jump-diffusive process, but that time-varying volatility, jumps and leverage effects are all present and must be accommodated if the fundamental arbitrage-free semi-martingale characterization is to be sustained.

A long line of studies, dating back to the seminal work of Mandelbrot (1963) and Fama (1965), documents that the unconditional distributions of day-to-day and longer horizon stock returns exhibit fatter tails than the normal distribution. Correspondingly, a large literature seeks to describe and explain this empirical regularity through alternative non-normal distributions, often inspired

---

* Correspondence to: Tim Bollerslev, Department of Economics, Duke University, Durham, NC 27708, USA.
E-mail: boller@econ.duke.edu

by the Mixture-of-Distributions Hypothesis (MDH) originally proposed by Clark (1973). The basic MDH stipulates that prices only move in response to new information, or 'news'. While the basic MDH treats the mixing variable as latent, Clark (1973), Epps and Epps (1976), and Tauchen and Pitts (1983) also relate it with trading volume.

Early studies focus on the unconditional distributional implications of the MDH. However, it is now well established that key features of the conditional return distribution, and the conditional variance in particular, are highly predictable (e.g., Engle, 2004). The pronounced predictability in volatility motivated empirical studies exploring the relationship between return variability and fundamental mixing variable(s) within the MDH context (e.g., Gallant *et al.*, 1992; Andersen, 1996; Liesenfeld, 1998; Bollerslev and Jubinski, 1999; Ane and Geman, 2000).[1]

In spite of the presence of such structured MDH approaches, the more ad hoc (G)ARCH class of models arguably ranks supreme for empirically characterizing conditional inter-daily return distributions (see, e.g., Andersen *et al.*, 2006a). Beyond providing a parsimonious and tractable approach to the time-varying return volatility, this literature has also uncovered a striking asymmetry between equity returns and volatility; i.e., large negative returns tend to be associated with higher future volatility than positive returns of the same magnitude. This asymmetry, forcefully documented by Nelson (1991), is generically labeled a leverage effect, although it is widely agreed that the effect has little to do with financial leverage.[2]

In contrast to the discrete-time formulations employed in the empirical MDH and (G)ARCH literatures, many important developments in theoretical asset pricing, and derivatives pricing in particular, are based on continuous-time models. For instance, the Black–Scholes option pricing formula assumes that prices evolve according to a homogeneous diffusion process. This assumption is obviously at odds with the leptokurtic unconditional daily return distributions, the pronounced volatility clustering, and the leverage effects discussed above, and much recent progress has been made in terms developing more empirically realistic continuous-time formulations. In particular, while the early contributions by Merton (1976) and Hull and White (1987) argue for the need to incorporate jumps and time-varying diffusive volatility in the pricing of options, respectively, recent studies document the need to simultaneously allow for both effects in order to satisfactorily represent observed security prices (e.g., Andersen *et al.*, 2002; Chernov *et al.*, 2003).

In this paper we combine insights from these separate strands of the literature by providing a framework for analyzing the distributional properties of discrete-time daily returns implied by a broad class of jump-diffusive models. Our approach is distinctly nonparametric and relies critically on the availability of high-frequency data for the construction of realized volatility measures. High-frequency, or tick-by-tick, prices have recently become available for a host of different financial instruments and markets, and the analysis of the corresponding realized variation measures have already provided important new empirical insights concerning the distributional properties and dynamic dependencies in financial market volatilities (see, e.g., Andersen and Bollerslev, 1998a; Andersen *et al.*, (2001a, 2003); Barndorff-Nielsen and Shephard, 2002). Pushing this analysis one step further, we show how the realized volatility measures may be used in the formulation of direct distributional tests for continuous-time models.

---

[1] The robustness of the empirical findings in Ane and Geman (2000) have recently been called into question by Gillemot *et al.* (2005) and Murphy and Izzeldin (2006).

[2] In fact, as discussed in more detail below, the asymmetry hitherto documented with daily and lower frequency data tends to be much more pronounced for aggregate equity index returns as opposed to individual stock returns, indirectly casting doubt on the financial leverage explanation.

Our empirical analysis provides the first comprehensive documentation that a broad set of individual equity return series may be converted into i.i.d. Gaussian series through a sequence of simple, theoretically motivated, nonparametric transformations. It may be seen as a logical extension of the earlier empirical investigations of Andersen *et al.* (2001b), who find the unconditional distributions of raw daily equity returns to have fat tails, but when standardizing these daily returns by the corresponding realized volatilities, constructed from the summation of high-frequency intraday squared returns, the distributions appear close to Gaussian. Nonetheless, it remains an approximate result as the null hypothesis of i.i.d. normality is rejected decisively if subjected to powerful statistical tests. From a theoretical perspective, this is not surprising. The (true) realized volatility standardized returns should be indistinguishable from a Gaussian if the true price process belongs to a certain class of pure diffusive processes and market microstructure frictions are negligible. However, various relevant market features may invalidate this result. First, there are inevitable errors in realized volatility measures due to discretization and noise. Second, it is likely there are discontinuities in the price path so the returns are not generated from a pure diffusion. Third, price and volatility innovations may be correlated, inducing asymmetry in the standardized return distribution.

To potentially obtain normality, each source of error warrants careful attention. We introduce a new set of diagnostics for guiding the choice of sampling frequency. These generalized volatility signature plots are designed to display the effects of microstructure noise as well as price jumps. Second, in addition to standard realized volatility measures we rely on the bipower variation measures of Barndorff-Nielsen and Shephard (2004) for separately measuring the continuous sample path variability and the variation due to jumps. Moreover, we extend the test for the occurrence of at least one jump per day in Barndorff-Nielsen and Shephard (2006) and Huang and Tauchen (2005) to a sequential jump detection scheme, directly identifying and estimating the within-day times and sizes of price jumps.[3] This allows us to construct jump-adjusted daily return series, while the extracted jump characteristics enable us to more directly gauge the impact and distributional implications of jumps.[4] Third, to alleviate the impact of return-volatility asymmetries, e.g., the leverage effect, we exploit a new financial-time sampling scheme in which we measure returns in event time, as defined by equidistant increments to the realized volatility of the jump-adjusted returns. In the diffusive semi-martingale setting, this realized volatility time-change should undo the impact of leverage style effects so that the financial-time return distributions become Gaussian. Again, all involved measures are obtained nonparametrically and the distributional implications are based strictly on probabilistic arguments, so that implied tests are applicable across the full range of standard jump diffusive models for asset returns.

Our approach is related to Peters and de Vilder (2006) and Andersen *et al.* (2007b) as they rely on a similar financial-time sampling and also undertake normality tests for the standardized return distributions.[5] However, they explore only a return series generated from futures contracts on the S&P500 equity index. These futures are near ideal in terms of having minimal microstructure distortions and high liquidity. We focus on the much broader set of 30 individual equity return series for the companies in the Dow Jones index. As a result, our series are subject to more noise, have more idiosyncratic return and volatility movements and have much higher volatility in

---

[3] Alternative non-parametric high-frequency data-based tests for jumps have recently been developed by Jiang and Oomen (2008), Mancini (2005), Christensen and Podolski (2007), and Lee and Mykland (2008).

[4] In concurrent and independent work, Fleming and Paye (2006) have studied the properties of daily returns scaled by realized bipower variation, but without any adjustments for leverage effects.

[5] See also Zhou (1998) for more informal empirical evidence along these lines for exchange rates.

general. Hence, we are able to shed light on the question of whether the findings from the benign setting studied previously carry over to a wider range of important return series. Furthermore, we obtain important evidence regarding the robustness of the prior studies. From a methodological point of view there are also a number of differences. Most strikingly, Peters and de Vilder (2006) make no adjustments for jumps and rely on tests with much less power than is the case in the current paper. Compared to Andersen *et al.* (2007b) we accommodate the issue of microstructure noise more directly through the generalized signature plots and rely on the new sequential jump detection technique. In fact, our identification of jump days is justified through the asymptotic distribution of the standard test statistic under a general diffusive null hypothesis as developed in Barndorff-Nielsen and Shephard (2004, 2006) and Huang and Tauchen (2005). The jump identification approach in Andersen *et al.* (2007b) appears to fare well in simulation settings but it is formally justified only under constant volatility across the trading day and the jump test turns more conservative under intraday variation in volatility. Moreover, given the potential importance of microstructure frictions, we explore whether the change of tick size for the Dow Jones stocks around January 2001 impacts the number of rejections of normality across the series. Such evidence is only meaningful on the basis of a large number of return series and the issue was not explored previously in this context. Finally, the empirical results related to the strength of the jump intensities and sizes, and the significance and magnitude of the leverage effects, for the individual stocks are of direct interest in their own right for a range of issues within financial economics.

The plan for the rest of the paper is as follows. The theoretical arguments for Gaussianity of the transformed return distributions are outlined in the next section. The realized variation measures and jump detection tests used in the practical implementation of the distributional tests are presented in Section 3. In Section 4 we discuss the data sources and issues related to the construction of the high-frequency returns and realized volatility measures, including generalized volatility signature plots designed to assess the adverse effects of market microstructure biases at the very highest sampling frequencies. Section 5 discusses preliminary summary statistics related to the importance of jumps and leverage effects. The outcomes of the distributional tests are summarized in Section 6. Section 7 concludes. More detailed evidence for each individual stock is available in a supplementary appendix on the journal's website.

2. THEORETICAL FRAMEWORK

Jump-diffusion models represent the asset price as a sum of a continuous sample path component and occasional discontinuous jumps. The class encompasses the leading parametric models in the asset pricing and, especially, the derivatives pricing literature.[6]

In particular, let $p(t)$ denote the continuous-time log-price process. The generic jump-diffusion model may then be expressed in stochastic differential equation form as

$$\mathrm{d}p(t) = \sigma(t)\,\mathrm{d}w\,(t) + \kappa(t)\,\mathrm{d}q\,(t), \quad t \geq 0 \tag{1}$$

where the instantaneous volatility process $\sigma(\cdot) > 0$ is càdlàg, $w(\cdot)$ denotes a standard Brownian motion independent of the drift, the counting process $q(t)$ is normalized so that $\mathrm{d}q(t) = 1$ represents

---

[6] Although this formulation, as given by equation (1), allows for both time-varying jump sizes and intensities, it rules out infinite activity Lévy processes; see, e.g., Cont and Tankov (2004) for a discussion of such processes, Todorov (2009) for an application involving jump-driven stochastic volatility models, and Barndorff-Nielsen *et al.* (2006) on using realized variation measures for certain infinite activity jump processes.

a jump at time t, and $dq(t) = 0$ otherwise, and $\kappa(t)$ denotes the jump size if a jump occurs at time t. For notational simplicity we exclude a drift term, $\mu(t)dt$, in equation (1), but the theoretical results can readily be extended to allow for a drift, $\mu(\cdot) \neq 0$.[7]

While asset pricing arguments often are cast in continuous time, empirical investigations are invariably based on discretely sampled prices. We denote the one-period continuously compounded discrete-time returns implied by the jump diffusion in (1) as

$$r_t \equiv p(t) - p(t-1), \quad t = 1, 2, \ldots \tag{2}$$

and we refer to the unit time interval as a 'day'. The distributional characteristics of the discrete-time returns obviously depend directly on the underlying continuous-time model. We next consider three sets of increasingly general modeling assumptions, and discuss how appropriately standardized and adjusted returns should be i.i.d. standard normal under each, thus providing theoretical guidance for empirical analysis into the importance of different model features.

