

Continuous-Time Models, Realized Volatilities, and Testable Distributional Implications for Daily Stock Returns¹

Torben G. Andersen²
Northwestern University

Tim Bollerslev³
Duke University

Per Frederiksen⁴
Nordea Markets

Morten Ørregaard Nielsen⁵
Cornell University

January 21, 2008

¹We are grateful to Asger Lunde for help with data extraction and for making his Gauss codes available, and Mark Kamstra for detailed comments on an earlier draft. We also thank Frank Diebold for numerous discussions on closely related ideas, as well as three anonymous referees and conference and seminar participants at the November 2005 SAMSI web conference on Lévy Processes in Finance and Econometrics, the December 2005 Thiele Symposium in Copenhagen, the March 2006 Risk Management Conference at Mont Tremblant, the April 2006 Realized Volatility Conference in Montreal, London School of Economics, University of Lausanne, the Swiss National Bank, and the Department of Statistics at Northwestern University. The work of Andersen and Bollerslev was supported by a grant from the NSF to the NBER, and CREATES funded by the Danish National Research Foundation. The work of Nielsen by a grant from the Danish Social Sciences Research Council (grant no. FSE 275-05-0220), and CREATES funded by the Danish National Research Foundation.

²Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, NBER, and CREATES; phone: 847-467-1285; e-mail: t-andersen@kellogg.northwestern.edu

³Corresponding author. Department of Economics, Duke University, Durham, NC 27708, NBER, and CREATES; phone: 919-660-1846; e-mail: boller@econ.duke.edu

⁴Equity Trading and Derivatives, Nordea Markets, 1401 Copenhagen C, Denmark; phone: +45-3333-6683; email: per.frederiksen@nordea.com

⁵Department of Economics, Cornell University, 486 Uris Hall, Ithaca, NY 14853, and CREATES; phone: 607-255-6338; e-mail: mon2@cornell.edu

Abstract

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 intraday 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 thirty 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.

JEL Classifications: C1, G1.

Keywords: Return distributions, continuous-time models, mixture-of-distributions hypothesis, financial-time sampling, high-frequency data, volatility signature plots, realized volatilities, jumps, leverage and volatility feedback effects.

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 thirty 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 semimartingale 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 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 tried to relate it to trading volume. Other early attempts at directly linking the mixture with observable economic variables include Epps & Epps (1976) and Tauchen & Pitts (1983), both of whom argue for an association with the market activity as captured by the cumulative trading volume over daily or inter-daily horizons.

These earlier studies focused primarily 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; see, e.g., Engle (2004). The pronounced predictability in volatility has, in turn, provided the impetus for a number of more recent empirical studies exploring the relationship between return variability and the fundamental mixing variable(s) within the MDH context; see, e.g., Gallant, Rossi & Tauchen (1992), Andersen (1996), Liesenfeld (1998), Bollerslev & Jubinski (1999), and Ane & 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, Bollerslev, Christoffersen & Diebold (2006). 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, in the sense that large negative returns tend to be associated with higher future volatility than positive returns of the same magnitude. This asymmetry, first forcefully documented by Nelson (1991), is now

¹The robustness of the empirical findings in Ane & Geman (2000) have recently been called into question by Gillemot, Farmer & Lillo (2005) and Murphy & Izzeldin (2006).

generically referred to as a leverage effect, although it is widely agreed that the effect has little, if anything, 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, have been based on continuous-time methods and models. For instance, the Black-Scholes option pricing formula assumes that prices evolve continuously according to a homogeneous, or time-invariant, 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 empirically more realistic continuous-time formulations. In particular, while the early contributions by Merton (1976) and Hull & White (1987) argued for the need to incorporate jumps and time-varying diffusive volatility in the pricing of options, respectively, a number of more recent studies have demonstrated the need to simultaneously allow for both effects in order to satisfactorily represent actually observed daily and weekly prices; see, e.g., Andersen, Benzoni & Lund (2002) and Chernov, Gallant, Ghysels & Tauchen (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 crucially 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 & Bollerslev (1998*a*), Andersen, Bollerslev, Diebold & Labys (2001, 2003), and Barndorff-Nielsen & 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.

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 the logical extension of the earlier empirical investigations of Andersen, Bollerslev, Diebold & Ebens (2001) who report that the unconditional distributions of the raw daily stock returns in the Dow-Jones index have fat tails, but when standardizing the daily returns by the corresponding realized volatilities, constructed from the summation of high-frequency intra-day 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

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

theoretical perspective, this is actually 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 empirically relevant features may invalidate this result. First, there are inevitable errors in the realized volatility measures due to discretization and noise, as we do not sample the high-frequency returns continuously and they are subject to market microstructure frictions. Second, it is likely there are jumps, or 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.

Each potential source of error warrants careful attention in the implementation phase. 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 & Shephard (2004), defined by the summation of properly scaled adjacent absolute high-frequency returns, 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 & Shephard (2006) and Huang & 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 as well as the distributional implications of jumps.⁴ Third, to alleviate the impact of return-volatility asymmetries, e.g., the so-called 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 semimartingale 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 distributional tests are applicable across the full range of standard jump diffusive models for asset returns.

Our approach is related to Peters & de Vilder (2006) and Andersen, Bollerslev & Dobrev (2007) 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

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

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

⁵See also Zhou (1998) for more informal empirical evidence along these lines for the Deutschemark/Dollar exchange rate.

minimal microstructure distortions and high liquidity. We focus on the much broader set of thirty 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 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, by investigating a much larger set of assets, we obtain important evidence regarding the robustness of the findings which is not feasible on the basis of the prior studies. From a methodological point there are also a number of differences. Most strikingly, Peters & 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, Bollerslev & Dobrev (2007) we accommodate the issue of microstructure noise more directly through the generalized signature plots and we rely on the new sequential jump detection technique which is related to well-known procedures from the literature. In fact, our identification of jump days is directly justified through the asymptotic distribution of the standard test statistic under a general diffusive null hypothesis as developed in Barndorff-Nielsen & Shephard (2004, 2006) and Huang & Tauchen (2005). The jump identification approach in Andersen, Bollerslev & Dobrev (2007) appears to fare well in simulation settings but it is only formally justified under constant volatility across the trading day and the jumps 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 has the impact of reducing the number of rejections of normality across the series. Such evidence is only truly 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 broad range of issues within financial economics. We further expect the new financial-time scale, dictated by the realized variation measure, to be useful in future work in terms of refining the popular MDH by providing a directly observable candidate for the otherwise latent mixing variable.

The plan for the rest of the paper is as follows. The formal theoretical arguments underlying the 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 different 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 and guard against the adverse effects of market microstructure biases at the very highest sampling frequencies. Section 5 discusses some preliminary summary statistics related to the importance and timing of jumps and leverage effects. The outcomes of the different distributional tests are summarized in Section 6, along with various robustness checks and a general discussion of the implications for

continuous-time stock price modeling. Section 7 concludes. More detailed information for each of the individual stocks are provided in tabular and graphical form 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,

$$dp(t) = \sigma(t) dw(t) + \kappa(t) dq(t), \quad t \geq 0, \quad (1)$$

where the instantaneous, or diffusive, volatility process $\sigma(\cdot) > 0$ is càdlàg, $w(\cdot)$ denotes a standard Brownian motion assumed to be independent of the drift, the counting process $q(t)$ is normalized such that $dq(t) = 1$ corresponds to a jump at time t , and $dq(t) = 0$ otherwise, with the $\kappa(t)$ process describing the size of the jump if a jump actually occurs at time t . For notational simplicity we explicitly exclude a drift term, $\mu(t) dt$, in equation (1), but all of the theoretical results below could easily be extended to allow for a non-zero drift, $\mu(\cdot) \neq 0$.⁷

While theoretical pricing arguments often most conveniently are expressed within a continuous-time framework, empirical investigations are invariably based on discretely sampled prices, or returns. We denote the one-period continuously compounded discrete-time returns implied by the jump-diffusion in (1) as,

$$r_t \equiv p(t) - p(t-1), \quad t = 1, 2, \dots, \quad (2)$$

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.

