Journal of Econometrics 150 (2009) 151–166

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

# A discrete-time model for daily S & P500 returns and realized variations: Jumps and leverage effects

Tim Bollerslev [a], Uta Kretschmer [b], Christian Pigorsch [c], George Tauchen [a,*]

[a] Department of Economics, Duke University, Box 90097, Durham, NC 27708, United States

[b] Department of Economics, University of Mannheim, L7, 3-5, D-68131 Mannheim, Germany

[c] Department of Economics, University of Bonn, Adenauerallee 24-42, D-53113 Bonn, Germany

ARTICLE INFO

*Article history:*
Available online 25 December 2008

*JEL classification:*
C1
C3
C5
G1

*Keywords:*
Realized volatility
Bipower variation
Jumps
Leverage effect
Simultaneous equation model

ABSTRACT

We develop an empirically highly accurate discrete-time daily stochastic volatility model that explicitly distinguishes between the jump and continuous-time components of price movements using nonparametric realized variation and Bipower variation measures constructed from high-frequency intraday data. The model setup allows us to directly assess the structural inter-dependencies among the shocks to returns and the two different volatility components. The model estimates suggest that the leverage effect, or asymmetry between returns and volatility, works primarily through the continuous volatility component. The excellent fit of the model makes it an ideal candidate for an easy-to-implement auxiliary model in the context of indirect estimation of empirically more realistic continuous-time jump diffusion and Lévy-driven stochastic volatility models, effectively incorporating the interdaily dependencies inherent in the high-frequency intraday data.



## 1. Introduction

Modeling of financial market volatility has been one of the most active areas of research in empirical finance and time series econometrics over the past two decades. A current theme in this literature concerns the question of whether financial prices, and equity prices in particular, may be adequately described by continuous sample path processes, or whether the price movements exhibit discontinuities, or jumps.[1]

One strand of the literature has sought to answer the question through the estimation of specific parametric continuous time models. This literature dates back to the early work of Merton (1976), with more recent contributions allowing for both jumps and time-varying stochastic volatility including Andersen et al. (2002), Bates (2000), Chernov et al. (2003), Eraker (2004), Eraker et al. (2003), and Pan (2002), among others. Still, the estimation of parametric jump diffusion models remains difficult, and the existing empirical results based on daily or coarser frequency data typically do not allow for a very clear distinction between pure diffusion multi-factor stochastic volatility models and lower-order models with jumps. Of course, given the often large within-day price movements, the daily data most often used in the estimation of the models may simply not be informative enough to provide a firm answer. At the same time, the direct estimation of specific parametric volatility models with large samples of high-frequency intraday data remains extremely challenging from a computational perspective and moreover requires that all of the market microstructure complications inherent in the high-frequency data be properly incorporated into the model.

This in turn has motivated a second more recent strand of the literature in which the intraday data is used in the construction of lower-frequency nonparametric daily volatility measurements. This literature, beginning with the work of Andersen and Bollerslev (1998), Andersen et al. (2001b) and Barndorff-Nielsen and Shephard (2002b) builds on the general result that under ideal conditions the sum of successively finer sampled high-frequency squared returns converges to the quadratic variation of the price process.[2] The quadratic variation, of course, includes both the continuous sample path variation and the jumps. However, combining

* Corresponding author. Tel.: +1 919 660 1812; fax: +1 919 684 8974.
*E-mail address:* george.tauchen@duke.edu (G. Tauchen).

[1] In addition to the implications for the direct modeling of the price process pursued in the present paper, the answer to the question has important implications for risk management and asset pricing more generally.

[2] Earlier important related contributions include the work by Dacorogna et al. (1993) and Müller et al. (1990).

0304-4076/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2008.12.001

the realized variation with the realized Bipower variation measure first introduced by Barndorff-Nielsen and Shephard (2004, 2005), allows for a direct nonparametric decomposition of the total price variation into its two separate components. Utilizing these ideas, Andersen et al. (2007) and Huang and Tauchen (2005) both report empirical evidence in support of non-trivial contributions to the overall daily price variation coming from the jump component.

