SCHOOL OF ECONOMICS AND MANAGEMENT
FACULTY OF SOCIAL SCIENCES
AARHUS UNIVERSITY

CREATES
Center for Research in Econometric Analysis of Time Series

CREATES Research Paper 2009-49

Realized Volatility and Multipower Variation

Torben G. Andersen and Viktor Todorov

School of Economics and Management
Aarhus University
Bartholins Allé 10, Building 1322, DK-8000 Aarhus C
Denmark

Realized Volatility and Multipower Variation*

Torben G. Andersen[†‡] and Viktor Todorov [§]

May 2009

Abstract

This paper reviews basic notions of return variation in the context of a continuous-time arbitrage-free asset pricing model and discusses some of their applications. We first define return variation in the infeasible continuous-sampling case. Then we introduce realized measures obtained from high-frequency observations which provide consistent and asymptotically normal estimates of the underlying return variation. The paper discusses applications of these measures for reduced-form volatility modeling and forecasting as well as testing for the presence of jumps.

Keywords: realized volatility, multipower variation, jumps, quadratic variation, volatility estimation, volatility forecasting, jump testing, continuous-time stochastic volatility model.

JEL classification: C22, C51, C52, G12.

*An abbreviated version of this manuscript is forthcoming in "Encyclopedia of Quantitative Finance", John Wiley & Sons.

†Andersen's work was supported by a grant from the NSF to the NBER and CREATES funded by the Danish National Research Foundation.

‡Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: t-andersen@kellogg.northwestern.edu

§Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: v-todorov@kellogg.northwestern.edu.

Many financial markets effectively operate in continuous time with multiple transaction prices and quotes recorded each second. Although such "ultra high-frequency" data are not useful for assessing the expected mean return of the underlying asset, it is highly informative regarding the strength of another key financial characteristic, namely return volatility. In particular, it is feasible to estimate the realized return volatility over a fixed time period directly from high-frequency return with good precision without imposing a specific parametric structure on the return dynamics. Hence, as access to tick-by-tick data became more commonplace and the first empirical links between cumulative absolute return measures and the underlying volatility were established in the mid 1990's, the literature on extracting time-varying return variation measures from high-frequency data has grown dramatically. The initial developments focused on measures reflecting the actual squared return variation, loosely denoted realized volatility, over daily and weekly frequencies. However, extended measures capturing different return moments and/or providing more robust inference regarding the continuous versus jump components of the return variation process are now an integral part of this literature. The potential applications to areas such as risk management, derivatives pricing, portfolio choice, market microstructure and general asset pricing are basically unlimited. This entry briefly reviews the main developments in this literature.

A general no-arbitrage setting for a continuously evolving logarithmic asset price is given by a jump-diffusive process. Formally, on an appropriate probability space, the stochastic logarithmic price process X_t is defined via[1]

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + \sum_{0 \leq s \leq t} \Delta X_s, \tag{1}$$

where the predictable drift component, b_s, signifies the instantaneous (continuously compounded) mean return, the diffusive coefficient, σ_s, reflects the strength of the diffusive volatility, W_s denotes a standard Brownian motion, and $\Delta X_s := X_s - X_{s-}$ indicates the jump size and is non-zero only if the price jumps exactly at time s. In addition, we require that the drift, diffusive and jump processes satisfy some (mostly) regularity conditions, see, e.g., Jacod (2008)[2]. Equation (1) is a generic representation of the jump diffusive stochastic volatility model that serves as a work horse for continuous time finance.

We are interested in inference regarding the actual volatility displayed by X over a given interval of time $[0, T]$. In the (infeasible) scenario where we have access to a continuous record of X_t we can directly determine the actual

[1] In this definition of X_t we implicitly assume that the jumps are of finite variation (so that we do not need to compensate them). Most results below apply to more general jump specifications, see Barndorff-Nielsen et al. (2006), Jacod (2008), Ait-Sahalia and Jacod (2009), Jacod and Todorov (2009).