⁶Although this formulation, as formally defined in equation (1), is quite general in allowing for both time-varying jump sizes and jump intensities, it does rule out infinite activity Lévy processes with infinitely many jumps over finite time intervals; see, e.g., Cont & Tankov (2004) for a general discussion of such processes, Todorov (2007) for an application involving jump-driven stochastic volatility models, and Barndorff-Nielsen, Shephard & Winkel (2006) for a discussion of realized variation measures in the context of certain infinite activity jump processes.

⁷Empirically the inclusion of a drift term would simply require the subtraction of a mean from the daily returns. In the empirical analysis we explicitly look at both raw and mean-adjusted returns. The results, reported below, are virtually identical.

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,

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

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 & Bollerslev (1998a), Andersen, Bollerslev, Diebold & Labys (2001), and Barndorff-Nielsen & 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. We provide a more in-depth discussion of these ideas in the context of our empirical implementation of equation (3) in Section 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 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, Johannes & Polson (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). However, suppose the jumps were

⁸The integrated volatility also plays a central role in option pricing models allowing for time-varying volatility; see, e.g., the aforementioned paper by Hull & White (1987).

known, and let the corresponding jump-adjusted returns be denoted by,

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

where the sum is to be interpreted as consisting of the non-zero jumps that occurred over the $(t-1, t]$ time-interval. Since all of the variation in the jump-adjusted returns comes from the diffusion component, standardizing by the integrated volatility should again result in a normal distribution,

$$\tilde{r}_t \left(\int_{t-1}^t \sigma^2(\tau) d\tau \right)^{-1/2} \sim N(0, 1), \quad t = 1, 2, \dots. \quad (5)$$

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

2.3 Leverage and Volatility Feedback Effects

The distributional results in the preceding two sections rule out so-called leverage and volatility feedback effects by assuming that 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 this assumption is unrealistic as the return-volatility relation is conditionally asymmetric in the sense that large negative returns are associated with larger volatilities than are positive returns of the same magnitude; see, e.g., Black (1976), Christie (1982), and Bollerslev, Litvinova & Tauchen (2006). Here, leverage effect is defined as correlation between volatility and past returns and volatility feedback as correlation between volatility and future returns.⁹ This implies 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. Intuitively, a priori fixing the sampling interval to encompass identical increments of realized return variation breaks the link between the numerator and denominator in equations (3) and (5).

Formally, let the event-time, or financial-time, sampling scheme be defined by,

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

⁹A similar leverage or volatility feedback effect could in principle work through the jump component. However, the related empirical evidence in Bollerslev, Kretschmer, Pigorsch & Tauchen (2008) suggests that the asymmetry works almost exclusively through the diffusive component.

where $t_0 \equiv 0$, and τ^* denotes the fixed financial-time unit spanned by each of the returns.¹⁰ For ease of comparison with the one-period, or daily, return distributions discussed above, we focus our empirical investigations on the case in which τ^* equals the unconditionally expected one-period, or daily, 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,

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

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

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

Importantly, this last theoretical 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. In contrast to traditional procedures for analyzing continuous-time models, our approach does not depend upon the validity of any particular parametric model. Yet, at the same time, the approach provides direct guidance for the specification of more realistic parametric models within the general class of jump-diffusions defined by (1).

3 Empirical Return and Variation Measures

Our empirical analysis of various transformed daily return distributions relies importantly on the availability of intra-day data. We let T denote the number of days for which intra-day price observations are available. The daily return series is then given by the increment in the observed log-prices over each trading day, i.e.,

$$R_t = p_{t,M} - p_{t,0}, \quad t = 1, \dots, T, \quad (10)$$

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 explicitly excludes the part of the daily variation associated

¹⁰A corresponding "business-time" sampling scheme for pure jump processes has previously been used in the analysis of Oomen (2006), while Zhou (1998) refers to similarly sampled returns as de-volatilized returns. This is also related to the ϑ -time sampling scheme advocated by Dacorogna, Gencay, Müller, Pictet & Olsen (2001) although they use a different underlying realized power variation scale.

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

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 are naturally thought of as deterministically occurring jumps. We treat them accordingly, and our trading day returns are then simply equal to the daily returns adjusted for the (observed) overnight jump. This approach allows us to avoid dealing with market closures where reliable high-frequency data generally are not available. At the same time, this feature implies that applications of the current results for predicting the distribution of future daily or lower frequency 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 this study, but the concurrent work by Andersen, Bollerslev & Huang (2006) exemplify 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 used in the construction of evenly spaced prices, say $M + 1$ per day, where preferably many more price observations are available each day. Denote the corresponding j 'th intra-daily log-price for day t by $p_{t,j}$, where $j = 0, 1, \dots, M$ and $t = 1, \dots, T$. The continuously compounded M intra-daily returns for day t are then similarly denoted by,

$$r_{t,j} = p_{t,j} - p_{t,j-1}, \quad j = 1, \dots, M, \quad t = 1, \dots, T. \quad (11)$$

The precision of the resulting nonparametric realized volatility and jump measures, as discussed further below, depends directly upon 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 a practical 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. Ways in which 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; see, e.g., Nielsen & Frederiksen (2008), Ait-Sahalia, Mykland & Zhang (2005), Bandi & Russell (2007), Barndorff-Nielsen, Hansen, Lunde & Shephard (2006), and Hansen & Lunde (2006), among many others. One complicating feature is that none of these contributions allow for the presence of jumps which is a major contributor to the price variability in the series we explore. In the empirical results reported below, we follow much of the literature in the use of a 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 new volatility signature type plots, as detailed in Section 4.

¹²The estimates reported in Hansen & 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 & Bollerslev (1998a), Andersen, Bollerslev, Diebold & Labys (2001) and Barndorff-Nielsen & Shephard (2002), we define the realized volatility for day t by,¹³

$$RV_t \equiv \sum_{j=1}^M r_{t,j}^2, \quad t = 1, \dots, T. \quad (12)$$

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 $M \rightarrow \infty$,

$$RV_t \rightarrow_p \int_{t-1}^t \sigma^2(s) ds + \sum_{s=q(t-1)+1}^{q(t)} \kappa^2(s), \quad t = 1, \dots, T. \quad (13)$$

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, in a series of recent papers Barndorff-Nielsen & 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 lag-one realized bipower variation is defined by,

$$BV_t \equiv \mu_1^{-2} \sum_{j=2}^M |r_{t,j}| |r_{t,j-1}|, \quad t = 1, \dots, T, \quad (14)$$

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

$$BV_t \rightarrow_p \int_{t-1}^t \sigma^2(s) ds, \quad t = 1, \dots, T. \quad (15)$$

Intuitively, for very large values of M , there will be at most one jump in any two adjacent time interval of length $1/M$. Since the contribution of each absolute return associated with the diffusion component in the limit (over an infinitesimal interval) is negligible, the product of a squared jump return with the adjacent absolute return is also vanishingly small asymptotically. Meanwhile, the scaling of the bipower variation term ensures that the consistency for the diffusive return variation is maintained. Now, combining equations (13) and (15), it follows readily that

