An Unbiased Measure of Realized Variance

Peter Reinhard Hansen*
Brown University
Department of Economics, Box B
Providence, RI 02912
Phone: (401) 863-9864
Email: Peter_Hansen@brown.edu

Asger Lunde
The Aarhus School of Business
Department of Information Science
Fuglesangs Allé 4 DK-8210 Aarhus V
Phone (+45) 89486688
Email: alunde@asb.dk

March 31, 2004

Abstract

The realized variance (*RV*) is known to be biased because intraday returns are contaminated with market microstructure noise, in particular if intraday returns are sampled at high frequencies. In this paper, we characterize the bias under a general specification for the market microstructure noise, where the noise may be autocorrelated and need not be independent of the latent price process. Within this framework, we propose a simple Newey-West type correction of the *RV* that yields an unbiased measure of volatility, and we characterize the optimal unbiased *RV* in terms of the mean squared error criterion. Our empirical analysis of the 30 stocks of the Dow Jones Industrial Average index shows the necessity of our general assumptions about the noise process. Further, the empirical results show that the modified *RV* is unbiased even if intraday returns are sampled every second.

JEL Classification: C10; C22; C80.

Keywords: Realized Variance; High-Frequency Data; Integrated Variance.

*Financial support from the Danish Research Agency, grant no. 24-00-0363 is gratefully acknowledged. We thank Torben Andersen, Tim Bollerslev, and Nour Meddahi for insightful comments. All errors remain our responsibility.

1 Introduction

The realized variance (*RV*) has become a popular empirical measure of volatility, because it is consistent for the integrated variance (*IV*) under ideal circumstances, see Andersen & Bollerslev (1998), Barndorff-Nielsen & Shephard (2002*b*) and Meddahi (2002). However, the consistency is established under assumptions that do not hold in practice and it is well known that the *RV* is biased and inconsistent for the *IV*, see e.g. Andreou & Ghysels (2002) and Oomen (2002).

The key to the *RV*'s shortcomings is that of an errors-in-variables problem. Ideally, the *RV* would be based on high-frequency intraday returns of the true price process, p^*. However, in practice the *RV* is derived from observed trades and/or quotes that are contaminated with market microstructure noise, while p^* is latent. The measurement errors cause the intraday returns to be autocorrelated and this makes the *RV* biased and inconsistent for the *IV*. So in practice it is common to use moderate-frequency intraday returns, such as 5-minute returns, because this can partially offset the bias, see e.g. Andersen & Bollerslev (1997). A theoretical justification for this approach has recently been established by Bandi & Russell (2003) who derived the optimal choice of sampling frequency in terms of the mean squared error criterion. However, sampling at a low frequency has the obvious drawback that it does not incorporate all data points (whereby information is lost). So this approach may result in an inefficient measure of volatility. This aspect is emphasized by Zhang, Mykland & Aït-Sahalia (2003) who show that an alternative estimator, which is based on subsampling techniques, is consistent for the *IV*. Further, Hansen & Lunde (2004*b*) have shown that a simple bias correction allows one to sample at very high frequencies, which leads to a more efficient estimator of *IV* than the standard *RV*. The results of Zhang et al. (2003), Bandi & Russell (2003), and Hansen & Lunde (2004*b*) provide important insight about the bias and the trade-off between bias and variance, however their results are derived for a particular type of market microstructure noise (*iid noise*), which does not hold in practice. In fact, our empirical analysis shows that the noise process has a time-dependence that persists for up to about two minutes.

The *RV*'s bias is caused by the autocorrelation in intraday returns, and the main idea of this paper is to remove the bias by incorporating the empirical autocovariances. This will allow us to sample intraday returns at high frequencies, and the hope is that this approach will preserve almost all information, such that the resulting *RV*-measure is a precise estimator of *IV*. We analyze the problem in a general framework where the observed price process equals the latent price process plus a noise term, $p(t) = p^*(t) + u(t)$. The framework is general in the sense that $u(t)$ may be autocorrelated and need not be independent of $p^*(t)$. The paper makes five main contributions. First, we characterize the bias of the *RV* for a very general type of market microstructure noise. The case where u is an iid process that is independent of p^*, (referred to as *iid noise*) is a special case of this framework. Second, we show that a simple autocorrelation correction of the *RV* yields an unbiased measure of the volatility. The bias-corrected realized variance, RV_{AC}, has a form that is similar to the robust covariance estimators, such as those of Newey & West (1987) and Andrews & Monahan (1992). Third, we consider the problem of identifying the optimal measure of volatility in the MSE sense. This problem is simplified considerably by restricting the attention to estimators that are conditionally unbiased, and within this class of estimators we show that the optimal *RV* is that with the smallest unconditional variance. Fourth, we undertake an extensive empirical analysis of the 30 stocks of the Dow Jones Industrial Average (DJIA) index, and we construct five years of daily unbiased volatilities for these equities. The empirical results are encouraging, because the RV_{AC} is found to be unbiased even if based on 1-second intraday returns! The observed price process can be constructed from both transaction data and mid-quotations and the two methods lead to very different types of market microstructure noise. In fact, the biases have the opposite sign, so the two sampling methods lead to quite different volatility signature plots. Fifth, we compare two methods for constructing artificial price processes in continuous time from tick-by-tick data. These are the previous-tick and the linear-interpolation methods that are discussed in Dacorogna, Gencay, Müller, Olsen & Pictet (2001), and we conclude that the latter is not suitable for sampling at high frequencies.

We choose to analyzed the realized variance under a general specification for the noise process, because this has several advantages. The generalization does not harm the clarity of our results that nest several results of the existing literature as special cases. However, the generality is more than aesthetically pleasing, because our empirical analysis shows that the noise process is both dependent across time and correlated with the innovations of the true latent price process. So a general specification for the noise process is vital for an understanding of the effects that market microstructure noise has on the *RV*.

We have argued that the bias of the *RV* is due to a measurement error problem. In the light the existing literature on measurement errors, it is somewhat remarkable that our bias correction is functional without imposing much structure on the noise process (measurement errors). For example, Mahajan (2003) analyzed binary choice models under a general specification for the measurement errors. While this framework is quite different from the present setting, a key insight from Mahajan's results is that consistent estimation requires identifying assumptions about the measurement errors. In our framework, we need not impose much (identifying) structure on the measurement errors, because the martingale property of the latent price process, p^*, makes it possible to identity the magnitude of the bias.

Previous approaches to bias-reduction methods include the filtering techniques used by Andersen, Bollerslev, Diebold & Ebens (2001) (moving average) and Bollen & Inder (2002) (autoregressive); and the approach by French, Schwert & Stambaugh (1987), Zhou (1996), and Hansen & Lunde (2004*b*) that corrects for first-order autocorrelation. Our approach is to correct for the first q_m autocorrelations, where q_m is increased as the sampling frequency, m, increases, such that the autocorrelation terms span a fixed period of time (and not our choice for m). The motivation for this is that the autocorrelation in intraday returns is specific to a period of time. For example, it may take 'pricing errors' two minutes to vanish, say, in which case we would choose q_m such that the q_m terms cover at least a two-minute window for any m.

There are several arguments for using high-frequency data to construct measures of volatility. For example, when the objective is to evaluate the predictive abilities of ARCH-type models, the RV can be substituted as a proxy for the unobserved conditional variance. This has several advantages over the use of squared daily returns as a proxy for the conditional variance. Andersen & Bollerslev (1998) showed that out-of-sample evaluation of ARCH-type models based on squared daily returns can be misleading (too poor), whereas an evaluation that is based on the RV is more accurate. So the RV is particularly useful for evaluation and comparison of multiple volatility models, because the RV makes it easier to tell which models that are better than others. More importantly, an empirical ranking of volatility models that is based on squared daily returns can be inconsistent for the true ranking, which is less likely to be the case if the RV is used, see Hansen & Lunde (2003). Time series of the RV provide valuable information about the dynamics of volatility, and this idea is explored in Andersen, Bollerslev, Diebold & Labys (2003) who proposed a dynamic model for the RV. The predictive ability of such models can be very good, in fact, Andersen, Bollerslev & Meddahi (2002) have shown that such models may perform almost as well as the true data generating process. Andersen, Bollerslev, Diebold & Labys (2001) provide additional arguments on the merits of the RV.

This paper is organized as follows. Section 2 contains definitions and descriptions of the basic properties of *RV*. Section 3 contains the theoretical results that include: A characterization of the *RV*'s bias in a general setting; the unbiased measure of the *RV*; and the characterization of the optimal unbiased *RV*. Section 4 contains our empirical analysis of the 30 DJIA equities and we propose a particular choice of *RV*-measure that is motivated by our findings. Section 5 summarizes and contains concluding remarks. Appendix A describes the method we use to construct confidence intervals for the average *RV* in our figures, and Appendix B contains all the proofs of the paper.

2 Definitions and Notation

We let $\{p(t)\}_{t\in I}$ be a logarithmic price process over a time interval I, and let $[a,b]\subset I$ be a compact interval that is partitioned into m intervals of equal length $\Delta_m \equiv (b-a)/m$. The interval, $[a,b]$, will typically span a trading day, so we refer to $y_{i,m} \equiv p(a+i\Delta_m) - p(a+i\Delta_m - \Delta_m)$, $i=1,\ldots,m$, as *intraday returns.* The *realized variance at frequency m* is defined by,

$$RV_{[a,b]}^{(m)} \equiv \sum_{i=1}^{m} y_{i,m}^2. \tag{1}$$

When $p(t)$ is a semimartingale[1] the *RV* is by definition a consistent estimator of the *quadratic variation*, (*QV*), of $\{p(t)\}_{t\in[a,b]}$, see Andersen & Bollerslev (1998) and Barndorff-Nielsen & Shephard (2002*b*). The stochastic volatility models define a particular class of semimartingales. These satisfy a stochastic difference equation of the form $dp(t) = \mu(t)dt + \sigma(t)dw(t)$, where $\mu(t)$ and $\sigma(t)$ are time-varying random functions and $w(t)$ is a standard Brownian motion. The *integrated variance* for such processes is defined by, $IV_{[a,b]} \equiv \int_a^b \sigma^2(t)dt$, and equals the *QV* for this class of semimartingales. The *IV* is fundamental for the pricing of derivative securities, see e.g. Hull & White (1987), which makes it a natural population measure of volatility.