[2] Most importantly, X must be an Ito semimartingale which essentially requires the drift, diffusion and jump compensator to be absolutely continuous with respect to time.

realization of the so-called quadratic variation of X,

$$QV_{[0,T]} = \int_0^T \sigma_s^2 ds + \sum_{0 \le s \le T} |\Delta X_s|^2. \tag{2}$$

Importantly, in the continuous-record case, we may also perfectly identify the jumps in the sample path and we thus "observe" the continuous and discontinuous parts of the quadratic variation and obtain the decomposition,

$$QV_{[0,T]}^c = \int_0^T \sigma_s^2 ds \quad \text{and} \quad QV_{[0,T]}^d = \sum_{0 \le s \le T} |\Delta X_s|^2. \tag{3}$$

The quadratic variation is of particular interest as it is closely related to the (variance) risk of the asset. This is most transparent in a simplified setting. For a pure diffusion, if the (stochastic) volatility process is independent from the return innovations, and ignoring the (negligible) return variation induced by innovations to the mean drift, the returns are conditionally Gaussian with variance governed by $QV_{[0,T]}^c$. If the return interval is small, such as a day, we may safely ignore the mean drift altogether and for simplicity equate it to zero. We then have the distributional result,[3]

$$(X_T - X_0)\Big|\{\sigma_s\}_{0 \le s \le T} \;\sim\; N\left(0, \int_0^T \sigma_s^2 ds\right). \tag{4}$$

Notice that the Gaussian distributional result is conditional on the realization of the so-called integrated variance which a priori is stochastic. In other words, the returns follow a Gaussian mixture distribution with the realization of the (diffusive) quadratic variation governing the actual return variation. If this integrated variance process is persistent, i.e., displays pronounced positive serial correlation, then the return distribution is symmetric, but displays both unconditional and conditional leptokurtosis along with volatility clustering. The main features which may modify this characterization and induce both asymmetries and more extreme outliers in the conditional return distribution are correlation between the return innovations and the volatility, the so-called leverage effect, and the presence of jumps. Nonetheless, the interpretation of $QV_{[0,T]}$ as the relevant return variation measure remains intact. Moreover, the multivariate version of equation (4) applies as stated in Andersen et al. (2003).

In practice we do not observe financial series continuously but rather obtain transaction price and quote data referring to specific points in time. Hence, the challenge is to design estimators from the discrete observations of X that can estimate continuous-time quantities like the quadratic variation and its decomposition into continuous and jump components. A closely related issue is the development of tests for the presence of jumps in the (partially) observed path of X. Suppose we observe X at $0, ..., i\Delta_n, ...[T/\Delta_n]\Delta_n$, where Δ_n is a

[3]Detailed discussion about the associated conditions and the interpretation of these relationships may be found in Andersen et al. (2007a).

small positive sampling interval and $[y]$ denotes the integer part of y. Since the time interval, T, is fixed, we cannot exploit standard "large (sample length) T " asymptotics for developing estimation and testing procedures but instead appeal to continuous record asymptotics, i.e., $\Delta_n \to 0$. For example, we may think of $[0, T]$ as one trading day and assume we observe the price process every 10 minutes, 5 minutes, 1 minute, etc. The exposition is split in three parts. We first define realized volatility as a feasible estimator of quadratic variation and discuss its relation to the underlying return distribution. Next, the notion of multipower variation is introduced and we outline its use for robust inference regarding critical aspects of the quadratic return variation. In the following section we discuss the application of realized multipower variation in formal testing for the presence of jumps. Finally, we briefly touch on procedures designed to mitigate the impact of market microstructure noise, which potentially allow for more efficient inference, as well as some multivariate generalizations of the results discussed in this survey.

1 Realized Volatility

Realized Volatility, also interchangeably referred to as Realized Quadratic Variation, is defined as

$$RV_T = \sum_{i=1}^{[T/\Delta_n]} |\Delta_i^n X|^2, \tag{5}$$

where $\Delta_i^n X := X_{i\Delta_n} - X_{(i-1)\Delta_n}$[4]. Hence, RV_T is simply a cumulative sum of squared high-frequency observations. Its usefulness stems from the fact that it consistently estimates the (unobservable) quadratic variation as $\Delta_n \to 0$, i.e.

$$RV_T \xrightarrow{\mathbb{P}} QV_T. \tag{6}$$

A few comments are in order. First, the estimator is model-free and applies very broadly independently of parametric model assumptions. Moreover, it remains valid in the multivariate case as well. Second, the measure speaks to the return variation over an interval and not to the underlying spot volatility at any given point. In fact, inference regarding spot volatility is very challenging without restrictive assumptions due to the lack of a continuous record of price observations and the impact of microstructure frictions at ultra-high frequencies. Third, RV_T resembles a rolling sample variance estimator. This statistic, labeled "historical volatility" has precedents in the literature[5]. Fourth, Merton (1980) notes that, in theory, high-frequency data can provide perfect inference whenever volatility is locally constant. Direct linkage of an empirical realized volatility measure based on intraday returns to an underlying (stochastic) quadratic return variation appears first in Andersen and Bollerslev (1998a). Fifth, one trading day

[4]Some authors label RV_T the Realized Variance, and refer to its square root as the realized volatility. However, in this entry we follow the early literature and term RV_T realized volatility.