2.1. No Jumps, Leverage, or Volatility Feedback Effects

The simplest and most commonly used continuous-time models are based on the dual assumptions of no jumps, or $q(t) \equiv 0$, along with no leverage and volatility feedback effects, or $\sigma(t)$ and $w(\tau)$ independent for all $t \geq 0$ and $\tau \geq 0$. In this situation it follows by standard arguments that

$$r_t \left(\int_{t-1}^{t} \sigma^2(\tau)d\tau \right)^{-1/2} \sim N(0, 1), \qquad t = 1, 2, \ldots \tag{3}$$

The *integrated volatility* normalizing the returns has the interpretation of the *ex post* return variability conditional on the sample path realization of the $\sigma(\tau)$ process over the corresponding discrete-time return interval, $(t-1, t]$.[8] Of course, the integrated volatility is not directly observable. However, starting with the work of Andersen and Bollerslev (1998a), Andersen *et al.* (2001a), and Barndorff-Nielsen and Shephard (2002), ways in which to accurately measure the integrated volatility on the basis of high-frequency data have received increasing attention in the literature. In Section 3 we provide a more in-depth discussion of these ideas in the context of our empirical implementation of equation (3).

Meanwhile, the popular GARCH and discrete-time stochastic volatility models in essence provide particular parametric approximations to the expectation of the integrated volatility conditional on the time $t - 1$ information set:

$$\sigma_{t|t-1}^2 = E_{t-1} \left(\int_{t-1}^{t} \sigma^2(\tau)d\tau \right)$$

Hence, from equation (3), only if the integrated volatility process is perfectly predictable will the GARCH standardized returns, $r_t \sigma_{t|t-1}^{-1}$, be normally distributed. In general, of course, the diffusive volatility process varies non-trivially over the $(t-1, t]$ interval, resulting in a mixture-of-normals

---

[7] The inclusion of a drift term simply requires subtraction of a mean from the daily returns. In the empirical analysis we consider both raw and mean-adjusted returns. The results, reported below, are virtually identical.
[8] The integrated volatility also plays a central role in option pricing models allowing for time-varying volatility; see, e.g., the aforementioned paper by Hull and White (1987).

distribution for the corresponding GARCH standardized returns, with the mixture dictated by the distribution of the integrated volatility forecast errors; see also the reasoning behind the use of conditional fat-tailed GARCH error distributions in Bollerslev (1987).

2.2. Jumps

A number of recent studies have argued for the importance of explicitly allowing for jumps, or $q(t) \neq 0$, when modeling speculative rates of return; see, e.g., Andersen *et al.* (2002), Bates (1996, 2000), Chernov *et al.* (2003), Eraker *et al.* (2003), Eraker (2004), and Johannes (2004). This adds an additional component to the *ex post* price variation process, and also invalidates the Gaussianity of the standardized returns in (3). Suppose the jumps were known, and let the corresponding jump-adjusted returns be denoted by

$$\tilde{r}_t \equiv p(t) - p(t-1) - \sum_{s=q(t-1)}^{q(t)} \kappa(s), \qquad t = 1, 2, \ldots \tag{4}$$

where the sum comprises all of the non-zero jumps over the $(t-1, t]$ time interval, and we assume that the jump process is independent of the Brownian process $w(t)$ in equation (1). All of the variation in the jump-adjusted returns now originates from the diffusion component, so standardizing by the integrated volatility should again result in a normal distribution:

$$\tilde{r}_t \left(\int_{t-1}^{t} \sigma^2(\tau) \mathrm{d}\tau \right)^{-1/2} \sim N(0, 1), \qquad t = 1, 2, \ldots \tag{5}$$

In practice, of course, the timing and magnitude of jumps are not known for sure, so the result in (5) is not directly testable. To circumvent this, we rely on two new non-parametric jump-detection procedures for disentangling the continuous and discontinuous sample path components, in turn providing an operational approximation to (5).

2.3. Leverage and Volatility Feedback Effects

The distributional results of the preceding sections rule out so-called leverage and volatility feedback effects by assuming the Brownian motion driving the diffusive price innovations, $w(\tau)$, and the volatility process, $\sigma(t)$, are independent for all $\tau, t \geq 0$. A number of studies argue in contrast that the return–volatility relation is conditionally asymmetric as large negative returns are associated with larger volatilities than are positive returns of the same magnitude; e.g., Black (1976), Christie (1982), and Bollerslev *et al.* (2006). Here, leverage effect is defined as correlation between volatility and past returns and volatility feedback as correlation between volatility and future returns.[9] The leverage effect may be induced by contemporaneous negative correlation between the diffusive price innovations and the volatility innovations in the underlying continuous time model. Likewise, the feedback effect will arise from a positive correlation between volatility

---

[9] A similar leverage or volatility feedback effect could in principle work through the jump component. However, the related empirical evidence in Bollerslev *et al.* (2008) suggests that the asymmetry works almost exclusively through the diffusive component.

innovations and the drift in the price process. This feature involves compensation via the mean return for an increase in return volatility. Often, this effect is seen initiated through a negative price reaction to a volatility shock, thus also involving negative correlation between volatility and price movements. From discretely observed data, this latter effect is hard to separate from the leverage effect. In either case, these interactions imply that the *ex post* integrated volatility in the denominator on the left-hand-side of (3) and (5) are informative about both the sign and magnitude of the corresponding returns, so the standardized distributions are no longer Gaussian, let alone mean zero. However, by measuring returns over equal increments of integrated volatility instead of calendar time, the resulting time-changed returns remain Gaussian, even in the presence of leverage and volatility feedback effects.

Formally, let the event-time, or financial-time, sampling scheme be defined by $t_0 \equiv 0$ and

$$t_k \equiv \inf_{t > t_{k-1}} \left(\int_{t_{k-1}}^{t} \sigma^2(\tau) d\tau > \tau^* \right), \quad k = 1, 2, \ldots \quad (6)$$

where τ^* denotes the fixed financial-time unit spanned by each return.[10] For ease of comparison with the daily return distributions discussed above, we focus on the case in which τ^* equals the unconditionally expected one-period integrated variance:

$$\tau^* \equiv E \left(\int_{t-1}^{t} \sigma^2(\tau) d\tau \right) \quad (7)$$

Denote the corresponding jump-adjusted financial-time sampled returns by

$$\tilde{r}_k^* \equiv p(t_k) - p(t_{k-1}) - \sum_{s=q(t_{k-1})}^{q(t_k)} \kappa(s), \quad k = 1, 2, \ldots \quad (8)$$

It follows then by the Time-Change for Martingales Theorem (Dambis, 1965; Dubins and Schwartz, 1965) that

$$\tilde{r}_k^* \tau^{*-1/2} \sim N(0, 1), \quad k = 1, 2, \ldots \quad (9)$$

Importantly, this result establishes normality of the appropriately adjusted and standardized returns for *any* jump-diffusion model.[11]

We next discuss the nonparametric high-frequency data-based procedures used in implementing and testing each of the distributional results presented above. Our approach does not depend upon the validity of any particular parametric model. Nonetheless, the approach provides guidance for the specification of more realistic parametric models within the general class of jump diffusions defined by (1).

---

[10] A corresponding 'business-time' sampling scheme for pure jump processes has previously been used by Oomen (2006), while Zhou (1998) refers to similarly sampled returns as de-volatized. It is also reminiscent of the ϑ-time sampling scheme advocated by Dacorogna *et al.* (2001), although they employ a different realized power variation scale.

[11] This is also related to the earlier work of Lai and Siegmund (1983), and the idea of sampling autoregressive processes in equal increments of Fisher information.

3. EMPIRICAL RETURN AND VARIATION MEASURES

Our empirical analysis of transformed daily return distributions relies on the availability of intra-day data. If such data are available for T trading days, the return series is given by the increment of the observed log-price over each trading day, i.e.

$$R_t = p_{t,M} - p_{t,0}, \qquad t = 1, \ldots, T \tag{10}$$

where $p_{t,0}$ denotes the opening, or first, log-price on day t, and $p_{t,M}$ refers to the closing, or last, price on day t. This definition excludes the part of the daily variation associated with the overnight return, as the closing price on day $t-1$, $p_{t-1,M}$, typically differs from the opening price on the following day t, $p_{t,0}$.[12] However, the overnight returns may naturally be labeled deterministically occurring jumps. We treat them accordingly, so our trading day returns simply equal the daily returns adjusted for the (observed) overnight jump. Of course, this implies that applications of the current results for predicting the distribution of future returns must incorporate explicit corrections, not only for jumps within the trading periods but also for the price variability associated with market closures. These additional issues fall outside the scope of the present study, but the concurrent work by Andersen *et al.* (2009) exemplifies how this may be implemented in practice.

To avoid the problem of irregularly spaced high-frequency return observations, an imputation scheme (see, e.g., Dacorogna *et al.*, 2001) is usually applied to construct evenly spaced prices, say $M+1$ per day, where preferably many more observations are available each day. Denote the jth intra-daily log-price for day t by $p_{t,j}$, where $j = 0, 1, \ldots, M$ and $t = 1, \ldots, T$. The M continuously compounded intra-daily returns for day t are similarly denoted:

$$r_{t,j} = p_{t,j} - p_{t,j-1}, \quad j = 1, \ldots, M, \quad t = 1, \ldots, T \tag{11}$$

The precision of the resulting nonparametric realized volatility and jump measures depends on the value of M. In theory, the larger the number of intra-day returns, the higher the precision of the estimators. At the same time, from an empirical perspective, the larger the value of M, the more sensitive the estimates are to the influences of market microstructure 'noise' not contemplated within the theoretical model in equation (1), including price discreteness, bid–ask spreads, and non-synchronous trading effects. How to best account for these frictions and the practical choice of M in the construction of realized volatility measures have recently been the subject of intensive research efforts; e.g., Nielsen and Frederiksen (2008), Ait-Sahalia *et al.* (2005), Bandi and Russell (2008), Barndorff-Nielsen *et al.* (2008), and Hansen and Lunde (2006), among many others. In the empirical results reported below, we instead follow much of the early literature in the use of a relatively sparse fixed 5-minute, or $M = 78$, sampling frequency. However, we explicitly justify this particular choice of M for each of the stocks through the use of volatility signature type plots, as detailed in Section 4.

---

[12] The estimates reported in Hansen and Lunde (2005) suggest that about twenty percent of the total daily return variation is attributable to the overnight period.

