

Jump-Robust Volatility Estimation using Nearest Neighbor Truncation

Torben G. Andersen, Dobrislav Dobrev, Ernst Schaumburg*

Abstract

We propose two new jump-robust estimators of integrated variance that allow for an asymptotic limit theory in the presence of jumps. Specifically, our MedRV estimator has better efficiency properties than the tripower variation measure and displays better finite-sample robustness to jumps and small (“zero”) returns. We stress the benefits of local volatility measures using short return blocks, as this greatly alleviates the downward biases stemming from rapid fluctuations in volatility, including diurnal (intraday) U-shape patterns. An empirical investigation of the Dow Jones 30 stocks and extensive simulations corroborate the robustness and efficiency properties of our nearest neighbor truncation estimators.

JEL classification: C14; C15; C22; C80; G10.

Keywords: High-frequency data; Integrated variance; Finite activity jumps; Realized volatility; Jump robustness; Nearest neighbor truncation; Intraday U-shape patterns.

*Torben G. Andersen: Northwestern University; NBER; CREATES; t-andersen@kellogg.northwestern.edu
Dobrislav Dobrev: Federal Reserve Board of Governors, Dobrislav.P.Dobrev@frb.gov
Ernst Schaumburg: Federal Reserve Bank of New York, Ernst.Schaumburg@ny.frb.org

Andersen gratefully acknowledges financial support from the NSF through a grant to the NBER, from the *Zell Center for Risk* at the Kellogg School and from CREATES funded by the Danish National Research Foundation. Furthermore, he appreciates the CME Group for providing S&P 500 futures data from CME DataMine.

We thank Federico Bandi, Luca Benzoni, Jean Jacod, Per Mykland, Roel Oomen, Roberto Renò, Neil Shephard, Viktor Todorov, Lan Zhang, and, in particular Mark Podolskij and Kevin Sheppard for their insights. We also thank participants at the Singapore Management University Conference in Honor of P.C.B. Phillips, July 2008, the CREATES Volatility Symposium, Aarhus, August 2008, the Chicago/London Conference on Financial Markets, December 2008, the Humboldt-Copenhagen Conference on Recent Developments in Financial Econometrics, Berlin, March 2009, the North American Summer Meeting of the Econometric Society, Boston, June 2009, the Conference of the Society for Financial Econometrics, Geneva, June 2009, and seminar participants at the Federal Reserve Board and the U.S. Commodity Futures Trading Commission for their comments.

The views in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, the Federal Reserve Bank of New York or of any other person associated with the Federal Reserve System.

1. Introduction

In the ubiquitous continuous-time no-arbitrage semimartingale framework for modeling asset prices, it is often useful to obtain separate estimates of the continuous part of the volatility process versus the return variation induced by discontinuities or jumps. This is, for instance, the case in the context of risk management, option pricing and volatility forecasting. Thus far, the dominant approach for delivering jump-robust volatility estimates from intradaily return observations has been the so-called realized bipower variation measure, introduced by Barndorff-Nielsen and Shephard (2004), in which volatility is estimated by the cumulative sum of products of adjacent absolute returns.¹ While the bipower variation, by construction, ensures that (finite activity) jumps will not impact the consistency of the volatility estimate, it does not allow for a feasible asymptotic theory under the jump alternative and is subject to a fairly significant finite sample jump distortion (upward bias) that may be of concern in applications.² To obtain improved finite sample jump robustness and an asymptotic theory under the jump alternative, the bipower variation has been generalized in subsequent work to tripower and higher order multipower variation (MPV) measures, which employ products of powers of three or more adjacent absolute returns.³ Tripower variation is theoretically the most efficient among those. However, it is also more susceptible to market microstructure contamination of the high-frequency return series than the bipower variation. In particular, multipower variation measures are sensitive to the presence of very small (zero) returns arising from stale quotes and rounding to a discrete price grid. In applications, the prevalence of zero returns is often substantial, thus introducing a separate source of potential bias.

We propose two simple alternatives to the prevailing bipower and tripower variation measures that provide additional robustness to jumps and/or market microstructure noise by using nearest neighbor truncation. The first estimator obtains jump robustness by appropriately scaling the square of the minimum of two consecutive intraday absolute returns. If one of these returns is large, e.g., due to the presence of a jump during the interval, this return is automatically discarded and all weight in the computation falls on the adjacent diffusive returns. Asymptotically, under finite jump activity, we never encounter two adjacent jumps so, like bipower, the measure retains consistency for the underlying integrated diffusive variance. However, this “minimum” or “MinRV” estimator suffers from a similar exposure to small (zero) returns as the traditional MPV estimators. In addition, large (absolute) returns are inherently more informative of the underlying volatility than small returns, so our minimum estimator is not particularly efficient. Consequently, we introduce another variant which considers three consecutive intraday returns and simply squares the median absolute return among the three. This estimator also, asymptotically, avoids including the impact of a jump in the measure while reducing the sensitivity to the smallest absolute returns within the trading day and enhancing efficiency.

The unifying theme behind our new estimators is that the absolute returns are truncated at a level controlled by the neighboring returns. “MinRV” uses one-sided truncation, comparing each intraday return with the subsequent absolute return, while “MedRV” uses two-sided truncation, picking the median of three adjacent absolute returns. Hence, these estimators exploit *adaptive* truncation which endogenously controls for the local level of volatility. This avoids the potentially delicate choice of an ex-ante threshold required for the asymptotically more efficient

¹Alternative jump-robust estimators include, among others: Mancini (2006), Andersen et al. (2008), Christensen and Podolskij (2007), Dobrev (2007).

²Studies of the finite sample behavior of the bipower statistic include Barndorff-Nielsen and Shephard (2004), Huang and Tauchen (2005), Lee and Mykland (2007), Andersen et al. (2007).

³See Barndorff-Nielsen et al. (2006a) and Barndorff-Nielsen et al. (2006c).

truncated RV approaches of Mancini (2006) and Aït-Sahalia and Jacod (2007) or the truncated bipower variation of Corsi et al. (2010).

The endogenous “nearest neighbor” truncation enhances the robustness of our estimators and allows for the development of an asymptotic distribution theory covering both the “no-jump” null hypothesis and the “jump” alternative, facilitating inference about the presence of jumps. Specifically, the MedRV estimator has better theoretical efficiency properties than the tripower variation measure and displays better finite-sample robustness to jumps and the occurrence of “zero” returns in the sample. Moreover, it is trivial to modify the MinRV/MedRV estimators to gain robustness to so-called “bounce backs” (due to isolated and errant outliers in the price record) either by using staggered rather than adjacent returns, as originally suggested by Huang and Tauchen (2005) for MPV estimators, or, even more effectively, by pre-averaging the returns, as first proposed by Podolskij and Vetter (2009).

We define the MinRV estimator as arising from the sequential use of the $\min(\cdot)$ operator on blocks of two returns and the MedRV estimator from applying the $\text{med}(\cdot)$ operator on blocks of three returns. Increasing the block size over this minimum length leads to a gradual efficiency loss, analogous to that observed for higher order MPV measures. Instead, as a theoretically attractive avenue for more efficient jump-robust volatility estimation exploiting larger block sizes, we consider the recent quantile realized volatility (QRV) estimator of Christensen et al. (2010) based on optimally combining observations of extreme quantiles within blocks of twenty or more data points.⁴ However, the reliance on larger blocks has a non-trivial practical cost in finite samples, in the form of a bias. A critical assumption is that the returns within each block are i.i.d. Gaussian and thus, in particular, that volatility is constant across the block. Although asymptotically valid, this assumption becomes progressively harder to maintain in practice as the block size increases to encompass a wider calendar interval. This is due to the pronounced variation in volatility across the trading day, which renders the underlying returns within a longer block non-homogeneous. In addition, although trade and quote arrivals are correlated with increments in volatility they are not well enough aligned to ensure homogeneity (i.e., constant volatility) of the observed log-price increments in either calendar or tick time. Consequently, the gap between the finite sample and asymptotic properties of such estimators tends to be substantially wider than for our “local” MinRV and MedRV estimators. We provide extensive evidence on the finite sample properties of the alternative estimators in simulations as well as for individual stocks in the Dow-Jones 30 index between January 2005 and July 2009.

The remainder of the paper progresses as follows: Section 2 lays out the basic setup and introduces several popular jump-robust measures of integrated volatility along with our MinRV and MedRV estimators. The asymptotic properties of the new estimators are laid out in a series of propositions. Section 3 provides an empirical application to the set of stocks in the Dow Jones 30 index. Section 4 presents extensive simulation evidence exploring the impact of a variety of features on the performance of the alternative estimators. Section 5 provides concluding observations, while all formal proofs are relegated to the appendix.

2. Jump-Robust IV Estimation

We consider the univariate logarithmic price process $Y = \{Y_t\}_{0 \leq t \leq 1}$ of an asset defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P)$ so that Y is adapted to the filtration and evolves in

⁴In fact, our “minimum” and “median” estimators share features of both the multipower and quantile estimators. MinRV and MedRV rely on functions of (small) overlapping blocks of adjacent returns like the former, while they exploit the squared quantiles of the (absolute) returns over a (short) block, thus mimicking qualitative aspects of the latter. Another recent estimator of this type is the Realized Outlyingness Weighted Quadratic Covariation (ROWQCov) estimator of Boudt et al. (2008).

continuous time as described by the following jump-diffusive representation,

$$dY_t = \mu_t dt + \sigma_t dB_t + dJ_t \quad (1)$$

where μ is a locally bounded and predictable process, while σ is cadlag and bounded away from zero almost surely. The price process is observed at the $N+1$ discrete points in time $0 \leq t_0 < t_1 < \dots < t_N \leq 1$ over a given period which we refer to as a trading day. The corresponding returns and time intervals are denoted $\Delta Y_i = Y_{t_i} - Y_{t_{i-1}}$ and $\Delta t_i = t_i - t_{i-1}$, $i = 1, \dots, N$. Finally, J denotes a finite activity jump process and dJ_t is either zero (no jump) or a real number indicating the size of the jump at time t . Our finite activity jump assumption implies that there are only a finite number of jumps over the trading day. The subsequent analysis exploits standard continuous record in-fill asymptotics where the time increment between successive return observations uniformly shrinks toward zero as N increases.

The object of interest is the continuous part of the quadratic variation, or the integrated variance (IV), defined as

$$IV = \int_0^1 \sigma_u^2 du$$

This rules out using the popular realized volatility measure which estimates the total quadratic variation of a semimartingale, including the contribution from (squared) jumps.⁵

2.1. Multipower Variation (MPV) Measures

The initial, and by far most widely used, estimator of IV in the presence of jumps is the bipower variation (BV) measure of Barndorff-Nielsen and Shephard (2004). It can be shown to be consistent for IV in the absence of market microstructure noise but under otherwise very general conditions. It is given as,

$$BV_N = \frac{\pi}{2} \left(\frac{N}{N-1} \right) \sum_{i=1}^{N-1} |\Delta Y_i| |\Delta Y_{i+1}|. \quad (2)$$

The intuition for the consistency and jump robustness of the BV estimator is straightforward: If $\Delta Y_i, \Delta Y_{i+1} \sim i.i.d.N(0, \frac{\sigma^2}{N})$ then $\mathbb{E}[|\Delta Y_i| |\Delta Y_{i+1}|] = \frac{2}{\pi} \frac{\sigma^2}{N}$, and $\left(\frac{N}{N-1} \right)$ is a required finite sample correction factor. As such, each term of the bipower variation measure delivers an unbiased estimate of the underlying local (spot) variance. Moreover, asymptotically, as the returns span near infinitesimal intervals, there will at most be a single jump within two adjacent intervals. This isolated jump will be dampened by the multiplication by a small adjacent (diffusive) return of order $(\frac{1}{\sqrt{N}})$. As N grows this is sufficient to render the jump contribution asymptotically negligible. Nonetheless, in practical applications there will be an upward (finite sample) bias due to large jumps as the adjacent return is not vanishing, but fixed in size, reflecting the underlying choice of sampling frequency. The latter is typically governed by market conventions as well as the liquidity and microstructure features of the market. Another drawback of the BV estimator is that the jumps only vanish at the rate of \sqrt{N} which is not sufficient to deliver a continuous-record central limit theory in the presence of jumps in the price path. The desire to obtain an operational asymptotic theory under jump alternatives was a major reason for the introduction of MPV statistics which are analyzed thoroughly in, e.g., Barndorff-Nielsen et al. (2006c).

⁵See, e.g., Andersen et al. (2009), Bandi and Russell (2007), Barndorff-Nielsen and Shephard (2007), McAleer and Medeiros (2008) for surveys of the realized volatility literature.

The requisite extension of the bipower variation statistic is particularly intuitive in the case of equally spaced sampling, i.e., $\Delta t = t_i - t_{i-1} = 1/N$, for all $i = 1, \dots, N$. The class of MPV statistics is then defined via the cumulative sum of m products of adjacent absolute returns raised to the (r/m) 'th order, where m is a positive integer and r a positive real number, usually an integer. Hence, the cumulative power of the adjacent products equals r . These statistics provide consistent estimators for the corresponding integrated power of volatility,

$$MPV_N(m; r) = d_{m,r} \left(\frac{N}{N-m+1} \right) (N)^{r/2-1} \sum_{i=1}^{N-m+1} |\Delta Y_i|^{\frac{r}{m}} \dots |\Delta Y_{i+m-1}|^{\frac{r}{m}} \xrightarrow{P} \int_0^1 \sigma_u^r du \quad (3)$$

where $d_{m,r}$ is a known constant dependent only on m and r , while $\left(\frac{N}{N-m+1} \right)$ is a finite sample correction factor.⁶ If the adjacent returns are i.i.d. Gaussian, each summand in (3) delivers an unbiased estimate of the power of spot volatility. The sum therefore provides a (converging) Riemann approximation to the integrated power of the volatility process.

This MPV measure generalizes the entire first generation of estimators in the realized volatility literature, as one obtains the standard realized volatility measure as $RV_N = MPV_N(1; 2)$, while $BV_N = MPV_N(2; 2)$, and additional oft-applied measures include the tripower variation $TPV_N = MPV_N(3; 2)$ the quadpower variation $QPV_N = MPV_N(4; 2)$, and the fourth order power variation $PV_N(4) = MPV_N(1; 4)$. In the presence of a finite activity jump process, the RV estimator is not consistent for the integrated variance, the BV statistic is consistent but does not allow for an asymptotic theory under the jump alternative, while the realized tripower and quadpower measures both provide consistency and allow an associated asymptotic mixed normal limit theory. This property is maintained for $MPV_N(m; 2)$ for $m \geq 3$. Likewise, the fourth order power variation is consistent for the integrated fourth power of the volatility process, the so-called integrated quarticity, but allows for an asymptotic theory only in the absence of jumps. Robust alternatives, which provide both consistency and asymptotic theory under finite activity jumps, are given by $MPV_N(m; 4)$ for $m \geq 5$.⁷

The existence of numerous estimators begs the question of which one is preferable. Not surprisingly, this cannot be answered in general. However, using the ideal setting of no microstructure noise, near infinitely frequent sampling and no jumps, the $MPV_N(m; r)$ measure of lowest order in m delivering the desired feature, whether consistency or a mixed normal limit theory, is the more efficient estimator. Specifically, for estimating the integrated variance in the absence of jumps, the realized volatility estimator is most efficient. Analogously, bipower variation is the preferred consistent jump-robust estimator for IV while tripower variation is the estimator with minimal asymptotic variance among this class which allows for the development of an asymptotic theory under the jump alternative.

Of course, the frictionless setting is not representative of actual market conditions. In particular, various market features limit the sampling frequency so that the impact of jumps cannot be fully neutralized. Thus, it may be desirable to apply higher order MPV measures as they provide better (finite sample) dampening of the jump component. In fact, Veraart (2008) finds from an extensive simulation exercise that the finite sample jump distortion is sufficiently influential to render $MPV_N(10; 2)$ and $MPV_N(10; 4)$ preferable to lower order MPV statistics. One caveat is that the simulations assume a very smooth evolution of the diffusive volatility process and that

⁶In the case of Gaussian i.i.d. price changes $d_{m,r} = \mu_r^{-m}$, where $\mu_p = E|U|^p = 2^{p/2} \frac{\Gamma(\frac{1}{2}(p+1))}{\Gamma(\frac{1}{2})}$, $U \sim N(0, 1)$, see, e.g., Barndorff-Nielsen and Shephard (2004) and Barndorff-Nielsen et al. (2006a).