Another population measure of volatility is the *conditional variance* (*CV*) that plays a pivotal role in ARCH-type models. The *CV* is defined by $CV_{[a,b]} \equiv \mathrm{var}(r_{[a,b]}|\mathcal{F}_a)$, where $r_{[a,b]} \equiv p(b) - p(a)$ and $\mathcal{F}_a$ denotes the information set at time a. So $CV_{[a,b]}$ is the variance of the innovation in $p(t)$ over the interval, $[a,b]$, conditional on the information at time a. The

[1]For the theory of semimartingales in relation to the present context, see e.g., Foster & Nelson (1996), Andersen, Bollerslev, Diebold & Labys (2001), and Barndorff-Nielsen & Shephard (2002*b*).

relations between the various quantities are the following: RV is generally consistent for QV, which (sometimes) equals the IV. Further, $E(IV|\mathcal{F}_a) \simeq CV$ with equality if $\{\mu(t)\}_{t=a}^{b}$ is $\mathcal{F}_a$-measurable. Thus the RV can be used to approximate these population measures under various assumptions. The empirical properties of the RV have been studied in various settings by Andersen, Bollerslev, Diebold & Ebens (2001) and Andersen, Bollerslev, Diebold & Labys (2001, 2003). A theoretical comparison between the IV and the RV is given in Meddahi (2002), and an asymptotic distribution theory of the RV (in relation to the IV) is established in Barndorff-Nielsen & Shephard (2002*a*).

To simplify the notation in the empirical analysis, we sometimes write $RV^{(x\,\text{min})}$ if the realized variance is based on x-minute intraday returns ($RV^{(x\,\text{min})}$ in place of $RV^{(m)}$, where $x = (b-a)/m$). We also use subscript-t to refer to day t and write $RV_t^{(m)}$ in place of $RV_{[a,b]}^{(m)}$ where $[a, b]$ represents the hours of day t that the market is open. A similar notation is used for QV, IV, and CV. Finally, we use $1_{\{\cdot\}}$ to denote the usual indicator function.

3 Realized Variance: Sampling Methods and Bias Issues

In this section, we discuss the empirical methods for constructing intraday returns from observed transactions and quotes, and propose a bias-adjustment that yields an unbiased RV. We also characterize the optimal unbiased RV in terms of the mean squared error criterion. Our bias correction is based on the following observation: Market microstructure noise causes intraday returns to be autocorrelated – in particular when these are sampled at a high frequency – and the autocorrelation is the reason that the RV is biased. So the empirical autocovariance function of intraday returns contains information that makes it possible to correct for the bias.

The RV is defined from intraday returns that require the value of the price process to be known at particular points in time. In practice, the price process is latent and prices must be interpolated from transaction and quotation data. These interpolated prices need not equal the true prices for a number of reasons that relate to market microstructure effects and aspects of the interpolation method. First, lack of liquidity could cause the observed price to differ from the true price, for example during short periods of time where large trades are being executed. Second, structural aspects of the market, such as the bid-ask spread and the discrete nature of price data that implies rounding errors. A third source of pricing errors can arise from the econometric method that is used to construct the artificial price data. The method is not unique and involves several choices such as: should prices be inferred from transaction data or mid-quotes; how to construct prices at points in time where no transaction or quotation occurred (at the exact same point in time).[2] A fourth source of pricing errors relates to the quality of the data. For example, tick-by-tick data sets contain misrecorded prices, such as transaction prices that are recorded to be zero. While zero-prices are easy to identify and remove from the data set, other misrecorded prices need not be.

3.1 Methods for Calculating the Realized Variance

The most common measure of *RV* requires equidistant price data. Equidistant price data must typically be constructed artificially from raw (irregularly spaced) price data. Let $t_0 < \cdots < t_N$ be distinct times at which raw prices, p_{t_j}, $j = 0, \ldots, N$, are observed. For example, p_{t_j} may be an actual transaction price or a mid-quote at time t_j. For any point in time, $\tau \in [t_0, t_N)$, we define the two artificial continuous time processes,

$$\begin{aligned} p(\tau) &\equiv p_{t_j}, && \tau \in [t_j, t_{j+1}), \\ \tilde{p}(\tau) &\equiv p_{t_j} + \frac{\tau - t_j}{t_{j+1} - t_j}(p_{t_{j+1}} - p_{t_j}), && \tau \in [t_j, t_{j+1}). \end{aligned}$$

The former is the *previous-tick* method that was proposed by Wasserfallen & Zimmermann (1985), and the latter is the *linear interpolation* method, see Andersen & Bollerslev (1997). Both methods are discussion in Dacorogna et al. (2001, sec. 3.2.1).[3] The artificial continuous time processes make it straight forward to construct equidistant intraday returns, such that the *RV* can be calculated for any frequency.

[Figure 1 about here]

In Figure 1 we have plotted the *volatility signature plots* for 4 of the 30 stocks that currently make up the Dow Jones Industrial Average (DJIA) index. Andersen, Bollerslev, Diebold & Labys (2000*b*) introduced the signature plot, which plots average realized variance, $\overline{RV}^{(m)} \equiv n^{-1}\sum_{t=1}^{n} RV_t^{(m)}$ against the sampling frequency, m, where the average is taken over n periods (days). A signature plot yields valuable information about the *RV*'s bias and can uncover important properties of the noise process. The signature plots in Figure 1 are based on five

[2] For example, Hansen & Lunde (2004*a*) fitted cubic splines to mid-quotes and noted that this method may 'over-smoothen' the price process, which would create a positive autocorrelation in the intraday returns.

[3] An alternative method for constructing a measure of *RV* is the Fourier method that was proposed by Malliavin & Mancino (2002), see also Barucci & Reno (2002).

years of daily RVs, where $RV_t^{(m)}$ is calculated from intraday returns that spans the period from 9:30 AM to 16:00 PM (the hours that the exchanges are open). We calculate $RV_t^{(m)}$ using both the previous-tick method and the linear interpolation method. The horizontal line represents $\bar{\sigma}^2 \equiv \overline{RV}^{(30\,\text{min})}$ that is a natural 'target' for an average RV, because 30-minute returns are expected to be almost uncorrelated. So $\overline{RV}^{(30\,\text{min})}$ should be (almost) unbiased for the average IV. The shaded area about $\bar{\sigma}^2$ represents an approximate 95% confidence interval for the average volatility. These confidence intervals are computed using a method that is described in Appendix A.

From Figure 1 we see that the RVs that are based on low and moderate frequencies appear to be approximately unbiased. However, at higher frequencies (more frequent than 10-minute sampling) the market microstructure effects are pronounced and the average RV diverges from $\bar{\sigma}^2$ as m is increased. Even the RV that is based on 5-minute returns is biased in some cases. This is particularly the case for the GE and PG equities when transaction prices are used. This is an important observation because 5-minute intraday returns is the most commonly used sampling frequency, and it seems that one would need to use 20-minute or 30-minute returns to obtain an almost unbiased RV, unless a bias reduction technique is employed.

There are two other important observations to be made from Figure 1. One is the difference between the signature plots of the previous-tick method and the linear-interpolation method. The latter is always below the former, and much below when intraday returns are sampled at high frequencies. This reveals an unfortunate property of the linear-interpolation method.

Lemma 1 *Let N be fixed and consider the RV based on the linear-interpolation method. It holds that $RV_{[a,b]}^{(m)} \xrightarrow{p} 0$ as $m \to \infty$.*

The important implication of Lemma 1 is that the RV should not be constructed from an artificial price process that is based on the linear-interpolation method – at least not if intraday returns are sampled at a high frequency. The property that $RV_{[a,b]}^{(m)} \xrightarrow{p} 0$ as $m \to \infty$ is also evident from most of the plots in Figure 1. However, the plots for the MSFT equity do not reveal this property, because this equity is traded/quoted very frequently, such that the drop off occurs for a larger m, than those displayed in the signature plots. Given the result of Lemma 1 and its empirical relevance, we shall entirely use the previous-tick method to construct the RVs in the remainder of the paper.

Another important observation of Figure 1 is that the transaction-based RVs have a positive bias whereas the quotation-based RVs have a negative bias, (previous-tick based RVs). This yields important information about certain properties of the market microstructure noise, as we show in the next subsection.

3.2 Characterizing the Bias of the Realized Variance

We define the noise process as $u(t) \equiv p(t) - p^*(t)$ for $t \in [a, b]$, where p^* is the true latent price process, and p is the observed price process. The noise process, u, can also be viewed as a measurement error or a pricing error, as discussed earlier. We shall refer to u as the noise process, and initially we make the following assumptions about u.

Assumption 1 *The noise process, u, is covariance stationary with mean zero, such that its autocorrelation function is defined by $\pi(s) \equiv E[u(t)u(t+s)]$.*

A simple example of a noise process that satisfies Assumption 1 is the iid noise process that has $\pi(s) = 0$ for all $s \neq 0$ and the Ornstein–Uhlenbeck specification that Aït-Sahalia, Mykland & Zhang (2003) used to analyze estimation of diffusion processes, in the presence of market microstructure noise.