3.1. Realized Volatility and Jumps

Following Andersen and Bollerslev (1998a), Andersen *et al*. (2001a) and Barndorff-Nielsen and Shephard (2002), we define the realized volatility for day t by[13]

$$RV_t \equiv \sum_{j=1}^{M} r_{t,j}^2, \qquad t = 1, \ldots, T \tag{12}$$

From the theory of quadratic variation, RV_t generally provides a consistent (in probability and uniformly in t) estimator of the daily increment to the quadratic variation for the underlying log-price process $p(\cdot)$ defined in (1). Specifically, for $M \to \infty$:

$$RV_t \longrightarrow_p \int_{t-1}^{t} \sigma^2(s) \, ds + \sum_{s=q(t-1)+1}^{q(t)} \kappa^2(s), \qquad t = 1, \ldots, T \tag{13}$$

Absent jumps, the second term vanishes and the realized volatility consistently estimates the integrated volatility which provides the contemporaneous standardization factor for the daily returns in the previous section. In general, however, the realized volatility measure includes the contribution to the total variation stemming from the squared jumps, and as such will not afford a consistent estimator of the requisite continuous sample path variation.

Meanwhile, Barndorff-Nielsen and Shephard (2004, 2006) show that separate nonparametric identification of the terms on the right-hand-side of equation (13) is possible through the use of so-called bipower variation measures. Specifically, the realized bipower variation is defined by

$$BV_t \equiv \mu_1^{-2} \sum_{j=2}^{M} |r_{t,j}||r_{t,j-1}|, \qquad t = 1, \ldots, T \tag{14}$$

where $\mu_1 = \sqrt{2/\pi}$. It can be shown that, even in the presence of jumps, for $M \to \infty$:

$$BV_t \longrightarrow_p \int_{t-1}^{t} \sigma^2(s) ds, \qquad t = 1, \ldots, T \tag{15}$$

Intuitively, for very large values of M, there is at most one jump in any two adjacent intervals of length $1/M$. Since the contribution of each absolute return associated with the diffusion component in the limit is negligible, any product involving a jump return will also be vanishingly small asymptotically. Moreover, the scaling factor for bipower variation ensures that it is consistent for the diffusive return variation. Hence, combining equations (13) and (15), the contribution to the total return variation stemming from the jump component is consistently estimated by the difference between the two. That is, for $M \to \infty$:

$$RV_t - BV_t \longrightarrow_p \sum_{s=q(t-1)}^{q(t)} \kappa^2(s), \qquad t = 1, \ldots, T \tag{16}$$

---

[13] We will refer interchangeably to this estimator as the realized volatility, the realized variation, or simply the variance. The exact meaning will be clear from the context.

Although formally consistent for the squared jumps, nothing prevents $\mathrm{RV}_t - \mathrm{BV}_t$ from becoming negative for finite values of M, especially when no jumps occur on day t. Similarly, part of the continuous price movements will invariably be attributed to the jump component due to sampling variation, resulting in small positive values of $\mathrm{RV}_t - \mathrm{BV}_t$ for finite M, even if there are no jumps, or $q(t) = q(t-1)$. Hence, following the empirical analysis in Andersen *et al.* (2007a), we refine our empirical analysis by considering the notion of *significant jumps*, only associating the most extreme price moves with the discontinuous jump component.

In particular, based on the asymptotic distribution theory in Barndorff-Nielsen and Shephard (2004, 2006) and the extensive simulation evidence in Huang and Tauchen (2005), we assess the significance of the daily jump component via the feasible logarithmic test statistic:

$$Z_t \equiv \sqrt{M} \frac{\ln \mathrm{RV}_t - \ln \mathrm{BV}_t}{((\mu_1^{-4} + 2\mu_1^{-2} - 5)\mathrm{TQ}_t \mathrm{BV}_t^{-2})^{1/2}} \longrightarrow_d N(0, 1) \tag{17}$$

where the realized tripower quarticity measure in the denominator is defined by

$$\mathrm{TQ}_t \equiv \frac{1}{M} \mu_{4/3}^{-3} \sum_{j=3}^{M} |r_{t,j}|^{4/3} |r_{t,j-1}|^{4/3} |r_{t,j-2}|^{4/3}, \quad t = 1, \ldots, T \tag{18}$$

and $\mu_{4/3} = 2^{2/3}\Gamma(7/6)/\Gamma(1/2)$, with $\Gamma(\cdot)$ denoting the gamma function. Thus, only (statistically) extreme positive values of $\mathrm{RV}_t - \mathrm{BV}_t$ are attributed to the jump component, i.e.

$$\mathrm{JV}_t \equiv I_{\{Z_t > \Phi_{1-\alpha}\}}(\mathrm{RV}_t - \mathrm{BV}_t), \quad t = 1, \ldots, T \tag{19}$$

where $I_{\{\cdot\}}$ denotes the indicator function, $\Phi_{1-\alpha}$ refers to the $(1 - \alpha)$ fractile of the standard normal distribution, and α denotes the chosen significance level.

Given our estimator for the squared jumps, an estimator for the continuous sample path variability, or integrated volatility, component is naturally obtained by the residual variation:

$$\mathrm{CV}_t \equiv \mathrm{RV}_t - \mathrm{JV}_t = I_{\{Z_t \leq \Phi_{1-\alpha}\}}\mathrm{RV}_t + I_{\{Z_t > \Phi_{1-\alpha}\}}\mathrm{BV}_t, \quad t = 1, \ldots, T \tag{20}$$

That is, we estimate the continuous volatility component by realized volatility on days with no significant jumps and by realized bipower variation on days with significant jump(s). The empirical results reported below rely on a significance level of $\alpha = 1\%$, but we also experimented with $\alpha = 5\%$ and 0.1%, resulting in qualitatively similar conclusions.[14]

The procedure discussed above provides a practical approach for identifying the jump contribution to the daily return variation. It does not, however, identify the individual jumps themselves. We next discuss two different methods for doing so.

---

[14] The use of standard significance levels automatically ensures that both JV_t and CV_t are non-negative, as $\Phi_{1-\alpha} > 0$ for $\alpha < 1/2$, while consistent estimation of the continuous and jump components would formally require that the significance level approaches zero with the sample size (see Barndorff-Nielsen and Shephard, 2006). Our choice of a low $\alpha = 1\%$ for each stage of the sequential testing scheme reflects our desire for a conservative approach to jump detection, so that only highly significant returns are removed from the continuous part of the return variation, which is the critical component for the subsequent distributional tests.

3.2. Jump-Adjusted Returns

In the absence of jumps and leverage effects, the daily returns should be approximately normally distributed when standardized by the corresponding integrated volatility, or an empirical estimate thereof. In general, however, the daily returns defined by the model in (1) may be comprised of both continuous price movements and discontinuous jumps. Building on the realized volatility measures defined above, we consider two different nonparametric procedures for directly identifying and estimating the intra-day jumps and the corresponding jump-adjusted returns.

Simple Jump Adjustments
Our first estimation scheme is based on the premise that jumps are relatively rare events. In particular, assume that there is at most one jump each day. It then follows from the arguments above that $JV_t \to_p \kappa_t^2$. Of course, this still leaves the sign of the jump undetermined. Appealing to the intuitive idea of signing the single day t jump on the basis of the largest (absolute) intra-day return, this estimation scheme defines the daily time series of jumps by[15]

$$\tilde{\kappa}_t \equiv \mathrm{sgn}(\{r_{t,k} : |r_{t,k}| = \max_{j \in \{1,\dots,M\}} |r_{t,j}|\}) \sqrt{JV_t}, \quad t = 1, \dots, T \tag{21}$$

where $\mathrm{sgn}(\cdot)$ is equal to 1 or -1 depending upon the sign of the argument. Accordingly, we denote the corresponding jump-adjusted daily returns by

$$\tilde{R}_t \equiv R_t - \tilde{\kappa}_t, \qquad t = 1, \dots, T \tag{22}$$

where $R_t = p_{t,M} - p_{t,0}$ denotes the daily return. As we move from sampling returns in calendar time to financial time, as defined by equal increments of integrated variance, knowing the exact jump times becomes essential in defining the new timescale. Of course, it is also possible that multiple jumps occur on certain days, violating the assumption underlying the simple procedure in (21). Hence, we next introduce a sequential jump identification scheme designed to facilitate inference regarding *all* significant jumps along with their *timing* within the trading day.

Sequential Jump Adjustments
A significant Z_t statistic, as defined in (17), only indicates the presence of one or more jumps. Our more detailed jump detection scheme applies this same statistic sequentially to identify potentially multiple significant jumps over the same day.

Intuitively, in the absence of any jumps, so that $RV_t - BV_t \to_p 0$, the average contribution of each squared intra-day return to the continuous sample path component is simply $M^{-1} \sum_{k=1}^{M} r_{t,k}^2$. Now assuming only a single jump on day t, this suggests the following alternative estimator for the day t contribution to the volatility coming from that jump:

$$I_{\{Z_t > \Phi_{1-\alpha}\}} \left(\max_{j \in \{1,\dots,M\}} r_{t,j}^2 - \frac{1}{M-1} \sum_{\substack{k=1 \\ k \neq j}}^{M} r_{t,k}^2 \right), \qquad t = 1, \dots, T$$

---

[15] We also experimented with signing the jumps on the basis of the total daily returns, resulting in very similar findings to those reported below.

This, of course, also directly identifies the time of the jump by the value of j that achieves the maximum. Now, eliminating this particular intra-day return in the calculation of a new jump-corrected realized volatility measure allows for the construction of a modified jump statistic to test for the presence of additional (smaller) jumps.

More precisely, in identifying the first jump, RV_t is based on the summation of all the squared intra-day returns. If the corresponding test in (17) rejects, we conclude that there is at least one jump during day t, and in turn identify its contribution to the total daily variation as the difference between the largest squared intra-day return and the average of the remaining $M - 1$ squared returns. Then, in identifying a possible second jump we define the day t realized volatility corrected for one jump as the summation of the squared returns, where the squared return containing the first jump is replaced by the average of the remaining $M - 1$ squared returns. If the new test statistic obtained by replacing RV_t in (17) with this jump-corrected realized volatility measure does not reject, we conclude that there is exactly one jump on day t, and we stop the sequential procedure. If, on the other hand, the test still rejects, we conclude that there are at least two jumps, and associate the contribution to the total variation coming from the second jump with the second largest squared intra-day return less the average of the remaining squared returns. More generally, after having identified i jumps, we calculate the jump-corrected realized volatility using the remaining $M - i$ returns scaled by $M/(M - i)$, continuing this sequential procedure until the corresponding test in (17) no longer rejects.[16]

Thus, having identified the total number of jumps, say J, during day t, as well as the magnitude of each of the jumps by the corresponding high-frequency returns:

$$\hat{\kappa}_{t,i} \equiv r_{t,j_i}, \qquad i = 1, \ldots, J, \qquad t = 1, \ldots, T \qquad (23)$$

where j_i denotes the exact time interval of the intra-day return associated with the ith jump, we calculate the jump-adjusted daily return as[17]

$$\hat{R}_t \equiv R_t - \sum_{i=1}^{J} \hat{\kappa}_{t,i}, \qquad t = 1, \ldots, T \qquad (24)$$

Similarly, we define the total variation on day t due to jumps as

$$JVS_t \equiv \sum_{i=1}^{J} JVS_{t,i}, \qquad t = 1, \ldots, T \qquad (25)$$

where $JVS_{t,i}$ gives the contribution from the ith jump, defined as the difference between the ith largest intra-day squared return and the average of the $M - J$ squared returns that are not

---

[16] We do not remove any returns in the computation of the bipower variation statistics. First, this has obvious asymptotic justification as the bipower variation statistic is consistent for the integrated variance in the presence of jumps. Second, removing one intraday return from the realized volatility computation does not alter the contribution from the remaining terms. In contrast, the realized bipower variation is not immune to this operation as it alters the two adjacent terms, often significantly. In view of this feature, the conservative nature of our jump detection scheme is best preserved by not sequentially adjusting the bipower variation statistic.