The nonparametric volatility measures have also inspired the development of a series of new and simple-to-implement reduced form volatility forecasting models in which the realized volatilities are modeled by standard discrete-time time series procedures, examples of which include Andersen et al. (2003, 2007), Corsi (2004), Corsi et al. (2008), Deo et al. (2006), Koopman et al. (2005) and Martens et al. (2004), among others. By effectively incorporating the high-frequency data into the volatility measurements, these simple discrete-time models generally out-perform existing more complicated parametric volatility models based on the corresponding return observations only. The simplicity of these methods, however, comes at the cost of disregarding information about the different volatility components. With the exception of Andersen et al. (2007), who simply included lagged measures for the jump component into a univariate linear forecasting model for the total realized variation, none of the above listed studies have made use of the decomposition of the total variation into its separate continuous and jump components. Meanwhile, the apparent relevance of jumps along with the distinctly different distributional features of the continuous and jump components, supports the idea of a more structured approach to realized volatility modeling.

This is the main theme of the present paper, in which we develop an empirically highly accurate multivariate discrete-time volatility model for the returns and the realized continuous sample path and jump variation measures. Our joint modeling of the returns and the two volatility components in turn allows us to directly assess the importance of the often documented asymmetric relationship between returns and volatility, and whether the observed so-called "leverage effect" is caused by a negative correlation of the lagged returns with the current continuous volatility component and/or current jumps.[3] We initially estimate the model equation-by-equation under the implicit assumption that the disturbances are independent across the three equations. However, our univariate estimation results reveal important nonlinear contemporaneous dependencies among the residuals, and we go on to account for these in a general recursive simultaneous equation system, explicitly allowing for contemporaneous nonlinear interdependencies. Despite the general and very flexible structure of the model, full information maximum likelihood estimation remains relatively straightforward. The recursive structure also makes simulations from the model easy to implement, which we use in checking different aspects of our final preferred specification. Our model estimates are based on daily realized volatilities and returns constructed from high-frequency five-minute S&P500 index futures over the 1985 to 2004 sample period. As part of our analysis, we also highlight some of the key distributional features of the corresponding daily Bipower variation and relative jump measures that any satisfactory fully specified continuous or discrete-time model will have to account for.

The remainder of the paper is organized as follows. Section 2 provides a short review of the relevant theory and construction of the pertinent volatility measures. Section 3 discusses the data and summary statistics for the different measures motivating the specification of our empirical model. The formulation of the three basic model equations for the returns, Bipower variation and relative jump series is detailed in Section 4. The resulting equation-by-equation estimates are presented in Section 5, along with an assessment of the cross-equation dependencies in the disturbances. Section 6 describes the joint recursive model and corresponding maximum likelihood estimates. Simulations from the model are used in Section 7 to further investigate the adequacy of the fit. Section 8 concludes with a brief summary and some suggestions for future research.

[3] The recent empirical analysis in Bollerslev et al. (2006b) also points toward the existence of a contemporaneous leverage effect in the form of cross-correlated high-frequency returns and absolute returns.

## 2. Realized volatility, bipower variation and jumps

We begin by a brief review of the relevant theory underlying the different variation measures employed in our empirical model. A more thorough theoretical treatment can be found in Andersen et al. (2001b), Barndorff-Nielsen and Shephard (2002a) and Protter (2004).