Analogous to our definition of the intraday returns, $y_{i,m}$, we define the true intraday returns, $y^*_{i,m} \equiv p^*(a + i\Delta_m) - p^*(a + i\Delta_m - \Delta_m)$, and the noise increments, $e_{i,m} \equiv u(a + i\Delta_m) - u(a + i\Delta_m - \Delta_m)$, where $\Delta_m \equiv (b-a)/m$. Note that $e_{i,m} = y_{i,m} - y^*_{i,m}$ for all $i = 1, \ldots, m$. Below, we analyze the bias that is defined by,

$$\text{bias}(RV^{(m)}_{[a,b]}) \equiv E[RV^{(m)}_{[a,b]} - \sum_{i=1}^{m}(y^*_{i,m})^2],$$

and we show that $\text{bias}(RV^{(m)}_{[a,b]})$ is closely related to the properties of $\pi(s)$ in the neighborhood of $s = 0$.[4]

An important aspect of our analysis is that our assumptions allow for a dependence between u and p^*. This is a generalization of the assumptions made in the existing literature, and our empirical analysis shows that this generalization is needed, in particular if prices are sampled

[4] Alternatively, one can define the bias to be $E[RV^{(m)}_{[a,b]} - IV^*_{[a,b]}]$, where IV^* is the integrated variance of p^*. However, this quantity decomposes into $E[RV^{(m)}_{[a,b]} - \sum_{i=1}^{m}(y^*_{i,m})^2] + E[\sum_{i=1}^{m}(y^*_{i,m})^2 - IV^*_{[a,b]}]$, where the latter is specific to the data generating process for p^* and does not depend on the pricing error, u. So it is without loss of generality that we focus on the first term, because it captures all the effects from the noise process, u.

from mid-quotations. The dependence between u and p^*, which is relevant for our analysis, is give in the form of the following dependence parameters.

Definition 1 *We define* $\gamma_m^{(h)} \equiv E(y_{i,m}^* e_{i+h,m})$ *for* $h = 0, \pm 1, \pm 2, \ldots$.

The special case where u and p^* are independent corresponds to $\gamma_m^{(h)} = 0$. The parameters $\gamma_m^{(h)}$ for $h \neq 0$ are only used in our proofs, whereas $\gamma_m^{(0)}$ is the key dependence parameter that appears in our bias expressions. We are now ready to characterized the *RV*'s bias in this general framework.

Theorem 2 *Given Assumption 1 the bias of the realized variance is given by*

$$\text{bias}(RV_{[a,b]}^{(m)}) = 2m\gamma_m^{(0)} + 2m[\pi(0) - \pi(\Delta_m)].$$

Remark 1 *An interesting observation of Theorem 2 is that the bias is always positive when the noise process is independent of* p^* $(\gamma_m^{(0)} = 0)$*. However, the total bias can be negative if* $\gamma_m^{(0)} < -[\pi(0) - \pi(\Delta_m)]$*. This corresponds to the case where the downwards bias, which is due to the negative correlation between* $e_{i,m}$ *and* $y_{i,m}^*$*, exceeds the upwards bias that is caused by the* u*. This appears to be the case for the RVs that are based on mid-quotes, see Figure 1.*

The last term of the bias expression of Theorem 2, $\pi(0) - \pi(\Delta_m)$, shows that the bias is tied to the properties of $\pi(s)$ in the neighborhood of zero. Since $\Delta_m \to 0$ as $m \to \infty$, it is natural to analyze the asymptotic bias that is defined by

$$\text{bias}(RV_{[a,b]}^{(\infty)}) \equiv \lim_{m\to\infty} E[RV_{[a,b]}^{(m)} - \sum_{i=1}^{m}(y_{i,m}^*)^2].$$

Corollary 3 *Suppose that the assumptions of Theorem 2 hold and that* $\gamma \equiv \lim_{m\to\infty}[\gamma_m^{(0)}]/\Delta_m$ *and* $\pi'(0) \equiv \lim_{m\to\infty}[\pi(\Delta_m) - \pi(0)]/\Delta_m$ *exist and are finite. The asymptotic bias is given by,* $\text{bias}(RV_{[a,b]}^{(\infty)}) = 2(b-a)(\gamma - \pi'(0))$.

Next, we formulate the result for the special case with iid noise, where $\pi'(0)$ is undefined (or $\pi'(0) = -\infty$). This case has previously been analyzed by Corsi, Zumbach, Müller & Dacorogna (2001), Zhang et al. (2003), Bandi & Russell (2003), and Hansen & Lunde (2004*b*).

Corollary 4 (iid Noise) *Suppose that* u *is an iid process that is independent of* p^**, i.e.* $\pi(s) = 0$ *for all* $s \neq 0$ *and* $\gamma_m^{(0)} = 0$*. Then* $\text{bias}(RV_{[a,b]}^{(m)}) = 2m\pi(0)$.

So unlike the situation in Corollary 3, where the asymptotic bias is finite, the bias diverges to infinity for the case with iid noise.

The standard approach to reduce the bias, is to choose m sufficiently small, such that the autocorrelation is relatively small. However, a drawback of sampling at low frequencies is that this approach does not incorporate all the information in the data, and may result in an inefficient estimator of volatility. This suggests that it might be better to sample at a higher frequency and use a device for bias reduction. That is the agenda of the next subsection.

3.3 Bias Correcting the Realized Variance

We need some additional assumptions before we can establish a bias-correction method for the RV. The following assumption concerns the time-dependence of the noise process.

Assumption 2 *The noise process has finite dependence in the following sense. For some finite $\rho_0 \geq 0$ it holds that $\pi(s) = 0$ for all $s > \rho_0$ and for any $t_0 < t_1$ we have $E[u(s)(p^*(t_0) - p^*(t_1))] = 0$ for all $s \notin [t_0 - \rho_0, t_1 + \rho_0]$. Finally, $E[y^*_{i,m} y^*_{j,m}] = 0$ for $i \neq j$.*

The assumption is trivially satisfied under the iid noise assumption, while a more interesting class of noise processes with finite dependence are those of the moving average type, $u(t) = \int_0^{\rho_0} \theta(s) dB(s)$, where $B(s)$ represents a standard Brownian motion and $\theta(s)$ is a bounded (non-random) function on $[0, \rho_0]$. The autocorrelation function for a process of this kind is given by $\pi(s) = \int_s^{\rho_0} \theta(t)\theta(t-s)dt$, for $s \in [0, \rho_0]$. The last part of Assumption 2 requires that the increments in p^* are uncorrelated, which holds if $dp^* = \sigma(t)dw$, where w is a standard Brownian motion. In the more general case where $dp^* = \mu(t)dt + \sigma(t)dw$, the assumption would approximately hold for large m, under suitable smoothness conditions on $\mu(t)$.

Theorem 5 *Suppose that Assumptions 1 and 2 hold and let q_m be such that $q_m/m > \rho_0$. Then* $\text{bias}(RV^{(m)}_{\text{AC}[a,b]}) = 0$, *where*

$$RV^{(m)}_{\text{AC}[a,b]} \equiv \sum_{i=1}^{m} y_{i,m}^2 + 2\sum_{h=1}^{q_m} \frac{m}{m-h} \sum_{i=1}^{m-h} y_{i,m} y_{i+h,m}.$$

The important implication of Theorem 5 is that it is possible to bias-correct the RV when the noise process satisfies the mild requirements of Assumption 1 and 2.

It is worth to point out that our assumptions about the noise process are substantially weaker than the iid noise assumption. In the following we use the notation, $RV^{(m)}_{\text{AC}_{q_m}}$ to make the number of terms in the bias correction explicit. The following corollary formulates our result for the special case with iid noise, where $q_m = 1$ is sufficient for obtaining an unbiased RV.

Corollary 6 *Let Assumptions 1 and 2 hold and suppose that $\pi(s) = 0$ for all $s \neq 0$, (iid noise). Then $RV^{(m)}_{\mathrm{AC}_1[a,b]} \equiv \sum_{i=1}^{m} y_{i,m}^2 + 2\frac{m}{m-1}\sum_{i=1}^{m-1} y_{i,m}y_{i+1,m}$ is unbiased.*

The corollary generalizes a result in Hansen & Lunde (2004*b*) who made the additional assumption that p^* and u are independent. Interestingly, Zhang et al. (2003) have proposed an estimator that is consistent under the same set of assumptions as Hansen & Lunde (2004*b*). This estimator is based on an average over K_m subsample-estimators, where the noise terms in each of the subsamples are independent (across subsamples). Thus the variance of this estimator is proportional to m/K_m^2, and the consistency is achieved by increasing K_m such that $m/K_m^2 \to 0$ as $m \to \infty$, and Zhang et al. (2003) show that $K_m = m^{2/3}$ minimizes the asymptotic MSE of their estimator. While our estimator is unbiased it is not consistent. However, this could be achieved by averaging over an increasing number of $RV^{(m)}_{\mathrm{AC}}$s, where the different ms are all primes, such that every error term only appears in one $RV^{(m)}_{\mathrm{AC}}$. However, from a practical viewpoint this estimator is not very interesting, because the iid noise assumption is not valid in practice. While $RV^{(m)}_{\mathrm{AC}}$ is generally unbiased, the properties of the subsample estimator are unknown in the general case, where the noise process is autocorrelated and correlated with the true returns. The consistency of the subsample estimator is driven by the fact that it is an average of quantities that contain independent error terms, which relies on the unrealistic iid noise assumption.

A drawback of RV_{AC} is that it, in theory, could result in a negative estimate of volatility, because the covariance terms are scaled upwards by $\frac{m}{m-h}$. In our empirical analysis, we shall see that this problem has some practical relevance when intraday returns are constructed from transaction data at high frequencies. When mid-quotes are used the problem does not seem to have much practical relevance. However, to fully rule out the possibility of getting a negative estimate of volatility, one could use a different bias-reduction method that employs a different kernel, for example the Bartlett kernel that is used in the Newey & West (1987) covariance estimator. While a different kernel will not be entirely unbiased, it may result is a smaller mean squared error than that of the RV_{AC}, so this is an interesting avenue for future research.