[17] This definition effectively assigns zero diffusive returns to the jump intervals. A natural alternative is to define the jump returns as the mean of the non-jump returns over the trading day. Our results are not materially affected by this choice.

associated with jump(s). That is:

$$\text{JVS}_{t,i} \equiv I_{\{Z_{t,i} > \Phi_{1-\alpha}\}} \left(\max_{j_i \in \{1,\ldots,M\} \setminus \{j_1,\ldots,j_{i-1}\}} r_{t,j_i}^2 - \frac{1}{M-J} \sum_{k \in \{1,\ldots,M\} \setminus \{j_1,\ldots,j_J\}} r_{t,k}^2 \right) \qquad (26)$$

where $Z_{t,i}$ denotes the ith sequential jump statistic, as discussed above. Lastly, the corresponding continuous volatility component is simply defined by

$$\text{CVS}_t \equiv \text{RV}_t - \text{JVS}_t, \qquad t = 1, \ldots, T \qquad (27)$$

which, in line with the earlier definition in (20), guarantees that each of the two daily time series are non-negative, and add up to the total daily realized variation.

The definition of $\tilde{\kappa}_t$ in (21) provides a rough estimate of $\sum_{i=1}^{J} \hat{\kappa}_{t,i}$. This suggests that the two procedures should produce similar jump adjustments for days in which there is only one jump. However, the ability of the sequential procedure to identify multiple (significant) jumps as well as their timing is important for the construction of jump-adjusted intra-day return series and these, in turn, constitute a critical input to the empirical analysis below.

4. DATA DESCRIPTION

4.1. Data Sources and Construction

Our data are extracted from the Trade And Quotation (TAQ) database, and consist of all recorded trades and quotes for the Dow Jones Industrial Average (DJIA) stocks for the 5-year period spanning 2 January 1998 to 31 December 2002. The ticker symbols and names of the stocks are listed in Table A1 of the supplementary appendix.[18] We only use the prices from the New York Stock Exchange (NYSE), with the exception of Intel and Microsoft, both of which are more actively traded on the National Association of Security Dealers Automated Quotation (NASDAQ) system. Mirroring the data-cleaning procedures of Hansen and Lunde (2006), we filter the series to remove price observations equal to zero, prices occurring outside the 9:30 a.m. to 4:00 p.m. official trading day, as well as extreme outliers or misrecorded price observations. This leaves us with 2–4 million prices for each stock, except for Intel and Microsoft, which both have around 26 million prices recorded over the sample. Finally, we delete days of early closing or late opening of the exchange and days in which trading in a particular stock was suspended for an extended period, resulting in approximately 1255 'intact' days for each stock.

To minimize market microstructure effects, we rely exclusively on mid-quotes and an imputation scheme involving the last quote preceding each 5-minute mark, in the construction of equally spaced 5-minute returns; i.e., $M = 78$ observations per day.[19] The choice of a 5-minute return interval is in line with the existing empirical literature and, as argued in Bandi and Russell (2008), it is also generally close to (mean-squared-error) 'optimal' for the standard realized variation measure and the TAQ data analyzed here. Importantly, however, our use of a 5-minute sampling

---

[18] All of the tables in the supplementary appendix are available on the journal's website.

[19] As argued in Hansen and Lunde (2006), using mid-quotes reduces the spurious serial correlation in the high-frequency returns due to bid–ask bounce and non-synchronous trading effects.

scheme in the present context, explicitly allowing for jumps, is further corroborated by the volatility signature plots discussed next.

4.2. Volatility Signature Plots

The conventional realized volatility signature plot popularized by Andersen *et al*. (2000b) provides a simple informal framework for gauging the impact of market microstructure frictions by plotting the average sample mean of RV_t over a long time-span as a function of the sampling frequency of the underlying intra-day returns, or M. In the absence of any frictions and dynamic dependencies in the returns, the realized volatilities are all consistent for the same total variation and hence, in practice, the signature plot should flatten out at the frequencies for which the microstructure frictions cease to have a distorting influence.

The signature plots in Figure A1 of the supplementary appendix for each of the individual stocks extend this idea by plotting the average realized bipower variation measures together with the standard realized variation for different sampling frequencies. Cursory inspection reveals a close similarity in the general shape across the individual stocks. We summarize the results in Figure 1 by plotting the median values (over the 30 stocks) of the average realized variation measures for each sampling frequency (measured in seconds).[20]

By reproducing the average (across days) RV_t and BV_t measures as a function of $1/M$ in the same graph, the volatility signature plot affords an informal way to gauge the importance of
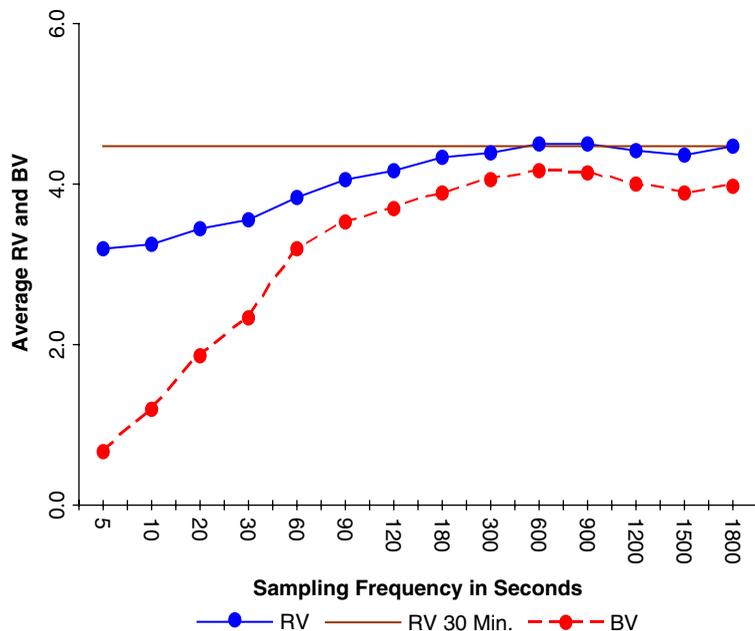

Figure 1. Median generalized volatility signature plots. This figure is available in color online at www.interscience.wiley.com/journal/jae

---

[20] Both of the variation measures have been converted to percent by multiplication by 10,000.

jumps. In particular, it follows from equation (16) that, under ideal conditions and for $1/M \to 0$, the distance between the two lines provides a consistent estimate of the total variation due to jumps.[21] In practice, of course, this theoretical prediction will be obscured by market microstructure 'noise', as directly evidenced by the systematic decline in both lines in Figure 1 in the range of 2–5 minutes, or 120–300 seconds. At the same time, the difference between the lines tends to stabilize at a sampling frequency of only 2 minutes, or 120 seconds. These effects are also in line with the extensive simulation results for the two measures based on empirically relevant continuous-time processes subject to 'noise' reported in Nielsen and Frederiksen (2008).[22] This suggests that both realized volatility and bipower variation measures are adversely affected by microstructure frictions at lower frequencies but the impact is correlated and tends to cancel so the gap between them, and hence the estimate of the jump component, remains remarkably stable for $1/M$ in excess of 120 seconds. Overall, this supports our use of a 5-minute return interval as a reasonable, albeit for some stocks somewhat conservative, uniform sampling scheme.

5. PRELIMINARY DATA ANALYSIS

As highlighted in the theoretical discussion, the presence of jumps and volatility feedback or leverage effects will cause the distribution of returns standardized by realized volatility to be non-Gaussian. Hence, we first present a set of summary statistics speaking to the importance of each of these features.

5.1. Jumps

We begin by considering jumps. We first report results based on the simple jump-detection procedure, followed by the more involved sequential jump-detection scheme.

Simple Jump Detection
Relying on the simple jump-detection method and a significance level of $\alpha = 1\%$, Table I displays the mean duration between significant jumps, the relative contribution of jumps to the realized variation, i.e., JV_t/RV_t, the mean size of the jump component for significant jump days, and lastly the corresponding mean (absolute) jump size, i.e., $|\tilde{\kappa}_t|$ as defined in equation (21). For ease of interpretation, we summarize the results in terms of the mean, standard deviation, minimum, and maximum of the statistics over all 30 DJIA stocks, with detailed results for each individual stock deferred to Table A2 in the supplementary appendix.

The mean duration between jumps ranges from a low of 4.1 days (HON) to a high of 10.1 days (GE), with an average across all stocks of 6.3 days. This intensity, of almost one jump per week, is much higher than typically estimated from parametric models based on daily or coarser frequency return observations.[23] These initial summary statistics suggest that important additional insights may be obtained from the use of higher-frequency data in terms of disentangling the price process into continuous and jump components. This is also consistent with the accumulating evidence that price jumps associated with the release of macroeconomic announcements are much more readily

---

[21] Related volatility signature plots, including plots for various integrated quarticity measures, have recently been explored by Andersen *et al.* (2006b).
[22] The theoretical framework in Rosenbaum (2007) may also help provide an explanation for these patterns.
[23] See, e.g., the GARCH-jump model estimates for individual stocks in Maheu and McCurdy (2004).

Table I. Jumps: simple method

	Mean duration	Rel. jump contribution JV_t/RV_t	Mean size of jump component ($\times 10,000$)	Mean size of actual jumps ($\times 100$)
Mean across stocks	6.3201	0.0476	1.2119	0.9812
SD across stocks	1.6068	0.0133	0.3283	0.1233
Min. across stocks	4.1325	0.0256	0.6247	0.7352
Max. across stocks	10.0976	0.0746	2.0825	1.3121

Note: This table reports the mean, standard deviation, minimum, and maximum over the 30 DJIA stocks for the mean duration between jumps, the relative jump contribution to the realized volatility, the mean size of the jump component ($\times 10,000$), as well as the mean size (in percent) of the square-root jump component (i.e., the absolute value of the actual jumps). For further details, see Table A2 in the supplementary appendix.

analyzed on the basis of intra-day data rather than the traditional daily return series (see, e.g., Andersen *et al.*, 2003b).