### 2.1. Quadratic variation

Our analysis builds on the theory of quadratic variation. Let $p_t$ denote the logarithmic price of a financial asset. Assume that $p_t$ follows the continuous-time semimartingale jump diffusion process:

$$p_t = \int_0^t \mu(s)\mathrm{d}s + \int_0^t \sigma(s)\mathrm{d}W(s) + \sum_{j=1}^{N(t)} \kappa(s_j), \tag{1}$$

where the mean process $\mu(t)$ is continuous and of finite variation, $\sigma(t) > 0$ denotes the càdlàg instantaneous volatility, $W(t)$ is a standard Brownian Motion, and the $N(t)$ process counts the number of jumps occurring with possibly time-varying intensity $\lambda(t)$ and jump size $\kappa(s_j)$. The theory of quadratic variation then permits the derivation of nonparametric volatility measures that allow us to decompose the total price variation into its continuous and jump part. In particular, the quadratic variation process of (1),

$$[p]_t = \operatorname{plim} \sum_{j=0}^{n-1} (p_{\tau_{j+1}} - p_{\tau_j})^2, \tag{2}$$

where $\tau_0 = 0 \leq \tau_1 \leq \cdots \leq \tau_n = t$ denotes a sequence of partitions with $\sup_j\{\tau_{j+1} - \tau_j\} \to 0$ for $n \to \infty$, may be expressed as,

$$[p]_t = \int_0^t \sigma^2(s)\mathrm{d}s + \sum_{j=1}^{N(t)} \kappa^2(s_j), \tag{3}$$

that is, the *integrated variance* and the sum of the *squared jumps*. Of course, in the popular pure diffusion case where the $N(t)$ counting process is identically equal to zero, the second term disappears and the quadratic variation is simply equal to the integrated variance.

### 2.2. Realized variation

Most of our analysis will be focused on daily returns and volatilities. Hence, for notational simplicity we normalize the daily time interval to unity, denoting the corresponding daily returns by:

$$r_t = p_t - p_{t-1}, \quad t = 1, \ldots . \tag{4}$$

To formally define our empirical volatility measures, denote the day $t$, $j$th within-day return by:

$$r_{t,j} = p_{t-1+\frac{j}{M}} - p_{t-1+\frac{(j-1)}{M}}, \quad j = 1, \ldots, M, \tag{5}$$

where $M$ refers to the number of returns per day. The sum of the corresponding squared intradaily returns:

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2, \tag{6}$$

then affords a natural estimator of the *realized quadratic variation*. Following the recent literature we will interchangeably refer to this quantity as the *realized variance* or the *realized volatility*. The idea of measuring the ex-post variation of asset prices by summing over more frequently sampled squared returns dates back at least to Merton (1980), and was also applied by French et al. (1987), Hsieh (1991) and Poterba and Summers (1986), and more recently by Taylor and Xu (1997), inter alia. Meanwhile, the notion of realized variation was first formally related to the theory of quadratic variation within the context of finance and time-varying volatility modeling by Andersen and Bollerslev (1998), Andersen et al. (2001b), Barndorff-Nielsen and Shephard (2002b) and Comte and Renault (1998).

In particular, it follows from the theory discussed above that the realized variance will generally converge uniformly in probability to the quadratic variation as the sampling frequency, $M$, of the underlying returns approaches infinity:

$$RV_t \to \int_{t-1}^{t} \sigma^2(s)\mathrm{d}s + \sum_{j=N(t-1)+1}^{N(t)} \kappa^2(s_j). \tag{7}$$

In other words, the realized variance affords an ex-post measure of the true total price variation, including the discontinuous jump part.

In order to distinguish the continuous variation from the jump component, Barndorff-Nielsen and Shephard (2004) first proposed the so-called *Bipower variation* measure, defined by:

$$BV_t = \frac{\pi}{2} \sum_{j=2}^{M} |r_{t,j}||r_{t,j-1}|. \tag{8}$$

Importantly, for increasingly finely sampled returns the Bipower variation measure becomes immune to jumps and consistently (for increasing values of $M$) estimates the integrated variance:

$$BV_t \to \int_{t-1}^{t} \sigma^2(s)\mathrm{d}s. \tag{9}$$

Consequently, the difference between the realized variance and the Bipower variation affords a simple nonparametric estimator of the contribution to total price variation coming from the jump component.