[5]It is, e.g., employed to construct annual and monthly return variation measures from monthly and daily squared returns by Officer (1973) and French et al. (1987), respectively.

is the natural measurement period due to the pronounced intraday volatility patterns and significant news announcement effects. These features render interpretation of realized return variation over intraday periods much more complex, see, e.g., Andersen and Bollerslev (1998b)[6]. Finally, it is important to recognize that RV_T provides an ex-post estimate of the return variation which is conceptually distinct from the ex-ante conditional return variance.

In order to assess the accuracy with which RV_T measures QV_T, a Central Limit Theorem (CLT) - characterizing its behavior as the sampling frequency increases - is useful. In the simplest case without jumps[7] it follows from Barndorff-Nielsen and Shephard (2002, 2004b) and Barndorff-Nielsen et al. (2005) that

$$\frac{1}{\sqrt{\Delta_n}}\left(RV_T - QV_T\right) \xrightarrow{\mathcal{L}-s} \epsilon \times \sqrt{2\int_0^T \sigma_s^4 ds}, \tag{7}$$

where $\xrightarrow{\mathcal{L}-s}$ means stable convergence in law[8] (see, e.g., Jacod and Shiryaev (2003) for exposition and references) and ϵ is a standard normal variable, defined on an extension of the original probability space[9]. Moreover, for two non-overlapping intervals, the measurement errors for the quadratic variation are serially uncorrelated. Evidently, the precision of RV_T can be gauged formally only if we can obtain a measure of the so-called "Integrated Quarticity," $IQ_T = \int_0^t \sigma_s^4 ds$, from a given set of discrete price observations. This is, in fact, feasible through the general realized multipower variation statistics introduced below. Another reason to consider the general multipower variation statistics is the desire to decompose the quadratic variation into continuous and jump parts. The realized variation provides a consistent estimate of the combined quadratic variation. The statistics introduced in the next section enable us to isolate the effect of jumps, thus allowing for direct identification of the continuous part, QV_T^c.

2 Multipower Variation

The initial example of a multipower variation process is the so-called bipower variation proposed by Barndorff-Nielsen and Shephard (2004c, 2006a) and de-

[6]Other early uses of high-frequency data for volatility estimation include work by Olsen & Associates summarized in Dacorogna et al. (2001) as well as Hsieh (1991) and Zhou (1996). Moreover, Comte and Renault (1998) mention the possibility of estimating spot volatility from increments to quadratic variation approximated over short horizons.

[7]More generally, the "observed" path may contain jumps and the associated CLT is derived in Jacod (2008). Veraart (2007) discusses different feasible estimators of the asymptotic variance in the special case when the jumps in price and volatility do not arrive together.

[8]Stable convergence in law strengthens the usual convergence in law by guaranteeing that the convergence in (7) is joint with any random variable defined on the original probability space. This property is critical for feasible asymptotic inference as well as application of the delta method.

[9]Bootstrapping is an alternative approach for inference and may well improve precision obtained from a fairly large, but finite, number of intraday observations, see Goncalves and Meddahi (2009).

fined as

$$BV_T = \frac{\pi}{2} \sum_{i=2}^{[T/\Delta_n]} |\Delta_{i-1}^n X||\Delta_i^n X|. \tag{8}$$

Barndorff-Nielsen and Shephard (2004c, 2006a), Barndorff-Nielsen et al. (2006) prove that, even in the presence of a jump process[10],

$$BV_T \overset{\mathbb{P}}{\longrightarrow} QV_T^c. \tag{9}$$

Hence, it provides a consistent estimator for the integrated variance and we may thus further consistently estimate the jump contribution to quadratic variation via the difference between RV_T and BV_T.