The potential importance of jumps is also evident from the last three columns of the table. In particular, estimates of the relative contribution of the jump component range from 2.6% (GE) to 7.5% (MO), with an average value of 4.8%.[24] The more detailed results in Table A2 of the supplementary appendix also point towards a negative association between jump durations and relative jump contributions. Further, the mean size of the jump component (multiplied by 10,000) on days with significant jumps is estimated between 0.62 (JNJ) and 2.08 (HPQ), which compares to a typical daily realized variation (multiplied by 10,000) of around 3–4. In other words, on days identified to have a jump, about a third of the return variation is attributed to jumps. Finally, the mean absolute size of the 'simple' jumps, i.e., $|\tilde{\kappa}_t|$, ranges from 0.74% to 1.31%, with a mean across all stocks of 0.98%.

Sequential Jump Detection
The sequential jump-detection procedure accommodates the presence of multiple jumps on a given trading day. It follows from the detailed results for the individual stocks in Figure A2 in the supplementary appendix that the median (across stocks) estimated (unconditional) probability of a single jump for the 'typical' stock is roughly 14%, while there is a 2% probability of two jumps. Meanwhile, the probability of three or more jumps in one day is very small, but not zero. This illustrates the potential importance of the sequential jump detection procedure, as most stocks have many days with multiple jumps.

At the same time, comparing the summary statistics in Table II for the sequential jump detection method to the corresponding statistics for the simple method in the last three columns of Table I, the numbers are generally fairly close. The relative contribution of the jump component for the sequential procedure ranges from 2.1% (GE) to 5.8% (MO), just slightly lower than the numbers for the simple method. Similarly, the mean size of the sequential jump component averaged across the stocks equals 1.83, compared to 2.08 in Table I, and the mean absolute jump size ranges from

---

[24] Our jump contribution measure is downward biased due to the conservative jump test. An asymptotically unbiased estimate of the overall jump contribution is given by the average across all stocks of the ratio $(\overline{RV} - \overline{BV})/\overline{RV}$, where the bar denotes average across all days. This value is 7.7%, indicating that there may be a fair amount of return variation within the (so classified) continuous component which actually stems from relatively smaller jumps. Of course, such misclassification would tend to render it harder to obtain Gaussian distributions for the standardized returns in the empirical analysis below.

Table II. Jumps: sequential method

	Rel. jump contribution JVS_t/RV_t	Mean size of jump component ($\times 10,000$)	Mean size of actual jumps ($\times 100$)
Mean across stocks	0.0373	1.0394	0.9282
SD across stocks	0.0101	0.3050	0.1177
Min. across stocks	0.0212	0.5065	0.6828
Max. across stocks	0.0575	1.8309	1.2464

Note: This table reports the mean, standard deviation, minimum, and maximum over the 30 DJIA stocks for the relative jump contribution to the realized volatility, the mean size of the jump component ($\times 10,000$), as well as the mean size (in percent) of the absolute value of the actual jumps. For further details, see Table A3 in the supplementary appendix.

Table III. Simple and sequential jump correlations

	Correlation	RMSE	Theil's U
Mean across stocks	0.9450	0.0062	0.2999
SD across stocks	0.0332	0.0033	0.1036
Min. across stocks	0.8722	0.0030	0.1086
Max. across stocks	0.9945	0.0200	0.5508

Note: This table reports the mean, standard deviation, minimum, and maximum over the 30 DJIA stocks for the correlation, root mean squared error (RMSE), and Theil's U statistic for the two daily jump series based on the simple and sequential methods, respectively. Observations where both series are zero have been removed. For further details, see Table A4 in the supplementary appendix.

a low of 0.68% (JNJ) to a high of 1.25% (HPQ), with the overall absolute mean jump size of 0.93% again being slightly below that in Table I.

The close coherence between the two daily jump component series, JV_t and JVS_t in equations (19) and (25) is further underscored by Table III, which presents various correlation measures between the two. To focus on the relation between the jump series, all common no-jump (zero) observations were excluded from the computations. The first column reports the standard sample correlation coefficient, the second the root mean squared error (RMSE) calculated as the square root of the sum of the squared differences between the series, and the third Theil's scale invariant U-statistic. As above, the results are summarized through the mean, standard deviation, minimum, and maximum across the 30 stocks, with detailed results for each stock deferred to Table A4 of the supplementary appendix. It is evident that the two differently estimated jump components are close. For instance, the lowest sample correlation equals 0.87 (WMT) and the average value is 0.95. Also, the RMSEs and Theil's U-statistic are generally low across the stocks. Hence the sequential procedure retains the information regarding jump occurrence and relative size on a day-to-day basis but, importantly, also identifies the intra-day timing of *all* significant jumps, which is critical for the subsequent analysis.

5.2. Leverage and Volatility Feedback Effects

The second key assumption underlying the normality of the integrated volatility standardized returns concerns the lack of correlation between the diffusive volatility process and the Brownian motion innovations to the price process.

In order to assess the validity of this assumption, Figure A3 of the supplementary appendix graphs the 5-minute cross-correlations for each of the stocks, i.e., corr($|r_j|, r_{j+i}$), where for notational simplicity r_j for $j = 1, \ldots, J$ refers to time series of approximately $J = 1,255 \times 78 = 97,890$ demeaned 5-minute returns available for each stock. An initial cursory look suggests a broadly similar shape across stocks, although the idiosyncratic noise inherent in the individual estimates makes it hard to draw sharp conclusions. Hence, we summarize the evidence in Figure 2 by plotting the median value, across the stocks, of each of the high-frequency cross-correlations. Figure 2 reveals a clear tendency for the correlations between $|r_j|$ and r_{j+i} to be negative for negative i, while the correlations typically are positive or near zero for positive values of i. Of course, there is a striking spike around $i = 0$, which is also present for most individual stocks. As such, this points to the existence of a potentially distorting high-frequency leverage effect for at least some of the stocks, but not much of a volatility feedback effect.[25]

Table IV provides summary statistics related to the leverage and volatility feedback type effects. Specifically, the table reports estimates of each individual effect as well as the difference between the two; the more detailed findings for each individual stock are again reported in the supplementary
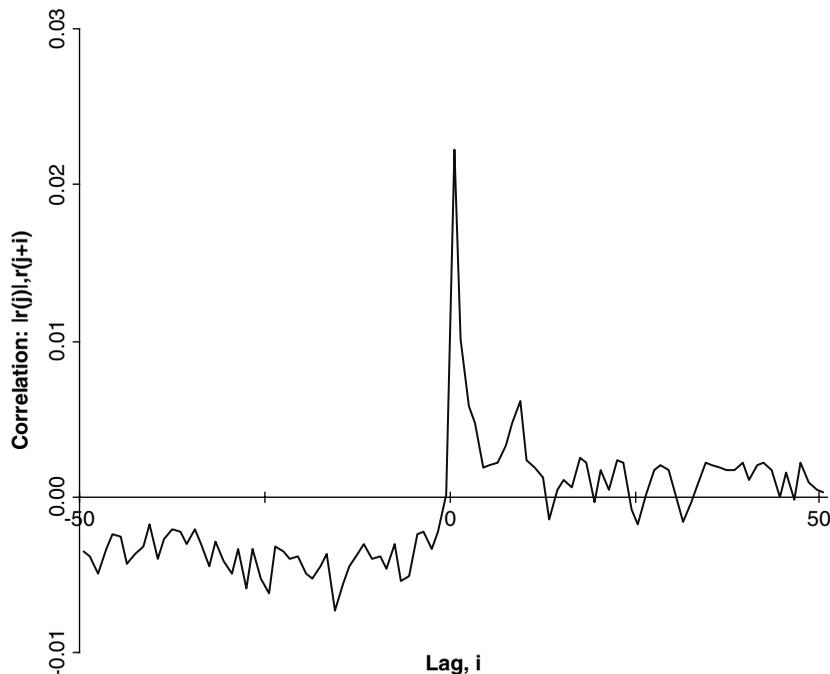

Figure 2. Median high-frequency leverage and volatility feedback effects

---

[25] This is consistent with the corresponding plots for high-frequency S&P500 futures returns in Bollerslev *et al.* (2006), which show even more pronounced negative cross-correlations for negative lags along with cross-correlations close to zero for positive lags.

Table IV. Leverage and volatility feedback effect estimates

	Leverage	Feedback	Difference
Mean across stocks	−0.0166	0.0076	−0.0243
SD across stocks	0.0151	0.0087	0.0145
Min. across stocks	−0.0560	−0.0155	−0.0658
Max. across stocks	0.0053	0.0226	−0.0035
Significance at 5% level	9	10	20
Significance at 1% level	6	5	14

Note: This table reports the mean, standard deviation, minimum, and maximum over the 30 DJIA stocks for the leverage and volatility feedback effect estimates along with their numerical difference, as described in the main text. The last two rows report the number of stocks (out of 30) for which the corresponding *t*-statistics, based on a heteroskedasticity and autocorrelation consistent Newey–West type covariance matrix estimator, are significantly different from zero at the 5% and 1% levels. For further details, see Table A5 in the supplementary appendix.

appendix, Table A5. The average *leverage effect* for an individual stock is estimated by

$$\frac{1}{K-2} \sum_{i=2}^{K-1} \frac{1}{J-K+1} \sum_{j=K}^{J} |r_j| r_{j-i}$$

while the *volatility feedback effect* is calculated as

$$\frac{1}{K-2} \sum_{i=2}^{K-1} \frac{1}{J-K+1} \sum_{j=1}^{J-K+1} |r_j| r_{j+i}$$

That is, the leverage effect is measured as the (unweighted) mean of the sample cross-covariances between the absolute returns and the lagged $2, \ldots, (K-1)$ period returns, corresponding to the $K-2$ cross-correlations immediately to the left of negative one in the figures. Similarly, the volatility feedback effect is measured as the mean of the sample cross-covariances between the absolute returns and the returns $2, \ldots, (K-1)$ periods into the future, corresponding to the sum of the first $K-2$ cross-correlations immediately to the right of one in the figures. For conciseness, we focus on $K = 30$, but identical qualitative findings are obtained for other values of K. Also, to guard against spurious non-synchronous trading effects, we explicitly exclude the first (positive and negative) cross-covariance but including these does not materially affect the results.[26]

More formal tests generally confirm the visual impression. The auto-covariances corresponding to the leverage effect are negative, while the volatility feedback auto-covariances are close to zero and, if anything, positive, on average. Interestingly, although the effects are statistically insignificant for most stocks, there is considerable cross-sectional variation in the magnitude of the leverage effect, and for some stocks the cross-covariances are quite significant.[27] We also note that the difference between the two effects is negative for all stocks, and significantly so at the 5% level for 20 of the 30.

[26] We also calculated the same statistics for the jump-adjusted returns, resulting in very similar numbers to those reported in the tables. These results are available upon request.

[27] These high-frequency based findings are corroborated by conventional EGARCH models for the daily returns which produce most significant volatility asymmetries for the stocks for which the leverage effects in Table A5 in the supplementary appendix are the largest.