Meanwhile, the extensive simulation evidence in Huang and Tauchen (2005) suggests that an empirically more robust measure is provided by the relative jump statistic, $RJ_t = (RV_t - BV_t)RV_t^{-1}$, or the (approximate) logarithmic version[4]:

$$J_t = \log RV_t - \log BV_t. \tag{10}$$

Hence, in the empirical results reported on here, we will rely on a joint model for $BV_t$ and $J_t$ as a way of capturing the distinct components accounting for the total daily price variation. The $J_t$ measure is in theory restricted to be non-negative. However, in practice for finite values of $M$, $BV_t$ may exceed $RV_t$ so that $J_t$ becomes negative. In the approach adopted here, we will simply treat these "measurement errors" as part of the $J_t$ process. Alternatively, building on the asymptotic (for increasing $M$) distribution theory in Barndorff-Nielsen and Shephard (2004), it would be possible to truncate the $J_t$ process, and only associate the values beyond a certain threshold with the jump component. This is the approach adopted in Andersen et al. (2007), who rely on a large critical value for identifying only the most significant jumps entering a reduced form univariate forecasting model for $RV_t$. In contrast, by jointly modeling the returns, the relative jump measure and the Bipower variation, we avoid the arbitrary choice of any pre-specified significance level affecting the selection of "significant" jumps.

## 3. Data and stylized facts

The theory discussed in the preceding section underlying the consistency of the $BV_t$ and $J_t$ measures formally hinges on the notion of increasingly finer sampled high-frequency returns. In practice, however, the sampling frequency is invariably limited by the actual quotation, or transaction frequency. Moreover, the observed high-frequency prices are further "contaminated" by a host of market microstructure frictions, including price discreteness and bid-ask spreads. These effects combine to render the basic assumption of a semimartingale price process invalid at the tick-by-tick level. In response to this, a number of authors, including Andersen et al. (2001a,b, 2007), have advocated the use of coarser sampling frequencies as a simple way to alleviate these contaminating effects, while maintaining most of the relevant information in the high-frequency data. This is also the approach adopted here.[5]

Specifically, while our primary data consists of tick-by-tick transaction prices for the S&P500 Index futures contracts traded on the Chicago Mercantile Exchange, ranging from January 1, 1985 to December 31, 2004, we follow Andersen et al. (2007) in computing our daily realized variance and jump measures from five-minute returns constructed using the nearest prices to each five-minute mark for the most actively traded contracts.[6] We also exclude all overnight returns.

The resulting daily series are displayed in Fig. 1. All of the series exhibit the widely-documented volatility clustering effect. Also, the variance of the logarithmic realized variance exceeds that of the logarithmic Bipower variation series. Consistent with this, the jump series depicted in the last panel exhibits many, mostly positive, small values. These small observations, including the small negative values, may be attributed to measurement, or discretization, errors due to the use of finitely many returns in the construction of the underlying measures. At the same time, the series also contains a number of more extreme observations indicative of genuine large-sized jumps on those days.

4 The empirical evidence in Huang and Tauchen (2005) for the S&P500 index also suggests that the relative contribution of jumps to the total price variation based upon the $RJ_t$ measure amounts to roughly 7%.

5 Several recent studies have proposed alternative procedures to more effectively make use of all the tick-by-tick data including the notion of an optimal sampling frequency, $M$, in the sense of minimizing the MSE of the resulting realized volatility measure as suggested by Bandi and Russell (2005) and Aït-Sahalia et al. (2005), business type sampling schemes dictated by the activity of the market, as in, e.g., Oomen (2005), the use of various pre-filtering or kernel type procedures to explicitly account for the serial correlation induced by the market microstructure noise as in, e.g., Andersen et al. (2001a), Bollen and Inder (2002), Corsi et al. (2001), Hansen and Lunde (2006), and Zhou (1996), along with sub-sampling schemes designed to adjust for the bias and inconsistency in the simple realized volatility estimator for increasing values of $M$, as first developed by Zhang et al. (2005). The recent paper by Barndorff-Nielsen et al. (2008) provides a unified theoretical framework for analyzing most of these estimators within a kernel based representation, along with a discussion of optimal kernel and bandwidth choices. Meanwhile, to the best of our knowledge none of these ideas has yet been formally extended to allow for similar measurements of the integrated variance in form of robust to market microstructure noise modified realized Bipower variation measures. Hence, in the empirical results reported on below, we simply rely on the same coarse sampling interval in the construction of both measures.