More generally we may define Realized Multipower Variation by raising m successive absolute price increments to arbitrary powers, indexed by the vector $\mathbf{p} = (p_1, ..., p_m)$, where $p_j \geq 0$, $j = 1, ..., m$.

$$V(X; \mathbf{p}; \Delta_n)_T = (\mu_{p_1} \dots \mu_{p_m})^{-1} \sum_{i=m}^{[T/\Delta_n]} |\Delta_{i-m+1}^n X|^{p_m} \dots |\Delta_i^n X|^{p_1}, \tag{10}$$

for μ_a denoting the a-th absolute moment of a standard normal. By appropriate scaling of the realized multipower variation, *in the absence of jumps*, we obtain corresponding convergence result

$$\Delta_n^{1-p/2}\, V(X; \mathbf{p}; \Delta_n)_T \overset{\mathbb{P}}{\longrightarrow} \int_0^T \sigma_s^p ds, \tag{11}$$

where $p = p_1 + ... + p_m$. In analogy to the bipower variation case, this result is robust to the presence of jumps provided $\max\{p_i\}_{i=1,\dots,m} < 2$, see Barndorff-Nielsen et al. (2006). Further generalizations and asymptotic properties are explored in Barndorff-Nielsen et al. (2005) for X continuous[11].

A prominent special case of multipower variation arises when $\mathbf{p}$ is a scalar. Then $V(X; p; \Delta_n)_T$ is labeled Realized Power Variation of order p. Obviously, $V(X; 2; \Delta_n)_T$ coincides with the realized volatility estimator. In addition, a natural estimator of integrated quarticity under the null hypothesis of no jumps may be based on the realized fourth power variation, $V(X; 4; \Delta_n)_T$. While this statistic is inconsistent under the jump alternative, many realized multipower variation statistics are consistent in the presence of jumps, including the so-called realized tripower variation with $\mathbf{p} = (4/3, 4/3, 4/3)$ and the realized quadpower variation with $\mathbf{p} = (1, 1, 1, 1)$.

An alternative way of estimating the continuous part of the quadratic variation is to use the truncated squared power variation proposed by Mancini (2001,

[10]Unfortunately, the CLT for the Bipower Variation derived under the assumption that X *does not have jumps* on the observed path does not hold in presence of jumps.

[11]Some of these properties are extended to the case when X contains jumps in Barndorff-Nielsen et al. (2006)

2009) and also analyzed by Jacod (2008). It is formally defined as

$$\widetilde{V}(X,2,\Delta_n)_T = \sum_{i=1}^{[T/\Delta_n]} |\Delta_i^n X|^2 1_{\{|\Delta_i^n X| \leq \alpha(\Delta_n)^{\xi}\}}, \tag{12}$$

for arbitrary $\alpha > 0$ and $\xi \in (0, 1/2)$. The intuition is simple - we discard increments higher than a given threshold in the summation of the squared returns, thus effectively discarding the impact of jumps.

Given their consistency for QV_T and QV_T^c and the absence of serial correlation in the associated measurement errors, RV_T and BV_T serve as a natural basis for measuring and forecasting return volatility. The literature is too extensive for a thorough review so we only provide an account of a few established findings. First, Andersen et al. (2003) demonstrate that reduced form time series models for realized volatility within the ARFIMA class generate forecasts which dominate traditional stochastic volatility and GARCH models estimated from daily return data. The long memory dependence incorporated in the ARFIMA setting is invariably highly significant and an important factor in forecast performance[12]. Another key feature in the superior forecast performance is the ability of RV_T to adapt quickly to shifts in the underlying level of volatility. This allows forecasts to be conditioned on a more accurate assessment of the current volatility state compared to models based on daily returns. Second, Andersen et al. (2007b) find that forecast performance may be boosted further through a decomposition of RV_T into the continuous and jump part. The diffusive volatility is the main source of long range dependence so separating out the jumps allows for a more refined measure of the relevant volatility state and improves forecast precision correspondingly[13]. Third, Andersen et al. (2006) document that this type of reduced form time series forecasts, at a minimum, perform on par with market based forecasts such as implied volatility forecasts extracted from option prices. Andersen et al. (2004) provides an analytic framework for gauging the performance of reduced form realized volatility forecasts across a broad class of popular diffusive stochastic volatility specifications, finding that the loss of predictive ability is minor relative to using (infeasible) forecasts based on the true underlying model. The usefulness of reduced form realized volatility and (in the multivariate extension) covariation forecast models from practical perspectives is documented, e.g., in Fleming et al. (2003) and Bandi et al. (2008) for portfolio allocation, in Andersen et al. (2006) for dynamic estimation of systematic (beta) market risk, in Andersen and Benzoni (2009) for specification testing of term structure models and in Bandi et al. (2008) for option pricing and associated trading strategies.

Realized multipower variation also facilitates estimating general parametric continuous-time models which is challenging due to the latency of the stochastic

[12] Typical estimates for the parameter indicating the degree of fractional integration fall in the range of 0.30 to 0.47, suggesting a strictly stationary but highly persistent volatility dynamics.