The results suggest that financial-time sampling is necessary to restore normality of the standardized return distributions, at least for some stocks. Of course, whether the high-frequency leverage and volatility feedback effects are large enough to cause noticeable distortions in the standardized return distributions remains an empirical question, to which we now turn.

6. DAILY RETURN DISTRIBUTIONS

6.1. Unconditional Return Distributions

It is well established that the unconditional distributions of daily stock returns are fat-tailed. At the same time, our theory predicts that suitably jump-adjusted and standardized returns should be i.i.d. Gaussian. Hence, as a natural benchmark, we first provide a summary of the raw unconditional return distributions for the DJIA stocks. The first row of Table V confirms the above-mentioned stylized facts. Using the normality test of Andersen *et al.* (2007b) involving the joint distribution of the first four sample moments, the null hypothesis that the unconditional return distribution, or $R_t/\sqrt{\text{var}(R_t)}$, is standard normal is rejected at the 1% level for all stocks.[28] Table A6 in the supplementary appendix indicates that the overwhelming rejections are due primarily to excess kurtosis.

These results are as expected if the underlying return volatility is time-varying since this induces a mixture type distribution. We next look at the unconditional distributions obtained by standardizing the daily returns with the one-day-ahead conditional volatility forecasts from a conventional GARCH(1,1) model.

Table V. Daily return distributions

Series	Raw returns Significance		Demeaned returns Significance	
	5% level	1% level	5% level	1% level
$R_t/\sqrt{\text{var}(R_t)}$	30	30	30	30
$R_t/\sqrt{\text{GARCH}(1,1)}$	30	30	30	30
$R_t/\sqrt{\text{RV}_t}$	21	12	18	9
$\tilde{R}_t/\sqrt{\text{CV}_t}$	18	10	15	9
$\hat{R}_t/\sqrt{\text{CVS}_t}$	20	12	18	11
$\hat{R}_k^*/\sqrt{E(\text{CVS}_t)}$	13	6	11	5
$\hat{R}_{5k,5}^*/\sqrt{5E(\text{CVS}_t)}$	6	3	2	2

Note: This table reports the number of stocks (out of 30) for which the hypothesis of normality is rejected based on the joint test for the first four moments. The results in the last two columns are based on subtracting the sample mean from the return series in the numerator. R_t refers to the daily return, while $\tilde{R}_t$ and $\hat{R}_t$ denote the daily jump-adjusted returns calculated according to the simple and sequential procedures, respectively. RV_t gives the total realized variation. The continuous variation based on the simple and sequential jump-adjustment procedures are denoted by CV_t and CVS_t, respectively. $\hat{R}_k^*$ refers to the financial-time return series constructing from the sequential jump-adjusted intra-day returns spanning $E(\text{CVS}_t)$ time-units. Lastly, $\hat{R}_{5k,5}^* \equiv \hat{R}_{5k}^* + \hat{R}_{5k-1}^* + \hat{R}_{5k-2}^* + \hat{R}_{5k-3}^* + \hat{R}_{5k-4}^*$ defines the financial-time return series spanning $5E(\text{CVS}_t)$ time-units. For further details regarding each of the individual stocks, see Table A6 in the supplementary appendix.

---

[28] Ignoring potential complications arising from correcting for jumps, this procedure is equivalent to testing that the first four orthogonal Hermite polynomials are equal to zero. As such, the test is a special case of the general class of normality tests developed by Bontemps and Meddahi (2005a,b) based on the so-called Stein equation.

6.2. GARCH Standardized Returns

The results for GARCH standardized returns, $R_t/\sqrt{\text{GARCH}(1,1)}$, in the second row of Table V, are again fully consistent with the existing literature. Although the mass in the tails of the GARCH standardized return distributions shrinks relative to that of the unconditional distributions, they remain significantly leptokurtic for all stocks; see Bollerslev (1987), Baillie and Bollerslev (1989), and Hsieh (1989) for early related evidence.[29]

Of course, if the underlying price and volatility processes evolve stochastically within the trading day, the GARCH volatilities, at best, represent the one-day-ahead conditional expectations of the corresponding (latent) integrated volatilities. As argued in Section 2, the GARCH standardized returns should therefore follow a fat-tailed mixture-of-normals distribution, with the mixture determined by the distribution of the GARCH volatility forecast errors *vis-à-vis* the true integrated volatilities. We next explore the distributions obtained by standardizing returns by realized volatilities. Since the latter provide more accurate *ex post* estimates of the integrated volatility realizations than *ex ante* GARCH forecasts, we expect these distributions to be closer to normal.

6.3. Realized Volatility Standardized Returns

We now focus on the distribution of realized volatility standardized returns, $R_t/\sqrt{\text{RV}_t}$. From the density and QQ-plots for the individual stocks in Figures A8 and A9 of the supplementary appendix, it is evident that the RV standardized distributions are much closer to the reference Gaussian distributions than the raw and GARCH standardized returns. In particular, the tails of the QQ-plots have improved considerably and mostly feature only slight deviations from the straight 45-degree line. These findings are also in accord with earlier studies by Andersen *et al.* (2000a, 2001b), arguing through similar informal graphical tools and summary statistics that the sample distributions of the RV_t standardized returns are close to Gaussian.

Complementing this informal evidence, the third row in Table V reports results from applying our formal moment-based test to the realized volatility standardized return distributions. Importantly, as shown in Andersen *et al.* (2007b), under the null hypothesis of a time-invariant, or homogeneous, diffusion, the fourth population moment of the RV_t standardized returns equals $m_4 = 3\frac{M}{M+2}$, rather than the standard normal value of three, and we use this value in implementing the test. Given the $M = 78$ 5-minute returns per day, this translates into a value of 2.925. The results confirm that the first four sample moments of $R_t/\sqrt{\text{RV}_t}$ generally adhere fairly closely to those of the slightly modified Gaussian distribution. Specifically, the implicit null of an underlying continuous-time diffusion is *not* rejected for nine of the 30 stocks at the 5% significance level, while the tests are insignificant for 18 stocks at the 1% level.

Nonetheless, looking at the more detailed statistics in Table A6 in the supplementary appendix, we find that the sample kurtosis for $R_t/\sqrt{\text{RV}_t}$ remains significantly different from the theoretical value of $m_4 = 2.925$ that should obtain for a homogeneous diffusion in many cases. Of course, many studies argue for the importance of allowing for jumps of stock returns, and the empirical results in Section 5.1 support this notion. The presence of a few large jumps tends to imply that the RV_t standardized distribution has thinner tails than the (modified) normal because the jumps inflate the denominator realized volatility disproportionately. More generally, however, the

---

[29] We also experimented with alternative EGARCH-M models, allowing for volatility asymmetry resulting in very similar findings; see Kim and Kon (1994) for related evidence.

presence of jumps simply obfuscates the asymptotic normality of the $R_t/\sqrt{\mathrm{RV}_t}$ distribution. Indeed, even though the majority of the rejections in Table V arise from exceedingly low sample values of m_4, for a few stocks the empirical values are significantly larger than 2.925. In an attempt to clarify these issues, we next consider the distribution of jump-adjusted returns standardized by an estimate of the corresponding continuous sample path variation.

6.4. Jump-Adjusted Realized Volatility Standardized Returns

Following Sections 2.2 and 3.2, we consider jump-adjusted return distributions using both the simple and sequential jump-detection schemes. Summary results of the normality tests for these distributions, labeled $\tilde{R}_t/\sqrt{\mathrm{CV}_t}$ and $\hat{R}_t/\sqrt{\mathrm{CVS}_t}$, respectively, are given in rows four and five of Table V. Perhaps surprisingly, the results indicate that neither of the jump-adjusted standardized series are systematically closer to Gaussian than the $R_t/\sqrt{\mathrm{RV}_t}$ non-adjusted realized volatility standardized returns. The hypothesis of normality is rejected for 18 stocks at the 5% level using the simple method and 20 stocks using the sequential procedure, compared to 21 stocks for the non-adjusted returns. Similarly, at the 1% level, 10 and 12 stocks reject for the jump-adjusted returns, while 12 stocks reject for the unadjusted returns.

Although jumps appear important and, according to Section 5.1, account for about a third of the return variation on jump days, adjusting for jumps fails to restore normality to the standardized returns. One reason is that jumps largely self-standardize: a large jump tends to inflate the (absolute) value of both the return (numerator) and the realized volatility (denominator) of standardized returns, so the impact is muted. Thus, even if jumps impact the raw return distribution significantly they exert much less influence on the realized volatility standardized distribution. In sum, the remaining, still appreciable, deviations from normality likely stem from a different source. One potential factor is systematic dependencies between the numerator and denominator of the standardized returns, as indicated in Section 2.3. Moreover, the empirical correlation-based measures discussed in Section 5.2 also suggest that a leverage type effect may be at work. To explore this possibility, we now consider the properties of jump-adjusted standardized returns sampled in financial time, i.e., equal increments of integrated volatility.

6.5. Jump-Adjusted Financial-Time Standardized Returns

If realized daily integrated volatility conveys information about the corresponding daily returns, or vice versa, as implied by the leverage and volatility feedback effects, discretely sampled returns from a diffusive process, standardized by integrated volatility, are generally not Gaussian. However, as discussed in Section 2.3, the dependence between the numerator and denominator of the standardized returns may be broken by sampling in so-called event, or financial, time. The new sequential jump-adjustment procedure, which identifies the timing of the jumps within the day, permits the construction of such financial-time returns by accumulating the jump-adjusted intra-day returns until they span identical increments of the CVS_t process, but time-varying calendar-time intervals. To compute $\hat{R}_k^*$ in practice, we include intra-day returns until the cumulative squared returns exceeds τ^*; i.e., the average daily (when $\tau^* = E(\mathrm{CVS}_t)$) respectively weekly ($\tau^* = 5E(\mathrm{CVS}_t)$) realized volatility in calendar time. Importantly, only non-jump returns as identified by the sequential jump detection scheme were included, since the simple jump-adjustment method does not identify the timing of all jumps and so is less appropriate for this purpose.

The second to last row of Table V, labeled $\hat{R}_k^*/\sqrt{E(\mathrm{CVS}_t)}$, reports results from applying the moment-based tests to jump-adjusted financial time returns where, for ease of comparison, the financial time unit is calibrated to an average trading day; i.e., $\tau^* = E(\mathrm{CVS}_t)$. Interestingly, the move to financial-time sampling results in a marked reduction in the number of stocks for which normality is rejected, with only six (five for demeaned returns) stocks now being significantly non-Gaussian at the 1% level. The quality of the approximation afforded by the normal distribution is also evident from the density and QQ-plots in Figures A14 and A15 in the supplementary appendix which, except for a few stocks, display a remarkably close coherence between the empirical and theoretical distributions.