6 The volatility signature plot for the same data depicted in Corsi et al. (2008) suggests that the returns are largely immune to the contaminating influences of the market microstructure noise at that frequency. In particular, the ratios of the sample means of the five-minute based realized measures to the ones based on 15- and 30-min sampling, equal 0.9936 and 0.9746 for the realized variance, and 0.9732 and 0.9660 for the Bipower variation, respectively.

**Table 1**
Descriptive statistics.

| Series | Mean | Std. dev. | Median | Skewness | Exc. kurt. | Ljung − Box(10) |
|---|---|---|---|---|---|---|
| $\sqrt{RV_t}$ | 0.8627 | 0.5935 | 0.7586 | 15.3509 | 496.7651 | 10155.72 |
| $\log RV_t$ | −0.5139 | 0.8775 | −0.5527 | 0.5950 | 1.7981 | 22023.20 |
| $\sqrt{BV_t}$ | 0.8340 | 0.5359 | 0.7348 | 11.1561 | 288.4633 | 12223.28 |
| $\log BV_t$ | −0.5817 | 0.8845 | −0.6163 | 0.5418 | 1.4807 | 21715.55 |
| $\log\left(\frac{RV_t}{BV_t}\right)$ | 0.0678 | 0.1263 | 0.0538 | 1.7766 | 12.2675 | 51.44 |
| $r_t$ | 0.0254 | 1.0946 | 0.0511 | −2.1655 | 96.2483 | 117.29 |
| $r_t/\sqrt{RV_t}$ | 0.0866 | 1.0027 | 0.0739 | 0.0503 | −0.1497 | 14.86 |