[13] An alternative approach exploiting a mix of different realized power variation statistics for forecasting is the so-called MIDAS regressions proposed by Ghysels et al. (2006).

volatility and the presence of price jumps. Bollerslev and Zhou (2002) estimate affine jump-diffusion models by treating the realized volatility as the unobservable Quadratic Variation in a method-of-moments procedure[14]. They show via simulations that the impact of the measurement error on the parameter estimates is small[15]. Corradi and Distaso (2006) provide formal justification for the estimator of Bollerslev and Zhou (2002). They consider joint asymptotics, both long time span and continuous record, and show that (in the general case) when the intra-period sampling frequency increases faster than sample length, the measurement error of realized volatility (and in some cases bipower variation) is asymptotically negligible. Todorov (2009) further generalizes these results. He proves the uniform integrability of realized multipower variation statistics under general specifications for the price jump process and demonstrates that the realized multipower variation can be used effectively in making semi-parametric inference about general classes of continuous-time models.

3 Jump Testing

Another important application of realized multipower variation is to determine whether, *on a given path*, the process X contains jumps or not. This may again be accomplished in entirely model-independent fashion, i.e., without imposing additional structure on X besides specification (1). Barndorff-Nielsen and Shephard (2004c, 2006a) propose testing for jumps by determining whether $QV^d_{[0,T]}$ is statistically different from zero. For this purpose they derive the joint asymptotic distribution of RV and BV conditional on the path of X not containing jumps.

Huang and Tauchen (2005) discusses various transformations of these tests which often lead to significant improvements in finite sample test performance. They also document that these procedures generally provide reliable tests. Alternative tests for the *null hypothesis of no jumps* can be derived by replacing the bipower variation with other realized multipower variation estimates of QV^c_T or with the truncated squared power variation[16].

Recently Ait-Sahalia and Jacod (2009) propose an attractive alternative to these tests. They construct a test statistic as a ratio of power variation of order four computed over two different frequencies,

$$\Phi_n^{(j)} = \frac{V(X;4;k\Delta_n)_T}{V(X;4;\Delta_n)_T}, \tag{13}$$

[14]A similar approach is adopted in Garcia et al. (2006) in the estimation of objective and risk-neutral distributions.

[15]Barndorff-Nielsen and Shephard (2002, 2006b) consider also estimation of continuous-time models by using realized volatility and Todorov (2006) adds also the realized fourth power variation in the inference. These papers compute moments in closed-form. Similarly, Meddahi (2002) derives moments in closed form for the general eigenfunction stochastic volatility models and uses them to study the difference between quadratic variation and realized volatility (in the case of no jumps). Moments of realized multipower variation statistics (e.g., bipower variation) are not known in general, unlike the moments of their continuous-time counterparts.

[16]See also Jiang and Oomen (2008) for an alternative but related jump test.

where $k \geq 2$ is an integer (typically 2 or 3). This test statistic behaves very differently depending on whether X contains jumps on the (partially) observed path or not, since,

$$\Phi_n^{(j)} \xrightarrow{\mathbb{P}} \begin{cases} 1 & \text{if } X \text{ contains jumps} \\ k & \text{if } X \text{ does not contain jumps,} \end{cases} \tag{14}$$

Ait-Sahalia and Jacod (2009) derive the CLT for $\Phi_n^{(j)}$ both when the given path of X contains jumps and when it does not. This allows for tests of both the null of no jumps (like above) and the null of jumps.

Finally, Lee and Mykland (2008) pursue a different strategy as they test for jumps at each single high-frequency observation. The motivation is that, under the null hypothesis of no jumps, the high-frequency increments, standardized by the estimated local volatility, are asymptotically normally distributed. It has the convenient feature that the exact timing of significant jumps within the trading period is identified[17].

4 Some Important Extensions

An important feature ignored in our exposition is the impact of so-called market microstructure noise. This refers to the patterns induced in high-frequency returns due to features such as a discrete price grid, the absence of continuous trading, and the presence of a bid-ask spread. Hence, the observed price may be seen as a noisy indicator of an underlying ideal price contaminated by a (small) noise process. This noise component induces a bias in the empirical RV measures which tends to grow with increasing sampling frequency. As such, when realized volatility is computed from ultra high-frequency returns, it is critical to adjust for the impact of microstructure noise. Andersen et al. (2000) suggest gauging the highest sampling frequency at which the systematic bias induced by the noise is negligible, via a so-called volatility signature plot, and use that frequency for realized volatility computations. However, it is potentially more efficient to sample more frequently and adjust for the noise component. This approach is pursued in a number of papers in recent years, see, e.g., Zhang et al. (2005), Bandi and Russell (2006), Hansen and Lunde (2006), Barndorff-Nielsen et al. (2008) and Zhang (2006). While the asymptotic theory has developed for the case of a pure diffusive price process, the properties of the various procedures in the presence of jumps remain largely unknown. Recently, Podolskij and Vetter (2009) propose to modify the bipower variation in a way which is robust both to the microstructure noise and price jumps (of finite activity).