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

the contribution to the total variation coming from the jump component is consistently estimated by the difference between the two. That is, for $M \rightarrow \infty$,

$$RV_t - BV_t \xrightarrow{p} \sum_{s=q(t-1)}^{q(t)} \kappa^2(s), \quad t = 1, \dots, T. \quad (16)$$

Although formally consistent for the squared jumps, nothing prevents $RV_t - BV_t$ 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 $RV_t - BV_t$ for finite M , even if there are no jumps, or $q(t) = q(t-1)$. Hence, following the empirical analysis in Andersen, Bollerslev & Diebold (2007), 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 & Shephard (2004, 2006), coupled with the extensive simulation evidence in Huang & Tauchen (2005), we assess the significance of the daily jump component according to the feasible logarithmic test statistic,

$$Z_t \equiv \sqrt{M} \frac{\ln RV_t - \ln BV_t}{((\mu_1^{-4} + 2\mu_1^{-2} - 5) TQ_t BV_t^{-2})^{1/2}} \xrightarrow{d} N(0, 1), \quad (17)$$

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

$$TQ_t \equiv \frac{1}{M} \mu_{4/3}^{-3} \sum_{j=3}^M |r_{t,j}|^{4/3} |r_{t,j-1}|^{4/3} |r_{t,j-2}|^{4/3}, \quad t = 1, \dots, T, \quad (18)$$

and $\mu_{4/3} = 2^{2/3} \Gamma(7/6) / \Gamma(1/2)$ with $\Gamma(\cdot)$ denoting the gamma function. Thus, only extreme (in a formal statistical sense) positive values of $RV_t - BV_t$ will be attributed to the squared jump component, i.e.,

$$JV_t \equiv I_{\{Z_t > \Phi_{1-\alpha}\}} (RV_t - BV_t), \quad t = 1, \dots, T, \quad (19)$$

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,

$$CV_t \equiv RV_t - JV_t = I_{\{Z_t \leq \Phi_{1-\alpha}\}} RV_t + I_{\{Z_t > \Phi_{1-\alpha}\}} BV_t, \quad t = 1, \dots, T. \quad (20)$$

That is, we estimate the day t continuous volatility component by the realized volatility if there are no significant jumps on day t , and by the realized bipower variation on days with significant jump(s). In the empirical results reported below we rely on a significance level of $\alpha = 1\%$,

but we also experimented with $\alpha = 5\%$ and 0.1% , resulting in qualitatively very similar overall conclusions.¹⁴

The high-frequency data based measures discussed above afford a simple-to-implement approach for estimating the daily integrated volatility and the sum of squared jumps over a given day. The approach does not, however, identify the individual jumps themselves. We next discuss two different procedures for doing so.

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.

3.2.1 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 $JV_t \rightarrow_p \kappa_t^2$. 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

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

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

$$\tilde{R}_t \equiv R_t - \tilde{\kappa}_t, \quad t = 1, \dots, T, \quad (22)$$

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 time scale. Of course, it is also possible that

¹⁴Note that 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 & Shephard (2006). Our choice of a low $\alpha = 1\%$ (and $\alpha = 0.1\%$ in our robustness checks) also helps alleviate concerns about sequential testing biases in the sense of providing a low Bonferroni-type bound for the more advanced sequential jump detection procedure discussed below.

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

multiple jumps may occur on certain days, violating the basic assumption underlying the simple procedure in (21). Hence, we next introduce a more advanced jump identification scheme designed to facilitate inference regarding *all* significant jumps along with their exact *timing* within the trading day.

3.2.2 Sequential Jump Adjustments

The Z_t statistic defined in (17) only signifies whether the difference between the realized volatility and the bipower variation measures is large enough to indicate the presence of one or more jumps on day t . 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 $RV_t - BV_t \rightarrow_p 0$, 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,

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

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.

Thus, having identified the total number of jumps, say J , during day t , as well as the magni-

tude of each of the jumps by the corresponding high-frequency returns,

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

where j_i denotes the exact time-interval of the intra-day return associated with the i 'th jump, we calculate the jump-adjusted daily return as,

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

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

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

where $JVS_{t,i}$ gives the contribution from the i 'th jump, defined as the difference between the i 'th largest intra-day squared return and the average of the $M - J$ squared returns that are not associated with jump(s). That is,

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

where $Z_{t,i}$ denotes the i 'th sequential jump statistic, as discussed above. Lastly, the corresponding continuous volatility component is simply defined by,

$$CVS_t \equiv RV_t - JVS_t, \quad t = 1, \dots, T, \quad (27)$$

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) may in effect be interpreted as an estimate of $\sum_{i=1}^J \hat{\kappa}_{t,i}$. This suggests that the two procedures should give rise to similar jump-adjustments for days in which there is only one jump. However, the ability of the sequential procedure to identify multiple jumps within a single day and the timing of jumps is crucial for the construction of jump-adjusted intra-day return series and these series, in turn, serve as a critical input into the empirical analysis below.

4 Data Description

4.1 Data Sources and Construction

Our data is extracted from the Trade And Quotation (TAQ) database, and consist of all recorded trades and quotes for all the 30 Dow Jones Industrial Average (DJIA) stocks for the five-year

period spanning January 2, 1998 through December 31, 2002. The ticker symbols and names for each of the stocks are listed in Table A1 of the supplementary appendix.¹⁶ We only use the prices from the New York Stock Exchange (NYSE), with the exception of Intel and Microsoft, both of which are most actively traded on the National Association of Security Dealers Automated Quotation (NASDAQ) system. Mirroring the data cleaning procedures of Hansen & Lunde (2006), we filter the raw price series to remove price observations equal to zero, prices occurring outside the 9:30 AM to 4:00 PM official trading day, as well as extreme outliers or mis-recorded price observations. This leaves us with somewhere between 2-4 million prices for each of the stocks, with the exception of Intel and Microsoft, both of which have around 26 million prices over the five-year sample. Further, we also delete days of early closing or late opening of the exchange, as well as days in which trading in a particular stock was suspended for an extended period of time, resulting in approximately 1, 255 “intact” days for each stock.

To minimize the impact of 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.¹⁷ The choice of a 5-minute return interval is in line with most of the existing empirical literature and, as argued in Bandi & Russell (2007), this sampling frequency 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 scheme in the present context, explicitly allowing for jumps, is further supported by the volatility signature plots discussed next.

4.2 Volatility Signature Plots


The conventional realized volatility signature plot popularized by Andersen, Bollerslev, Diebold & Labys (2000*b*) provides a simple and 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 begin to flatten out at the frequencies for which the market microstructure frictions cease to have a distorting influence.

The signature plots in Figure A1 in the supplementary appendix for each of the individual stocks generalize this idea, by plotting the average realized bipower variation measures together with the standard realized variation for different sampling frequencies. A cursory look at these graphs reveals a close similarity in the general shape across the different stocks. In order to effectively summarize the results we plot in Figure 1 the median values (over the 30 stocks) of

¹⁶All of the tables in the supplementary appendix are available from the authors upon request.

¹⁷As argued in Hansen & Lunde (2006), the use of mid-quotes tends to reduce spurious serial correlation in the high-frequency returns due to bid-ask bounce and non-synchronous trading effects.

Figure 1: Median generalized volatility signature plots



the average realized variation measures for each sampling frequency (measured in seconds).¹⁸

By reproducing the average (across days) RV_t and BV_t measures as a function of $1/M$ within the same graph, the volatility signature plot affords an informal way to gauge the importance of jumps. In particular, it follows from equation (16) that, under ideal conditions and for $1/M \rightarrow 0$, the distance between the two lines provides a consistent estimate of the total variation due to jumps.¹⁹ 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 two lines shows a tendency to stabilize at a sampling frequency of only two minutes, or 120 seconds. These effects are also directly in line with the extensive simulation based evidence for the two measures based on empirically realistic continuous-time processes subject to “noise” reported in Nielsen & Frederiksen (2008).²⁰ As such, this suggests that even though the individual realized volatility and bipower variation measures can be adversely affected by market microstructure frictions at

¹⁸Both of the variation measures have been converted to percent by multiplication with 10,000.