In the time-series literature the lag-length, q_m, is typically chosen such that $q_m/m \to 0$ as $m \to \infty$, e.g., $q_m = \mathrm{int}[4(m/100)^{2/9}]$. However, in the present context this would be inappropriate because this would lead to $q_m = 3$ when a typical trading day (390 minutes) is divided into 78 intraday returns (5-minute returns), and $q_m = 6$ if the day was divided into 78 intraday returns (30-second returns). So the former q would cover 15 minutes whereas the latter would cover 3 minutes (6×30 seconds) and, in fact, the period shrinks to zero as $m \to \infty$. The autocorrelation in intraday returns is specific to a period of time, which does not depend on m. So it is more appropriate to keep the width of the 'autocorrelation-window', $\frac{q_m}{m}$, constant. This also makes $RV_{\mathrm{AC}}^{(m)}$ more comparable across different frequencies, m. Thus we set $q_m = \mathrm{ceil}(\frac{w}{(b-a)/m})$, [5] where w is the desired length of the lag window and $b - a$ is the length of the sampling period (both in units of time), such that $(b - a)/m$ is the period covered by each intraday return. In this case we will write $RV_{\mathrm{AC}_w}^{(m)}$ in place of $RV_{\mathrm{AC}_{q_m}}^{(m)}$. So if $w = 15\,\mathrm{min}$ and $b - a = 390\,\mathrm{min}$, we would include $q_m = \mathrm{ceil}[m/26]$ autocorrelation terms, and use the notation $RV_{\mathrm{AC}_{15\,\mathrm{min}}}^{(m)}$ for any m.

The fact that q_m must be such that $q_m/m > \rho_0 \geq 0$ has the implications that the modified realized variance, $RV_{\mathrm{AC}[a,b]}^{(m)}$, cannot be consistent for $QV_{[a,b]}$.[6] This property is common for covariance estimators in the standard time-series setting, whenever q_m/m does not converge to zero sufficiently fast, see e.g., Kiefer, Vogelsang & Bunzel (2000) and Jansson (2002).

3.4 Optimal Sampling Frequency for Intraday Returns

In the following we let σ_t^2 denote a theoretical measure of volatility for day t, such as QV_t, IV_t, or CV_t. Our objective is to determine the 'best' estimator of σ_t^2 amongst a set of possible RV-measures, which we denote by $\{\hat{\sigma}_{\lambda,t}^2, \lambda \in \Lambda\}$. So λ indexes both the frequency (m) and the window, w, we use to correct for the autocorrelation in intraday returns. A natural objective for choosing λ, is to minimize the mean squared error and define the optimal RV-measure to be the solution to

$$\hat{\sigma}_{\lambda^*,t}^2 \equiv \arg\min_{\lambda \in \Lambda} E[(\hat{\sigma}_{\lambda,t}^2 - \sigma_t^2)^2]. \tag{2}$$

An immediate obstacle for an empirical implementing of this objective is that σ_t^2 is unobserved. So it is not possible to calculate the sample average of $(\hat{\sigma}_{\lambda,t}^2 - \sigma_t^2)^2$ for different λs, and rank

[5] $\mathrm{ceil}(x)$ denotes the smallest integer that is greater than or equal to x.

[6] In the absence of the market microstructure noise we have that $\mathrm{var}(y_{i,m}^2) \approx 2\sigma_{i,m}^4$ and $\mathrm{var}(y_{i,m}y_{i+h,m}) \approx \sigma_{i,m}^2\sigma_{i+h,m}^2$, where $\sigma_{i,m}^2 \equiv \mathrm{var}(y_{i,m})$. So we note that $\mathrm{var}[RV_{\mathrm{AC}[0,1]}^{(m)}] \approx 2\sum_{i=1}^m \sigma_{i,m}^4 + \sum_{h=1}^{q_m}(2\frac{m}{m-h})^2 \sum_{i=1}^{m-h} \sigma_{i,m}^2\sigma_{i+h,m}^2$ which approximately is of the same order as $2(1+2q_m)\int_0^1 \sigma^4(s)ds$. So the variance is $O(m)$ times the *integrated quarticity*, see Barndorff-Nielsen & Shephard (2002*a*).

these to determine λ^*. However, the problem simplifies considerably by restricting our attention to estimators that are conditionally unbiased. We denote the set of conditionally unbiased estimators by $\{\hat{\sigma}^2_{\lambda,t}, \lambda \in \Lambda_{\text{ub}}\}$, where

$$E[\hat{\sigma}^2_{\lambda,t}|\sigma^2_t] = \sigma^2_t \qquad \text{for all} \quad \lambda \in \Lambda_{\text{ub}}. \tag{3}$$

For this class of RV-measures we obtain the following result.

Theorem 7 *It holds that*

$$\arg\min_{\lambda\in\Lambda_{\text{ub}}} E[(\hat{\sigma}^2_{\lambda,t} - \sigma^2_t)^2] = \arg\min_{\lambda\in\Lambda_{\text{ub}}} \text{var}[\hat{\sigma}^2_{\lambda,t}]. \tag{4}$$

Thus, within a class of conditionally unbiased estimators, the optimal estimator of σ^2_t is the one that has the smallest unconditional variance. So rather than attempting to solve (2), which requires knowledge about σ^2_t, we shall determine the empirical solution to (4). In practice, this will be done in two steps. First, we eliminate the estimators that are found to be biased and then we compare the estimators in terms of their variance.

4 Empirical Comparison of Bias-Corrected *RV*s

We put our theoretical results to work in an empirical analysis of the 30 stocks of the DJIA. These are listed in Table 1.

[Table 1 about here]

Our sample period spans the five years from January 2, 1998 to December 31, 2002, which results in $n = 1,255$ trading days. The data are transaction prices and quotes from the Trade and Quote (TAQ) database. For each of the 30 equities we extract data from the exchange where the equity is most actively traded. Table 1 lists the selected exchange for each of the equities. The raw data were filtered for outliers and we discarded transactions outside the period from 9:30am to 4:00pm, and the intraday returns were constructed using the previous-tick method. Additional details about the data are given in a separate technical appendix. We have derived results for all 30 equities, but to conserve space we only include the results for four equities in the figures and tables. The four equities were selected because we found them to be representative for the 30 equities. The corresponding figures and tables for the other 26 equities are given in the technical appendix.

[Figure 2 about here]

Figure 2 presents the volatility signature plots for RV_{AC}, for various values of w. The figure shows that the bias correction yields an average RV that is almost constant across frequencies, and very close to $\bar{\sigma}^2 \equiv \overline{RV}^{(30\,\mathrm{min})}$ which is an appropriate target for average volatility.

The signature plot of $RV_{\mathrm{AC}_1}^{(m)}$ provide valuable information about the time-dependence in u and identifies the appropriate values for w (those for which $RV_{\mathrm{AC}_w}^{(m)}$ is unbiased). The $RV_{\mathrm{AC}_1}^{(m)}$ only corrects for the first-order autocorrelation, so if $RV_{\mathrm{AC}_1}^{(m)}$ differs from $\bar{\sigma}^2$ it strongly suggests that one lag is not sufficient to capture the time-dependence in the noise process. When transaction data are used it seems that the minimum requirement for w is about 30, 15, 90, and 20 (300) seconds for the four equities, whereas the corresponding numbers for the quotation data are 120, 30, 90, and 30 (300) seconds. The numbers in brackets reflect that the volatility signature plots for PG are more ambiguous about the appropriate value for w, and the dependence may persist for as long as five minutes for the PG stock. The RVs that are based on a smaller w appear to be biased.

Next, we apply the result of Theorem 7 to identify the optimal unbiased RV. Since Theorem 7 applies to conditionally unbiased estimators only, we must discard the RVs that appear to be biased and the remaining problem is to determine the m and w (for w sufficiently large) that lead to the smallest variance of RV_{AC}.

[Figure 3 about here]

The plots in Figure 3 are signature plots for the standard deviations of RV_t (and not the sample average $\overline{RV}$), so these signature plots are quite different from those in Figures 1 and 2. More precisely, Figure 3 contains plots of the square root of $n^{-1}\sum_{t=1}^{n}(RV_{\mathrm{AC}_w,t}^{(m)} - \overline{RV}_{\mathrm{AC}_w}^{(m)})^2$ against m, for different values of w. So these plots contain sample information about which of the RVs that has the smallest variance.

The are three interesting observations that can be made from Figure 3. The first one is that the standard deviation is increasing in w. This is not too surprising, because it simply reflects that the estimation of additional covariance terms reduces the precision of the estimator, when the covariances are truly zero. The second observation is that the standard deviation is decreasing in m, although the gains are fairly modest once the sampling frequency increases beyond sampling every ten seconds. So the conclusion that can be drawn from Figure 3 is that the best unbiased RV is achieved by using the smallest possible value for w for which the RV is unbiased, and the largest possible value for m. Finally, the third observation that can be made from Figure 3 is that mid-quote data typically yields a smaller standard deviation than does transaction data, and this suggests that mid-quotes yield the most precise RV.

[Table 2 about here]

As discussed earlier, RV_{AC} need not be positive. We evaluate the practical relevance of this issue by counting the occurrences of negative values for various combinations of m and w. Table 2 shows that negative values do occur if either w is large or m is small. However, the configuration with $w = 2\,\text{min}$ and 1-second sampling, which appears to be a suitable configuration of RV_{AC}, according to the results of Figure 2, did not result in negative estimates except for the transaction data for MSFT. For this case we observed 20 negative RV_{AC}s out of the 1,255 days. Bid-ask bounces creates a sharp negative autocorrelation in intraday returns, which is more likely to yield a negative estimate, see West (1997) for a discussion on this issue. This is particular the case for an actively traded equity such as MSFT, where transaction prices jump between the bid and ask prices several times every minute.

When mid-quotes are used the possibility of getting a negative value of RV_{AC} does not seem to be of practical relevance. This provides another argument for using mid-quotes to construct the RV rather than transaction data.

5 Summary and Concluding Remarks

We have analyzed the standard RV and an alternative unbiased RV within a general framework for the market microstructure noise. Our assumptions allow for time-dependence in the noise process and allow the noise process to be correlated with the innovations in the latent price process. The framework generalizes the assumptions that have previously been made in the literature, and our empirical results clearly show that both generalizations are needed in practice.