[17] Andersen et al. (2007) explore a similar strategy determining the exact location of, potentially multiple, jumps from uniform test statistics across all high-frequency returns over each trading day. Another alternative is to invoke a sequential test where, conditional on the jump statistic being significant at a given level, the most dominant jump is identified and removed and the jump test procedure then repeated until the jump statistic no longer is significant at the given test level. This strategy is employed by Andersen et al. (2007).

In addition, our exposition has focused on the univariate case. If X is continuous, the multivariate extension is conceptually straightforward and has been done in Barndorff-Nielsen and Shephard (2004b), Barndorff-Nielsen et al. (2005) (although the practical application can be challenging). If X contains jumps matters are much more complicated, as realized multipower variation can behave very differently depending on whether the jumps in individual series arrive jointly or not. This is exploited in Jacod and Todorov (2009) who propose tests for both whether the jumps in the individual components arrive together or not without any restriction on the possible jump dependence (Barndorff-Nielsen and Shephard (2004a) consider measuring whether the covariation induced by jumps is zero as a way to test for common jumps). Jacod and Todorov (2009) document a non-trivial number of days with common arrival of jumps in exchange rates. Todorov and Bollerslev (2007) consider linear dependence between individual stock and market jumps and show that the associated discontinuous market betas can behave very differently from the continuous ones.

In summary, realized volatility and the related realized multipower variation literature is vibrant. Given the large efficiency gains obtained from volatility measurements exploiting high-frequency data relative to daily data and the increasing availability of tick-by-tick data, both the theoretical and empirical research within this area will surely continue to grow in coming years.

References

Ait-Sahalia, Y. and J. Jacod (2009). Testing for Jumps in a Discretely Observed Process. *Annals of Statistics 37*, 184–222.

Andersen, T. and L. Benzoni (2009). Do Bonds Span Volatility Risk in the U.S. Treasury Market? A Specification Test for Affine Term Structure Models. *Journal of Finance*, forthcoming.

Andersen, T. and T. Bollerslev (1998a). Answering the Critics: Yes, Standard Volatility Models Do Provide Accurate Forecasts. *International Economic Review 39*, 885–905.

Andersen, T. and T. Bollerslev (1998b). Deutsche Mark- Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies. *Journal of Finance 53*, 219–265.

Andersen, T., T. Bollerslev, and F. Diebold (2007a). Parametric and Nonparametric Volatility Measurement. In Y. Ait-Sahalia and L. Hansen (Eds.), *Handbook of Financial Econometrics*. North Holland.

Andersen, T., T. Bollerslev, and F. Diebold (2007b). Roughing It Up: Including Jump Components in the Measurement, Modeling, and Forecasting of Return Volatility. *Review of Economics and Statistcis 89*.

Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2000). Great Realizations. *Risk 13*, 105–108.

Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2003). Modeling and Forecasting Realized Volatility. *Econometrica 71*, 579–625.

Andersen, T., T. Bollerslev, F. Diebold, and G. Wu (2006). Realized Beta: Persistence and Predictability. In T. Fomby and D. Terrell (Eds.), *Advances in Econometrics: Econometric Analysis of Economic and Financial Time Series*, Volume 20. Elsevier Science.

Andersen, T., T. Bollerslev, and D. Dobrev (2007). No-Arbitrage Semimartingale Restrictions for Continuous-Time Volatility Models subject to Leverage Effects, Jumps and i.i.d. Noise: Theory and Testable Distributional Assumptions. *Journal of Econometrics 138*, 125–180.

Andersen, T., T. Bollerslev, P. Frederiksen, and M. Nielsen (2007). Continuous-time Models, Realized Volatilities, and Testable Distributional Implications for Daily Stock Returns. Working paper, Kellogg School, Northwestern University.

Andersen, T., T. Bollerslev, and N. Meddahi (2004). Analytic Evaluation of Volatility Forecasts. *International Economic Review 45*, 1079–1110.

Andersen, T., P. Frederiksen, and A. Staal (2006). The Information Content of Realized Volatility Forecasts. Working paper, Kellogg School, Northwestern University.