⁷See Barndorff-Nielsen et al. (2006c) for the behavior of MPV estimators under both the no-jump null hypothesis and under the jump alternative in more general scenarios.

equally spaced ultra high-frequency returns are available at near arbitrary sampling intervals. These features render volatility near constant across sequences of ten adjacent high-frequency returns, thus ensuring that the returns within each block truly are approximately i.i.d. Gaussian distributed. In reality, however, most return volatility series feature a pronounced U-shape across the trading day, resulting in sharp movements in volatility over short periods of time. Moreover, fresh quote or trade observations are often not available at the very highest sampling frequencies as both trade and quote intensities undergo significant intraday fluctuations as well. The result is that a block of adjacent (non-stale) price or quote observations often spans a non-trivial time interval, and thus the assumption of homogenous returns and, in particular, constant volatility is questionable. Moreover, the extent of this problem is proportional to the number of adjacent returns used by a given estimator. As such, it may be important to explore the impact of the “locality” of the estimator in practical applications. We present evidence from both actual equity data and an extensive simulation design that this feature, indeed, is a major determinant of the performance of such estimators.

In summary, the class of MPV measures embodies a tradeoff between efficiency and locality on the one hand and jump robustness on the other. This motivates our introduction of alternative estimators based on nearest neighbor truncation that retain the local nature of the bipower and tripower variation measures while providing better finite sample robustness and asymptotic efficiency than any MPV estimator allowing for a feasible asymptotic theory in the presence of (finite activity) jumps.

2.2. The MinRV and MedRV Estimators

We propose the following MinRV and MedRV estimators of integrated variance,

$$\text{MinRV}_N = \frac{\pi}{\pi - 2} \left(\frac{N}{N - 1} \right) \sum_{i=1}^{N-1} \min (|\Delta Y_i|, |\Delta Y_{i+1}|)^2 \quad (4)$$

$$\text{MedRV}_N = \frac{\pi}{6 - 4\sqrt{3} + \pi} \left(\frac{N}{N - 2} \right) \sum_{i=2}^{N-1} \text{med} (|\Delta Y_{i-1}|, |\Delta Y_i|, |\Delta Y_{i+1}|)^2$$

The scaling factors ensure that every summand on the right hand side of equation (4) provides an unbiased estimate of the underlying spot variance if the corresponding block of returns are i.i.d. Gaussian.⁸ The robustness of the MinRV and MedRV estimators compared, e.g., to the BV estimator in equation (2) stems from the fact that returns contaminated by a large jump are fully eliminated by the min/med operators. That is, if a (large) jump occurs within one of the two terms for the MinRV estimator, the min operator simply squares the adjacent (small) diffusive return. There will be an induced bias as we now effectively compute the square of a single return rather than of the minimum of two independent returns, but the bias is an order of magnitude smaller than for bipower. For illustration, assume there is a (large) price jump of size Δj_i in the interval $[t_{i-1}; t_i]$ but no jump in the adjacent intervals so that $|\Delta Y_{i-1}| \ll |\Delta j_i|$ and $|\Delta Y_{i+1}| \ll |\Delta j_i|$. The distortion to the overall BV from the interval containing the jump return clearly dominates the corresponding maximal distortion to the MinRV measure, i.e.,

$$\frac{\pi}{2} |\Delta j_i| (|\Delta Y_{i-1}| + |\Delta Y_{i+1}|) \gg \frac{\pi}{\pi - 2} [|\Delta Y_{i-1}|^2 + |\Delta Y_{i+1}|^2] \quad (5)$$

⁸The derivation of these scaling factors is a simple exercise in integration and is available from the authors upon request.

as the left hand side is of order $(1/\sqrt{N})$ versus $(1/N)$ on the right hand side. More generally, the upward bias due to jumps for any MPV statistic, $MPV_N(m; 2)$, $m \geq 2$, is of order $(1/N^{1-\frac{1}{m}})$, thus approaching $(1/N)$ from above for large m . The expression (5) also reflects the important fact that only the number of jumps, *not* their actual size, influences the bias of the MinRV and MedRV estimators. Hence, by construction, these measures provide better finite sample jump robustness than the MPV statistics.⁹

As indicated, the MinRV and MedRV estimators are consistent for the integrated variance.

Proposition 1 *Let the log-price process Y_t be given by the jump-diffusion (1) with finite jump activity. Assume further that μ_t is adapted and locally bounded, σ_t is adapted, cadlag and $\inf_{t \geq 0} \sigma_t > 0$ almost surely. Then we have, as $N \rightarrow \infty$,*

$$\text{MinRV}_N \xrightarrow{P} \int_0^1 \sigma_u^2 du \quad \text{and} \quad \text{MedRV}_N \xrightarrow{P} \int_0^1 \sigma_u^2 du$$

Under slightly stronger assumptions about the volatility process we obtain a corresponding asymptotic distribution theory.

Proposition 2 *Let the log-price process Y_t be given by the jump-diffusion (1) with finite jump activity. Assume further that μ_t is adapted and locally bounded, σ_t is bounded away from zero almost surely and follows an Ito process of the form (A1) given in the appendix, then*

$$\begin{aligned} \sqrt{N}(\text{MinRV}_N - IV) &\xrightarrow{\text{stable } \mathcal{D}} MN \left(0, 3.81 \int_0^1 \sigma_u^4 du \right) \\ \sqrt{N}(\text{MedRV}_N - IV) &\xrightarrow{\text{stable } \mathcal{D}} MN \left(0, 2.96 \int_0^1 \sigma_u^4 du \right) \end{aligned}$$

Moreover, these results remain valid for finite as well as infinite activity jumps in the volatility process subject only to the weak regularity conditions of Barndorff-Nielsen et al. (2006a), stipulating only that the resulting general Ito semimartingale, extending Assumption (A1), has jump characteristics that are locally bounded.

The distributional convergence is stable and the notation MN denotes a mixed Gaussian limiting distribution, i.e., a normal distribution conditional on the realization of the integrated quarticity, $IQ = \int_0^1 \sigma_u^4 du$, where, importantly, the limiting normal variate is independent of the (random) IQ process.

Proposition 2 resembles existing limit theories for MPV statistics in scenarios without jumps. In this case, MinRV and MedRV are less efficient than the optimal MPV statistic (bipower). Of course, the advantage is the near elimination of finite sample jump distortion.

Arguably, it is more meaningful to compare the MinRV and MedRV estimators to higher order MPV measures as the latter also allow for a feasible asymptotic theory in the presence of jumps. Table C.1 tabulates the relevant efficiency (asymptotic variance) factors. MinRV is on par with the sixth order MPV estimator, while MedRV is more efficient than even the best (tripower) estimator within this set of MPV measures.

⁹Likewise, the MinRV and MedRV estimators are simpler than the truncated RV estimator of Mancini (2006) which annihilates the impact of jumps beyond a pre-specified threshold. The choice of threshold can be delicate when (latent) volatility is time varying. Promising progress, in terms of practically implementing this style of estimator, is reported in concurrent work by Corsi et al. (2010) using a combination of multipower and threshold techniques. The performance of the latter approach in comparison to various neighbor-truncation estimators is thoroughly investigated in Andersen et al. (2011).

To further illustrate the relationship between the estimators, we derive the joint asymptotic distribution between the MinRV, MedRV and the standard RV, BV and TV estimators under the no-jump null:

Proposition 3 *Assume the conditions of Proposition 2 hold. If jumps are absent from the price process, then the following joint asymptotic stable distributional result holds,*

$$\sqrt{N} \begin{bmatrix} RV_N - IV \\ BV_N - IV \\ TV_N - IV \\ MinRV_N - IV \\ MedRV_N - IV \end{bmatrix} \xrightarrow{\text{stable } \mathcal{D}} MN \left(0, \begin{bmatrix} 2 & 2 & 2 & 2 & 2 \\ & 2.61 & 2.74 & 2.98 & 2.53 \\ & & 3.06 & 3.14 & 2.61 \\ & & & 3.81 & 3.05 \\ & & & & 2.96 \end{bmatrix} IQ \right)$$

If there is finite jump activity in the price process, then the distributional result associated with the lower trivariate system, involving the TV, MinRV and MedRV measures, remain valid.

The proposition implies that both MinRV and MedRV are highly correlated with BV and TV as well as with each other.¹⁰ This suggests that the scope for applying a jump-robust GMM procedure by combining BV, TV, MinRV and MedRV estimators may be somewhat limited. More importantly, Proposition 3 can be used to construct jump tests based on MedRV or MinRV analogous to those proposed by Barndorff-Nielsen and Shephard (2004) based on BV. For example, it follows that the jump test statistic given by the difference $RV_N - MedRV_N$ has the following asymptotic distribution:¹¹

$$\sqrt{N}(RV_N - MedRV_N) \xrightarrow{\text{stable } \mathcal{D}} MN(0, 0.96 IQ) \quad (6)$$

To recapitulate, Proposition 2 implies that feasible inference regarding the realized latent integrated variance is possible based on the MinRV and MedRV estimators, even in the presence of finitely many price jumps. Proposition 3 implies that MinRV and MedRV can also be used for feasible inference regarding the realized latent jump variance. However, as for the MPV measures, this requires a consistent estimator for the integrated quarticity. Such jump-robust estimators are readily constructed from higher order MPV statistics. An alternative is to construct an estimator for the quarticity in a direct extension of the principles behind the MinRV and MedRV estimators:

$$\begin{aligned} \text{MinRQ}_N &= \frac{\pi N}{3\pi - 8} \left(\frac{N}{N-1} \right) \sum_{i=1}^{N-1} \min(|\Delta Y_i|, |\Delta Y_{i+1}|)^4 \\ \text{MedRQ}_N &= \frac{3\pi N}{9\pi + 72 - 52\sqrt{3}} \left(\frac{N}{N-2} \right) \sum_{i=2}^{N-1} \text{med}(|\Delta Y_{i-1}|, |\Delta Y_i|, |\Delta Y_{i+1}|)^4 \end{aligned}$$

In fact, the asymptotic theory is entirely analogous and results similar to Propositions 1-2 hold for these quarticity estimators.¹²

¹⁰The implied correlations between IV estimators are: $\text{Corr}(BV, \text{MinRV}) = 94.5\%$, $\text{Corr}(BV, \text{MedRV}) = 91.0\%$, $\text{Corr}(TV, \text{MinRV}) = 92.0\%$, $\text{Corr}(TV, \text{MedRV}) = 86.7\%$, $\text{Corr}(\text{MinRV}, \text{MedRV}) = 92.0\%$.

¹¹Likewise, it is straightforward to derive the asymptotic distribution of the corresponding ratio- and log-based version of this jump test via the so-called “delta-rule.” These tests tend to be better behaved in finite samples, see, e.g., the discussion in Huang and Tauchen (2005).

¹²The theory and practice of IQ estimation is explored in Andersen et al. (2011), generalizing the neighborhood truncation approach in directions that facilitate practical robust and efficient IQ estimation.

3. Finite Sample Evidence - Dow Jones 30 Stocks

In this section we gauge the empirical performance of the MinRV and MedRV estimators on the set of Dow Jones 30 stocks using NYSE/TAQ data from January 1, 2005 through May 31, 2007. Ignoring days with abbreviated trading hours around major holidays we obtain a sample of 601 days with intraday observations based on mid-quotes. Prior to analysis, we remove bid-ask quote pairs that represent spread outliers (less than 0.1% of the data) and then apply a standard error filter advocated by Barndorff-Nielsen et al. (2009) for raw tick data. Finally, we invoke a very mild functional filtering procedure, introduced in Andersen et al. (2011), for the jump-robust estimators that guards against dramatic violations of the null hypothesis of iid Gaussian price innovations over the local return block as manifested in extreme outliers for the functional values associated with the IV estimate. Since the outlier filter is applied at a nominal significance level of 10^{-6} , under the null hypothesis, it typically removes at most a single intraday observation for each stock across the entire sample. Consequently, this functional filter is helpful in eliminating gross distortions caused by highly irregular observations that survive the mild initial filtering of the raw data, thereby providing additional robustification without any noticeable impact on the estimates in the absence of violations of the local null hypothesis.

We compare the finite sample efficiency of sub-sampled MinRV and MedRV to sub-sampled MPV measures. Sub-sampling is a simple way of enhancing the efficiency of an estimator and was originally proposed for RV by Zhang et al. (2005). It involves taking the average of an estimator across all possible sub-samples, at a given sampling frequency, obtained by starting at different offsets and scaling up to match the full length of the trading day. For comparison, we also consider a sub-sampled version of the recently developed Quantile RV (QRV) estimator of Christensen et al. (2010). There is little evidence regarding the preferred construction of QRV, so we follow their specific implementation.¹³

To keep the exposition manageable we present results for daily measures averaged over the full sample and across all stocks only. This conveys the systematic differences among estimators while reducing the impact of idiosyncratic features of individual stocks or specific days.

3.1. “Tick” Time and Calendar Time Sampling

An important aspect of the implementation of any IV estimator is whether to sample in “tick” or “calendar” time. These schemes represent alternative extreme views of the dependence between observation times and price moves. Tick time sampling can be justified, from standard limiting theory, when volatility is constant in tick time, i.e., the quadratic variation of the price process is perfectly correlated with the observation times. Calendar time sampling, on the other hand, can be justified when observation times are exogenous to the price process, i.e., inference may be carried out conditional on the observation times. In particular, correlation between

¹³QRV exploits the fact that, in a large block of zero mean i.i.d. Gaussian returns, the largest absolute returns contain the most information about the variance. Hence, QRV exploits fairly extreme quantiles, necessitating the use of long blocks. We follow Christensen et al. (2010) in optimally weighting the quantiles {0.05, 0.10, 0.15, 0.20, 0.80, 0.85, 0.90, 0.95} on rolling blocks of twenty returns based on the finite sample values of the associated scaling factors, exploiting sample sizes up to 23,400 (one per second):

$$QRV = \frac{1}{\nu} \sum_{\lambda \in \{0.80, 0.85, 0.9, 0.95\}} \omega_{\lambda} [g^2(\lambda) + g^2(1 - \lambda)]$$

where $g(\lambda)$ is the order statistic of the λ^{th} return percentile and ω_{λ} provides the optimal weighting. The constant ν is a normalizing constant whose (finite sample) value must be determined by numerical integration or simulation. Importantly, it is not possible to reliably interpolate the QRV scaling factors across nearby sample sizes because they display a pronounced oscillation.

quote/trade arrivals and changes in volatility is ruled out.¹⁴ Neither assumption is likely to be fully accurate, so we study our estimators under both sampling schemes.

Figure C.1 presents a set of “signature plots,” depicting the average values of the various integrated variance measures across stocks and trading days as a function of the (effective) sampling frequency.¹⁵ Ideally, all the jump-robust measures, obtained via distinct procedures and sampling frequencies, should coincide and form a horizontal line, reflecting the average value of the integrated variance across the stocks in the sample. However, as is evident from Figure C.1, nonlinear and non-overlapping signature plots are commonplace. This indicates the presence of biases, stemming from different sensitivities among the estimators to market microstructure noise and other finite sample distortions, including jumps. This raises the question of what combination of estimator and sampling frequency ensures reasonably unbiased measures.

Focusing initially on calendar versus tick time, we note from Figure C.1, Panels 1.A and 2.A, that the measures generally produce similar values and patterns in the IV estimates as a function of frequency under the alternative sampling schemes although there is an improvement in alleviating the extent of the dip at the very highest frequencies from tick based sampling. Once we reach 90 seconds, the plots are nearly identical for most estimators. In particular, the tendency towards producing an inverted U-shape, rather than a flat line, across the frequency range is apparent for both sampling schemes. The lower volatility estimates for calendar time sampling at high frequencies is consistent with the presence of additional zero returns, arising from the application of the “previous tick rule” in calendar time. This rule stipulates that, for any interval without a quote arrival, one uses the last previously observed mid-quote as the end-of-interval price, thus mechanically producing zero returns in these instances. Table C.2 below shows that this effect is likely to be substantial at the highest frequencies.

A couple of other features are striking. First, the RV estimator attains a higher and more stable value than the others. Since RV, per construction, includes the contribution from jumps, the discrepancy between RV and the other measures provides an indication of the overall jump contribution to the return variation across the Dow Jones stocks. The jump-robust IV estimators appear more susceptible to distortions as they fluctuate more significantly across the frequencies, with the pronounced dip for the shortest return intervals being the most glaring.¹⁶ Almost universally, the maximum value is obtained around the two-minute frequency. Our subsequent explorations do, in fact, document fairly significant downward biases at the highest and lowest frequencies, suggesting that IV estimates based on sampling in the range of 90-180 seconds are the least biased. In particular, we study the effects of stale quotes, unevenly spaced data, iid microstructure noise and strong intraday volatility patterns in controlled Monte Carlo experiments below. Second, the jump-robust IV estimators differ substantially from each other: while MinRV, MedRV and BV lead to roughly similar average estimates the tripower (TV) and quantile RV (QRV) estimators yield much lower average values across many of the frequencies. The analysis in the following sections sheds further light on the reasons for this disparate behavior of the alternative IV measures.