We established several important results. First, we characterized the bias of the RV under the general assumptions. Then we showed that it is simple to bias-correct the RV and derived the unbiased RV_{AC}. The bias-correction is based on the empirical autocovariance function and RV_{AC} has a form that is similar to well-known robust covariance estimators. Next, we showed that the MSE-optimal unbiased RV can be identified empirically, because the unconditional variance is proportional to the MSE for conditionally unbiased RVs.

We applied our results to five years of high-frequency data for the 30 stocks of the DJIA index. Our analysis revealed several interesting qualitative and quantitative features of market microstructure noise. An important result is that market microstructure noise is time-dependent and when mid-quote data are use, the noise is clearly correlated with the innovations in the true price process. The empirical analysis also showed that the RV_{AC} does an excellent job at removing the bias, which the standard measure of RV suffers from, and we discussed how signature plots can be used to choose the parameters in RV_{AC}'s specification. The bandwidth parameter, w, which defines q_m, should be sufficiently large to fully capture the time-dependence in u. The volatility signature plot of $RV_{\mathrm{AC}_1}^{(m)}$ provides guidance about a suitable choice for w, and our plots suggested that the minimum value for w for the 30 DJIA stocks is in the range from 30 seconds to 2 minutes. The empirical results also showed that the larger is m the more precise is $RV_{\mathrm{AC}}^{(m)}$. This is quite intuitive, because the larger is m the more information is preserved in the estimator. The analysis also suggested that mid-quote data yields a more accurate RV-measure than do transaction data.

Based on our empirical results, we make the following recommendation for constructing an unbiased realized variance: Use mid-quotes data; construct intraday returns using the previous-tick method; use ultra high-frequency intraday returns (1-second); correct for the autocorrelation using a window-length that is sufficiently long to capture the persistence in the noise process (typically about 1-minute); use signature plots to make informed decisions about a good configuration for $RV_{\mathrm{AC}}^{(m)}$ in terms of w and m.

Our objectives of this paper were to characterize the bias and derive an unbiased RV under assumptions that are relevant for practical applications. While we have established several important results in this setting, our analysis also uncovered several new issues that should be addressed in future research. For example, the trade-off between bias and variance should be explored in future research, and it is quite plausible that the use of a different kernel in the bias-correction would result in a smaller mean squared error. Another important issue is to establish optimality results for a broader class of bias-corrected RVs, similar to the results that Bandi & Russell (2003) and Hansen & Lunde (2004*b*) have derived under the iid noise assumption, for the standard RV and RV_{AC_1}, respectively. In future research, we plan to derive the asymptotic distribution of RV_{AC}, analogous to the results that Barndorff-Nielsen & Shephard (2002*a*), Bandi & Russell (2003), and Hansen & Lunde (2004*b*) have derived for various RVs using more restrictive assumptions. Similar methods may also be useful for the estimation of realized covariances, see Andersen, Bollerslev, Diebold & Labys (2001), and future research should also include a comparison of RV_{AC} to model-based bias corrections and RVs that are based on different kernels. The literature on robust covariance estimation contains many estimators that can be used for bias reduction in this framework. So perhaps there is a second 'life' for the heteroskedasticity and autocorrelation consistent covariance estimators in this context.

A Appendix: Confidence Interval for Volatility Signature Plots

In our empirical analysis, we seek a measure that is informative about the precision of our bias corrected RVs. A minimal requirement is that the sample average of RV_{AC} is in a neighborhood of average daily volatility, and the latter is well approximated by the sample average of a low-frequency RV, which is presumed to be almost unbiased. In our empirical analysis we used the RV that is based on 30-minute intraday returns.

Let RV_t represent an unbiased RV-measure of the volatility for day t, and define the sample average, $\hat{\mu} \equiv n^{-1} \sum_{t=1}^{n} RV_t$, and consider the problem of constructing a confidence interval for $\mu \equiv E(RV_t)$, based on a sample RV_t, $t = 1, \ldots, n$. We follow the results of Andersen, Bollerslev, Diebold & Labys (2000*a*) and Andersen et al. (2003) and assume that the realized variance is log-normally distributed: $\log RV_t \sim N(\xi, \omega^2)$. Then it holds that $E[RV_t] = \mu = \exp(\xi + \omega^2/2)$ and $\text{var}(RV_t) = [\exp(\omega^2) - 1] \exp(\omega^2 + 2\xi)$.

Now, $\hat{\xi} \pm \sigma_{\hat{\xi}} c_{1-\alpha/2}$ is a $(1-\alpha)$-confidence interval for ξ, where $\hat{\xi} = n^{-1} \sum_{t=1}^{n} \log RV_t$, $\sigma_{\hat{\xi}}^2 \equiv \text{var}(\hat{\xi})$, and $c_{1-\alpha/2}$ is the appropriate quantile from the standard normal distribution. Since $\xi = \log(\mu) - \omega/2$ it follows that,

$$\begin{aligned} P[l \leq \xi \leq u] &= P[l + \omega/2 \leq \log(\mu) \leq u + \omega/2] \\ &= P[\exp(l + \omega/2) \leq \mu \leq \exp(u + \omega/2)], \end{aligned}$$

such that the confidence interval for ξ can be converted into one for μ. These calculations require ω^2 to be known, while in practice we must estimate ω^2. So we define $\eta_t \equiv \log RV_t - \hat{\xi}$ and use the Newey-West estimator,

$$\hat{\omega}^2 \equiv \sum_{t=1}^{n} \eta_t^2 + 2 \sum_{h=1}^{q} \left(1 - \frac{h}{q+1}\right) \sum_{t=1}^{n-h} \eta_t \eta_{t+h}, \qquad \text{where} \quad q = \text{int}[4(n/100)^{2/9}],$$

and subsequently set $\hat{\sigma}_{\hat{\xi}}^2 \equiv \hat{\omega}^2/n$. An approximate confidence interval for μ is now given by

$$\text{CI}(\mu) \equiv \exp(\log(\hat{\mu}) \pm \hat{\sigma}_{\hat{\xi}} c_{1-\alpha/2}),$$

which we recenter about $\log(\hat{\mu})$ (rather than $\hat{\xi} + \hat{\omega}/2$), because this ensures that the sample average of RV_t belongs to the confidence interval, $\hat{\mu} \in \mathrm{CI}(\mu)$.

B Appendix of Proofs

Proof of Lemma 1. First we note that $\tilde{p}(\tau)$ is continuous and piece-wise linear on $[a, b]$. So $\tilde{p}(\tau)$ satisfies the Lipschitz condition: $\exists\delta > 0$ such that $|\tilde{p}(\tau) - \tilde{p}(\tau + \epsilon)| \leq \delta\epsilon$ for all $\epsilon > 0$ and all τ. With $\epsilon = (b - a)/m$ we have that

$$\sum_{i=1}^{m} y_{i,m}^2 \leq \sum_{i=1}^{m} \delta^2 (b-a)^2 m^{-2} = \delta^2 (b-a)^2/m,$$

which shows that $RV_{[a,b]}^{(m)} \xrightarrow{p} 0$ as $m \to \infty$, and that the QV of $\tilde{p}(\tau)$ is zero with probability one. ■

Proof of Theorem 2. From the identities, $RV_{[a,b]}^{(m)} = \sum_{i=1}^{m} [y_{i,m}^{*2} + e_{i,m}^2 + 2y_{i,m}^* e_{i,m}]$ and $E(y_{i,m}^* e_{i,m}) = \gamma_m^{(0)}$ it follows that the bias is given by $\sum_{i=1}^{m} E(e_{i,m}^2) + 2E(y_{i,m}^* e_{i,m}) = m[E(e_{i,m}^2) + 2\gamma_m^{(0)}]$.

The result of the theorem now follows from the identity

$$E(e_{i,m}^2) = E[u(i\Delta_m) - u((i-1)\Delta_m)]^2 = 2[\pi(0) - \pi(\Delta_m)],$$

given Assumption 1. ■

Proof of Corollary 3. Since $m = (b - a)/[(b - a)/m) = (b - a)/\Delta_m$ we have that

$$\lim_{m\to\infty} m[\pi(0) - \pi(\Delta_m)] = \lim_{m\to\infty} (b-a)\frac{\pi(0) - \pi(\Delta_m)}{\Delta_m} = -(b-a)\pi'(0),$$

under the assumption that $\pi'(0)$ is well defined. ■

Proof of Corollary 4. Follows directly from Theorem 2 since both $\gamma_m^{(0)} = 0$ and $\pi(\Delta_m) = 0$. ■

Proof of Theorem 5. First, we note that

$$\begin{aligned} E(e_{i,m} e_{i+h,m}) &= E[u(i\Delta_m) - u((i-1)\Delta_m)][u((i+h)\Delta_m) - u((i+h-1)\Delta_m)] \\ &= 2\pi(h\Delta_m) - \pi((h-1)\Delta_m) - \pi((h+1)\Delta_m) \\ &= [\pi(h\Delta_m) - \pi((h+1)\Delta_m)] - [\pi((h-1)\Delta_m) - \pi(h\Delta_m)], \end{aligned}$$

such that

$$\sum_{h=1}^{q_m} E(e_{i,m} e_{i+h,m}) = [\pi(q_m\Delta_m) - \pi((q_m+1)\Delta_m)] - [\pi(0) - \pi(\Delta_m)],$$

where the first term equals zero given Assumption 2. Further, by Assumption 2 we have that $y^*_{i,m}$ is uncorrelated with the noise process when separated in time by at least ρ_0. So it follows that $0 = E[y^*_{i,m}(e_{i-q_m,m} + \cdots + e_{i+q_m,m})]$, which implies that $\gamma^{(0)}_m = -\sum_{h=1}^{q_m}(\gamma^{(-h)}_m + \gamma^{(h)}_m)$.