¹⁹Related generalized volatility signature plots, including plots for various integrated quarticity measures, have recently been explored by Andersen, Bollerslev, Frederiksen & Nielsen (2006).

²⁰The theoretical framework recently developed by Rosenbaum (2007) may also be useful in providing a more formal theoretical explanation for the observed patterns.

lower frequencies, the impact of the noise tends to cancel so that the difference between the two measures, and hence the estimate of the jump component, remains remarkably stable for $1/M$ in excess of 120 seconds. Taken as whole, this therefore support our use of a 5-minute return interval as a reasonable, albeit for some stocks also somewhat conservative, uniform sampling scheme.

This completes our discussion of the data and empirical measures used in testing the main theoretical distributional implications. However, before presenting the outcome of our moment-based tests for (approximate) normality, it is useful to first consider some preliminary summary statistics.

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. Thus, in order to place the evidence emerging from the subsequent distributional tests within an appropriate context, 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. Following the discussion above, we first report results based on the simple jump-detection procedure, followed by the more involved sequential jump-detection scheme, explicitly identifying the time(s) of jump(s) within the day.

5.1.1 Simple Jump Detection

Relying on the simple jump-detection method and a significance level of $\alpha = 1\%$, Table 1 displays the resulting mean duration between significant jumps, the relative contribution of jumps to the total realized variation, i.e., JV_t/RV_t , the mean size of the jump component for significant jump days (multiplied by 10,000), and lastly the corresponding mean (absolute) jump size (measured in percent), 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 thirty 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.²¹ As such, these initial summary statistics suggest that

²¹See, e.g., the parametric daily GARCH-jump model estimates for individual stocks reported in Maheu &

Table 1: Jumps - Simple Method

	Mean duration	Rel. jump contribution JV_t/RV_t	Mean size of jump component (x10,000)	Mean size of actual jumps (x100)
Mean across stocks	6.3201	0.0476	1.2119	0.9812
Std. dev. 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: The 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 (x10,000), 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.

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 analyzed on the basis of intra-day data rather than the traditional daily return series, see, e.g., Andersen, Bollerslev, Diebold & Vega (2003).

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 percent (GE) to 7.5 percent (MO), with an average value of 4.7 percent. The more detailed results in Table A2 of the supplementary appendix also point towards a negative association between durations between jumps and the 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 on average. Finally, the mean absolute size of the “simple” jumps, i.e., $|\tilde{\kappa}_t|$, ranges from 0.74 to 1.31 percent, with a mean across all stocks of 0.98 percent.

5.1.2 Sequential Jump Detection

In contrast to the simple method, the sequential jump-detection procedure explicitly accommodates the presence of multiple jumps on a given trading day. In particular, it follows from the detailed results for each of the individual stocks reported 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 percent, while there is a two-percent probability of two jumps in one day. Meanwhile, the probability of three or more jumps in one day is close to, although not identically equal to, zero. These results illustrate the potential importance of the sequential jump detection procedure, as most of the stocks have many days with multiple jumps.

McCurdy (2004).

Table 2: Jumps - Sequential Method

	Rel. jump contribution JVS_t/RV_t	Mean size of jump component (x10,000)	Mean size of actual jumps (x100)
Mean across stocks	0.0373	1.0394	0.9282
Std. dev. 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: The 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 (x10,000), 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 3: Simple and Sequential Jump Correlations


	Correlation	RMSE	Theil's U
Mean across stocks	0.9450	0.0062	0.2999
Std. dev. 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: The 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.

At the same time, comparing the summary statistics in Table 2 for the sequential jump detection method to the corresponding statistics for the simple method in the last three columns of Table 1, the numbers are generally fairly close. The relative contribution of the jump component for the sequential procedure ranges from 2.1 percent (GE) to 5.8 percent (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 1, and the mean absolute jump size ranges from a low of 0.68 percent (JNJ) to a high of 1.25 percent (HPQ), with the overall absolute mean jump size of 0.93 percent again being slightly below that in Table 1.

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 3 which presents a summary of various correlation measures between the two. To directly focus on the relation between the jump series, all the common no-jump zero observations were not included in the computations. The first column reports the standard sample correlation coefficient, the second column the root mean squared error (RMSE) calculated as the square-root of the sum of the squared differences between the two series, and the third Theil's scale invariant U-statistic. As above, the results are summarized in the form of the mean, standard deviation, minimum, and maximum across the thirty stocks,

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



with detailed results for each stock provided in 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), with an average value of 0.95 across the thirty stocks. Also, the RMSEs and Theil’s U-statistics are generally fairly low across the stocks. Hence, the sequential procedure retains the information regarding jump occurrence and relative importance on a day-to-day basis but, importantly, it also directly identifies the intra-day timing of *all* the 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 provide a preliminary assessment of the validity of this assumption, Figure A3 of the supplementary appendix graphs the 5-minute cross-correlations for each of the stocks, i.e., $\text{corr}(|r_j|, r_{j+i})$, where for notational simplicity r_j for $j = 1, \dots, J$ refers to time series of approximately $J = 1,255 \times 78 = 97,890$ demeaned 5-minute returns available for each stock. A cursory look at the graphs suggests a broadly similar shape across stocks, although the idiosyncratic

Table 4: Leverage and Volatility Feedback Effect Estimates

	Leverage	Feedback	Difference
Mean across stocks	-0.0166	0.0076	-0.0243
Std. dev. 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: The 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, respectively. For further details, see Table A5 in the supplementary appendix.

noise inherent in the individual estimates makes it hard to draw sharp conclusions. Hence, in an effort to minimize the impact of the estimation error, 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 of the 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.²²

In order to more formally test the statistical significance of these features, Table 4 provides summary statistics directly related to the leverage and volatility feedback type effects depicted in the figure. This allows for more formal testing concerning the statistical significance of these features. 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 appendix, Table A5. The average *leverage effect* for an individual stock is estimated by,

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

²²This is also directly in line with the corresponding plots for the 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. As such, the change to financial-time sampling is likely to be even more important in restoring normality for the standardized index returns than it is for the individual stock returns analyzed here.

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 (un-weighted) mean of the sample cross-covariances between the absolute returns and the lagged $2, \dots, (K-1)$ 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 $2, \dots, (K-1)$ 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.²³

The 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 a considerable cross-sectional variation in the magnitude of the leverage effect, in particular, and for some stocks the cross-covariances appear quite significant.²⁴ We also note that the difference between the two effects is negative for all stocks, and significantly so at the 5% level for twenty out of the thirty.

These results suggest that for some stocks the use of financial-time sampling may be necessary to restore normality of the standardized return distributions. Of course, whether the high-frequency leverage and volatility feedback effects are large enough to cause any 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 asset returns, and stock returns in particular, are fat-tailed. At the same time, our theory predicts that suitably jump-adjusted and standardized stock returns should be i.i.d. Gaussian. Hence, as a natural benchmark, we

²³We also calculated the same statistics for the jump-adjusted returns, resulting in very similar numbers to the ones reported in the tables. These results are available upon request.

²⁴These high-frequency based findings are further corroborated by our estimation of conventional EGARCH models for the daily returns, which tend to result in the most significant volatility asymmetries for the stocks for which the leverage effects in Table A5 in the supplementary appendix are also the largest. These additional estimation results are available upon request.