Bandi, F. and J. Russell (2006). Separating Microstructure Noise from Volatility. *Journal of Financial Economics 79*, 655–692.

Bandi, F., J. Russell, and C. Yang (2008). Realized Volatility Forecasting and Option Pricing. *Journal of Econometrics 71*, 34–46.

Bandi, F., J. Russell, and J. Zhu (2008). Using High-Frequency Data in Dynamic Portfolio Choice. *Econometric Reviews 27*, 163–198.

Barndorff-Nielsen, O., S. Graversen, J. Jacod, M. Podolskij, and N. Shephard (2005). A Central Limit Theorem for Realised Power and Bipower Variations of Continuous Semimartingales. In Y. Kabanov and R. Lipster (Eds.), *From Stochastic Analysis to Mathematical Finance, Festschrift for Albert Shiryaev*. Springer.

Barndorff-Nielsen, O., P. Hansen, A. Lunde, and N. Shephard (2008). Designing Realised Kernels to Measure the ex-post Variation of Equity Prices in the Presence of Noise. *Econometrica 76*, 1481–1536.

Barndorff-Nielsen, O. and N. Shephard (2002). Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models. *Journal of the Royal Statistical Society: Series B 64*, 253–280.

Barndorff-Nielsen, O. and N. Shephard (2004a). Bipower Covariation: Measuring and Testing Jumps in Multivariate Price Processes. Working paper.

Barndorff-Nielsen, O. and N. Shephard (2004b). Econometric Analysis of Realised Covariation: High-frequency Covariance, Regression and Correlation in Financial Economics. *Econometrica 72*, 885–925.

Barndorff-Nielsen, O. and N. Shephard (2004c). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics 2*, 1–37.

Barndorff-Nielsen, O. and N. Shephard (2006a). Econometrics of Testing for Jumps in Financial Economics using Bipower Variation. *Journal of Financial Econometrics 4*, 1–30.

Barndorff-Nielsen, O. and N. Shephard (2006b). Impact of Jumps on Returns and Realised Variances: Econometric Analysis of Time-Deformed Lévy Processes. *Journal of Econometrics 131*, 217–252.

Barndorff-Nielsen, O., N. Shephard, and M. Winkel (2006). Limit Theorems for Multipower Variation in the Presence of Jumps in Financial Econometrics. *Stochastic Processes and Their Applications 116*, 796–806.

Bollerslev, T. and H. Zhou (2002). Estimating Stochastic Volatility Diffusionusing Conditional Moments of Integrated Volatility. *Journal of Econometrics 109*, 33–65.

Comte, F. and E. Renault (1998). Long Memory in Continuous-Time Stochastic Volatility Models. *Mathematical Finance 8*, 291–323.

Corradi, V. and W. Distaso (2006). Semiparametric Comparision of Stochastic Volatility Models Using Realized Measures. *Review of Economic Studies 73*, 635–667.

Dacorogna, M., R. Gencay, U. Muller, R. Olsen, and O. Pictet (2001). *An Introduction to High-Frequency Finance* (1st ed.). San Diego: Academic Press.

Fleming, J., C. Kirby, and B. Ostdiek (2003). The Economic Value of Volatility Timing using Realized Volatility. *Journal of Financial Economics 67*, 473–509.

French, K., G. Schwert, and R. Stambaugh (1987). Expected Stock Returns and Volatility. *Journal of Financial Economics 19*, 3–29.

Garcia, R., M. Lewis, S. Pastorello, and E. Renault (2006). Estimation of Objective and Risk-neutral Distributions based on Moments of Integrated Volatility. Working paper, Universite de Montreal.

Ghysels, E., P. Santa-Clara, and R. Valkanav (2006). Predicting Volatility: how to get most out of Returns Data Sampled at Different Frequencies. *Journal of Econometrics 131*, 59–95.

Goncalves, S. and N. Meddahi (2009). Bootstrapping Realized Volatility. *Econometrica 77*, 283–306.

Hansen, P. and A. Lunde (2006). Realized Variance and Market Microstructure Noise. *Journal of Business and Economic Statistics 24*, 127–161.

Hsieh, D. (1991). Chaos and Nonlinear Dynamics: Applications to Financial Markets. *Journal of Finance 46*, 1839–1877.

Huang, X. and G. Tauchen (2005). The Relative Contributions of Jumps to Total Variance. *Journal of Financial Econometrics 3*, 456–499.

Jacod, J. (2008). Asymptotic Properties of Power Variations and Associated Functionals of Semimartingales. *Stochastic Processes and their Applications 118*, 517–559.