For $h \neq 0$ we have that $E(y_{i,m}y_{i+h,m}) = (\gamma^{(-h)}_m + \gamma^{(h)}_m) + E(e_{i,m}e_{i+h,m})$, since $E(y^*_{i,m}y^*_{i+h,m}) = 0$ by Assumption 2. It now follows that

$$\begin{aligned} E[\sum_{h=1}^{q_m} \tfrac{m}{m-h} \sum_{i=1}^{m-h} y_{i,m}y_{i+h,m}] &= \sum_{h=1}^{q_m}(\gamma^{(-h)}_m + \gamma^{(h)}_m)m - m[\pi(0) - \pi(\Delta_m)] \\ &= -m(\gamma^{(0)}_m + [\pi(0) - \pi(\Delta_m)]), \end{aligned}$$

which proves that $RV^{(m)}_{\mathrm{AC}}$ is unbiased. ■

Proof of Corollary 6. From $\pi(s) = 0$ for $s \neq 0$ we have that $\rho_0 = 0$, and the result now follows from Theorem 5, because $q_m/m > 0$ for all m ($q_m = 1$ for all m). ■

Proof of Theorem 7. Note that

$$E[(\hat{\sigma}^2_{\lambda,t})^2] = E[(\hat{\sigma}^2_{\lambda,t} - \sigma^2_t + \sigma^2_t)^2] = E[(\hat{\sigma}^2_{\lambda,t} - \sigma^2_t)^2] + E[(\sigma^2_t)^2] + 2E[(\hat{\sigma}^2_{\lambda,t} - \sigma^2_t)\sigma^2_t],$$

where the last term equals zero, since $E[\hat{\sigma}^2_{\lambda,t}|\sigma^2_t] = \sigma^2_t$ by assumption. The same identity implies that $E[\hat{\sigma}^2_{\lambda,t}] = E[\sigma^2_t]$ and it follows that

$$\mathrm{var}[\hat{\sigma}^2_{\lambda,t}] - E[(\hat{\sigma}^2_{\lambda,t} - \sigma^2_t)^2] = E[(\sigma^2_t)^2] - [E(\hat{\sigma}^2_{\lambda,t})]^2 = E[(\sigma^2_t)^2] - [E(\sigma^2_t)]^2 = \mathrm{var}[\sigma^2_t].$$

This proves that $\mathrm{var}[\hat{\sigma}^2_{\lambda,t}] \propto E[(\hat{\sigma}^2_{\lambda,t} - \sigma^2_t)^2]$, since $\mathrm{var}[\sigma^2_t]$ does not depend on λ. ■

References

Aït-Sahalia, Y., Mykland, P. A. & Zhang, L. (2003), How often to sample a continuous-time process in the presence of market microstructure noise, Working Paper w9611, NBER.

Andersen, T. G. & Bollerslev, T. (1997), 'Intraday periodicity and volatility persistence in financial markets', *Journal of Empirical Finance* **4**, 115–158.

Andersen, T. G. & Bollerslev, T. (1998), 'Answering the skeptics: Yes, standard volatility models do provide accurate forecasts', *International Economic Review* **39**(4), 885–905.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Ebens, H. (2001), 'The distribution of realized stock return volatility', *Journal of Financial Economics* **61**(1), 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2000*a*), 'Exchange rate return standardized by realized volatility are (nearly) Gaussian', *Multinational Finance Journal* **4**(3&4), 159–179.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2000*b*), 'Great realizations', *Risk* **13**(3), 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**(453), 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), 'Modeling and forecasting realized volatility', *Econometrica* **71**(2), 579–625.

Andersen, T. G., Bollerslev, T. & Meddahi, N. (2002), 'Analytic evaluation of volatility forecasts'. forthcoming in International Economic Review.

Andreou, E. & Ghysels, E. (2002), 'Rolling-sample volatility estimators: Some new theoretical, simulation, and empirical results', *Journal of Business & Economic Statistics* **20**(3), 363–376.

Andrews, D. W. K. & Monahan, J. C. (1992), 'An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator', *Econometrica* **60**, 953–966.

Bandi, F. M. & Russell, J. R. (2003), Microstructure noise, realized volatility, and optimal sampling, Working paper, Graduate School of Business, The University of Chicago.

Barndorff-Nielsen, O. E. & Shephard, N. (2002*a*), 'Econometric analysis of realised volatility and its use in estimating stochastic volatility models', *Journal of the Royal Statistical Society* **B 64**, 253–280.

Barndorff-Nielsen, O. E. & Shephard, N. (2002*b*), 'Estimating quadratic variation using realised volatility', *Journal of Applied Econometrics* **17**, 457–477.

Barucci, E. & Reno, R. (2002), 'On measuring volatility of diffusion processes with high frequency data', *Economics Letters* **74**, 371–378.

Bollen, B. & Inder, B. (2002), 'Estimating daily volatility in financial markets utilizing intraday data', *Journal of Empirical Finance* **9**, 551–562.

Corsi, F., Zumbach, G., Müller, U. & Dacorogna, M. (2001), 'Consistent high-precision volatility from high-frequency data', *Economic Notes* **30**(2), 183–204.

Dacorogna, M. M., Gencay, R., Müller, U., Olsen, R. B. & Pictet, O. V. (2001), *An Introduction to High-Frequency Finance*, Academic Press, London.

Foster, D. P. & Nelson, D. B. (1996), 'Continuous record asymptotics for rolling sample variance estimators', *Econometrica* **64**(1), 139–174.

French, K. R., Schwert, G. W. & Stambaugh, R. F. (1987), 'Expected stock returns and volatility', *Journal of Financial Economics* **19**(1), 3–29.

Hansen, P. R. & Lunde, A. (2003), 'Consistent preordering with an estimated criterion function, with an application to the evaluation and comparison of volatility models'. Brown University Working Paper, 2003-01
`http://www.econ.brown.edu/fac/Peter_Hansen`.

Hansen, P. R. & Lunde, A. (2004*a*), 'A forecast comparison of volatility models: Does anything beat a GARCH(1,1)?', *Journal of Applied Econometrics* . Forthcoming.

Hansen, P. R. & Lunde, A. (2004*b*), 'Realized variance and IID market microstructure noise'. Brown University Working Paper, 2004
`http://www.econ.brown.edu/fac/Peter_Hansen`.

Hull, J. & White, A. (1987), 'The pricing of options on assets with stochastic volatility', *Journal of Finance* **42**(2), 281–300.

Jansson, M. (2002), 'Autocorrelation robust tests with good size and power', *UC Berkeley, Mimeo* .

Kiefer, N. M., Vogelsang, T. J. & Bunzel, H. (2000), 'Simple robust testing of regression hypotheses', *Econometrica* **68**, 695–714.

Mahajan, A. (2003), 'Misclassified regressors in binary choice models'. Mimeo, Stanford University.

Malliavin, P. & Mancino, M. E. (2002), 'Fourier series method for measurement of multivariate volatilities', *Finance and Stochastics* **6**(1), 49–61.

Meddahi, N. (2002), 'A theoretical comparison between integrated and realized volatility', *Journal of Applied Econometrics* **17**, 479–508.

Newey, W. & West, K. (1987), 'A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix', *Econometrica* **55**, 703–708.

Oomen, R. A. C. (2002), 'Modelling realized variance when returns are serially correlated'. manuscript, Warwick Business School, The University of Warwick.

Wasserfallen, W. & Zimmermann, H. (1985), 'The behavior of intraday exchange rates', *Journal of Banking and Finance* **9**, 55–72.

West, K. D. (1997), 'Another heteroskedasticity and autocorrelation consistent covariance matrix estimator', *Journal of Econometrics* **76**, 171–191.

Zhang, L., Mykland, P. A. & Aït-Sahalia, Y. (2003), A tale of two time scales: Determining integrated volatility with noisy high frequency data, Working Paper w10111, NBER.

Zhou, B. (1996), 'High-frequency data and volatility in foreign-exchange rates', *Journal of Business & Economic Statistics* **14**(1), 45–52.

Tables and Figures

Table 1: Equities included in our empirical analysis.

Symbol	Name	Exchange	Trans./day	Quotes/day
AA	ALCOA INC	NYSE	1112	1625
AXP	AMERICAN EXPRESS COMPANY	NYSE	1778	2467
BA	BOEING COMPANY	NYSE	1461	2171
C	CITIGROUP	NYSE	2747	3890
CAT	CATERPILLAR INC	NYSE	996	1574
DD	DU PONT DE NEMOURS E I CO	NYSE	1486	2496
DIS	WALT DISNEY COMPANY	NYSE	1793	2601
EK	EASTMAN KODAK CO	NYSE	926	1501
GE	GENERAL ELECTRIC CORPORATION	NYSE	3277	3721
GM	GENERAL MOTORS CORP	NYSE	1301	2259
HD	HOME DEPOT INC	NYSE	2111	2930
HON	HONEYWELL INTERNATIONAL INC	NYSE	1061	1783
HWP	HEWLETT-PACKARD COMPANY	NYSE	1922	2557
IBM	INTERNATIONAL BUSINESS MACHINES	NYSE	2643	4306
INTC	INTEL CORP	NASDAQ	37588	21185
IP	INTERNATIONAL PAPER CO	NYSE	1186	1853
JNJ	JOHNSON AND JOHNSON	NYSE	1650	2028
JPM	MORGAN J.P. CO INC	NYSE	1751	2477
KO	COCA-COLA CO	NYSE	1608	2074
MCD	MCDONALDS CORP	NYSE	1376	1977
MMM	MINNESOTA MNG MFG CO	NYSE	1156	1862
MO	PHILIP MORRIS COMPANIES INC	NYSE	1843	3133
MRK	MERCK CO INC	NYSE	1889	2256
MSFT	MICROSOFT CORP	NASDAQ	34252	22385
PG	PROCTER & GAMBLE CO	NYSE	1680	3053
SBC	SBC COMMUNICATIONS INC	NYSE	1761	2367
T	AT & T CORP	NYSE	1979	1966
UTX	UNITED TECHNOLOGIES CORP	NYSE	1051	1711
WMT	WAL-MART STORES INC	NYSE	2007	2695
XOM	EXXON MOBIL CORPORATION	NYSE	2028	2803

The table lists the equities used in our empirical analysis. For each equity, we extract data from the exchange where it is most actively traded (third column). The average number of transactions and quotes per day are given in the last two columns.