¹⁴In practice, the issue with calendar time sampling is greatly alleviated if one avoids sampling at ultra-high frequencies. For example, the effect appears negligible for actively traded securities if volatility is computed from one-minute returns. Nonetheless, much current research aims to reduce the impact of noise sufficiently that even the highest sampling frequencies may be used productively. As such, the dependence between observation times (quotes or trades) and the price process is an intriguing area for research; see, e.g., Mykland et al. (2008).

¹⁵This graphical tool was introduced in Andersen et al. (2000).

¹⁶Chaboud et al. (2007) document a similar dip of the bipower variation estimates of IV for foreign exchange and Treasury bond data.

3.2. Stale Quotes

Table C.2 summarizes basic descriptive statistics for the Dow Jones components, including all companies that were part of the index during January 1, 2005 and May 31, 2007. These stocks are generally very actively traded, with an average of one new quote arrival every 2 seconds throughout the normal trading hours from 9:30am to 4:00pm ET. However, a significant number of these quotes simply repeat previous quotes, as reported in the first two columns of the table: the average number of non-duplicate quotes is only about one quarter of the recorded quotes. Remarkably, the two most actively traded stocks (INTC and MSFT) have the smallest number of non-duplicate quotes, as also observed by Griffin and Oomen (2008). As a consequence, tick-time sampled estimators may well behave differently if we censor the data for duplicates. Furthermore, the last column of the table indicates that the duplicate quotes are clustered as most stocks experience long spells of repeated quotes. The median number of zero returns lasting longer than 30 seconds on each day is 161, those lasting longer than 1 minute are about 48 on each day, and those lasting longer than 2 minutes are about 11 per day. Importantly, this pattern is observed without much variation across stocks (reported) or trading days (not reported). The effect of censoring for duplicate quotes on tick-time sampled estimators is shown in Figure C.1, Panel 3.A. For all estimators, except QRV, the dip at higher frequencies is substantially alleviated, and the overall level of the jump-robust estimators is uniformly higher. We also note that the TV measure starts behaving like the majority of the other IV estimators, leaving QRV as a relative outlier. These effects are not unexpected as, by removing long sequences of duplicate quotes, we dramatically reduce the number of high-frequency returns and alter the tick-time sampling scheme to feature fewer zero returns, even for the medium frequencies. Hence, while the impact at the highest frequencies may reflect an alleviation of a strong downward bias induced by zero returns, the lower frequency results are likely upward biased due to the systematic removal of periods with flat pricing. Thus, while it is encouraging to observe the anticipated response to the censoring of zero quote changes, we deem the standard tick-time sampling less likely to be biased for the moderate frequencies in the range of 1-3 minutes.¹⁷

3.3. Microstructure Noise

The prevalent way of dealing with microstructure noise is to sample at a lower frequency than the available data (e.g., 2 minutes) to gain noise robustness and then compensate for the efficiency loss by sub-sampling the estimator. A complementary approach termed “pre-averaging,” recently introduced by Podolskij and Vetter (2009), exploits the data at the highest frequency available, but uses local “pre-averaging” via a kernel function to produce a set of non-overlapping (asymptotically) noise free (in reality, noise reduced) observations to which standard IV estimators may be applied. In practice, this necessitates a choice of bandwidth (and kernel function) and leads to a familiar bias-variance trade-off which we study for our set of estimators. To render the results comparable, we ensure that the window length used in constructing pre-averaged returns matches the frequencies applied to compute the corresponding estimators without pre-averaging.¹⁸ Finally, note that throughout we sub-sample the pre-averaged estimators to enhance efficiency, as we do for all estimators without pre-averaging.

¹⁷The censoring of zero returns has been adopted in, e.g., Ait-Sahalia et al. (2011) and Hautsch and Podolskij (2010), while it is explored in some detail by Griffin and Oomen (2008). Phillips and Yu (2008) provide an explicit model of flat pricing and analyze the implications for the asymptotic distribution of RV.

¹⁸In particular, note that if the pre-averaged log-prices, for a window of K observations, are defined as $\bar{Y}_{t_i} = \frac{1}{K} \sum_{j=0}^{K-1} Y_{t_{i+j}}$, $i = 0, 1, \dots, N-K+1$, then the associated pre-averaged returns, $\bar{Y}_{t_{i+K}} - \bar{Y}_{t_i}$, cover $2K$ terms of the original log-price series Y_{t_i} , $i = 0, 1, \dots, N$. Since price averaging induces return serial correlation corresponding

For brevity, we subsequently refer to estimators that are pre-averaged and sub-sampled as “pre-averaged” estimators. The combined pre-averaging and sub-sampling procedure is described in detail in Appendix B.

The pre-averaged (and sub-sampled) estimators are displayed in Figure C.1, Panels 1.B-3.B. Overall, they mirror the qualitative behavior of the original measures in Panels 1.A-3.A, although the estimated level of volatility is notably lower at the highest frequencies, where the tripower variation measure continues to be strongly downward biased. The drop in the estimated volatility level is smaller for the medium range of frequencies, around 90-180 seconds, as revealed by a close inspection of Panels 1.B-2.B versus Panels 1.A-2.A. This is consistent with the pre-averaging procedure eliminating a bias stemming from a mild degree of “bouncing” of the quote midpoints around the underlying efficient price, even if the estimators simultaneously are subject to an opposite, downward bias arising from the presence of an excessive number of zero returns at the highest frequencies. Finally, we note again that the QRV measure delivers decidedly lower values over the medium to low frequencies. The simulation section below exemplifies some of the features that may rationalize this distinctive behavior of QRV.

3.4. Robustness across Volatility Regimes

Based on newly available NYSE/TAQ data, we consider a second shorter sample covering June 1, 2007 to July 31, 2009. Ignoring abbreviated trading days around major holidays we obtain a sample of 540 days for this period. It allows us to consider the robustness and relative performance of the estimators across regimes as the volatility is significantly higher during the more recent period.

The resulting signature plots are displayed in Figure C.2. Along with the much higher estimates of IV there is also a wider gap between RV and the jump-robust measures, suggesting more significant jump activity in this period. The relative behavior of the individual estimators, however, follows roughly the same pattern as in the first sample, with notably lower estimates of QRV and TV compared to MinRV, MedRV, and BV. Specifically, Panels 2.A and 2.B indicate that the volatility estimates now are consistent for an even wider range of frequencies spanning 1-5 minutes. This may reflect a general reduction in the noise-to-signal ratio over time due to the growing quantity and quality of high-frequency quotations and diminished bid-ask spreads. Panels 3.A and 3.B may further suggest that, for this period, it is beneficial to censor for duplicates for the sake of producing volatility estimates that are consistent across a broad set of frequencies and pre-averaging windows. However, the flat curves at the highest frequencies are partly an artifact of the low number of non-duplicate observations.¹⁹ Whatever the merits of duplicate censoring, it is encouraging that the evidence is consistent across large sets of alternative sampling frequencies and estimators. Overall, the behavior of the IV estimators appears robust to the volatility regime as the qualitative differences observed over the prior sample remain intact.

to the window length, the pre-averaged return series are computed from non-overlapping blocks of $2K$ underlying prices. Therefore, we calibrate the window size, K , for each pre-averaged IV estimator so that the sampling step of the corresponding estimator without pre-averaging equals $2K$.

¹⁹For many day-stock combinations we simply do not have enough non-zero returns to create a complete tick-time sampled return series at the highest frequencies. In those cases, we defer to the values obtained from the highest frequency that produce a viable return series. Consequently, there is a mechanical duplication of estimates from the lower frequencies which biases the averages of the high frequency values towards those obtained at the medium frequencies.

3.5. Pervasiveness of Intraday Volatility Patterns

One potential advantage of tick-time sampling of quote midpoints is the ability to dampen the impact of intra-day variations in volatility to the extent this variation is correlated with quote arrivals. If tick-time sampling succeeds in compensating for systematic volatility movements, then the diurnal U-shape pattern in volatility cannot be the source of any downward bias in IV estimates at low frequencies across sampling schemes. Hence, we take a look at the intraday volatility patterns stemming from alternative sampling procedures.

In the first column of Figure C.3 (Panels 1.A-3.A) we plot the average diurnal variance factor for the DJ30 sample. Several facts emerge. First, the intra-day U-shape is very pronounced with volatility clearly elevated during the first hour of trading, and this finding is robust across RV and IV measures.²⁰ Second, tick-time sampling does *not* alleviate the volatility pattern. On the contrary, the relative variance factor at the open increases from about 6 for calendar time sampling to more than 10 for tick-time sampling. Moreover this effect is independent of whether the tick time sampling scheme uses *all* available ticks or first brings the data to a one second grid to guard against inaccuracies in the intra-second quote or trade chronology. Third, the very pronounced spike at the open is qualitatively consistent with our observations based on the signature plots in Figures C.1-C.2.

Although DJ30 stocks are heavily quoted and traded, they are illiquid by the standards of the most active derivatives markets. To elucidate the relationship between market structure/liquidity and the efficacy of tick-time sampling, the second column of Figure C.3, Panels 1.B-3.B, displays the diurnal variance factor for an extraordinarily active derivatives market, namely the S&P500 E-mini futures contract traded on the electronic Globex platform of the CME Group. Since trading begins Sunday through Thursday at 5pm local time (labeled Central Time, or CT) and lasts until 3:15pm CT the following day,²¹ the trading day is 1,335 minutes versus 390 for the individual stocks. The diurnal variance factor is calculated using the same method as for the DJ30 stocks.

The top right Panel 1.B conveys a dramatic diurnal pattern in calendar time. In particular, there are spikes at the release of regularly scheduled macroeconomic announcements at 7:30am (bucket number 870) and 9am CT (bucket number 960). Not surprisingly, the equity index is more sensitive to these events than individual stocks.²² Moreover, the U.S. trading day from 8:30am-3:00pm CT (buckets 930-1320) accounts for the vast majority of the daily return variation with the jump at 8:30am CT being particularly apparent. During the U.S. stock market trading hours, an approximate U-shape pattern is visible and, disregarding the last 15 minutes of trading following the closure of the U.S. equity markets, the E-mini volatility pattern seems qualitatively to match the average pattern from the individual stocks quite well.

Much more dramatic differences arise in tick time. Panel 2.B of Figure C.3 shows the result of tick time sampling after bringing observations to a one-second grid. The dramatic variation in volatility that was apparent in calendar time is almost entirely eliminated except for a slight hump around bucket number 430 which, for the median day, occurs around 8:30am CT and therefore matches the elevated volatility associated with the opening of the stock market. The ultra-high frequency of the futures market means that there is a noticeable difference when going to tick time sampling using all ticks (Panel 3.B), where even less volatility variation is visible.

²⁰For legibility only RV and MedRV is shown in Figure C.3.

²¹There is also a minor trading window at 3:30pm-4:30pm CT on Monday through Thursday but this period is not very liquid and typically not considered part of the official trading day.

²²Since individual stocks do not trade at 7:30am CT, the impact of those announcements will be largely absorbed into the overnight price movement, but there is likely to be some spill-over in the form of a more intense price discovery process following the market opening on such announcement days.

The only sign of elevated volatility occurs prior to bucket number 200 which, for the average day, corresponds to 8:30am CT and therefore is consistent with our findings above.

Thus, while the tick time sampling schemes are ineffective in rendering the average intraday volatility approximately constant for our application to quotes on the DJ30 stocks, this is not necessarily true for all markets. In particular, consistent with the prior literature, we do find that tick time sampling is helpful in rendering the average intraday returns nearly homogeneous for active index derivatives contracts, even if some of the market opening effects remain visible. However, beyond the diurnal volatility patterns, there are also unexpected random volatility shocks during the course of trading, and for these it is less likely that tick time sampling will be sufficient to homogenize the volatility series. As such, our findings support the intuition that estimators based on short blocks of returns will be beneficial in alleviating any potential bias that may be induced by a strong return variation over short time horizons. Our simulation results below reinforce these conclusions further.

3.6. Common Features Uncovered from the Empirical Study

We summarize some of the general features uncovered by our empirical study. First, the jump-robust estimators display an inverted U-shape as a function of the sampling frequency, with the maximum average volatility level obtained within the 1-5 minute range. The drop in the estimators at higher frequencies are particularly striking for some estimators over the initial sample period. Second, the tripower variation estimator displays the most severe bias at the highest frequencies, while the QRV estimator appears dramatically downward biased at the medium to lower frequencies. Third, our high-frequency return series obtained from quote mid-points display no sign of the extreme upward bias at the highest frequencies that is associated with i.i.d. type noise and commonly observed in estimators based on trade prices. Fourth, moving from calendar to tick-time sampling improves the coherence across the alternative measures. Fifth, moving from sub-sampled to pre-averaged measures produces almost identical average volatility estimates over the medium frequency range while the drop at the highest frequencies is larger and the drop at the lowest frequencies slightly less pronounced. Sixth, the qualitative behavior is similar across the two samples, except for the smaller drop in the volatility estimates at the highest frequencies and the broader range of roughly stable estimates for the medium frequencies in the latter period. Seventh, across the preferred range of frequencies and regardless of the sampling scheme, our MinRV and MedRV estimators tend to be below the less jump-robust, and thus likely more upward biased, BV estimator, while staying somewhat above the TV and QRV estimators that are more prone to downward biases due to zero or non-homogeneous returns on larger blocks. Finally, we caution that the findings surely reflect the specific structure and liquidity of the market for actively traded U.S. stocks. The results will be different for less liquid markets, such as those for smaller stocks or corporate bonds, and for more liquid assets, such as very active futures or foreign exchange markets.

It is natural to view the volatility measures for the medium range of frequencies as approximately unbiased, as we observe a better coherence across the alternative IV estimators and a closer resemblance to the pattern displayed by the RV estimator over this range. Moreover, there are a set of plausible explanations for the presence of significant downward biases at both the highest and lowest return frequencies. However, since we do not observe the actual day-by-day realization of IV we cannot directly assess the plausibility of this hypothesis. Consequently, we now turn to a thorough simulation study to shed additional light on the potential sources of bias and the comparative performance of the various estimators.

4. Finite Sample Simulation Evidence

We conduct a series of Monte Carlo experiments focusing on features of the data generating process that may affect the finite sample behavior of the various IV estimators. In particular, we compare the performance of sub-sampled/pre-averaged MinRV and MedRV estimators to sub-sampled/pre-averaged BV, TV, and QRV benchmarks for a set of models embodying distinct features. The emphasis is on the qualitative impact. In reality, such features are all present simultaneously and interact with each other, creating rather complex patterns in tick-by-tick data. Hence, the simulations are not designed to replicate the quantitative magnitude of statistics observed from actual data in all dimensions, but rather to help identify features which may be relevant to explain the systematic patterns in the empirical results.

We consider the following set of models which mostly deviate from the simple i.i.d. Gaussian benchmark in a single dimension to facilitate direct interpretation of the qualitative impact:

- *Model 1*: “BM”. This is our baseline Brownian motion model with sampling on an equi-spaced time grid. It provides the ideal setting under which the finite sample performance of the IV estimators should be closely in line with the underlying asymptotic theory.

- *Model 2*: “SV-U”. This is a stochastic volatility model (two-factor affine) with intraday U-shape volatility pattern and sampling on an equi-spaced time grid. It allows us to isolate potential finite sample biases of the estimators due to time variation in volatility.

- *Model 3*: “BM + Sparsity”. This is a Brownian motion model with sampling on a sparse (exogenously random) time grid. While not necessarily realistic, this model is helpful for studying the potential distortion of the estimators when applied on non-homogeneously sampled returns, effectively inducing spurious variations in their volatility.

- *Model 4*: “BM + 1 Jump”. This is a Brownian motion model with one jump on each day and sampling on an equispaced time grid. It serves to illustrate the degree of finite sample jump robustness of the alternative IV estimators.

- *Model 5*: “BM + 4 Jumps”. This is a Brownian motion model with four jumps on each day and sampling on an equispaced time grid. We use this jump specification to study the impact of multiple (potentially adjacent) jumps.

- *Model 6*: “BM + Noise”. This is a Brownian motion model with sampling on an equispaced time grid subject to i.i.d. microstructure noise. It allows us to shed light on potential distortions due to microstructure noise.

- *Model 7*: “BM + 1 Bounce Back”. This is a Brownian motion model with sampling on an equispaced time grid subject to a single large price outlier resulting in adjacent jumps of opposite sign. It serves to illustrate the effectiveness of suppressing this particular type of noise in the data by standard noise-reduction techniques such as pre-averaging.