Table 5: Daily Return Distributions

Series	Raw Returns		Demeaned Returns	
	Significance		Significance	
	5 % level	1 % level	5 % level	1 % level
$R_t/\sqrt{Var(R_t)}$	30	30	30	30
$R_t/\sqrt{GARCH(1,1)}$	30	30	30	30
$R_t/\sqrt{RV_t}$	21	12	18	9
$\tilde{R}_t/\sqrt{CV_t}$	18	10	15	9
$\hat{R}_t/\sqrt{CVS_t}$	20	12	18	11
$\hat{R}_k^*/\sqrt{E(CVS_t)}$	13	6	11	5
$\hat{R}_{5k,5}^*/\sqrt{5E(CVS_t)}$	6	3	2	2

Note: The 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 $E(CVS_t)$ 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 $5E(CVS_t)$ time-units. For further details regarding each of the individual stocks, see Table A6 in the supplementary appendix.

first provide a summary of the raw unconditional return distributions for the thirty DJIA stocks. The first row of Table 5 confirms the stylized facts regarding daily stock returns. Using the normality tests of Andersen, Bollerslev & Dobrev (2007) involving the joint distribution of the first four sample moments, the null hypothesis that the unconditional return distribution, or $R_t/\sqrt{Var(R_t)}$, is standard normal is rejected at the 1% level for all of the stocks.²⁵ Table A6 in the supplementary appendix indicates that the overwhelming rejections are due primarily to the sample kurtosis being significantly in excess of three.

As discussed at length above, these results are exactly what we should expect to find if the underlying return volatility is time-varying, as this naturally induces a mixture type distribution. Motivated by this line of reasoning, we next look at the unconditional return distributions obtained by standardizing the daily returns with the one-day-ahead conditional volatility forecasts from a conventional discrete-time GARCH(1,1) model.

6.2 GARCH Standardized Returns

The results for the GARCH standardized returns, $R_t/\sqrt{GARCH(1,1)}$, reported in the second row of Table 5, are again fully consistent with the existing literature. In particular, even though

²⁵This test is formally equivalent to testing that the first four orthogonal Hermite polynomials are equal to zero, and may be seen as a special case of the more general class of normality tests developed by Bontemps & Meddahi (2005a, 2005b) for testing the so-called Stein equation.

the mass in the tails of the GARCH standardized return distributions is reduced relative to that of the unconditional return distributions, the distributions remain significantly leptokurtic for all stocks; see Bollerslev (1987), Baillie & Bollerslev (1989), and Hsieh (1989) for early empirical evidence along these lines.²⁶

Of course, if the underlying price and volatility processes evolve stochastically within the trading day, the GARCH volatilities will 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-a-vis the true integrated volatility realizations. In order to explore this conjecture, we next turn to the distributions obtained by standardizing the daily returns with the corresponding realized volatilities. Intuitively, since the measured realized volatilities provide more accurate ex-post estimates of the integrated volatility realizations than the ex-ante GARCH forecasts, we would 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, $R_t/\sqrt{RV_t}$. From the density and QQ-plots for the individual stocks in Figures A8 and A9 of the supplementary appendix, respectively, it is evident that the RV standardized return distributions are much closer to the reference Gaussian distributions than are the raw and GARCH standardized returns. Most obvious, the tails in the QQ-plots have improved considerably and, with a few exceptions, show only slight deviations from the straight 45-degree line. As previously mentioned, these findings are in line with the earlier studies by Andersen, Bollerslev, Diebold & Labys (2000a) and Andersen, Bollerslev, Diebold & Ebens (2001), arguing through similar informal graphical tools and simple summary statistics that the sample distributions of the RV_t standardized returns tend to be well approximated by a normal distribution.

Complementing this informal evidence, the third row in Table 5 reports the results from applying our formal moment-based test to the realized volatility standardized return distributions. Importantly, however, as shown in Andersen, Bollerslev & Dobrev (2007), 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. In the present context with $M = 78$ 5-minute returns per day, this translates into a value of 2.925. The results confirm that the first four sample moments of $R_t/\sqrt{RV_t}$ 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

²⁶We also experimented with the estimation of alternative EGARCH-M formulations, explicitly allowing for volatility asymmetry, or leverage and/or volatility feedback effects, resulting in very similar findings to the ones reported here; see also Kim & Kon (1994) for earlier related empirical evidence.

out of the thirty stocks at the 5% significance level, while the tests turn out to be insignificant for eighteen of the thirty 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 $R_t/\sqrt{RV_t}$ remains significantly different from the theoretical value of $m_4 = 2.925$ that should obtain for a homogeneous diffusion in many cases. Of course, as noted above, several studies argue for the importance of allowing for jumps in the modeling of aggregate equity index returns, in particular, and the empirical results for the individual stocks reported in Section 5.1 support this notion. The presence of a few large jumps tends to imply that the RV_t standardized return distribution has thinner tails than the (modified) normal because the jumps inflate the denominator realized volatility disproportionately. However, more generally the presence of jumps simply serves to obfuscate the asymptotic, for increasing sampling frequency, normality of the $R_t/\sqrt{RV_t}$ distribution. Indeed, even though the majority of the rejections reported in Table 5 arise from exceedingly low sample values of m_4 , for a few of the 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. The summary results of the normality tests for the corresponding distributions labelled $\tilde{R}_t/\sqrt{CV_t}$ and $\hat{R}_t/\sqrt{CVS_t}$, respectively, are given in the fourth and fifth rows of Table 5. Perhaps somewhat surprisingly, the results indicate that neither of the jump-adjusted standardized return series are systematically closer to being normally distributed than are the $R_t/\sqrt{RV_t}$ non-adjusted realized volatility standardized returns. The hypothesis of normality is rejected for eighteen of the stocks at the 5% level using the simple method and twenty of the stocks using the more advanced sequential procedure, compared to twenty-one stocks for the non-adjusted returns. Similarly, at the 1% level, ten and twelve stocks reject for the jump-adjusted return distributions, while twelve stocks reject for the unadjusted returns.

Thus, even though jumps appear empirically important and, according to the estimates in Section 5.1, account for approximately a third of the total daily price variation on jump days, adjusting for jumps fails in terms of restoring normality of the standardized return distributions. One reason is that jumps are largely self-standardizing in the sense that a large jump tends to inflate the (absolute) value of both the daily return and the daily realized volatility, i.e., both the numerator and denominator of the standardized returns, so that they effectively cancel each other and the overall effect is muted. As a result, even if the jumps impact the raw return distribution significantly they exert much less influence on the realized volatility standardized

return distribution. In sum, the remaining, still appreciable, deviations from normality likely stem from a different source. According to our theoretical framework, the only other candidate is systematic dependencies between the numerator and denominator of the daily standardized returns. Of course, the simple empirical correlation-based measures discussed in Section 5.2 indeed suggest that a leverage type effect may be at work. To investigate this conjecture, our final set of results study 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 the realized value of daily integrated volatility is informative about the corresponding daily returns, or vice versa, as implied by the leverage and volatility feedback effects, the discretely sampled integrated volatility standardized returns from a continuous sample path diffusion are generally not normally distributed. Meanwhile, as discussed in Section 2.3, this dependence between the numerator and the denominator of the standardized return distribution may be broken by sampling the returns in so-called event time, or financial time. The new sequential jump-adjustment procedure, which explicitly identifies the timing of the jumps within the day, permits the construction of such financial-time returns by accumulating the jump-adjusted intraday returns so that they span identical increments in the corresponding CVS_t process, but time-varying calendar-time intervals. In practice, to calculate \hat{R}_k^* we included intraday returns until the sum of squared returns was equal to or exceeded τ^* ; i.e., the average daily (when $\tau^* = E(CVS_t)$) respectively weekly (when $\tau^* = 5E(CVS_t)$) realized volatility in calendar time. Importantly, only non-jump returns as identified by the sequential jump detection scheme were included.²⁷

The second to last row in Table 5, labelled $\hat{R}_k^*/\sqrt{E(CVS_t)}$, reports the results from applying the moment-based tests to the jump-adjusted returns spanning an identical amount of continuous sample path volatility where, for ease of comparison with the other entries in the table, we calibrate the unit of the financial time scale to one average trading day; i.e., $\tau^* = E(CVS_t)$. Interestingly, the move to financial-time sampling results in a marked reduction in the number of stocks for which the assumption of normality is rejected, with only six (five for the demeaned returns) stocks now rejecting at the 1% level. The close approximation afforded by the normal distribution is also evident from the detailed density and QQ-plots in Figures A14 and A15 in the supplementary appendix which, except for a few stocks, show a remarkably close coherence between the empirical and theoretical distributions.