Jacod, J. and A. N. Shiryaev (2003). *Limit Theorems For Stochastic Processes* (2nd ed.). Berlin: Springer-Verlag.

Jacod, J. and V. Todorov (2009). Testing for Common Arrivals of Jumps for Discretely Observed Multidimensional Processes. *Annals of Statistics 37*, 1792–1838.

Jiang, G. and R. Oomen (2008). A New Test for Jumps in Asset Prices. *Journal of Econometrics 144*, 352–370.

Lee, S. and P. Mykland (2008). Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics. *Review of Financial Studies 21*, 2535–2563.

Mancini, C. (2001). Disentangling the Jumps of the Diffusion in a Geometric Jumping Brownian Motion. *Giornale dell'Instituto Italiano degli Attuari LXIV*, 19–47.

Mancini, C. (2009). Nonparametric Threshold Estimation for Models with Stochastic Diffusion Coefficient and Jumps. *Scandinavian Journal of Statistics 36*.

Meddahi, N. (2002). Theoretical Comparision Between Integrated and Realized Volatility. *Journal of Applied Econometrics 17*, 479–508.

Merton, R. (1980). On Estimating the Expected Return on the Market. *Journal of Financial Economics 8*, 323–361.

Officer, R. (1973). The Variability of the Market Factor of the New York Stock Exchange. *Journal of Business 46*, 434–453.

Podolskij, M. and M. Vetter (2009). Estimation of Volatility Functionals in the Simultaneous Presence of Microstructure Noise and Jumps. *Bernoulli 15*, 634–658.

Todorov, V. (2006). Econometric Analysis of Jump-Driven Stochastic Volatility Models. Working paper, Duke University.

Todorov, V. (2009). Estimation of Coninuous-time Stochastic Volatility Models with Jumps using High-Frequency Data. *Journal of Econometrics 148*, 131–148.

Todorov, V. and T. Bollerslev (2007). Jumps and Betas: A New Theoretical Framework for Disentangling and Estimating Systematic Risks. Working paper, Northwestern University and Duke University.

Veraart, A. (2007). Feasible Inference for Realized Variance in the Presence of Jumps. Working paper.

Zhang, L. (2006). Efficient Estimation of Stochastic Volatility using Noisy Observations: A Multi-scale Approach. *Bernoulli 19*, 1019–1043.

Zhang, L., P. Mykland, and Y. Ait-Sahalia (2005). A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High Frequency Data. *Journal of the Americal Statistical Association 100*, 1394–1411.

Zhou, B. (1996). High-Frequency Data and Volatility in Foreign-Exchange Rates. *Journal of Business and Economic Statistics 8*, 45–52.

Research Papers 2009

2009-36: Tom Engsted and Thomas Q. Pedersen: The dividend-price ratio does predict dividend growth: International evidence

2009-37: Michael Jansson and Morten Ørregaard Nielsen: Nearly Efficient Likelihood Ratio Tests of the Unit Root Hypothesis

2009-38: Frank S. Nielsen: Local Whittle estimation of multivariate fractionally integrated processes

2009-39: Borus Jungbacker, Siem Jan Koopman and Michel van der Wel: Dynamic Factor Models with Smooth Loadings for Analyzing the Term Structure of Interest Rates

2009-40: Niels Haldrup, Antonio Montañés and Andreu Sansó: Detection of additive outliers in seasonal time series

2009-41: Dennis Kristensen: Pseudo-Maximum Likelihood Estimation in Two Classes of Semiparametric Diffusion Models

2009-42: Ole Eiler Barndorff-Nielsen and Robert Stelzer: The multivariate supOU stochastic volatility model

2009-43: Lasse Bork, Hans Dewachter and Romain Houssa: Identification of Macroeconomic Factors in Large Panels

2009-44: Dennis Kristensen: Semiparametric Modelling and Estimation: A Selective Overview

2009-45: Kim Christensen, Silja Kinnebrock and Mark Podolskij: Pre-averaging estimators of the ex-post covariance matrix

2009-46: Matias D. Cattaneo, Richard K. Crump and Michael Jansson: Robust Data-Driven Inference for Density-Weighted Average Derivatives

2009-47: Mark Podolskij and Mathias Vetter: Understanding limit theorems for semimartingales: a short survey

2009-48: Isabel Casas and Irene Gijbels: Unstable volatility functions: the break preserving local linear estimator

2009-49: Torben G. Andersen and Viktor Todorov: Realized Volatility and Multipower Variation