4.1. Simulation Design

In each model, the price process $\{Y_t\}$ follows a driftless Brownian motion with instantaneous volatility $\sigma(t)$:

$$dY(t) = \sigma(t) dW_1(t)$$

Across all model specifications, the unconditional IV of each day is calibrated to 0.000159 corresponding to an annualized volatility of 20% (assuming 252 trading days per year). This roughly matches the average level of volatility observed in our DJ 30 sample between January 2005 and May 2007. We also match the average sample frequency of 2 seconds resulting in 11,700 intraday observations. We adopt an equispaced time grid across all models except for Model 3, where we use (exogenously) random sampling.

In Model 2, the stochastic volatility model with intraday U-shape volatility pattern is given by

$$\begin{aligned} dY(t) &= \sigma_u(t) \sigma_{sv}(t) dW_1(t) \\ \sigma_{sv}^2(t) &= \sigma_1^2(t) + \sigma_2^2(t) \\ d\sigma_1^2(t) &= \kappa_1 [\theta_1 - \sigma_1^2(t)] dt + \eta_1 \sigma_1(t) dW_{21}(t) \\ d\sigma_2^2(t) &= \kappa_2 [\theta_2 - \sigma_2^2(t)] dt + \eta_2 \sigma_2(t) dW_{22}(t) \end{aligned}$$

with W_{21}, W_{22} independent and the leverage effect captured by the instantaneous correlations $\rho_1 = \text{corr}(dW_1(t), dW_{21}(t))$ and $\rho_2 = \text{corr}(dW_1(t), dW_{22}(t))$. The two factor model parameters are calibrated in line with Andersen et al. (2005) in percentage form as $\kappa_1 = 0.6$, $\kappa_2 = 0.1$, $\theta_1 = 1.0582$, $\theta_2 = 0.5291$, $\eta_1 = 0.2$, $\eta_2 = 0.1$, $\rho_1 = 0.9$, $\rho_2 = -0.4$. Following Hasbrouck (1999) we model the diurnal volatility U-shape as the sum of two exponentials:

$$\sigma_u(t) = C + A e^{-at} + B e^{-b(1-t)}, \quad t \in [0; 1]$$

where the constants $A = 0.75$, $B = 0.25$, $C = 0.88929198$, $a = 10$, $b = 10$ are calibrated to produce a strong asymmetric U-shape with variance at the open ($t = 0$) more than 3 times the midday variance ($t = 1/2$) and variance at the close about 1.5 times the midday variance.

In Model 3, non-homogeneous sampling is obtained by drawing a random sample of 11,701 points (without replacement) out of the full daily time grid (23,401 seconds from 9:30 am to 4:00 pm). The resulting sample size is identical to the one obtained with 2-second equidistant sampling, but the price observations are subject to (exogenous) random sparsity.

In Models 4-5, price jumps are introduced by extending the log-price process as

$$dY(t) = \sigma(t) dW_1(t) + dJ_t$$

where the Poisson jump process (J_t) is assumed independent of (W_1, W_2). In order to stress test the IV estimators, we calibrate the process to match moderate “jump days” on which one (Model 4) or four (Model 5) Gaussian jumps account for an average increase of 25% in realized volatility (i.e. the jump contribution JV is 25% of IV or, equivalently, 20% of QV = IV + JV, thereby generating many jumps that are not too obvious return outliers).

In Model 6, we simulate Gaussian i.i.d. noise with a moderate noise-to-signal ratio $\lambda = 0.25$, defined as the ratio of annualized error variance to annualized IV. Finally, in Model 7 we consider a particular kind of noise caused by errors in the recorded price that produces so-called “bounce backs”, i.e., the occurrence of errant adjacent return jumps of opposite sign. In order to highlight any bias arising from such data imperfections, we calibrate the model so that the single bounce back accounts for 25% of IV or 20% of QV.

4.2. Simulation Results

For each model specification, we simulate 20,000 trading days, roughly matching the combined stock-day sample size in our empirical study, and tabulate the relative bias and efficiency of the sub-sampled IV estimators at sampling frequencies 12, 60, and 300 seconds (Table 3) as well as their pre-averaged counterparts for pre-averaging windows 12, 60, and 300 seconds (Table 4).²³

²³The relative bias is computed as the sample mean of \widehat{IV}/IV , while the relative efficiency factor at the 60-sec sample frequency is calculated as the sample mean of $390(\widehat{IV} - IV)^2/IQ$, where IV and IQ are the true simulated integrated variance and integrated quarticity on each day, while \widehat{IV} is the IV estimate for the given day. For example, the MSE factor for 60-sec sub-sampled RV is about 1.33, and thus in line with the theoretical value derived in Zhang et al. (2005).

The results for Model 1, Tables 3-4, confirm that, in a frictionless setting with homogeneous returns, all IV estimators are unbiased and their relative efficiency is closely in line with the asymptotic theory. However, under the more realistic scenario of stochastic volatility and diurnal volatility patterns all IV estimators exhibit a pronounced downward bias if sub-sampled or pre-averaged sparsely (bottom panels, Tables 3-4, Model 2). Moreover, QRV remains biased even for relatively high sub-sampling frequencies or small pre-averaging windows (top panels, Tables 3-4, Model 2). In contrast, the remaining IV estimators become essentially unbiased as the sampling frequency grows since they are better equipped to handle intraday volatility fluctuations due to the “locality” achieved by using short rolling blocks of returns. Thus, in this regard, the MinRV/MedRV and MPV measures have a clear finite sample advantage over QRV.

At the same time, all estimators are downward biased at the highest frequencies if the returns are non-homogeneous due to (exogenous) random sparsity of the available observations (top panels, Tables 3-4, Model 3). Notice also that the biases are larger for the pre-averaged than for the sub-sampled estimators, potentially contributing to the relatively lower average volatility estimates at high frequencies for the pre-averaged estimators observed in Panels 2.A-2.B of Figures C.1 and C.2. In the sparsity case, pre-averaging over wider windows or sampling at sparser frequencies (bottom panels, Tables 3-4, Model 3) eliminates the bias for all estimators. On the other hand, the results for the jump scenarios (Models 4-5, Tables 3-4) provide evidence that MinRV, MedRV, and QRV are considerably less biased in the presence of jumps, especially multiple ones, compared to the multipower variation measures. In this scenario, the pre-averaged estimators are less upward biased than the sub-sampled estimators, providing some support for relying on the pre-averaging estimators at the medium to lower frequencies. The sensitivity of the estimators to microstructure noise appears quite similar, and both sub-sampling (middle and bottom panels, Table 3, Model 6) and pre-averaging (middle and bottom panels, Table 4, Model 6) offer a sensible solution for all estimators. In the presence of bounce backs (Model 7), and absent noise correction, all estimators except QRV display large biases for all sampling frequencies (Table 3), as is to be expected. However, importantly, standard noise reduction techniques effectively eliminate any bias stemming from bounce backs (Table 4-5). In particular, pre-averaging removes the bias at any frequency since the procedure averages across the adjacent jumps of opposite sign, while sub-sampling dramatically reduces the bias at all but the highest frequencies as only two of the subsamples will be contaminated by jumps. Hence, for sub-sampling the bias remains at the 12 second frequency (Table 5, Panel 1) while it is non-existent at 60 and 300 seconds (Table 5, Panel 2-3).²⁴

Finally, we note that, matching the pre-averaging window size to the sub-sampling frequency, we obtain a consistently lower MSE of the pre-averaged estimators relative to their sub-sampled counterparts (Table 3 vs Table 4), which is in line with the underlying asymptotic theory.

Our Monte Carlo experiments suggest that the MinRV/MedRV estimators combine the main advantages of the existing IV estimators in finite samples: superior robustness to jumps akin to QRV as well as reasonable robustness to time-varying volatility like the multipower variation measures. Sub-sampling or pre-averaging either too sparsely or too frequently, relative to the

²⁴We conclude that the extremely large biases induced by bounce backs documented in simulations by Christensen et al. (2010) stem exclusively from their failure to apply pre-averaging, even as they detail the benefits of pre-averaging elsewhere in the article. Of course, most bounce backs will also be eliminated by sensible data cleaning (pre-filtering) techniques. Moreover, all MPV estimators as well as MinRV/MedRV can be modified by employing blocks of staggered (“skip-one”) returns rather than adjacent returns in order to alleviate bounce backs, see Huang and Tauchen (2005). The asymptotic properties of such modified estimators are identical to the original ones although, in finite samples, they are less local and thus more susceptible to the downward bias induced by the diurnal volatility pattern.

one-minute frequency, results in downward biases for all IV estimators, but especially QRV, if there are pronounced intraday variations in volatility or instances of spurious sparcity of the price observations. In this regard, the evidence in Tables 3-4 is qualitatively consistent with the signature plots for the Dow Jones 30 stocks in Section 3 (Figures 1 and 3). As such, the estimates obtained around the 2-minute frequency appear most reliable as they avoid both the (downward) biases induced at the highest frequencies by sparcity, noise and rounding and those at lower frequencies due to stochastic volatility and diurnal patterns.

5. Conclusion

We introduce two new estimators of integrated variance based on high-frequency return observations. These MinRV and MedRV estimators rely on nearest neighbor truncation as a novel way to achieve jump robustness, while sharing a number of important features with existing estimators. First, they mirror the traditional RV measure in simply cumulating a sum of squared intraday returns. Second, they resemble the bipower and tripower variation estimators in exploiting two or three adjacent return observations to obtain each summand within the variation measure. An important distinction, however, is that the MinRV and MedRV estimators dampen the effect of jumps at a faster asymptotic rate than any multipower variation (MPV) estimator. Third, they provide an endogenous and locally adaptive truncation of outliers compared to the (asymptotically more efficient) truncation RV and MPV estimators that in practice necessitates careful calibration of the truncation threshold to the latent time-varying volatility level. Fourth, since the minimum and the median correspond to particular quantiles, the MinRV and MedRV estimators are conceptually related to Quantile RV (QRV) estimators although there are several important differences. QRV emphasizes the efficiency gains from using extreme quantiles based on long blocks of observations, thereby sacrificing the locality of the estimator. In contrast, MinRV and MedRV achieve good efficiency properties exploiting very short overlapping blocks. Moreover, the finite sample scaling factors and normalizing constants of MinRV and MedRV are known in closed form and do not require costly numerical evaluation.

The MinRV and MedRV estimators are designed to obtain a number of specific properties. One, the finite sample jump-robustness is, by construction, excellent as influential (large) jumps are systematically discarded. Two, the use of overlapping windows helps extract additional information from the observed data. Three, relying on functionals of short blocks of returns renders the estimators robust to the strong intraday variation in volatility that severely afflicts jump-robust IV estimators exploiting functionals of long return blocks. Four, the adaptive nature of the (implicit) threshold avoids any delicate calibration of the appropriate cut-off level. Five, there is no need for auxiliary procedures to determine a window length or bandwidth. Six, the MedRV estimator is robust also to the presence of spurious isolated zero returns (quote or trade price duplicates). Finally, both MinRV and MedRV are very simple and entail only minor modifications of the popular realized volatility, bipower and tripower variation measures. In particular, the asymptotic results for the estimators can be derived in a manner analogous to the techniques applied in the extant literature. From that perspective, our appendix illustrates how one may recast the features of a new candidate estimator in the format required for application of the established methodology for deriving the asymptotic distribution theory.

The evidence gleaned from our analysis of the Dow Jones 30 stocks as well as the simulation study confirms that the MinRV and MedRV estimators possess excellent jump robustness (only minor upward bias) while, with adequate choice of sampling frequency, they also may be designed to avoid significant downward biases from sparse data and/or zero returns. Moreover, the MedRV measure is theoretically more efficient than all existing jump-robust MPV estimators

which allow for an asymptotic limit theory in the presence of jumps, i.e., tripower and higher order MPV measures. In practice, MinRV and MedRV appear to perform on par with the sub-sampled bipower variation statistic, while clearly improving on the jump-robustness of the latter. Finally, the new estimators dominate the less local QRV estimators as, in practice, the latter suffer significantly from the pronounced intraday pattern in volatility and other factors inducing violations of return homogeneity across longer return blocks. In this regard, our study highlights the benefits of local volatility measures based on short blocks of returns, as this greatly alleviates the biases stemming from rapid variation in intraday volatility, including the pronounced diurnal U-shaped volatility pattern. While not included in this article, one may find a comparison of the performance of our nearest neighbor truncation estimators to recently introduced alternative truncation RV estimators in Andersen et al. (2011).

Finally, we stress that the robustness properties of the estimators studied here provide no substitute for sensible pre-filtering of the data in order to eliminate major errors in recorded prices and time stamps. Furthermore, our findings are based on, and thus support, the application of noise reduction techniques such as sub-sampling or pre-averaging as these procedures result in substantial efficiency gains while effectively eliminating the impact of market microstructure noise at our favored frequencies of 1-3 minutes.

In conclusion, the MinRV and MedRV measures are promising candidates for practical applications involving the estimation of integrated variance due to their combination of reasonable efficiency and good robustness properties. Such estimators may be particularly attractive for estimation and inference in settings where the presence of jumps cannot be ignored. Moreover, it is simple to generalize these estimators to obtain corresponding measures of integrated quarticity, and the associated asymptotic limit theory is readily derived using the proof strategy developed here, see Andersen et al. (2011). In addition, the more general neighborhood truncation estimators introduced in that work should prove useful in sharpening the inference regarding integrated quarticity and in testing for the presence of price jumps.

References

- Aït-Sahalia, Y., Jacod, J., 2007. Volatility estimators for discretely sampled lvy processes. *Annals of Statistics* 35 (1), 355–392.
- Aït-Sahalia, Y., Mykland, P. A., Zhang, L., 2011. Ultra high frequency volatility estimation with dependent microstructure noise. *Journal of Econometrics* 160, 190–203.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., 2009. Parametric and nonparametric volatility measurement. North Holland, in *Handbook of Financial Econometrics*, Yacine Aït-Sahalia, Lars P. Hansen, and Jose A. Scheinkman (Eds.).
- Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2000. Great realizations. *Risk Magazine* 13, 105–108.
- Andersen, T. G., Bollerslev, T., Dobrev, D., 2007. No-arbitrage semi-martingale restrictions for continuous-time volatility models subject to leverage effects, jumps and i.i.d. noise: Theory and testable distributional implications. *Journal of Econometrics* 138 (1), 125–80.
- Andersen, T. G., Bollerslev, T., Meddahi, N., 2005. Correcting the errors: Volatility forecast evaluation using high-frequency data and realized volatilities. *Econometrica* 73 (1), 279–96.
- Andersen, T. G., Dobrev, D. P., Schaumburg, E., 2008. Duration based volatility estimation. Manuscript, Northwestern University.
- Andersen, T. G., Dobrev, D. P., Schaumburg, E., Jun 2011. A functional filtering and neighborhood truncation approach to integrated quarticity estimation. NBER Working Paper (17152).
- Bandi, F. M., Russell, J. R., 2007. Volatility. Elsevier Science, New York, in *Handbook of Financial Engineering*, V. Linetski and J. Birge (Eds.).
- Barndorff-Nielsen, O. E., Graversen, S. E., Jacod, J., Podolskij, M., Shephard, N., 2006a. A central limit theorem for realized power and bipower variations of continuous semimartingales. in *From stochastic calculus to mathematical finance : the Shiryaev Festschrift*, Springer Verlag.
- Barndorff-Nielsen, O. E., Graversen, S. E., Jacod, J., Shephard, N., 2006b. Limit theorems for bipower variation in financial econometrics. *Econometric Theory* 22 (4), 677–719.
- Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., Shephard, N., 2009. Realised kernels in practice: Trades and quotes. *Econometrics Journal* 12, C1–C32.
- Barndorff-Nielsen, O. E., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2 (1), 1–37.
- Barndorff-Nielsen, O. E., Shephard, N., 2007. Variation, Jumps, Market Frictions and High Frequency Data in Financial Econometrics. Cambridge University Press, in *Advances in Economics and Econometrics. Theory and Applications*, Ninth World Congress, R. Blundell, T. Persson, and W. Newey (Eds.).
- Barndorff-Nielsen, O. E., Shephard, N., Winkel, M., 2006c. Limit theorems for multipower variation in the presence of jumps. *Stochastic Processes and Their Applications* 116, 796–806.
- Boudt, K., Croux, C., Laurent, S., 2008. Outlyingness weighted quadratic covariation Working paper.
- Chaboud, A., Chiquoine, B., Hjalmarsson, E., Loretan, M., 2007. Frequency of observation and the estimation of integrated volatility in deep and liquid markets. Board of Governors of the Federal Reserve System, International Finance Discussion Papers.
- Christensen, K., Oomen, R., Podolskij, M., 2010. Realized quantile-based estimation of integrated variance. *Journal of Econometrics* 159, 74–98.
- Christensen, K., Podolskij, M., 2007. Range-based estimation of quadratic variation. Working paper, Ruhr-Universität Bochum.
- Corsi, F., Pirino, D., Ren, R., 2010. Threshold bipower variation and the impact of jumps on volatility forecasting. *Journal of Econometrics* 159, 276–288.
- Dobrev, D., 2007. Capturing volatility from large price moves: generalized range theory and applications. Working paper, Northwestern University.
- Griffin, J. E., Oomen, R. C., 2008. Sampling returns for realized variance calculations: Tick time or transaction time? *Econometric Reviews* 27, 230–253.
- Hasbrouck, J., 1999. The dynamics of discrete bid and ask quotes. *Journal of Finance* 54 (6), 2109–42.
- Hautsch, N., Podolskij, M., 2010. Pre-averaging based estimation of quadratic variation in the presence of noise and jumps: Theory, implementation, and empirical evidence. Working paper, Humboldt University and ETH, Zurich.
- Huang, X., Tauchen, G., 2005. The relative contribution of jumps to total price variance. *Journal of Financial Econometrics* 3 (4), 456–99.
- Jacod, J., Li, Y., Mykland, P., Podolskij, M., Vetter, M., 2009. Mirostructure noise in the continuous case: the pre-averaging approach. *Stochastic Processes and Their Applications* 26, 2803–2831.
- Jacod, J., Shiryaev, A. N., 2003. Limit theorems for stochastic processes, 2nd Edition. *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin; New York.