Comparing the results of the tests for leverage and volatility feedback effects for each stock in Table A5, with the more detailed results of the normality tests in Table A6, there is generally also

²⁷Note that the simple jump-adjustment procedure only identifies the days with at least one jump, and not the actual within day time of the jump(s). Hence, on jump-days the approximation to equal units of CV_t time could be very poor.


a close relation between the significance of the former tests and 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 the normality of the standardized return series in calendar time is rejected at the 5% level. Meanwhile, the p -value for the joint normality test for the one-day financial-time returns equals 0.316. Conversely, for JPM, one of the two stocks for which normality of the one-day returns is rejected at the 5% level in financial-time but not in calendar-time, the leverage effect appears insignificant and the volatility feedback effect is only marginally significant at the 10% level.

The CVS_t time series used in constructing the new financial-time scale varies considerably over the sample for many of the stocks. Consequently, some of the corresponding “one-day” observations span intra-day returns over several calendar days while others are based on the summation of only a few squared 5-minute returns. As such, the underlying asymptotic theory, for the number of high-frequency return observations going to infinity, may provide a poor approximation in these situations. Hence, the last row in Table 5, labelled $\hat{R}_{5k,5}^*/\sqrt{5E(CVS_t)}$, reports the results of the normality test applied to returns spanning one financial “week,” or five average “days;” i.e., $\tau^* = 5E(CVS_t)$. Remarkably, the normality assumption for this longer return horizon, but shorter time series, is now only rejected at the 1% level for three (two for the demeaned returns) of the thirty stocks.²⁸

To more directly highlight the improvements in the accuracy of the normal approximation afforded by the sequential distributional adjustments discussed above, and in turn the importance of the corresponding qualitative features of the underlying continuous-time model, Figure 3 plots the p -values for the different tests for each of the thirty individual stocks and return transformations in Table 5. If the respective distributions are indeed (approximately) Gaussian, and the individual tests are independent, these 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 clearly improves the situation, but all the 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 very close to uniformly distributed between zero and one. Thus, the p -value plots further support the hypothesis that by moving to financial time, as measured through the jump-adjusted realized return variation, the normality of the (jump-adjusted) returns is restored. In essence this confirms that inter-daily stock prices may usefully be thought of as discretely sampled observations from an underlying continuous-time jump-diffusion model, but for many stocks it is essential to also accommodate leverage and/or

²⁸Of course, the powers of the tests based on fewer longer horizon return observations are likely to be lower. However, we also looked at the distribution of the standardized returns over longer calendar time periods, and we did not observe a similar dramatic reduction in the number of rejections for those series. These results are available upon request.

Figure 3: p -values for the 30 DJIA stocks, Jan. 1998 - Dec. 2002, 5-minute sampling




volatility feedback effects.

6.6 Alternative Sampling Frequencies

Our empirical results hinge on the use of high-frequency data for the construction of reliable realized variation measures and associated jump detection and financial-time sampling schemes. In particular, the new generalized volatility signature plots discussed in Section 4.2 led us towards a uniform 5-minute sampling frequency for each of the DJIA stocks. In order to confirm that a sampling frequency around this range provides a reasonable trade-off between the preference for very finely sampled returns and the potential contaminating market microstructure influences, we report here the results obtained from the identical distributional tests but based on both more and less frequently sampled intra-day returns.

Figure 4 reports the p -values for the different return transformations based on a coarser 30-minute, or half-hour, sampling frequency, corresponding to the right-most points in the median volatility signature plot in Figure 1. Under ideal conditions, the daily realized volatility measures and jump detection tests based on “only” $M = 13$ half-hourly intra-day observations are, of course, subject to much larger measurement errors than the 5-minute based measures and tests with $M = 78$. This effect directly manifests itself in the form of the noticeable deterioration in the dispersion of the p -values for the realized volatility standardized return distributions, which are now visible 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

Figure 4: p -values for the 30 DJIA stocks, Jan. 1998 - Dec. 2002, 30-minute sampling



close to uniform.

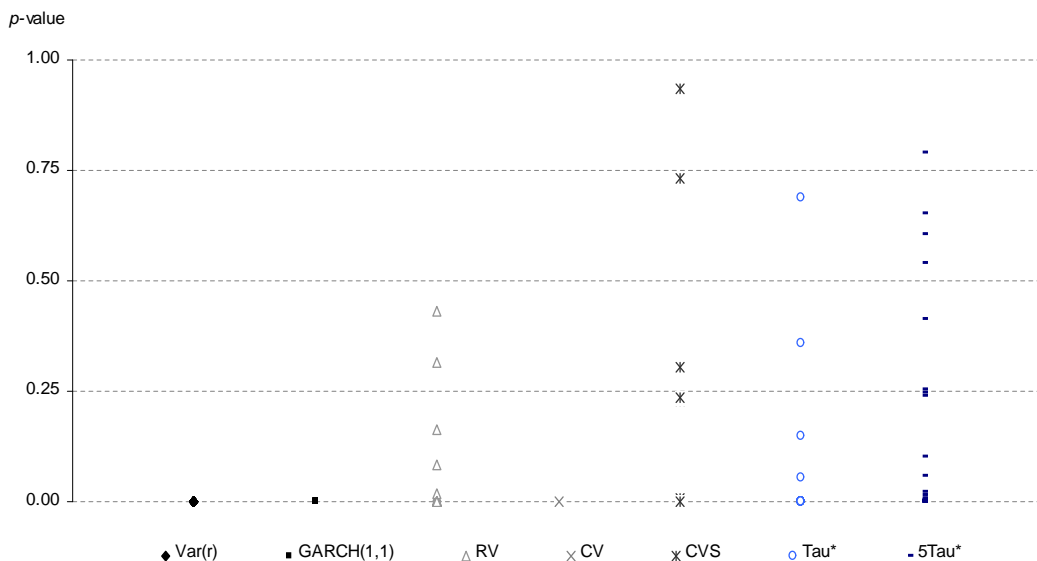
At the other end of the spectrum, Figure 5 reports the p -values obtained by using more finely sampled 30-second, or half-a-minute, returns; i.e., $M = 780$. This choice of M corresponds to the point in Figure 1 where the slope of the signature lines and the difference between the average realized variation and bipower variation measures start to diverge. Now a marked deterioration in the dispersion of the p -values for the financial-time returns becomes apparent. The contaminating influences from the market microstructure “noise” at this higher sampling frequency overwhelm the signal in the realized variation measures. Not surprisingly, a direct investigation of this very high-frequency series (not reported here) also revealed some quite dramatic violations of the basic arbitrage-free semi-martingale assumption for the price process.