Comparing the test results for leverage and volatility feedback effects for each stock in Table A5 with the normality tests in Table A6, there is generally also a close association between the significance of the former and the strength of the 'normality gains' obtained by moving to financial-time sampling. For instance, for IBM the leverage and volatility feedback effects are both significant, and consequently normality of the standardized return series in calendar time is rejected at the 5% level, while the p-value for the normality test for the one-day financial-time returns is 0.316. Conversely, for JPM, one of only two stocks for which normality of one-day returns is rejected at the 5% level in financial but not calendar time, the leverage effect appears insignificant and the volatility feedback effect is only marginally significant at the 10% level.

The CVS_t series used to construct the financial-time scale often varies considerably over the sample. Consequently, some 'one-day' observations span intra-day returns over several calendar days, while others are based on the sum of only a few squared 5-minute returns. In the latter case, the asymptotic theory, for the number of intra-day returns approaching infinity, provides a poor approximation. Hence, the last row of Table V, labeled $\hat{R}_{5k,5}^*/\sqrt{5E(\mathrm{CVS}_t)}$, reports on normality test applied to returns spanning one financial 'week', or five average 'days'; i.e., $\tau^* = 5E(\mathrm{CVS}_t)$. Remarkably, normality for this longer return horizon, but shorter time series, is now rejected at the 1% level for only three (two for demeaned returns) stocks.[30]

To highlight the improved accuracy of the normal approximation afforded by the sequential distributional adjustments, Figure 3 plots the p-values for the tests for each stock and return transformation underlying Table V. If these distributions are Gaussian and the individual tests independent, the p-values should be distributed uniformly on the unit interval. The raw and GARCH standardized daily return series invariably have p-values of zero, as indicated by the single point on the plot. Standardizing the returns by the realized volatilities improves the picture, but all p-values remain below 0.25, and the results for the standardized jump-adjusted returns do not fare any better. In contrast, the p-values for the 'daily' and 'weekly' financial-time returns appear close to uniformly distributed. Thus, the p-value plots further support the hypothesis that by moving to financial time normality of the (jump-adjusted) returns is restored. It is consistent with the notion that stock prices may be thought of as discretely sampled observations from a continuous-time jump-diffusion model, while also underscoring the impact of leverage and/or volatility feedback effects.

---

[30] Of course, the power of the tests based on fewer longer horizon returns is likely lower. However, we also studied the distribution of the standardized returns over longer calendar time periods, and did not observe a similar dramatic reduction in the number of rejections.
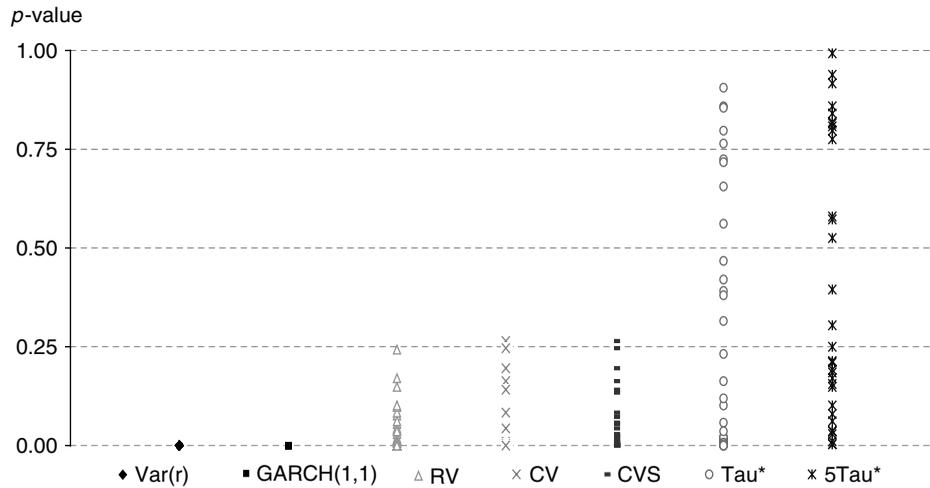
Figure 3. *p*-values for the 30 DJIA stocks, January 1998 to December 2002, 5-minute sampling

6.6. Alternative Sampling Frequencies

Our empirical results hinge on the use of high-frequency data for construction of reliable realized variation measures and associated jump detection and financial-time sampling schemes. In particular, the volatility signature plots introduced in Section 4.2 guided our selection of a 5-minute sampling frequency. To confirm that a sampling frequency in this range provides a reasonable trade-off between the preference for finely sampled returns and avoiding market microstructure contamination, we applied the same distributional tests to series based on both more and less frequently sampled intra-day returns.

Figure 4 reports *p*-values for the different return transformations based on a coarser 30-minute sampling frequency, corresponding to the right-most points in the median volatility signature plot in Figure 1. Under ideal conditions, the realized volatility measures and jump detection tests based on 'only' $M = 13$ half-hourly intra-day observations are subject to much larger measurement errors than the 5-minute based measures and tests. This effect manifests itself in a noticeable deterioration in the dispersion of the *p*-values for the realized volatility standardized returns, which are now visibly less consistent with a uniform distribution. Meanwhile, the distribution of the *p*-values for the financial-time returns, and the '5-day' returns in particular, still appear fairly close to uniform.

At the other end of the spectrum, Figure 5 displays *p*-values obtained using finely sampled 30-second returns; i.e., $M = 780$. This corresponds to the point in Figure 1 where the slope of the signature plots for the average realized volatility and bipower variation measures begin to diverge. A marked deterioration in the dispersion of the *p*-values for the financial-time returns is now apparent. The contaminating influences from the market microstructure 'noise' overwhelm the signal in the realized variation measures. Not surprisingly, direct investigation of this very high-frequency series (not reported here) reveals dramatic violations of the basic arbitrage-free semi-martingale assumption for the price process.
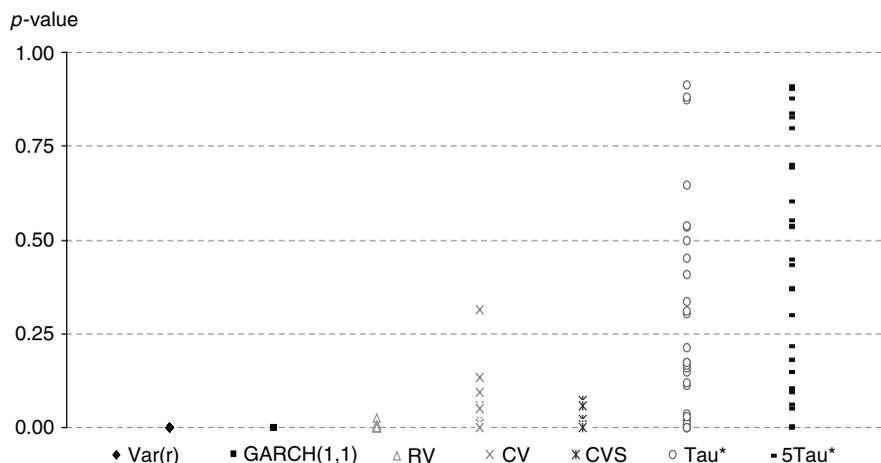
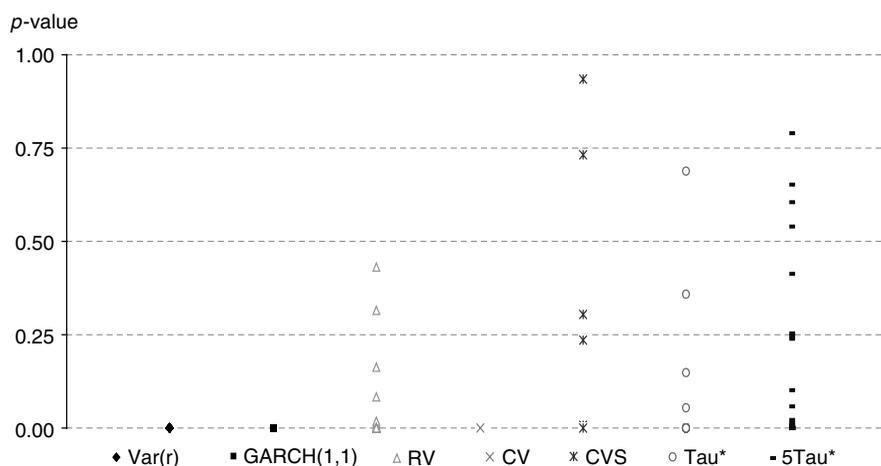Figure 4. *p*-values for the 30 DJIA stocks, January 1998 to December 2002, 30-minute sampling

Figure 5. *p*-values for the 30 DJIA stocks, January 1998 to December 2002, 30-second sampling

In sum, the 5-minute sampling frequency appears to be a reasonable choice for eliciting distributional information from the high-frequency data within this context.[31]

7. CONCLUDING REMARKS

We show how high-frequency intra-day data can be used to construct simple non-parametric realized variation measures and test statistics which shed light on the nature of daily or lower

---

[31] The minimum tick size on the NYSE was reduced to 1 cent on 29 January 2001. In the supplementary appendix we provide summary conclusions for the more recent time period February 2001 through December 2004, which mirror our more detailed empirical findings for the longer sample.

frequency return distributions. Each step in our sequential test procedure speaks directly to important qualitative features of the underlying return generating process. As such, the tests may serve as diagnostic tools in the specification of empirically realistic continuous-time models. In this regard, our empirical results for the set of DJIA stocks suggest that their price series may be satisfactorily described as discretely sampled observations from a jump-diffusion model, but only after allowing for leverage and/or volatility feedback effects.

Each step in the sequential procedure could be extended in a number of directions. As discussed, several recent studies argue for the use of new multi-scale or kernel-based realized volatility measures for more accurately measuring the true latent return variation (e.g., Bandi and Russell, 2008; Hansen and Lunde, 2006; Barndorff-Nielsen *et al*., 2008; Ait-Sahalia *et al*., 2005). Also, while the use of daily realized volatility measures conveniently circumvents complications associated with the strong intra-day volatility patterns (e.g., Andersen and Bollerslev, 1998b), the financial-time scale will invariably span different periods of the day, and it may prove beneficial to explicitly control for this feature. Moreover, a number of alternative jump detection procedures have recently been proposed (e.g., Jiang and Oomen, 2008; Mancini, 2005), and it would be interesting to compare and contrast the results obtained here to such alternative schemes.

It may also be informative to relate price jumps to news arrivals, either in the form of company specific news (e.g., Johannes, 2004), or macroeconomic announcements (e.g., Andersen *et al*., 2003b). Similarly, it might prove instructive to associate the financial-time scale defined by realized volatility to observable economic activity variables within the context of the MDH (e.g., Ane and Geman, 2000; Luu and Martens, 2003). From the reverse perspective, given that our realized volatility and jump transformations have a sound foundation in theory and appear to outperform prior MDH style models for the return distribution on the empirical dimension, it may be useful for MDH style models to relate their candidate economic mixing variables to the diffusive and jump return variation components estimated here.