- Lee, S. S., Mykland, P. A., December 2007. Jumps in financial markets: A new non-parametric test and jump dynamics. *Review of Financial Studies* 20.
- Mancini, C., 2006. Estimating the integrated volatility in stochastic volatility models with levy type jumps. Working paper, University of Firenze.
- McAleer, M., Medeiros, M. C., 2008. Realized volatility: A review. *Econometric Reviews* 27 (1-3), 10–45.
- Mykland, P. A., Renault, E., Zhang, L., 2008. Volatility estimation when trade times and returns are dependent. Working paper University of Chicago and UIC.
- Mykland, P. A., Zhang, L., 2009. Inference for continuous semimartingales observed at high frequency. *Econometrica* 77 (5), 1403–1445.
URL <http://econpapers.repec.org/RePEc:ecm:emetrp:v:77:y:2009:i:5:p:1403-1445>
- Phillips, P. C., Yu, J., 2008. Information loss in volatility measurement with flat price trading. Working paper, Yale University and Singapore Management University.
- Podolskij, M., Vetter, M., 2009. Estimation of volatility functionals in the simultaneous presence of microstructure noise and jumps. *Bernoulli* 15, 634–658.
- Revuz, D., Yor, M., 1999. Continuous martingales and Brownian motion, 3rd Edition. Grundlehren der mathematischen Wissenschaften. Springer, Berlin [u.a.].
- Veraart, A. E. D., 2008. Inference for the jump part of quadratic variation of itô semimartingales. working paper University of Aarhus.
- Zhang, L., Mykland, P. A., Ait-Sahalia, Y., 2005. A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association* 100 (472), 1394–1411.

Appendix A. Consistency and CLT

This appendix contains the proofs for the consistency and limiting distribution of the MinRV and MedRV estimators. The appendix is intended to provide a pedagogical step by step illustration of how one, where possible, may proceed to recast the features of a block based estimator into a format required for application of the existing asymptotic distribution theory results originally developed for MPV estimators.

Let Y_t be the log price process and assume that it follows a Brownian semimartingale

$$Y_t = Y_0 + \int_0^t a_u du + \int_0^t \sigma_{u-} dB_u \quad (\text{A.1})$$

where a is a locally bounded and predictable process and σ is adapted and cadlag and bounded away from zero. Without loss of generality, we further assume that the functions a, σ are uniformly bounded and that $\inf_{t>0} \sigma_t > 0$ a.s.²⁵ The extension allowing for finite activity jumps in Y_t is dealt with in Section Appendix A.3 below.

For the central limit theorem we require in addition that the volatility process follows a generalized Itô process:

$$\textbf{Assumption (A1)} : \sigma_t = \sigma_0 + \int_0^t \tilde{a}_u du + \int_0^t \tilde{\sigma}_{u-} dB_u + \int_0^t \tilde{v}_{u-} dW_u,$$

where \tilde{a} is locally bounded and predictable and $\tilde{\sigma}, \tilde{v}$ are cadlag and the Brownian motions B, W are uncorrelated. As before, we impose without loss of generality that the functions $\tilde{a}, \tilde{\sigma}$, and \tilde{v} are uniformly bounded and that $\inf_{t>0} \tilde{\sigma}_t > 0$ and $\inf_{t>0} \tilde{v}_t > 0$ a.s. In addition, as explained at the end of Section A.2, a general set of jump processes may be included in the volatility process specification without altering the results.

We assume that Y is observed at $N + 1$ evenly spaced time points spanning the interval $[0; 1]$. Below, we denote these observations by $Y_{i/N}$, $i = 0, \dots, N$, and the associated log-returns by $\Delta_i^N Y = Y_{i/N} - Y_{(i-1)/N}$, $i = 1, \dots, N$. The proofs involve sequences of standardized return observations and corresponding approximating sequences for which volatility is fixed across one or more returns. Hence, we introduce non-overlapping blocks of $M \geq 1$ returns for which the volatility process is constant. We assume we have $K = N/M$ such blocks in the sample. Consequently, we define the quantities,

²⁵As argued in Barndorff-Nielsen et al. (2006a), this follows from working with the stopped versions of the processes: $T_t^{(k)} = Y_{t \wedge T_k}$ and $\sigma_t^{(k)} = \sigma_{t \wedge T_k}$ where $T_k = \inf\{t | |a_t| + |\sigma_{t-}| \geq k\}$ and $T_k \nearrow \infty$ a.s.

$$\chi_i^N = \sqrt{N} \Delta_i^N Y \quad \text{and} \quad (\text{A.2})$$

$$\beta_i^{N,M} = \sqrt{N} \sigma_{\lfloor \frac{(i-1)/M \rfloor M}{N}} \Delta_i^N B = \sqrt{N} \sigma_{\lfloor \frac{(i-1)/M \rfloor}{K}} \Delta_i^N B, \quad (\text{A.3})$$

where $\lfloor \cdot \rfloor$ indicates the integer part of an expression. Hence, for each of the K return blocks, corresponding to $\beta_i^{N,M}$, the volatility remains fixed at the value it attains at the beginning of the block. We shall exploit that, for large N , $\chi_i^N \approx \beta_i^{N,M}$. The strategy of the proof is then, as in Barndorff-Nielsen et al. (2006a), henceforth BNGJPS, to first show convergence in probability and distribution for the approximate process and then argue that the difference is small.

Let $g : \mathbb{R}^2 \mapsto \mathbb{R}_+$ be given by,

$$g(x) = \frac{\pi}{\pi - 2} \min(|x_1|^2, |x_2|^2).$$

then, for any two bivariate vectors, $\mathbf{a} = (a_1, a_2)$ and $\mathbf{b} = (b_1, b_2)$, we have the bound

$$|g(\mathbf{a}) - g(\mathbf{b})| \leq (|a_1^2 - b_1^2| + |a_2^2 - b_2^2|). \quad (\text{A.4})$$

Furthermore, we also later use the fact that, except on the null set $\{(x, y) \in \mathbb{R}^2 | x = y\}$, we have

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} [\min(x_1^2, x_2^2 + \varepsilon z) - \min(x_1^2, x_2^2)] = \begin{cases} z & \text{if } x_2^2 < x_1^2 \\ 0 & \text{if } x_2^2 > x_1^2 \end{cases} \quad (\text{A.5})$$

Our MinRV estimator of IV may be written as,

$$\text{minRV}_N = \frac{\pi}{\pi - 2} \frac{1}{N - 1} \sum_{i=1}^{N-1} \min((\chi_i^N)^2, (\chi_{i+1}^N)^2) = \frac{1}{N - 1} \sum_{i=1}^{N-1} g(\chi_i^N, \chi_{i+1}^N).$$

The proof revolves around the sequences,

$$\mathbf{V}_N = \frac{1}{N} \sum_{i=1}^{N-1} g(\chi_i^N, \chi_{i+1}^N) \quad \text{and} \quad \mathbf{U}_N^M = \frac{1}{N} \sum_{i=1}^{N-1} g(\beta_i^{N,M}, \beta_{i+1}^{N,M}).$$

Since $\text{MinRV}_N = \frac{N}{N-1} V_N$, the V_N sequence is asymptotically equivalent to our MinRV estimator, while U_N^M is the approximating sequence.

We now introduce additional simplifying notation: for any adapted, integrable, d -dimensional cadlag process, Z , and for $N \geq j > i - 1 \geq 0$ we denote the expectation conditional on information at time $\frac{i-1}{N}$:

$$E_{i-1} \left[Z_{\frac{j}{N}} \right] = E \left[Z_{\frac{j}{N}} | \mathcal{F}_{\frac{i-1}{N}} \right] \quad (\text{A.6})$$

One useful implication of our ability to focus on the case with uniformly bounded drift and volatility functions is that, using the Burkholder-Davis-Grundy inequalities, we have,

$$E_{i-1} [|\chi_i^N|^p] \leq C \quad \text{and} \quad E_{i-1} [|\beta_i^{N,M}|^p] \leq C, \quad (\text{A.7})$$

where $p > 0$ and C denotes a generic positive constant which (with a slight abuse of notation) will take on disparate values in different places in what follows. We exploit this property of uniformly bounded moments repeatedly in the sequel.

We may now decompose our basic estimators for IV into a sum of conditional expectations and the associated martingale difference sequence: $V_N = V_{1N} + V_{2N}$ and $U_N^M = U_{1N}^M + U_{2N}^M$ where,

$$\mathbf{V}_{1N} = \frac{1}{N} \sum_{i=1}^{N-1} E_{i-1} [g(\chi_i^N, \chi_{i+1}^N)], \quad \mathbf{V}_{2N} = \frac{1}{N} \sum_{i=1}^{N-1} \{g(\chi_i^N, \chi_{i+1}^N) - E_{i-1} [g(\chi_i^N, \chi_{i+1}^N)]\}$$

$$\mathbf{U}_{1N}^M = \frac{1}{N} \sum_{i=1}^{N-1} E_{\lfloor \frac{i-1}{M} \rfloor M} \left[g \left(\beta_i^{N,M}, \beta_{i+1}^{N,M} \right) \right], \quad \mathbf{U}_{2N}^M = \frac{1}{N} \sum_{i=1}^{N-1} \left\{ g \left(\beta_i^{N,M}, \beta_{i+1}^{N,M} \right) - E_{\lfloor \frac{i-1}{M} \rfloor M} \left[g \left(\beta_i^{N,M}, \beta_{i+1}^{N,M} \right) \right] \right\}.$$

When $M = 1$ we will use the shorthand $\beta_i^N \equiv \beta_i^{N,1}$, $\mathbf{U}_N \equiv \mathbf{U}_N^1$ and similarly for the individual pieces \mathbf{U}_{1N} and \mathbf{U}_{2N} . These definitions allow us to decompose the main estimator:

$$\mathbf{V}_N = \mathbf{U}_{1N} + \mathbf{U}_{2N} + (\mathbf{V}_{1N} - \mathbf{U}_{1N}) + (\mathbf{V}_{2N} - \mathbf{U}_{2N}) \quad (\text{A.8})$$

Consistency of V_N can then be obtained by showing consistency of the estimator applied to the approximating Brownian path with piecewise constant volatility ($\mathbf{U}_N = \mathbf{U}_{1N} + \mathbf{U}_{2N}$) and then showing that the difference $V_N - U_N$ (the last two terms) is asymptotically negligible. This is what we do in Section Appendix A.1 below. To prove a CLT, we exploit a different decomposition (similar to Mykland and Zhang (2009)), in which we show the CLT for our estimator applied to an approximating Brownian motion in which volatility is held constant over a block of length M and then proceed to show that the difference between the original estimator and the estimator applied to the approximating process is negligible. This analysis is carried out in Section Appendix A.2 based on the decomposition:

$$\begin{aligned} \sqrt{N}(\mathbf{V}_N - IV) &= \sqrt{N}(\mathbf{V}_{1N} - \mathbf{U}_{1N}) + \sqrt{N}(\mathbf{V}_{2N} - \mathbf{U}_{2N}) + \sqrt{N}(\mathbf{U}_{1N} - IV) \\ &\quad + \sqrt{N}(\mathbf{U}_{2N} - \mathbf{U}_{2N}^M) + \sqrt{N}\mathbf{U}_{2N}^M \end{aligned} \quad (\text{A.9})$$

Appendix A.1. Proposition 1: Consistency

We proceed by analyzing equation (A.8) term by term through a series of lemmas. For brevity, we focus on the features that are specific to our estimator, while referring to proofs in the extant literature when feasible. This also serves to highlight the underlying structural similarities between our IV measure and some previously proposed IV estimators.

Lemma 4 *Under the maintained assumptions we have,*

$$\mathbf{U}_{1N} \xrightarrow{P} IV \quad (\text{A.10})$$

Moreover, if Assumption (A1) holds we obtain,

$$\sqrt{N}(\mathbf{U}_{1N} - IV) \xrightarrow{P} 0 \quad (\text{A.11})$$

Proof. First, note that

$$g(\beta_i^N, \beta_{i+1}^N) = \left[g(\beta_i^N, \beta_{i+1}^N) - g\left(\beta_i^N, \sqrt{N}\sigma_{\frac{i-1}{N}}\Delta_{i+1}^N B\right) \right] + g\left(\beta_i^N, \sqrt{N}\sigma_{\frac{i-1}{N}}\Delta_{i+1}^N B\right)$$

so we may write

$$\mathbf{U}_{1N} = \frac{1}{N} \sum_{i=1}^{N-1} \mathbb{E}_{i-1} \left[g(\beta_i^N, \beta_{i+1}^N) - g\left(\beta_i^N, \sqrt{N}\sigma_{\frac{i-1}{N}}\Delta_{i+1}^N B\right) \right] + \frac{1}{N} \sum_{i=1}^{N-1} \sigma_{\frac{i-1}{N}}^2 \quad (\text{A.12})$$

The first sum in (A.12) tends to zero in probability. To see this, note that the bound (A.4) implies the following limit in L_2 -norm:

$$\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^{N-1} \mathbb{E}_{i-1} \left[g(\beta_i^N, \beta_{i+1}^N) - g\left(\beta_i^N, \sqrt{N}\sigma_{\frac{i-1}{N}}\Delta_{i+1}^N B\right) \right] \right|^2 \leq \frac{C}{N} \mathbb{E} \sum_{i=1}^{N-1} |\sigma_{\frac{i}{N}}^2 - \sigma_{\frac{i-1}{N}}^2|^2 \rightarrow 0 \quad (\text{A.13})$$

where the convergence in (A.13) (which implies convergence in probability) follows from the fact that σ_t has finite quadratic variation. In addition, since $\{\sigma_t^2\}_{t \geq 0}$ is uniformly bounded and cadlag, the pointwise dominated convergence of $\sigma_u - \sigma_{\lfloor \frac{uN}{N} \rfloor} \rightarrow 0$ for $u \in [0; 1]$ follows, and Lebesgue's theorem yields

$$\sum_{i=1}^{N-1} \left[\int_{\frac{(i-1)}{N}}^{\frac{i}{N}} (\sigma_u^2 - \sigma_{\frac{(i-1)}{N}}^2) du \right] \xrightarrow{a.s.} 0 \quad (\text{A.14})$$

Together (A.13) and (A.14) imply $IV - U_{1N} \xrightarrow{P} 0$, which establishes (A.10). To show (A.11) we need the stronger assumption (A1). Define the sequence of independent standard normals $Z_i = \sqrt{N} \Delta_i^N B$, then Assumption (A1) yields

$$\mathbb{E}_{i-1} \left[(\sigma_{\frac{i}{N}}^2 - \sigma_{\frac{(i-1)}{N}}^2) Z_{i+1}^2 \mathbf{1}_{Z_{i+1}^2 < Z_i^2} \right] = \mathbb{E}_{i-1} \left[(\sigma_{\frac{i}{N}}^2 - \sigma_{\frac{(i-1)}{N}}^2) \varphi(Z_i^2) \right] = O_P(1/N) \quad (\text{A.15})$$

since $\varphi(Z_i^2) = \mathbb{E}_i[Z_{i+1}^2 \mathbf{1}_{Z_{i+1}^2 < Z_i^2}]$ is an even function of the Brownian path $\{B_t\}_{(i-1)/N < t < i/N}$. Now the property (A.5) yields

$$\mathbb{E}_{i-1} \left[g(\beta_i^N, \beta_{i+1}^N) - g\left(\beta_i^N, \sqrt{N} \sigma_{\frac{(i-1)}{N}} \Delta_{i+1}^N B\right) \right] = O_P(1/N) \quad (\text{A.16})$$