Table 2: Number of negative volatilities for the RV_{AC} estimator.

	Sampling frequency in seconds																							
	Trade data												Quote data											
Window	1	2	3	4	5	10	30	60	120	180	240	300	1	2	3	4	5	10	30	60	120	180	240	300
AA																								
$w=1$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=2$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=3$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=5$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=7$	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	0	0
$w=10$	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	1	1	0	0	0
GE																								
$w=1$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=2$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=3$	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=5$	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	1	0	1	0
$w=7$	0	0	0	0	0	0	0	0	3	3	1	6	0	0	0	0	0	0	0	1	1	1	1	2
$w=10$	1	1	1	1	1	2	2	4	1	5	2	6	1	1	1	1	1	1	1	1	1	2	3	2
MSFT																								
$w=1$	26	19	13	12	10	5	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=2$	20	10	7	3	3	2	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=3$	10	6	5	2	5	3	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
$w=5$	12	10	8	5	6	3	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	0
$w=7$	13	12	6	4	3	1	0	0	0	1	1	3	0	0	0	0	0	0	0	0	0	0	1	1
$w=10$	8	7	8	4	4	2	0	1	2	1	2	3	0	0	0	0	0	0	0	0	0	1	0	1
XOM																								
$w=1$	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0	0	0	0
$w=2$	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0	0	0	0
$w=3$	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0	0	0	0
$w=5$	0	0	0	0	0	0	0	0	1	1	3	3	0	0	0	0	0	0	0	0	0	0	2	0
$w=7$	0	0	0	0	0	0	1	0	2	11	3	13	0	0	0	0	0	0	0	0	2	4	2	7
$w=10$	3	3	4	4	4	4	3	5	9	9	8	13	2	2	2	3	2	3	4	5	6	5	9	7

This table reports the number of days that the RV_{AC} produced a negative estimate, as a function of w and m. The total number of trading days in our sample is 1,255.

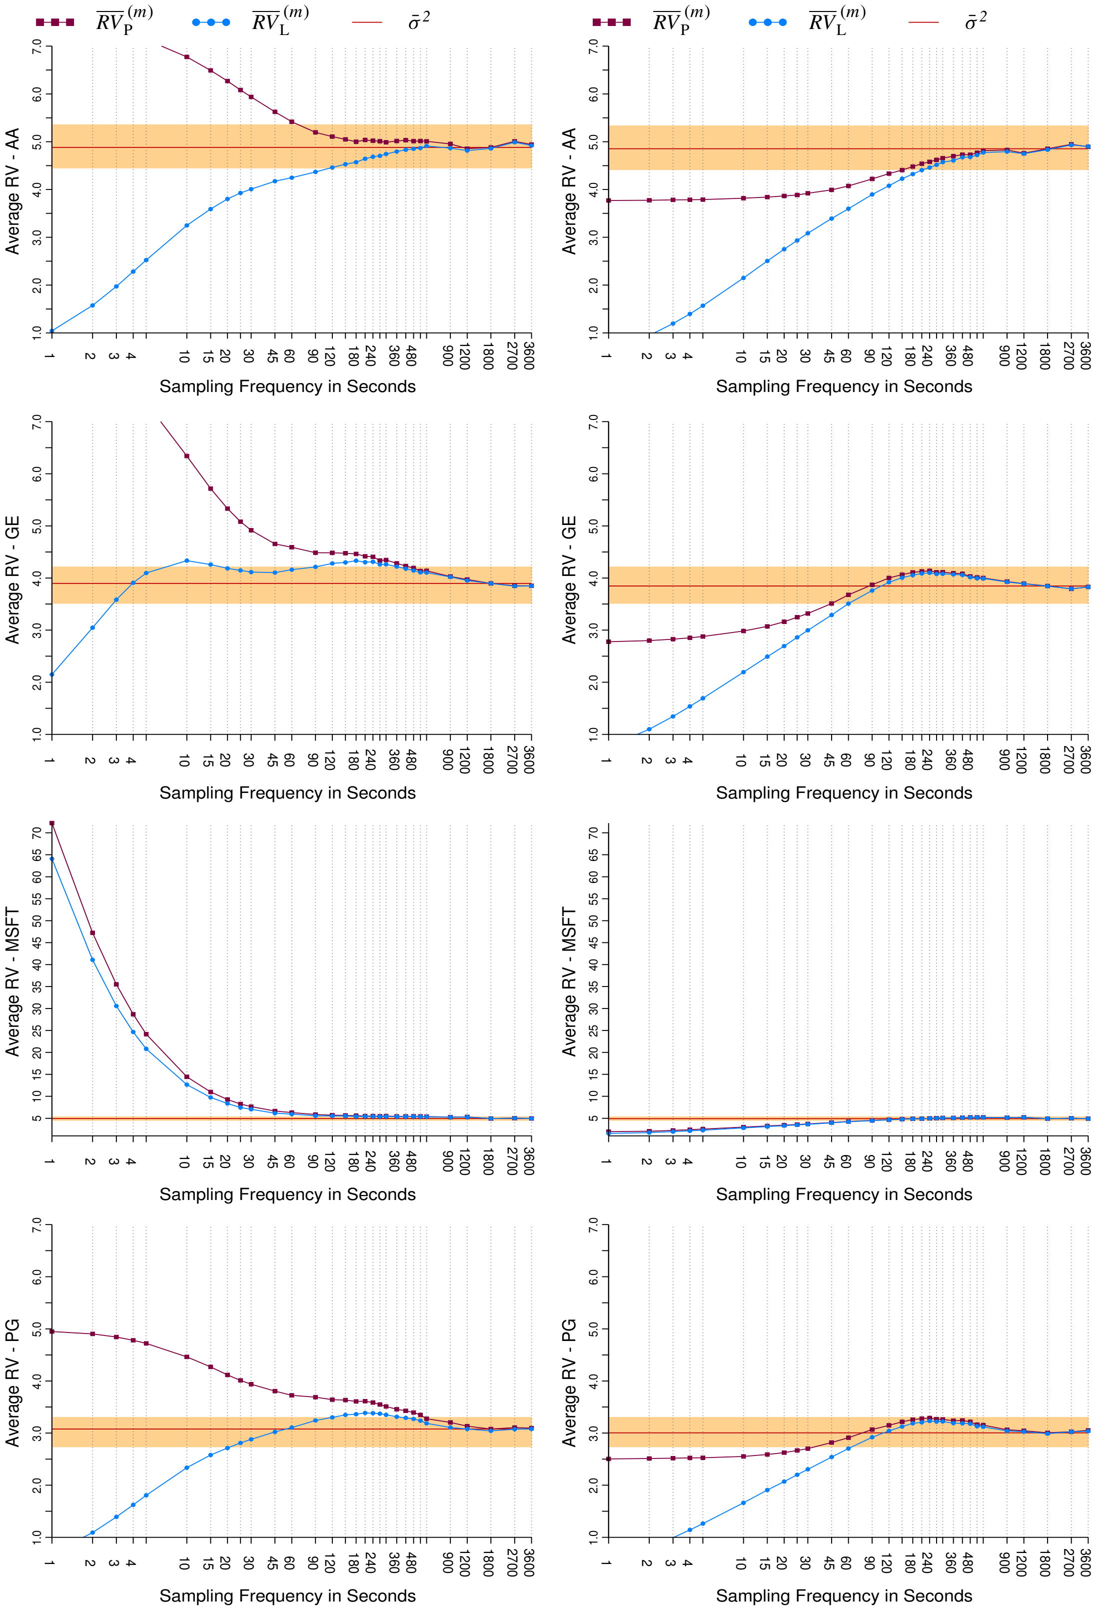

Figure 1: Signature plots of *RV* based on the previous-tick method, or the linear-interpolation method. Panels on the left are based on transaction data whereas panels on the right is based on mid-quotes. The shaded area about $\bar{\sigma}^2$ gives an approximate 95% confidence band for the average *IV*.