Another interesting question relates to the possible extension of the distributional results and test statistics derived here to a multivariate setting. Although the notion of realized covariation may be defined straightforwardly, practical issues related to the non-synchronicity of multiple high-frequency price series looms large, (e.g., de Pooter *et al*., 2008). The multivariate extension also presents challenges from a theoretical perspective in terms of the time deformation required to simultaneously guard against leverage and/or volatility feedback effects across multiple assets (e.g., Ploberger, 2005).

Last but not least, it is of interest to directly explore the usefulness of the return transformations and decompositions developed here for value-at-risk type calculations, volatility timing, and other related financial decisions (e.g., Fleming *et al*., 2003).

REFERENCES

Ait-Sahalia Y, Mykland PA, Zhang L. 2005. How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies* **18**: 351–416.

Andersen TG. 1996. Return volatility and trading volume: an information flow interpretation of stochastic volatility. *Journal of Finance* **51**: 169–204.

Andersen TG, Bollerslev T. 1998a. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* **39**: 885–905.

Andersen TG, Bollerslev T. 1998b. Deutchemark–dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance* **53**: 219–265.

Andersen TG, Bollerslev T, Diebold FX, Labys P. 2000a. Exchange rate returns standardized by realized volatility are (nearly) Gaussian. *Multinational Finance Journal* **4**: 159–179.

Andersen TG, Bollerslev T, Diebold FX, Labys P. 2000b. Great realizations. *Risk* **13**: 105–108.

Andersen TG, Bollerslev T, Diebold FX, Labys P. 2001a. The distribution of exchange rate volatility. *Journal of the American Statistical Association* **96**: 42–55.

Andersen TG, Bollerslev T, Diebold FX, Ebens H. 2001b. The distribution of realized stock return volatility. *Journal of Financial Economics* **61**: 43–76.

Andersen TG, Benzoni L, Lund J. 2002. An empirical investigation of continuous-time equity return models. *Journal of Finance* **57**: 1239–1284.

Andersen TG, Bollerslev T, Diebold FX, Labys P. 2003a. Modelling and forecasting realized volatility. *Econometrica* **71**: 579–625.

Andersen TG, Bollerslev T, Diebold FX, Vega C. 2003b. Micro effects of macro announcements: real-time price discovery in foreign exchange. *American Economic Review* **93**: 38–62.

Andersen TG, Bollerslev T, Christoffersen P, Diebold F. 2006a. Volatility forecasting. In *Handbook of Economic Forecasting*, Elliott G, Granger C, Timmermann A (eds). North-Holland: Amsterdam; 777–878.

Andersen TG, Bollerslev T, Frederiksen PH, Nielsen M. 2006b. Comment on Peter R. Hansen and Asger Lunde: realized variance and market microstructure noise. *Journal of Business and Economic Statistics* **24**: 173–178.

Andersen TG, Bollerslev T, Huang X. 2009. A semiparametric framework for modelling and forecasting jumps and volatility in speculative prices. *Journal of Econometrics* (forthcoming).

Andersen TG, Bollerslev T, Diebold FX. 2007a. Roughing it up: including jump components in the measurement, modeling, and forecasting of return volatility. *Review of Economics and Statistics* **89**: 701–720.

Andersen TG, Bollerslev T, Dobrev D. 2007b. 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**: 125–180.

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

Baillie RT, Bollerslev T. 1989. The message in daily exchange rates: a conditional variance tale. *Journal of Business and Economic Statistics* **7**: 297–305.

Bandi FM, Russell JR. 2008. Microstructure noise, realized variance, and optimal sampling. *Review of Economic Studies* **75**: 339–369.

Barndorff-Nielsen OE, Shephard N. 2002. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society Series B* **64**: 253–280.

Barndorff-Nielsen OE, Shephard N. 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* **2**: 1–48.

Barndorff-Nielsen OE, Shephard N. 2006. Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* **4**: 1–30.

Barndorff-Nielsen OE, Shephard N, Winkel M. 2006. Limit theorems for multipower variation in the presence of jumps. *Stochastic Processes and their Applications* **116**: 796–806.

Barndorff-Nielsen OE, Hansen PR, Lunde A, Shephard N. 2008. Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica* **76**: 1481–1536.

Bates D. 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutschemark options. *Review of Financial Studies* **9**: 69–107.

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

Black F. 1976. The pricing of commodity contracts. *Journal of Financial Economics* **3**: 167–179.

Bollerslev T. 1987. A conditional heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics* **69**: 542–547.

Bollerslev T, Jubinski D. 1999. Equity trading volume and volatility: latent information arrivals and common long-run dependencies. *Journal of Business and Economic Statistics* **17**: 9–21.

Bollerslev T, Litvinova J, Tauchen G. 2006. Leverage and volatility feedback effects in high-frequency data. *Journal of Financial Econometrics* **4**: 353–384.

Bollerslev T, Kretschmer U, Pigorsch C, Tauchen G. 2008. A discrete-time model for daily S&P500 returns and realized variations: jumps and leverage effects. *Journal of Econometrics* **150**: 151–166.

Bontemps C, Meddahi N. 2005a. Testing distributional assumptions: a GMM approach. Working paper, University of Montreal.

Bontemps C, Meddahi N. 2005b. Testing normality: a GMM approach. *Journal of Econometrics* **124**: 149–186.

Chernov M, Gallant AR, Ghysels E, Tauchen G. 2003. Alternative models of stock price dynamics. *Journal of Econometrics* **116**: 225–257.

Christensen K, Podolski M. 2007. Realized range-based estimation of integrated variance. *Journal of Econometrics* **141**: 323–349.

Christie AA. 1982. The stochastic behavior of common stock variances: value, leverage and interest rate effects. *Journal of Financial Economics* **3**: 145–166.

Clark PE. 1973. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* **41**: 135–155.

Cont R, Tankov P. 2004. *Financial Modeling with Jump Processes*. Chapman & Hall: London.

Dacorogna MM, Gencay R, Müller UA, Pictet OV, Olsen RB. 2001. *An Introduction to High-Frequency Finance*. Academic Press: San Diego, CA.

Dambis KE. 1965. On the decomposition of continuous submartingales. *Theory of Probability Applications* **10**: 401–410.

de Pooter M, Martens M, van Dijk D. 2008. Predicting the daily covariance matrix for S&P100 stocks using intraday data: but which frequency to use? *Econometric Reviews* **27**: 199–229.

Dubins L, Schwartz G. 1965. On continuous martingales. *Proceedings of the National Academy of Sciences* **53**: 913–916.

Engle RF. 2004. Nobel lecture. Risk and volatility: econometric models and financial practice. *American Economic Review* **94**: 405–420.

Epps TW, Epps ML. 1976. The stochastic dependence of security price changes and transaction volumes: implications for the mixture-of-distributions hypothesis. *Econometrica* **44**: 303–321.

Eraker B. 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* **59**: 1367–1403.

Eraker B, Johannes M, Polson N. 2003. The impact of jumps in volatility and returns. *Journal of Finance* **58**: 1269–1300.

Fama E. 1965. The behavior of stock prices. *Journal of Business* **38**: 34–105.

Fleming J, Paye BS. 2006. High-frequency returns, jumps and the mixture of normals hypothesis. Working paper, Jones Graduate School of Management, Rice University.

Fleming J, Kirby C, Ostdiek B. 2003. The economic value of volatility timing using realized volatility. *Journal of Financial Economics* **67**: 473–509.

Gallant AR, Rossi PE, Tauchen G. 1992. Stock-prices and volume. *Review of Financial Studies* **5**: 199–242.

Gillemot L, Farmer JD, Lillo F. 2005. Theres more to volatility than volume. Working paper, Santa Fe Institute.

Hansen L, Lunde A. 2005. A realized variance for the whole trading day based on intermittent high-frequency data. *Journal of Financial Econometrics* **3**: 525–554.

Hansen PR, Lunde A. 2006. Realized variance and market microstructure noise. *Journal of Business and Economic Statistics* **24**: 127–161.

Hsieh D. 1989. Modeling heteroskedasticity in foreign exchange rates. *Journal of Business and Economic Statistics* **7**: 307–317.

Huang X, Tauchen G. 2005. The relative contribution of jumps to total price variance. *Journal of Financial Econometrics* **3**: 456–499.

Hull JC, White A. 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance* **42**: 281–300.

Jiang GJ, Oomen RCA. 2008. Testing for jumps when asset prices are observed with noise: a 'swap variance' approach. *Journal of Econometrics* **144**: 352–370.

Johannes M. 2004. The statistical and economic role of jumps in interest rates. *Journal of Finance* **59**: 227–260.

Kim D, Kon SJ. 1994. Alternative models for the conditional heteroskedasticity of stock returns. *Journal of Business* **67**: 563–598.

Lai T, Siegmund D. 1983. Fixed accuracy estimation of an autoregressive parameter. *Annals of Statistics* **11**: 478–485.

Lee SS, Mykland PA. 2008. Jumps in financial markets: a new nonparametric test and jump dynamics. *Review of Financial Studies* **21**: 2535–2563.

Liesenfeld R. 1998. Dynamic bivariate mixture models: modeling the behavior of prices and trading volume. *Journal of Business and Economic Statistics* **16**: 101–109.

Luu JC, Martens M. 2003. Testing the mixture of distributions hypothesis using 'realized' volatility. *Journal of Futures Market* **23**: 661–679.

Maheu JM, McCurdy TH. 2004. News arrival, jump dynamics, and volatility components for individual stock returns. *Journal of Finance* **59**: 755–793.

Mancini C. 2005. Disentangling the jumps of the diffusion in a geometric jumping Brownian motion. Working paper, Department of Mathematics, University of Florence.

Mandelbrot B. 1963. The variation of certain speculative prices. *Journal of Business* **36**: 394–419.

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

Murphy A, Izzeldin M. 2006. Order flow, transaction clock, and normality of asset returns: Ane and Geman (2000) revisited. Working paper, Oxford and Lancaster Universities.

Nelson DB. 1991. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* **59**: 347–370.

Nielsen M, Frederiksen PH. 2008. Finite sample accuracy and choice of sampling frequency in integrated volatility estimation. *Journal of Empirical Finance* **15**: 265–286.

Oomen R. 2006. Properties of realized variance under alternative sampling schemes. *Journal of Business and Economic Statistics* **24**: 219–237.

Peters RT, de Vilder RG. 2006. Testing the continuous semimartingale hypothesis for the S&P 500. *Journal of Business and Economic Statistics* **24**: 444–454.

Ploberger W. 2005. A generalization of the Skorohod embedding to random elements in Polish spaces. Working paper, University of Rochester.

Rosenbaum M. 2007. A new microstructure noise index. Working paper, Université Paris-Est and CREST-ENSAE.

Tauchen G, Pitts M. 1983. The price variability–volume relationship on speculative markets. *Econometrica* **51**: 485–505.

Todorov V. 2009. Econometric analysis of jump-driven stochastic volatility models. *Journal of Econometrics* forthcoming.

Zhou B. 1998. Parametric and nonparametric volatility measurement. In *Nonlinear Modelling of High-Frequency Financial Time Series*, Dunis C, Zhou B (eds). Wiley: London; 109–129.