This ensures that the first term in (A.12) is asymptotically negligible, even when scaled up by \sqrt{N} . Hence, the remaining task is to show,

$$\sqrt{N} \left(\frac{1}{N} \sum_{i=1}^{N-1} \sigma_{\frac{(i-1)}{N}}^2 - IV \right) \xrightarrow{P} 0.$$

However, this is a common task in the proof of CLT for IV estimators and the method of proof is, by now, well established; see, e.g., BNGJPS where the result is shown for a general setting of which the current framework is a special case. A more intuitive and detailed exposition is provided by Barndorff-Nielsen et al. (2006b), henceforth BNGJS. ■

Lemma 5 *Under the maintained assumptions, we have*

$$U_{2N} \xrightarrow{P} 0 \quad (\text{A.17})$$

Proof. To simplify notation, define the martingale difference sequence $\left\{ \frac{1}{N} \eta_i^N, \mathcal{F}_{\frac{i}{N}} \right\}_{i \geq 0}$:

$$\eta_i^N = g(\beta_i^N, \beta_{i+1}^N) - E_{i-1} [g(\beta_i^N, \beta_{i+1}^N)]$$

Note that $\mathbb{E}[(\eta_i^N)^2 | \mathcal{F}_{\frac{i-1}{N}}] \leq C$, so applying the Cauchy-Schwartz inequality,

$$\mathbf{V} \left[\frac{1}{N} \sum_{i=1}^N \eta_i^N \right] = \frac{1}{N} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N ((\eta_i^N)^2 + 2\eta_i^N \eta_{i+1}^N) \right] \leq \frac{C}{N} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \mathbb{E}[(\eta_i^N)^2 | \mathcal{F}_{\frac{i-1}{N}}] \right] \leq \frac{C}{N} \rightarrow 0.$$

The L_2 convergence implies $\frac{1}{N} \sum_{i=1}^N \eta_i^N \xrightarrow{P} 0$. ■

Lemma 6 *Under the maintained assumptions, we have,*

$$(\mathbf{V}_{1N} - U_{1N}) \xrightarrow{P} 0. \quad (\text{A.18})$$

Under Assumption (A1), we obtain,

$$\sqrt{N} (\mathbf{V}_{1N} - U_{1N}) \xrightarrow{P} 0. \quad (\text{A.19})$$

Proof. To establish (A.19), and thus also (A.18), we must show,

$$\sqrt{N} (V_{1N} - U_{1N}) = \frac{1}{\sqrt{N}} \sum_{i=1}^{N-1} \mathbb{E}_{i-1} [(g(\chi_i^N, \chi_{i+1}^N) - g(\beta_i^N, \beta_{i+1}^N))] \rightarrow 0 \text{ as } N \rightarrow \infty. \quad (\text{A.20})$$

Using the bound (A.4), it follows that,

$$\begin{aligned} \sqrt{N} (V_{1N} - U_{1N}) &\leq \frac{1}{\sqrt{N}} \mathbb{E} \left[\sum_{i=1}^{N-1} |g(\chi_i^N, \chi_{i+1}^N) - g(\beta_i^N, \beta_{i+1}^N)| \right] \leq \frac{C}{\sqrt{N}} \mathbb{E} \left[\sum_{i=1}^N |(\chi_i^N)^2 - (\beta_i^N)^2| \right] \\ &= \frac{C}{\sqrt{N}} \sum_{i=1}^N \left(\mathbb{E}_{i-1} \left[h(\sqrt{N} \Delta_i^N Y) - \sigma_{\frac{i-1}{N}}^2 \right] \right) \end{aligned}$$

where we have defined the function $h(x) = x^2$. This formulation maps directly into the setting of BNGJPS where the results of this lemma are proven in a more general setting and for a generic $h(x)$ function subject to regularity conditions. In particular, our h function trivially satisfies the continuous differentiability and polynomial growth conditions necessary for the applicability of their analysis. An accessible, albeit lengthy, account of the steps of the argument may be found in BNGJS (2006, pp. 713-719). So while this proof is quite involved, the above reformulation of the relevant inequalities, as they arise within our specific setting, allows us to simply refer to previously published work for the result. ■

Lemma 7 *Under the maintained assumptions, we have,*

$$(\mathbf{V}_{2N} - \mathbf{U}_{2N}) \xrightarrow{P} 0. \quad (\text{A.21})$$

Moreover, we may strengthen this result further to obtain,

$$\sqrt{N} (\mathbf{V}_{2N} - \mathbf{U}_{2N}) \xrightarrow{P} 0. \quad (\text{A.22})$$

Proof. In order to demonstrate the second result of the lemma, which obviously implies the first, we define,

$$\xi_i^N = (1/\sqrt{N}) [g(\chi_i^N, \chi_{i+1}^N) - g(\beta_i^N, \beta_{i+1}^N)],$$

and we must then prove that,

$$\sum_{i=1}^{N-1} (\xi_i^N - \mathbb{E}_{i-1}[\xi_i^N]) \xrightarrow{P} 0$$

This expression constitutes a martingale difference sequence with respect to the filtration $\mathcal{F}_{\frac{i}{N}}$, so it suffices to show,

$$\sum_{i=1}^{N-1} \mathbb{E} [(\xi_i^N)^2] = \mathbb{E} \left[\sum_{i=1}^{N-1} \mathbb{E}_{i-1} [(\xi_i^N)^2] \right] \rightarrow 0 \text{ as } N \rightarrow \infty$$

Mimicking the type of steps undertaken in the proof of the previous lemma, including application of the uniform bound on moments of χ_i^N and β_i^N , we obtain,

$$\begin{aligned} \sum_{i=1}^{N-1} \mathbb{E} [(\xi_i^N)^2] &= \frac{1}{N} \mathbb{E} \left[\sum_{i=1}^{N-1} \mathbb{E}_{i-1} |g(\chi_i^N, \chi_{i+1}^N) - g(\beta_i^N, \beta_{i+1}^N)|^2 \right] \\ &\leq \frac{C}{N} \mathbb{E} \left[\sum_{i=1}^N \mathbb{E}_{i-1} [(h(\chi_i^N) - h(\beta_i^N))^2] \right]. \end{aligned}$$

As for the previous lemma, our reformulation of the task maps the problem into the corresponding task in BNGJPS (2006) who prove the current lemma in a more general setting. A detailed account of the requisite steps to complete the proof may again be gleaned from BNGJS (2006, pp. 704-706). ■

Taken together, Lemma 4 - 5 and the first parts of Lemma 6 - 7 imply the consistency of our estimator under the minimal maintained assumptions. The second part of Lemma 6 - 7 is critical for the proof of the central limit theorem below.

Appendix A.2. Proposition 2: The CLT

Lemma 8 Under assumption (A1), we have

$$\sqrt{N}U_{2N}^M \xrightarrow{\text{stable } \mathcal{D}} N \left(0, \nu \int_0^1 \sigma^4 du \right) \quad (\text{A.23})$$

where the constant $\nu = \text{Var}[g(U_0, U_1)] + 2 \text{Cov}[g(U_0, U_1), g(U_1, U_2)]$ for $U_0, U_1, U_2 \sim i.i.d.N(0, 1)$.

Proof. Consider splitting the N scaled return observations into K blocks, the k^{th} of which is the vector $\chi_k^M = \{\sqrt{N}\Delta_i^N Y\}_{i \in \{(k-1)M+1, \dots, kM\}}$. The corresponding vector of observations from the approximating Brownian motion where volatility is held constant over the block is $\beta_k^{N,M} = \{\beta_i^{N,M}\}_{i \in \{(k-1)M+1, \dots, kM\}}$. Next, define by $g_M(\cdot) : \mathbb{R}^M \mapsto \mathbb{R}$ the block estimator of volatility:

$$g_M(\beta_k^{N,M}) = \frac{1}{M} \sum_{i=(k-1)M+1}^{kM-1} g(\beta_i^{N,M}, \beta_{i+1}^{N,M}) \quad (\text{A.24})$$

We wish to apply Theorem IX.7.28 in Jacod and Shiryaev (2003) to $\sqrt{N}U_{2N}^M$. Defining the martingale difference sequence $\psi_k^{N,M} = \sqrt{M} \left(g_M(\beta_k^{N,M}) - \frac{M-1}{M} \sigma_{\frac{k-1}{K}}^2 \right)$ we can write

$$\begin{aligned} \sqrt{N}U_{2N}^M &= \frac{1}{\sqrt{K}} \sum_{k=1}^K \psi_k^{N,M} + \frac{1}{\sqrt{N}} \sum_{k=1}^{K-1} \left(g(\beta_{kM}^{N,M}, \beta_{kM+1}^{N,M}) - \mathbb{E}_{\frac{k-1}{K}} \left[g(\beta_{kM}^{N,M}, \beta_{kM+1}^{N,M}) \right] \right) \\ &= \frac{1}{\sqrt{K}} \sum_{k=1}^K \psi_k^{N,M} + o_P(1) \end{aligned} \quad (\text{A.25})$$

The last equality follows from the fact that each term in the second sum is centered and has bounded variance (given the uniform bound on σ_t). Thus the sum divided by \sqrt{N} will tend to zero provided $K = o_P(N)$.

We must now verify conditions (7.27)-(7.31) of Theorem IX.7.28. First note that $E[\psi_k^{N,M} | \mathcal{F}_{\frac{k-1}{K}}] = 0$ so that condition (7.27) is trivially satisfied. Condition (7.28) follows from the fact that

$$\frac{1}{K} \sum_{k=1}^K E \left[\left\{ \sqrt{M} \left(g_M(\beta_k^M) - \frac{M-1}{M} \sigma_{\frac{k-1}{K}}^2 \right) \right\}^2 \middle| \mathcal{F}_{\frac{k-1}{K}} \right] = \frac{\nu}{K} \sum_{k=1}^K \sigma_{\frac{k-1}{K}}^4 \xrightarrow{P} \nu \int_0^1 \sigma_u^4 du \quad (\text{A.26})$$

where the convergence in probability (and in fact a.s.) is a consequence of the volatility process being cadlag and uniformly bounded. Next, we turn to condition (7.29). Let $\Delta_k^M \mathbf{B} = \left(B_{\frac{k}{K}} - B_{\frac{k-1}{K}} \right)$, then $E \left[\psi_k^{N,M} \Delta_k^M \mathbf{B} \middle| \mathcal{F}_{\frac{k-1}{K}} \right] = 0$, which follows from the fact that the variables $\psi_k^{N,M}$ are centered and that g_M is an even function. Condition (7.30), stating that $E \left[(\psi_k^{N,M})^2 \mathbf{1}_{|\psi_k^{N,M}| > \varepsilon} \right] \xrightarrow{P} 0$, follows straightforwardly from the fact that σ is uniformly bounded.

Finally, let $\{N_t\}_{t \in [0;1]}$ be a bounded martingale orthogonal to B (i.e. the covariation $\langle B, N \rangle_t = 0$ a.s.). We want to show that, for each block k , $E[\psi_k^{N,M} \left(N_{\frac{k}{K}} - N_{\frac{k-1}{K}} \right) | \mathcal{F}_{\frac{k-1}{K}}] = 0$. For $t > \frac{k-1}{K}$ consider the martingale difference sequence $M_t = E \left[\psi_k^{N,M} \middle| \mathcal{F}_t \right]$. By the martingale representation theorem, $M_t = M_{\frac{k-1}{K}} + \int_{\frac{k-1}{K}}^t \varphi_u dB_u$ for some predictable process φ_u . Therefore the processes $\{M_t\}_{t > \frac{k-1}{K}}$ and $\{N_t - N_{\frac{k-1}{K}}\}_{t > \frac{k-1}{K}}$ are orthogonal and the product, $\{M_t(N_t - N_{\frac{k-1}{K}})\}$ is again a martingale which must then have mean zero. This verifies condition (7.31). Theorem IX.7.28 in Jacod and Shiryaev (2003) then implies that as N (and hence K and M) tend to infinity:

$$\sqrt{N}U_{2N}^M \xrightarrow{\text{stable}} N \left(0, \nu \int_0^1 \sigma^4 du \right) \quad (\text{A.27})$$

■

Lemma 9 *Under the maintained assumptions, we have*

$$\sqrt{N} \left(\mathbf{U}_{2N} - \mathbf{U}_{2N}^M \right) \xrightarrow{P} 0 \quad (\text{A.28})$$

Proof. Defining $\eta_i^{N,M} = g(\beta_i^{N,M}, \beta_{i+1}^{N,M}) - \mathbb{E}_{\lfloor (i-1)/M \rfloor M} \left[g(\beta_i^{N,M}, \beta_{i+1}^{N,M}) \right]$, we note that $\left\{ \frac{1}{\sqrt{N}} (\eta_i^N - \eta_i^{N,M}) \right\}_{i \geq 1}$ is a martingale difference sequence with respect to the filtration $\{\mathcal{F}_{i/N}\}$. To show that $\sqrt{N} \left(\mathbf{U}_{2N} - \mathbf{U}_{2N}^M \right) = \sum_{i=1}^{N-1} (\eta_i^N - \eta_i^{N,M}) / \sqrt{N} \rightarrow 0$ in probability, it therefore suffices (by Doob's inequality, e.g. Revuz and Yor (1999), p.54-55) to show that

$$\frac{1}{N} \mathbb{E} \left[\sum_{i=1}^{N-1} \left| g(\beta_i^N, \beta_{i+1}^N) - g(\beta_i^{N,M}, \beta_{i+1}^{N,M}) \right|^2 \right] \rightarrow 0 \quad (\text{A.29})$$

By the bound of $g(\cdot)$ we have

$$\begin{aligned} \frac{1}{N} \mathbb{E} \left[\sum_{i=1}^{N-1} \left| g(\beta_i^N, \beta_{i+1}^N) - g(\beta_i^{N,M}, \beta_{i+1}^{N,M}) \right|^2 \right] &\leq \frac{C}{N} \mathbb{E} \left[\sum_{i=1}^N \mathbb{E}_{i-1} \left| (\beta_i^N)^2 - (\beta_i^{N,M})^2 \right|^2 \right] \\ &\leq \frac{C}{N} \mathbb{E} \left[\sum_{i=1}^N \left| \sigma_{\frac{i-1}{N}}^2 - \sigma_{\frac{\lfloor (i-1)/M \rfloor M}{N}}^2 \right|^2 \right] = C \mathbb{E} \int_0^1 \left(\sigma_{\frac{uN}{N}}^2 - \sigma_{\frac{\lfloor uK \rfloor}{K}}^2 \right)^2 du = o_P(1) \end{aligned} \quad (\text{A.30})$$

where the last inequality follows from the uniform boundedness of σ_t and Lebesgue's theorem. ■

Importantly, the specification of the volatility process in Assumption (A1) can be extended to include finite as well as infinite activity jump processes subject only to very weak regularity conditions, stipulating that the volatility process evolves according to an Ito semimartingale where the jump components have locally bounded jump characteristics, as laid out in BNGJPS. This follows from the fact that the only terms in (A.9) affected by the inclusion of jumps are $\sqrt{N}(V_{1N} - U_{1N})$ and $\sqrt{N}(V_{2N} - U_{2N})$ which map into the corresponding terms in BNGJPS as outlined in the proofs above. As such, the distributional results of the paper cover a very wide range of underlying return generating processes.

Finally, we note that the proof for the MedRV estimator follows analogously by simply changing the g function accordingly. The proof of Proposition 3 is omitted, but it may be derived using the identical strategy, in which volatility is held constant over blocks of increasing size. In particular, the conditional covariance between the estimators can easily be calculated on each block and a stable convergence argument similar to Lemma 8 goes through.

Appendix A.3. The Asymptotic Distribution under Jump Alternatives

Suppose now the log price process is given as $X = Y + J$, where Y is a Brownian semi-martingale of the form (A.1) while J is a finite activity jump process. While the specification of the jump process is restrictive, it covers many cases of interest and it can be generalized to infinite activity jump processes along the lines of Barndorff-Nielsen et al. (2006c). Restating their Proposition 1 in our notation, we have $\sqrt{N} |Y_{\frac{i}{N}} - Y_{\frac{i-1}{N}}| = O_P(|\log(N)|^{1/2})$, which, as they show, follows from Levy's modulus of continuity theorem for Brownian motion. This immediately yields:

Proposition 10 *When J is a **finite activity** jump process, the asymptotic distribution of the MinRV and MedRV estimators applied to the processes $\{X_t\}$ and $\{Y_t\}$ are identical.*

Proof.