In sum, the 5-minute sampling frequency used in the construction of the realized variation measures in the preceding sections appears to provide a reasonable choice for eliciting important distributional information from the high-frequency price data in the present context.

6.7 More Recent Data

The minimum tick size for trades and quotes on the NYSE was reduced from a sixteenth of a dollar to one cent on January 29, 2001. The recent results in Hansen & Lunde (2006) suggest that this finer tick size has been accompanied by a substantial reduction in the magnitude and impact of the market microstructure noise for a variety of realized volatility measures. To investigate the robustness of our findings with respect to this change in market structure, we replicated the

Figure 5: p -values for the 30 DJIA stocks, Jan. 1998 - Dec. 2002, 30-second sampling



analysis with data spanning the more recent, but slightly shorter, time period from February 2001 through December 2004.

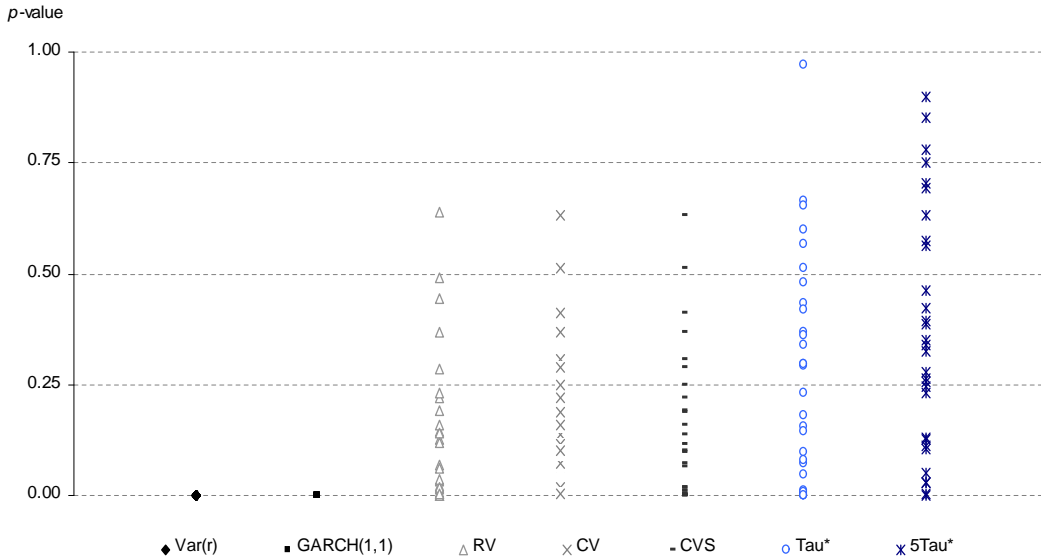
Without delving into details, a preliminary data analysis along the lines of Section 5 indicates that the jump intensities and the relative importance of jumps are somewhat smaller over the more recent period, but that jumps remain an important part of the price process for all stocks.²⁹ For instance, the average duration between jumps has increased by about fifty percent relative to the earlier period, and the relative jump contribution and the number of days with more than one jump have fallen by roughly half.

This diminished importance of jumps, coupled with the decreased magnitude of the market microstructure noise, suggest that the calendar-time standardized returns may be more closely approximated by a normal distribution. This is indeed the case. Comparing the p -value plots in Figures 6 and 3, the numbers for the $R_t/\sqrt{RV_t}$, $\tilde{R}_t/\sqrt{CV_t}$, and $\hat{R}_t/\sqrt{CVS_t}$ standardized distributions over the more recent time period in Figure 6 are clearly better dispersed than the values for the earlier period in Figure 3.³⁰ The closer coherence between the p -values for the RV_t , CV_t , and CVS_t standardizations and the p -values for the “daily” financial-time returns also indirectly suggests a diminished impact of the volatility asymmetry, or leverage effect, over the more recent time period. Still, with the exception of one or two stocks, all of the p -values for the

²⁹Complete results are available upon request.

³⁰Of course, the slightly shorter sample period is also likely to result in marginally lower powers of the different tests.

Figure 6: p -values for the 30 DJIA stocks, Feb. 2001 - Dec. 2004, 5-minute sampling



former standardizations are less than one-half. On the other hand, the p -values for the financial-time returns, and the “weekly” returns in particular, are again fairly close to being uniformly distributed over the entire unit interval.

As such, the broad conclusions for the more recent time period mirror our more detailed empirical findings for the longer earlier time period. Importantly, however, the new smaller tick sizes and apparent reduction in confounding market microstructure influences combine to present a less challenging environment for the realized volatility measures used in the analysis. Looking ahead, this suggests that the basic methodology and testing procedures developed here should provide even better guidance in future applications.

7 Concluding Remarks

We show how high-frequency intra-day data can be used in the construction of relatively simple-to-implement non-parametric realized variation measures and test statistics for better understanding the nature of the return distribution. Each step in the sequential test procedure speaks directly to important qualitative features of the underlying return generating process. As such, the tests may serve as a useful set of diagnostic tools in the specification of more accurate and empirically realistic continuous-time models. In this regard, our empirical results related to the individual DJIA stocks suggest that their price series may be satisfactorily described as discretely sampled observations from an underlying jump-diffusion model only after allowing for leverage and/or

volatility feedback effects.

Each step in the sequential procedure could be further extended and improved in a number of important directions. As discussed above, several recent studies argue for the use of new multi-scale or kernel-based realized volatility measures for more accurately measuring the true latent variation, e.g., Bandi & Russell (2007), Hansen & Lunde (2006), Barndorff-Nielsen, Hansen, Lunde & Shephard (2006), and Ait-Sahalia et al. (2005). Also, while the use of daily realized volatility measures conveniently circumvents complications associated with the strong intra-day patterns in volatility, e.g., Andersen & Bollerslev (1998*b*), the financial-time scale will invariably span different periods of the day, and it may prove beneficial to explicitly adjust and control for this feature. A number of alternative nonparametric jump detection procedures have recently been proposed in the literature, e.g., Jiang & Oomen (2005) and Mancini (2005), and it would be interesting to compare and contrast the results obtained here to some of these alternative schemes.

It would also be informative to directly relate the price jumps to different news arrivals, either in the form of company specific news, e.g. Johannes (2004), or more systematic macroeconomic news announcements, e.g. Andersen, Bollerslev, Diebold & Vega (2003). Similarly, it might prove instructive to associate the financial-time scale defined by the realized volatility to directly observable economic activity variables within the context of the MDH, e.g., Ane & Geman (2000) and Luu & Martens (2003). From the reverse perspective, given that our realized volatility and jump transformations have a sound theoretical basis and appear to significantly outperform any of the prior MDH style models for the return distribution on the empirical dimension, it is potentially very informative for future models within the MDH area to map their candidate economic mixing variables into the corresponding estimated diffusive and jump component series of return variation provided here.

Another very interesting question relates to the possible extension of the univariate distributional results and test statistics derived here to a multivariate setting. Although the notion of realized covariation may be defined in a straightforward manner, practical implementation issues related to the non-synchronicity of multiple high-frequency price series looms large, e.g., de Pooter, Martens & van Dijk (2007). The multivariate extension also presents formidable 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 would be interesting to more directly explore the return transformations and decompositions developed here in terms of their usefulness for value-at-risk type calculations, volatility timing, and other related financial decisions, e.g., Fleming, Kirby & Ostdiek (2003). We leave all of these issues for future research.