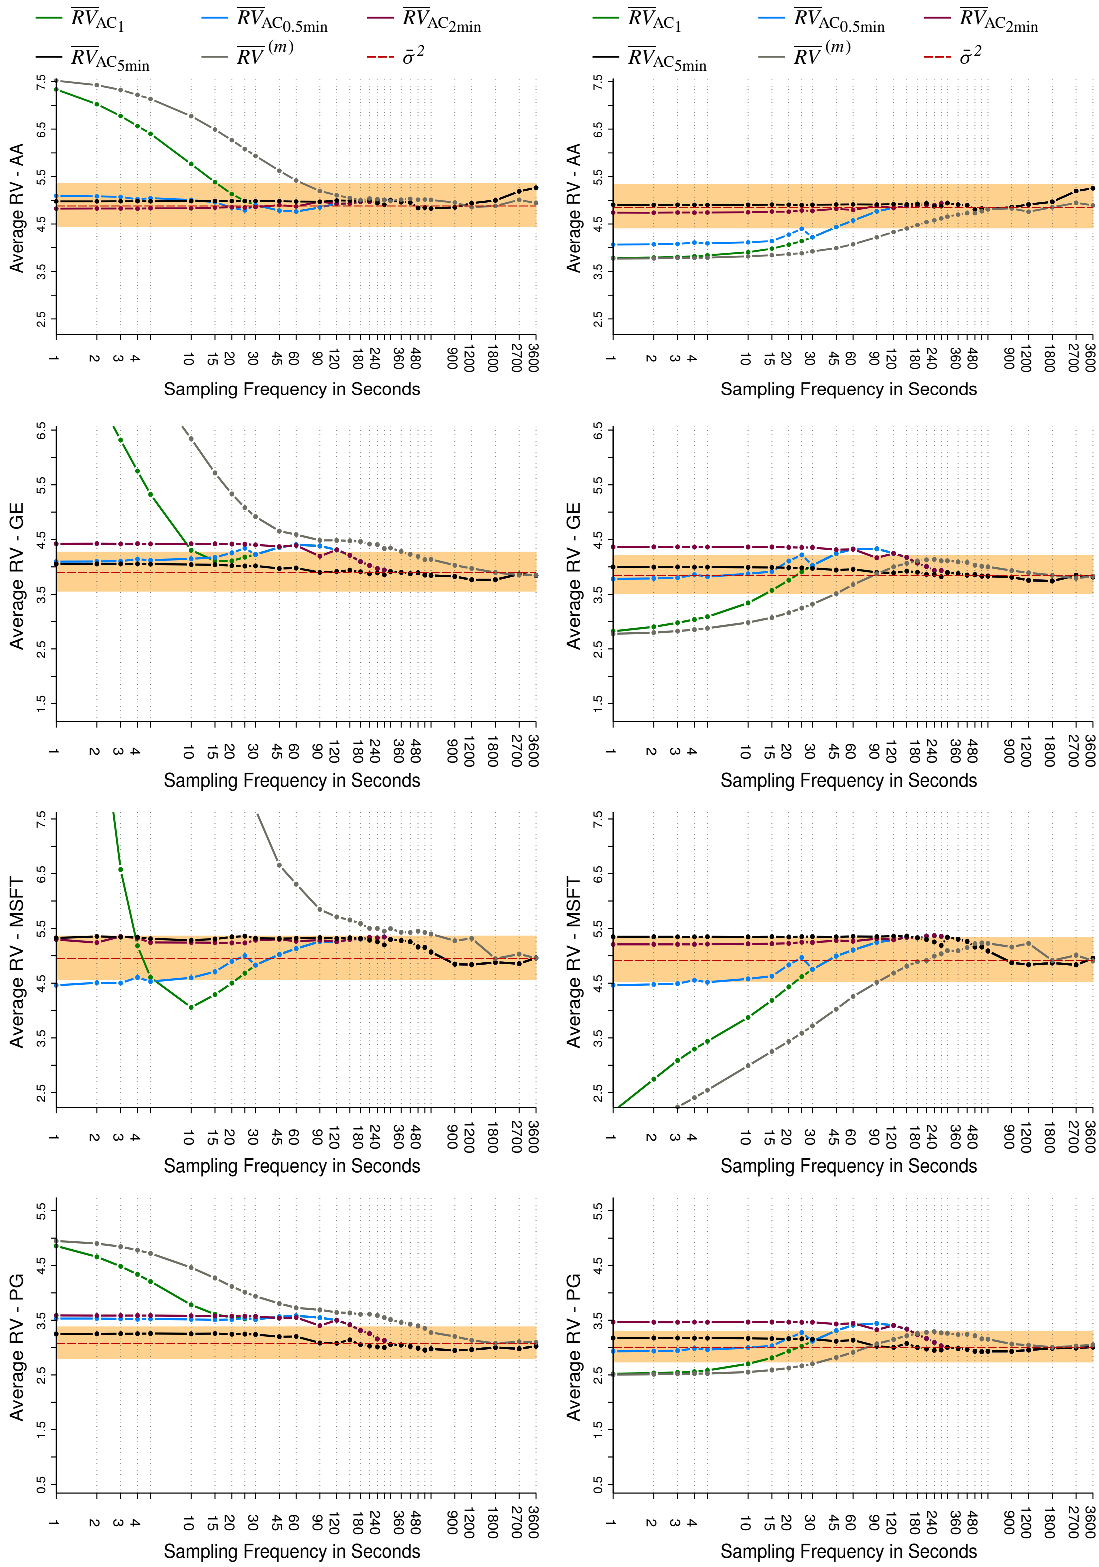

Figure 2: Volatility signature plots of $RV_{\mathrm{AC},t}$ for four different choices of w. Panels on the left are based on transaction data whereas panels on right are based on mid-quotes. The shaded area about $\bar{\sigma}^2$ represents an approximate 95% confidence band for the average IV.

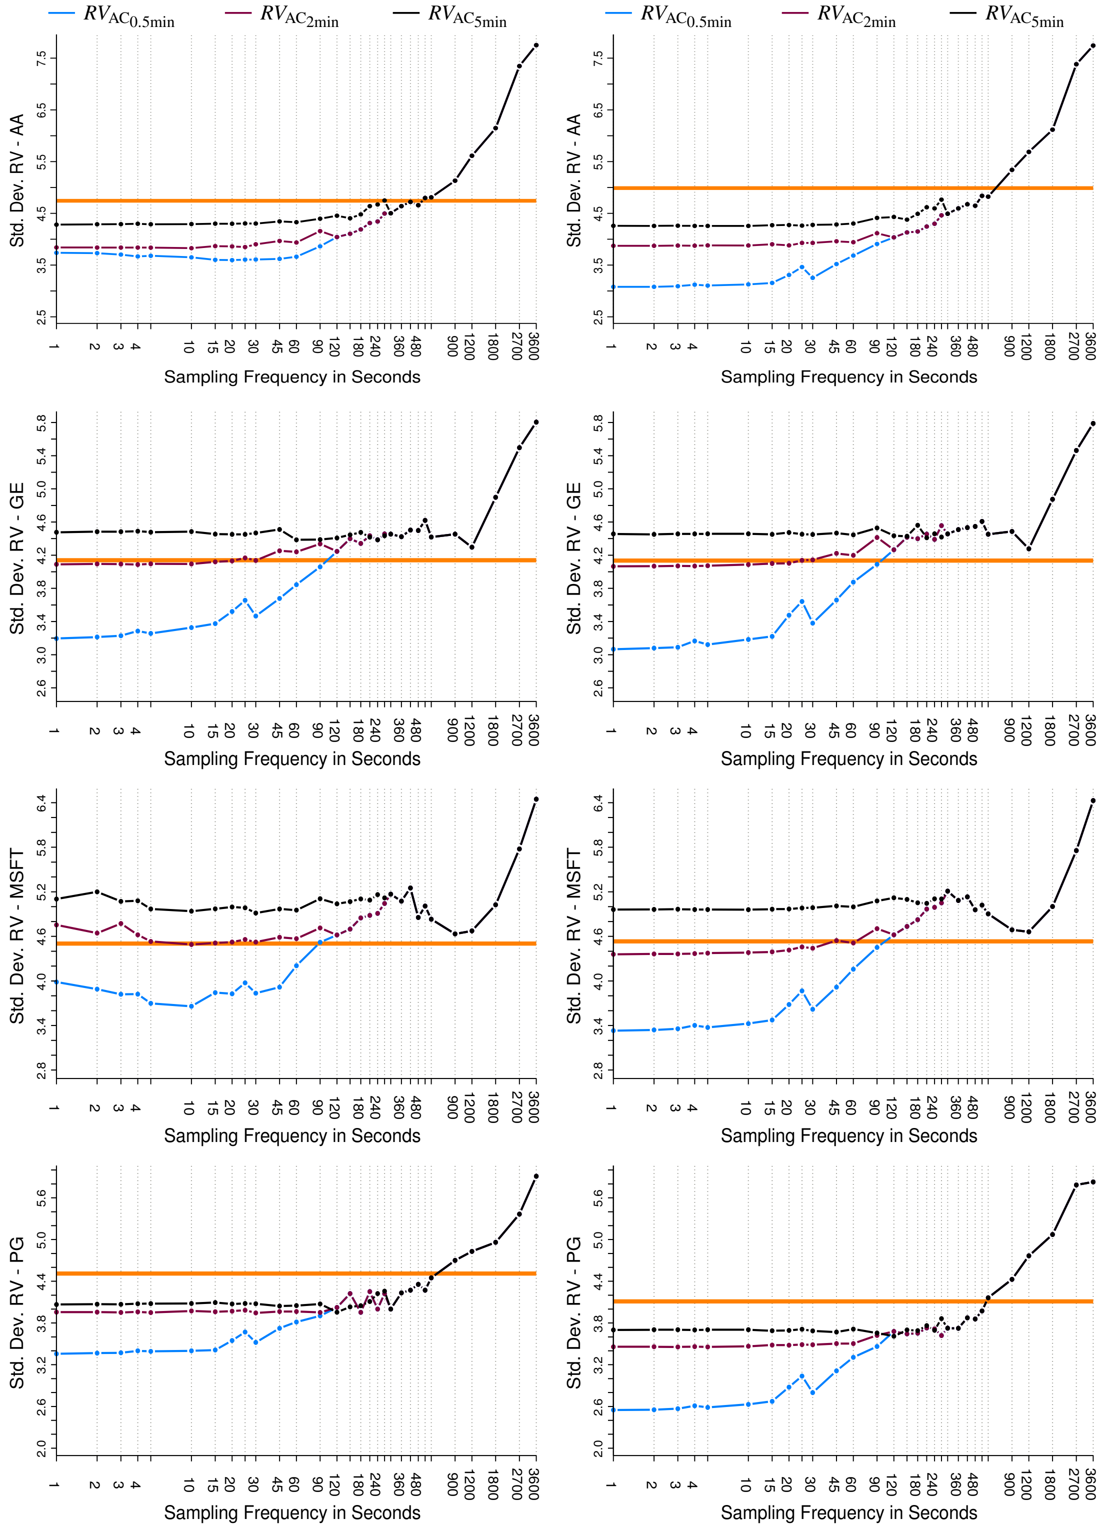

Figure 3: Signature plots of the sample standard deviation of $RV_{\mathrm{AC},t}$, which can be viewed as approximation of 'MSE signature plots', because the MSE of a conditionally unbiased RV is proportional to its unconditional variance. The horizontal line is the sample standard deviation of $RV^{(30min)}$. Plots for both transaction data (left) and mid-quotes (right) are displayed.