As before, we deal only with the MinRV case as the MedRV case is analogous. On a given realization of the path there is a finite number of jumps, so (asymptotically) at most one of the terms $|X_{\frac{i}{N}} - X_{\frac{i-1}{N}}|$ or $|X_{\frac{i+1}{N}} - X_{\frac{i}{N}}|$ will include a jump. Therefore, each term in the estimator (up to a normalizing constant) is

$$\min\left(|X_{\frac{i}{N}} - X_{\frac{i-1}{N}}|^2, |X_{\frac{i+1}{N}} - X_{\frac{i}{N}}|^2\right) = O_P\left(\frac{\log N}{N}\right)$$

regardless of whether a (single) jump occurred or not. Since only finitely many terms differ,

$$\begin{aligned} \sum_{j=1}^N \left[\min\left(|X_{\frac{j}{N}} - X_{\frac{j-1}{N}}|^2, |X_{\frac{j+1}{N}} - X_{\frac{j}{N}}|^2\right) - \min\left(|Y_{\frac{j}{N}} - Y_{\frac{j-1}{N}}|^2, |Y_{\frac{j+1}{N}} - Y_{\frac{j}{N}}|^2\right) \right] \\ = O_P\left(\frac{\log N}{N}\right) = o_P\left(\frac{1}{\sqrt{N}}\right) \end{aligned}$$

so neither the consistency nor the convergence in distribution is affected by the occurrence of finite activity jumps.

■

Appendix A.4. The Key Relations behind Proposition 3: Correlation Between Estimators

The correlation between the non-subsampled IV estimators can be calculated in a straightforward manner under the Brownian null. Denote by $U_1, \dots, U_5 \sim N(0, \sigma^2)$ a block of five adjacent returns and denote by $\mathcal{E}_1, \mathcal{E}_2 \in \{\text{BV}, \text{TV}, \text{MinRV}, \text{MedRV}\}$ two generic estimators using blocks of either two or three adjacent returns.

The covariance between two estimators based on 2-blocks of returns (e.g., $\mathcal{E}_1 = \text{BV}$ and $\mathcal{E}_2 = \text{MinRV}$) over a day with N returns is then given by

$$\text{Cov}[\mathcal{E}_1, \mathcal{E}_2] = \frac{N-1}{N} \text{Cov}[\mathcal{E}_1(U_1, U_2), \mathcal{E}_2(U_1, U_2)] + 2 \frac{N-2}{N} \text{Cov}[\mathcal{E}_1(U_1, U_2), \mathcal{E}_2(U_2, U_3)]$$

The covariance between an estimator based on a 2-block (e.g., $\mathcal{E}_1 = \text{BV}$) and an estimator based on a 3-block (e.g., $\mathcal{E}_2 = \text{TV}$) over a day with N returns is analogously given by

$$\text{Cov}[\mathcal{E}_1, \mathcal{E}_2] = 2 \frac{N-2}{N} \text{Cov}[\mathcal{E}_1(U_1, U_2), \mathcal{E}_2(U_1, U_2, U_3)] + 2 \frac{N-3}{N} \text{Cov}[\mathcal{E}_1(U_1, U_2), \mathcal{E}_2(U_2, U_3, U_4)]$$

Finally, the covariance between two estimators based on 3-blocks of returns (e.g., $\mathcal{E}_1 = \text{TV}$ and $\mathcal{E}_2 = \text{MedRV}$) over a day with N returns is

$$\begin{aligned} \text{Cov}[\mathcal{E}_1, \mathcal{E}_2] &= \frac{N-2}{N} \text{Cov}[\mathcal{E}_1(U_1, U_2, U_3), \mathcal{E}_2(U_1, U_2, U_3)] + 2 \frac{N-3}{N} \text{Cov}[\mathcal{E}_1(U_1, U_2, U_3), \mathcal{E}_2(U_2, U_3, U_4)] \\ &\quad + 2 \frac{N-4}{N} \text{Cov}[\mathcal{E}_1(U_1, U_2, U_3), \mathcal{E}_2(U_3, U_4, U_5)] \end{aligned}$$

The prerequisite covariances between functionals of random normals can easily be calculated (often in closed form) to the desired accuracy. Proposition 3 provides the asymptotic covariance matrix between RV, BV, TV, MinRV and MedRV, where the covariance terms with RV simply equals the variance of RV since the latter is the efficient estimator under the Brownian null.

Appendix B. IV estimators based on pre-averaged returns

Given N equally-spaced (log) returns, $\Delta_i^N Y = Y_{i/N} - Y_{(i-1)/N}$, $i = 1, \dots, N$, we define corresponding pre-averaged returns for any pre-averaging window size $2K \leq N$:

$$\Delta_i^N \bar{Y} = \frac{1}{K} \sum_{j=K}^{2K-1} Y_{(i+j)/N} - \frac{1}{K} \sum_{j=0}^{K-1} Y_{(i+j)/N}, \quad i = 1, \dots, N - 2K + 1. \quad (\text{B.1})$$

An equivalent definition with analytically more tractable expression is given by:

$$\Delta_i^N \bar{Y} = 2 \sum_{j=1}^{2K-1} g\left(\frac{j}{2K}\right) \Delta_{i+j}^N Y, \quad i = 1, \dots, N - 2K + 1, \quad (\text{B.2})$$

where $g(x) = x \wedge (1-x)$, $x \in [0, 1]$ is the pre-averaging kernel. We further define $\psi_K = \frac{1}{2K} \sum_{j=1}^{2K-1} 4g\left(\frac{j}{2K}\right)^2$ to be the finite sample analog of the variance scaling factor $\psi = \int_0^1 4g(u)^2 du = \frac{1}{3}$ induced by the pre-averaging kernel.

Consider the following $2K$ sub-samples of non-overlapping pre-averaged returns:

$$\begin{aligned} \bar{S}_1 &= \left\{ \Delta_{1+2(s-1)K}^N \bar{Y} : s = 1, \dots, \left\lfloor \frac{N}{2K} \right\rfloor \right\} \\ \bar{S}_2 &= \left\{ \Delta_{2+2(s-1)K}^N \bar{Y} : s = 1, \dots, \left\lfloor \frac{N-1}{2K} \right\rfloor \right\} \\ &\dots \\ \bar{S}_{2K} &= \left\{ \Delta_{2K+2(s-1)K}^N \bar{Y} : s = 1, \dots, \left\lfloor \frac{N-2K+1}{2K} \right\rfloor \right\} \end{aligned}$$

Let $\widehat{IV}[\bar{S}_k]$ denote the raw IV estimates obtained on each sub-sample of pre-averaged returns \bar{S}_k , $k = 1, \dots, 2K$. Then the pre-averaged (and sub-sampled) estimators $\widehat{IV}[\Delta_*^N \bar{Y}]$ for the full set of pre-averaged returns $\Delta_*^N \bar{Y} = \{\Delta_i^N \bar{Y}\}_{i=1}^{N-2K+1}$ can be defined as follows:

$$\widehat{IV}[\Delta_*^N \bar{Y}] = \frac{1}{2K} \sum_{k=1}^{2K} \frac{1}{\psi} \widehat{IV}[\bar{S}_k], \quad \widehat{IV}^{Adj}[\Delta_*^N \bar{Y}] = \frac{1}{2K} \sum_{k=1}^{2K} \frac{1}{\psi_K} \frac{\frac{N}{2K}}{\left\lfloor \frac{N-k+1}{2K} \right\rfloor} \widehat{IV}[\bar{S}_k],$$

where the latter expression incorporates a finite sample bias correction and is the form we use in our pre-averaged implementation of all estimators.

Consistency and asymptotic normality are clearly preserved by pre-averaging and sub-sampling, while noise-robustness improves. In particular, bounce-backs are near perfectly annihilated given that adjacent returns are subject to almost identical kernel weights. For further details on pre-averaging, please refer to Podolskij and Vetter (2009) and Jacod et al. (2009).

Appendix C. Figures and Tables

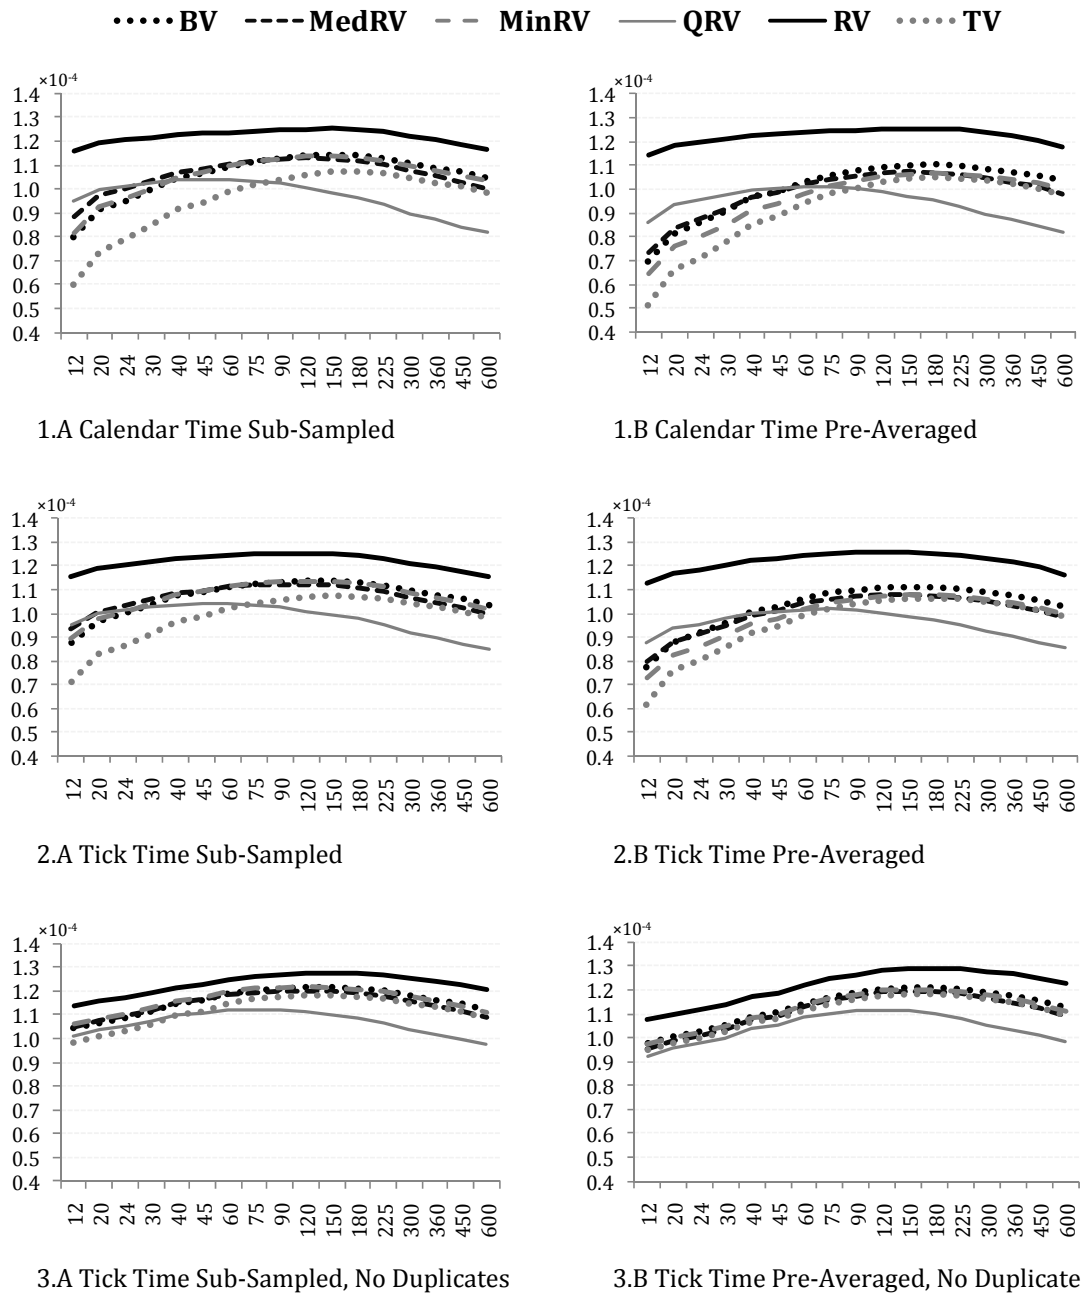


Figure C.1: **Average IV estimates across 33 DJ30 stocks between January 1 2005 and May 31, 2007.** The first column (Panels (1A)-(3A)) shows the average IV estimates produced by each sub-sampled estimator as a function of sampling frequency (measured in seconds on the x-axis). Panel (1A) contains the calendar time estimates, Panel (2A) the tick time estimates, and Panel (3A) the tick time estimates after filtering out duplicate quotes. The second column (Panels (1B)-(3B)) plots the corresponding estimates for the pre-averaged version of each estimator as a function of pre-averaging window (measured in seconds on the x-axis), where the pre-averaging is carried out at the highest frequency as described in section 3.3.

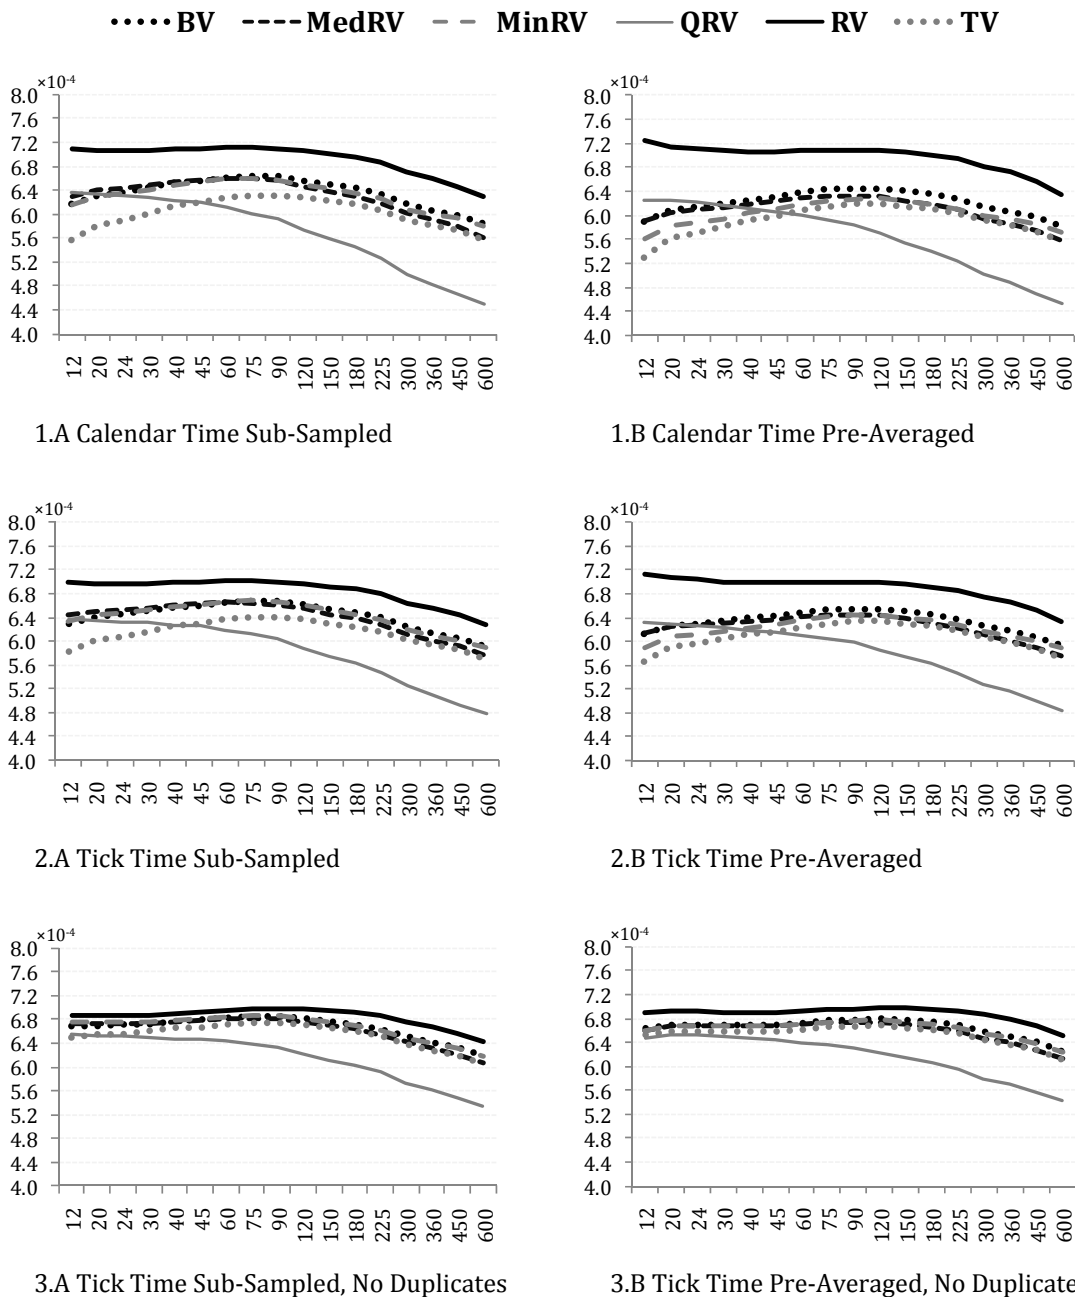
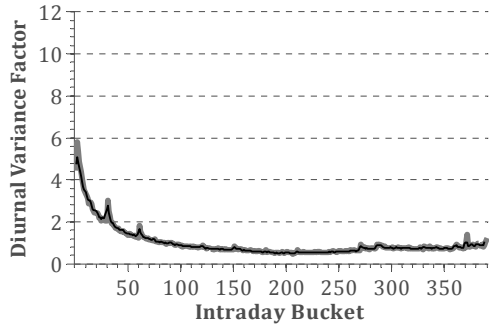
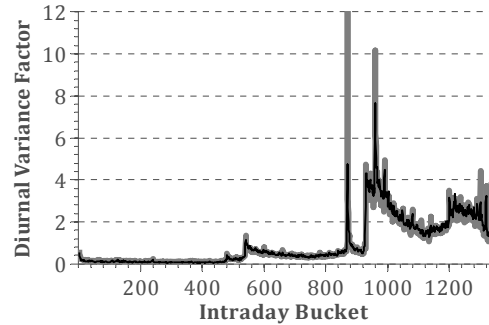


Figure C.2: **Average IV estimates across 33 DJ30 stocks between June 1 2007 and July 31, 2009.** The first column (Panels (1A)-(3A)) shows the average IV estimates produced by each sub-sampled estimator as a function of sampling frequency (measured in seconds on the x-axis). Panel (1A) contains the calendar time estimates, Panel (2A) the tick time estimates, and Panel (3A) the tick time estimates after filtering out duplicate quotes. The second column (Panels (1B)-(3B)) plots the corresponding estimates for the pre-averaged version of each estimator as a function of pre-averaging window (measured in seconds on the x-axis), where the pre-averaging is carried out at the highest frequency as described in section 3.3.