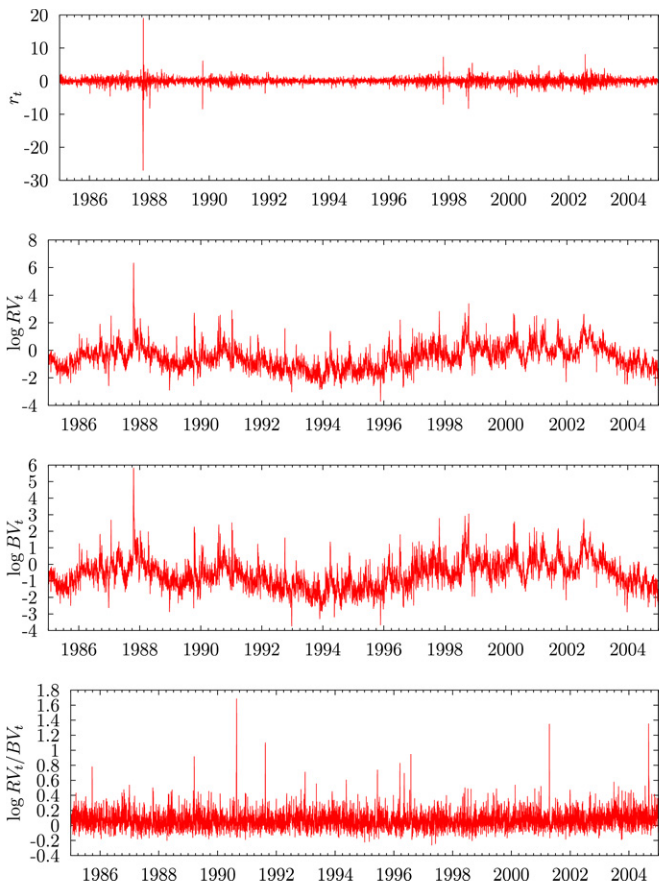


**Fig. 1.** Time Series of returns, logarithmic realized variance, logarithmic Bipower variation and jumps.

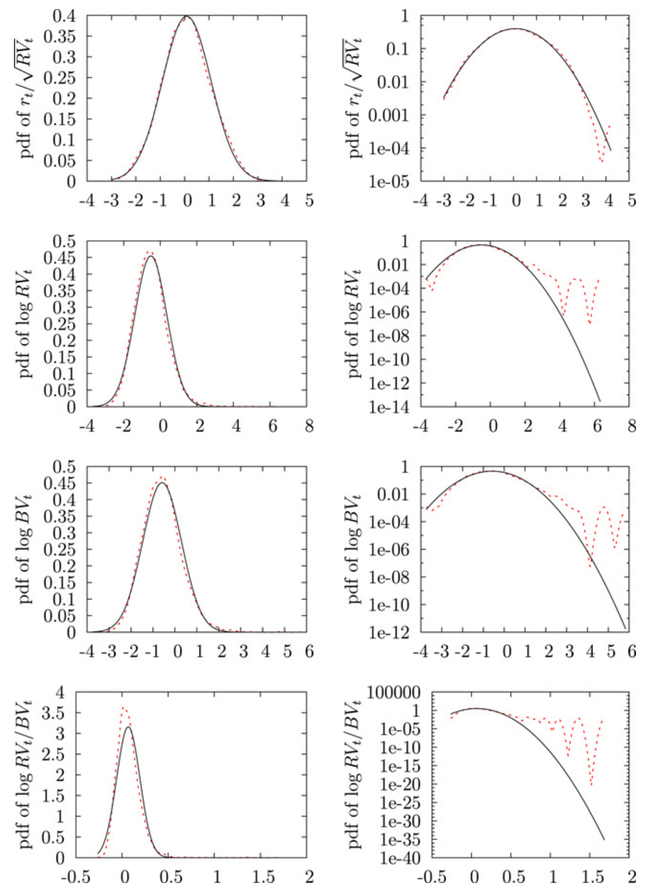


**Fig. 2.** Unconditional distributions of standardized returns, logarithmic realized variance, logarithmic Bipower variation and jumps. The left panel of the figure shows the kernel density estimates of the series (dashed line) and the normal density (solid line) for reference purposes. The right panel shows the same in log scale.

These visual impressions are confirmed by the summary statistics reported in Table 1. In particular, the mean and variance of the realized volatility both exceed the corresponding statistics for the square-root Bipower variation. It follows also from the table that the unconditional distribution of both volatility measures are highly skewed and leptokurtic. However, the logarithmic transform renders both approximately normal. This approximate log-normality is further supported by the kernel density plots presented in Fig. 2. Similar results for the realized volatility from other markets have previously been reported in Andersen et al. (2001a,b) among others. Meanwhile, the descriptive statistics and the corresponding kernel density plots for the relative jump measure, $J_t$, clearly indicate a positively skewed and leptokurtic distribution.[7] The unconditional distribution of the daily returns also show the expected excess kurtosis and negative skewness. At the same time, the distribution of the returns standardized by the realized volatility is surprisingly close to Gaussian, as previously documented by Andersen et al. (2001a).[8]

Turning to the last column in the table, all of the volatility measures exhibit highly significant own serial dependencies, as evidenced by the Ljung-Box test statistics for up to tenth order autocorrelation.[9] Furthermore, the sample autocorrelation functions in Fig. 3 for the two logarithmic volatility measures show the characteristic hyperbolic decay with autocorrelation coefficients being significant (compared to the conservative

[7] Note that the sign of the skewness is determined by the specific definition of our jump measure as the logarithmic ratio of $RV_t$ divided by $BV_t$. Barndorff-Nielsen and Shephard (2004) in contrast consider the inverse ratio resulting in a negatively skewed distribution.

[8] In the absence of jumps and independence between the innovation processes driving the returns and the volatility, the standardized returns defined by the stylized model in Eq. (1) should be normally distributed.

[9] The critical value of this test is 18.31.