References

- Ait-Sahalia, Y., Mykland, P. A. & 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, T., Bollerslev, T., Christoffersen, P. & Diebold, F. (2006), Volatility forecasting, in G. Elliott, C. Granger & A. Timmermann, eds, ‘Handbook of Economic Forecasting’, North-Holland, Amsterdam, pp. 777–878.
- Andersen, T. G. (1996), ‘Return volatility and trading volume: An information flow interpretation of stochastic volatility’, *Journal of Finance* **51**, 169–204.
- Andersen, T. G., Benzoni, L. & Lund, J. (2002), ‘An empirical investigation of continuous-time equity return models’, *Journal of Finance* **57**, 1239–1284.
- Andersen, T. G. & Bollerslev, T. (1998a), ‘Answering the skeptics: Yes, standard volatility models do provide accurate forecasts’, *International Economic Review* **39**, 885–905.
- Andersen, T. G. & Bollerslev, T. (1998b), ‘Deutchemark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies’, *Journal of Finance* **53**, 219–265.
- Andersen, T. G., Bollerslev, T. & Diebold, F. X. (2007), ‘Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility’, *Review of Economics and Statistics* **89**, 701–720.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Ebens, H. (2001), ‘The distribution of realized stock return volatility’, *Journal of Financial Economics* **61**, 43–76.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2000a), ‘Exchange rate returns standardized by realized volatility are (nearly) Gaussian’, *Multinational Finance Journal* **4**, 159–179.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2000b), ‘Great realizations’, *Risk* **13**, 105–108.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2001), ‘The distribution of exchange rate volatility’, *Journal of the American Statistical Association* **96**, 42–55.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), ‘Modelling and forecasting realized volatility’, *Econometrica* **71**, 579–625.

- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Vega, C. (2003), ‘Micro effects of macro announcements: Real-time price discovery in foreign exchange’, *American Economic Review* **93**, 38–62.
- Andersen, T. G., Bollerslev, T. & Dobrev, D. (2007), ‘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.
- Andersen, T. G., Bollerslev, T., Frederiksen, P. H. & Nielsen, M. Ø. (2006), ‘Comment on Peter R. Hansen and Asger Lunde: Realized variance and market microstructure noise’, *Journal of Business and Economic Statistics* **24**, 173–178.
- Andersen, T. G., Bollerslev, T. & Huang, X. (2006), ‘A semiparametric framework for modelling and forecasting jumps and volatility in speculative prices’, *Working paper, Duke and Northwestern Universities* .
- Ane, T. & Geman, H. (2000), ‘Order flow, transaction clock, and normality of asset returns’, *Journal of Finance* **55**, 2259–2284.
- Baillie, R. T. & Bollerslev, T. (1989), ‘The message in daily exchange rates: A conditional variance tale’, *Journal of Business and Economic Statistics* **7**, 297–305.
- Bandi, F. M. & Russell, J. R. (2007), ‘Microstructure noise, realized variance, and optimal sampling’, *Review of Economic Studies* **forthcoming**.
- Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A. & Shephard, N. (2006), ‘Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise’, *Working paper, Oxford and Stanford Universities* .
- Barndorff-Nielsen, O. E. & 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, O. E. & Shephard, N. (2004), ‘Power and bipower variation with stochastic volatility and jumps’, *Journal of Financial Econometrics* **2**, 1–48.
- Barndorff-Nielsen, O. E. & Shephard, N. (2006), ‘Econometrics of testing for jumps in financial economics using bipower variation’, *Journal of Financial Econometrics* **4**, 1–30.
- Barndorff-Nielsen, O. E., Shephard, N. & Winkel, M. (2006), ‘Limit theorems for multipower variation in the presence of jumps’, *Stochastic Processes and their Applications* **116**, 796–806.

- Bates, D. (1996), ‘Jumps and stochastic volatility: Exchange rate processes implicit in Deutschemark options’, *Review of Financial Studies* **9**, 69–107.
- Bates, D. S. (2000), ‘Post-’87 crash fears in the S&P 500 futures option market’, *Journal of Econometrics* **94**, 181–238.
- 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., 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* **forthcoming**.
- Bollerslev, T., Litvinova, J. & Tauchen, G. (2006), ‘Leverage and volatility feedback effects in high-frequency data’, *Journal of Financial Econometrics* **4**, 353–384.
- 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, A. R., 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, A. A. (1982), ‘The stochastic behavior of common stock variances - value, leverage and interest rate effects’, *Journal of Financial Economics* **3**, 145–166.
- Clark, P. E. (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 and Hall, London.

- Dacorogna, M. M., Gencay, R., Müller, U. A., Pictet, O. V. & Olsen, R. B. (2001), *An Introduction to High-Frequency Finance*, Academic Press, San Diego.
- Dambis, K. E. (1965), ‘On the decomposition of continuous submartingales’, *Theory of Probability Applications* **10**, 401–410.
- de Pooter, M., Martens, M. & van Dijk, D. (2007), ‘Predicting the daily covariance matrix for S&P100 stocks using intraday data - but which frequency to use?’, *Econometric Reviews* **forthcoming**.
- Dubins, L. & Schwartz, G. (1965), ‘On continuous martingales’, *Proceedings of the National Academy of Sciences* **53**, 913–916.
- Engle, R. F. (2004), ‘Nobel lecture. Risk and volatility: Econometric models and financial practice’, *American Economic Review* **94**, 405–420.
- Epps, T. W. & Epps, M. L. (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., Kirby, C. & Ostdiek, B. (2003), ‘The economic value of volatility timing using realized volatility’, *Journal of Financial Economics* **67**, 473–509.
- Fleming, J. & Paye, B. S. (2006), ‘High-frequency returns, jumps and the mixture of normals hypothesis’, *Working paper, Jones Graduate School of Management, Rice University*.
- Gallant, A. R., Rossi, P. E. & Tauchen, G. (1992), ‘Stock-prices and volume’, *Review of Financial Studies* **5**, 199–242.
- Gillemot, L., Farmer, J. D. & Lillo, F. (2005), ‘There’s 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, P. R. & 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, J. C. & White, A. (1987), ‘The pricing of options on assets with stochastic volatilities’, *Journal of Finance* **42**, 281–300.
- Jiang, G. J. & Oomen, R. C. A. (2005), ‘A new test for jumps in asset prices’, *Working Paper, Warwick Business School* .
- Johannes, M. (2004), ‘The statistical and economic role of jumps in interest rates’, *Journal of Finance* **59**, 227–260.
- Kim, D. & Kon, S. J. (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, S. S. & Mykland, P. A. (2008), ‘Jumps in real-time financial markets: a new nonparametric test and jump dynamics’, *Review of Financial Studies* **forthcoming**.
- 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, J. C. & Martens, M. (2003), ‘Testing the mixture of distributions hypothesis using ”realized” volatility’, *Journal of Futures Market* **23**, 661–679.
- Maheu, J. M. & McCurdy, T. H. (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, R. C. (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, D. B. (1991), ‘Conditional heteroskedasticity in asset returns: A new approach’, *Econometrica* **59**, 347–370.
- Nielsen, M. Ø. & Frederiksen, P. H. (2008), ‘Finite sample accuracy and choice of sampling frequency in integrated volatility estimation’, *Journal of Empirical Finance* **forthcoming**.
- Oomen, R. (2006), ‘Properties of realized variance under alternative sampling schemes’, *Journal of Business and Economic Statistics* **24**, 219–237.
- Peters, R. T. & de Vilder, R. G. (2006), ‘Testing the continuous semimartingale hypothesis for the S&P 500’, *Journal of Business & 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. (2007), ‘Econometric analysis of jump-driven stochastic volatility models’, *Working paper, Duke University*.
- Zhou, B. (1998), Parametric and nonparametric volatility measurement, *in* C. Dunis & B. Zhou, eds, ‘Nonlinear Modelling of High-Frequency Financial Time Series’, John Wiley and Sons Ltd., London.