— RV — MedRV



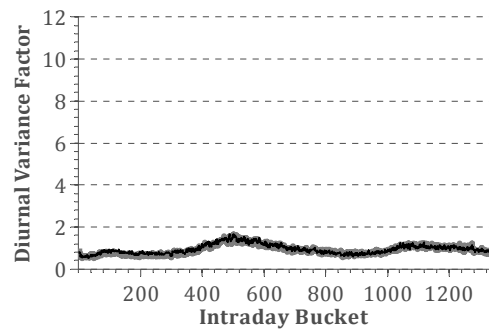
1.A DJ30 Stocks: Calendar Time Sample



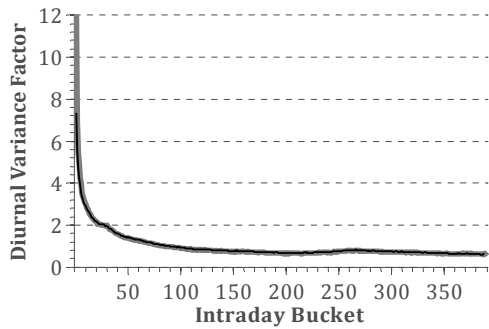
1.B E-Mini Futures: Calendar Time Sample



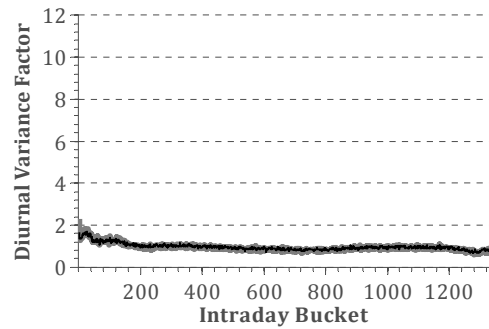
2.A DJ30 Stocks: Tick Time Sample, 1sec Ticks



2.B E-Mini Futures: Tick Time Sample, 1sec Ticks



3.A DJ30 Stocks: Tick Time Sample, All Ticks



3.B E-Mini Futures: Tick Time Sample, All Ticks

Figure C.3: **Diurnal volatility pattern across 33 DJ30 stocks and the E-mini S&P500 futures between January 1, 2005 and May 31, 2007.** We split each day into a number of equal sized buckets given by the number of minutes in a trading day and plot the average ratio across stock trading days between the realized variance contribution of each intraday bucket and the mean for all buckets on that same trading day, as measured by the RV and MedRV estimators. The plotted average ratio is therefore interpretable as the sample mean of the diurnal variance factor. The first column (Panels 1A-3A) shows the diurnal variance factor for the 33 DJ30 stocks throughout the trading day from 9:30am ET - 4:00pm ET, spanning 390 minutes. The second column (Panels 1.B-3.B) shows the diurnal variance factor for the ultra liquid S&P500 E-mini futures contract for the trading hours from 5:00pm CT the preceding day to 3:15pm CT, spanning 1335 minutes. Thus, for calendar time sampling (Panels 1.A and 1.B), the intraday buckets comprise 390 1-minute intervals for the DJ30 stocks and 1335 1-minute intervals for the E-Mini Futures. Accordingly, in the tick time sampling schemes each bucket represents $\frac{1}{390}$ of all observations for the DJ30 stocks and $\frac{1}{1335}$ of all observations for the E-Mini Futures. The tick time observation counts are computed by either first reducing the sample to one-second ticks based on the previous-tick method (Panels 2.A and 2.B) or directly taking all ticks without any reduction (Panels 3.A and 3.B).

Table C.1: Variance factors for multipower and MinRV and MedRV estimators. All estimators have an asymptotic variance of the form $\nu \int_0^1 \sigma_u^4 du$. The table displays the variance factor ν for each estimator.

	RV	BV	MedRV	MPV(3)	MPV(4)	MPV(5)	MPV(6)	MinRV	MPV(7)
Variance Factor	2.00	2.61	2.96	3.06	3.38	3.61	3.78	3.81	3.91

Table C.2: Descriptive statistics for the 33 stocks that were part of the Dow Jones 30 index between January 1, 2005 and May 31, 2007.

Symbol	Average # Quotes	Average # Non-Duplicate Quotes	Average Price (\$)	Average 2-min RV ($\times 10^{-4}$)	Average 2-min BV ($\times 10^{-3}$)	Average Log Spread ($\times 10^{-4}$)	Average Ratio Log Spread/Sigma	Average # zero returns per day lasting above 30s / 60s / 120s
AA	11,062	3,610	30	2.22	2.10	4.80	0.03	161 / 48 / 11
AIG	11,420	4,484	64	1.00	0.95	2.73	0.03	127 / 30 / 5
AXP	10,280	3,444	55	0.81	0.76	2.84	0.03	161 / 48 / 11
BA	10,481	4,919	74	1.26	1.19	2.97	0.03	113 / 21 / 3
BAC	12,364	3,496	48	0.68	0.64	2.87	0.04	161 / 49 / 11
C	13,201	4,073	49	0.77	0.73	2.70	0.03	144 / 42 / 9
CAT	9,660	4,361	71	1.62	1.54	3.01	0.02	131 / 27 / 4
CVX	13,353	6,169	63	1.64	1.58	2.88	0.02	82 / 15 / 2
DD	10,460	3,597	45	1.19	1.12	3.48	0.03	154 / 45 / 10
DIS	11,421	2,774	29	1.08	1.01	4.46	0.04	177 / 67 / 19
EK	7,645	1,769	26	2.19	1.93	5.48	0.04	189 / 81 / 28
GE	14,010	2,565	35	0.63	0.59	3.35	0.04	182 / 70 / 20
GM	10,225	3,211	30	3.70	3.49	5.09	0.03	162 / 51 / 13
HD	11,606	3,814	39	1.35	1.28	3.69	0.03	149 / 39 / 8
HPQ	12,284	3,324	32	1.61	1.50	4.32	0.04	163 / 55 / 14
IBM	11,626	5,531	86	0.84	0.80	2.43	0.03	89 / 14 / 2
INTC	17,285	1,638	22	1.67	1.51	4.63	0.04	206 / 99 / 35
IP	8,064	2,277	34	1.36	1.25	4.28	0.04	191 / 71 / 20
JNJ	11,722	4,091	64	0.52	0.49	2.47	0.04	130 / 32 / 6
JPM	12,263	3,453	42	0.88	0.83	3.27	0.04	163 / 52 / 12
KO	11,269	3,158	44	0.59	0.55	3.15	0.04	166 / 57 / 15
MCD	11,171	3,238	36	1.22	1.13	3.90	0.04	164 / 55 / 14
MMM	9,421	4,026	78	0.91	0.87	2.74	0.03	141 / 31 / 5
MO	10,544	4,195	74	1.09	0.93	2.38	0.02	142 / 33 / 5
MRK	11,424	3,315	37	1.58	1.38	3.84	0.03	165 / 54 / 13
MSFT	16,575	1,272	27	0.94	0.84	3.80	0.04	199 / 106 / 43
PFE	13,654	2,501	26	1.12	1.06	4.63	0.05	184 / 71 / 21
PG	11,663	4,101	58	0.72	0.69	2.69	0.03	144 / 38 / 7
T	11,970	2,472	27	1.08	1.00	4.94	0.05	177 / 77 / 27
UTX	10,201	4,432	68	1.07	1.01	3.16	0.03	127 / 27 / 4
VZ	12,339	3,056	35	0.98	0.91	3.81	0.04	165 / 61 / 18
WMT	12,475	4,345	48	0.91	0.88	3.01	0.03	130 / 32 / 6
XOM	14,831	6,133	64	1.39	1.34	2.27	0.02	90 / 17 / 2
Mean ALL:	11,757	3,601	47	1.23090	1.14720	0.00035	0.03	152 / 49 / 13
Median ALL	11,778	3,601	48	1.08398	1.00552	0.00033	0.03	161 / 48 / 11
Max ALL:	17,285	6,169	86	3.69647	3.48752	0.00055	0.05	206 / 106 / 43
Min ALL:	7,645	1,272	22	0.52107	0.49326	0.00023	0.02	82 / 14 / 2

Table C.3: **Relative bias and relative mean squared error (MSE) factors for sub-sampled IV estimators at the 12-second, 60-second, and 300-second sampling frequencies.** We report the bias and MSE factors for each sub-sampled estimator and each model based on sampling frequency 12 seconds (top panel), 60 seconds (middle panel), and 300 seconds (bottom panel). The MSE factor is computed as the sample mean of $390(IV - IV)^2/IQ$, where IV and IQ are the true simulated integrated variance and integrated quarticity on each day. The MSE factor for 60-sec sub-sampled RV is thus ≈ 1.33 which is the theoretical value shown by Zhang, Mykland, and Ait-Sahalia (2005). Each column corresponds to a given estimator: Realized volatility (RV), bipower variation (BV), tripower variation (TV), Realized Quantile RV (QRV), the MinRV and the MedRV. Each row corresponds to one of the models described in detail in Section 4.

Sub-Sampled Estimators	Relative Bias							Relative MSE						
	RV	BV	TV	QRV	MinRV	MedRV	RV	BV	TV	QRV	MinRV	MedRV		
12-sec frequency														
Model 1: BM	1.000	1.000	1.000	1.000	1.000	1.000	0.270	0.307	0.332	0.308	0.383	0.333		
Model 2: SV-U	0.999	0.998	0.998	0.941	0.998	0.998	0.271	0.310	0.335	1.480	0.386	0.336		
Model 3: BM + Sparsity	1.000	0.988	0.985	0.992	0.980	0.986	0.292	0.375	0.427	0.353	0.545	0.423		
Model 4: BM + 1 Jump	1.246	1.021	1.010	1.001	1.001	1.001	70.810	0.640	0.409	0.317	0.391	0.341		
Model 5: BM + 4 Jumps	1.251	1.042	1.025	1.007	1.007	1.008	48.769	1.279	0.634	0.327	0.475	0.469		
Model 6: BM + Noise	1.083	1.084	1.084	1.083	1.084	1.084	2.962	3.046	3.067	3.017	3.159	3.070		
Model 7: BM + 1 Bounceback	1.042	1.035	1.009	1.000	1.054	1.056	2.374	1.744	0.410	0.313	4.104	4.292		
60-sec frequency														
Model 1: BM	1.001	1.001	1.001	1.001	1.001	1.001	1.335	1.521	1.632	1.572	1.884	1.646		
Model 2: SV-U	0.997	0.994	0.991	0.940	0.994	0.991	1.334	1.495	1.596	2.509	1.842	1.606		
Model 3: BM + Sparsity	1.000	1.000	0.999	1.000	0.999	0.999	1.336	1.502	1.611	1.552	1.847	1.618		
Model 4: BM + 1 Jump	1.250	1.045	1.027	1.008	1.007	1.008	75.964	3.133	2.168	1.612	1.916	1.692		
Model 5: BM + 4 Jumps	1.249	1.084	1.061	1.030	1.030	1.035	50.385	5.833	3.708	2.058	2.770	2.868		
Model 6: BM + Noise	1.016	1.016	1.016	1.016	1.016	1.016	1.435	1.597	1.698	1.641	1.937	1.720		
Model 7: BM + 1 Bounceback	1.008	1.007	1.003	1.000	1.010	1.010	1.421	1.573	1.635	1.572	1.989	1.771		
300-sec frequency														
Model 1: BM	1.000	1.000	1.000	1.001	1.000	1.000	6.796	7.795	8.405	8.509	9.663	8.452		
Model 2: SV-U	0.988	0.976	0.966	0.881	0.976	0.965	6.529	7.186	7.613	10.520	8.810	7.667		
Model 3: BM + Sparsity	0.998	0.998	0.998	0.999	0.997	0.998	6.719	7.635	8.216	8.447	9.399	8.267		
Model 4: BM + 1 Jump	1.251	1.090	1.066	1.043	1.031	1.034	85.062	15.212	12.239	10.215	10.591	9.542		
Model 5: BM + 4 Jumps	1.244	1.150	1.127	1.115	1.098	1.110	56.661	23.918	19.528	18.068	17.665	18.560		
Model 6: BM + Noise	1.004	1.003	1.003	1.003	1.003	1.003	6.755	7.654	8.229	8.396	9.424	8.287		
Model 7: BM + 1 Bounceback	1.002	1.002	1.001	1.002	1.002	1.002	6.845	7.745	8.336	8.588	9.508	8.352		

Table C.4: **Relative bias and relative mean squared error (MSE) factors for pre-averaged IV estimators using 12-second, 60-second, and 300-second pre-averaging windows.** We report the bias and MSE factors for each pre-averaged estimator and each model based on pre-averaging window of 12 seconds (top panel), 60 seconds (middle panel), and 300 seconds (bottom panel). The MSE factor is computed as the sample mean of $390(\widehat{IV} - IV)^2/IQ$, where IV and IQ are the true simulated integrated variance and integrated quarticity on each day. Each column corresponds to a given estimator: Realized volatility (RV), bipower variation (BV), tripower variation (TV), Realized Quantile RV (QRV), the MinRV and the MedRV. Each row corresponds to one of the models described in detail in Section 4.

	Relative Bias							Relative MSE						
	RV	BV	TV	QRV	MinRV	MedRV		RV	BV	TV	QRV	MinRV	MedRV	
12-sec window	Model 1: BM	1.000	1.000	1.000	1.000	1.000		1.000	0.214	0.254	0.279	0.248	0.336	0.278
	Model 2: SV-U	0.999	0.999	0.998	0.941	0.999		0.998	0.213	0.255	0.278	1.422	0.338	0.279
	Model 3: BM + Sparsity	1.000	0.976	0.966	0.979	0.960		0.969	0.236	0.496	0.722	0.433	0.950	0.656
	Model 4: BM + 1 Jump	1.246	1.018	1.009	1.001	1.001		1.001	70.723	0.503	0.333	0.251	0.341	0.282
	Model 5: BM + 4 Jumps	1.251	1.036	1.021	1.006	1.006		1.007	48.703	0.973	0.492	0.259	0.376	0.332
	Model 6: BM + Noise	1.071	1.071	1.071	1.071	1.071		1.071	2.221	2.261	2.286	2.263	2.350	2.291
	Model 7: BM + 1 Bounceback	1.036	1.008	1.005	1.001	1.003		1.003	1.784	0.311	0.297	0.252	0.346	0.287
60-sec window	Model 1: BM	1.001	1.001	1.001	1.001	1.001		1.001	1.071	1.290	1.421	1.300	1.721	1.412
	Model 2: SV-U	0.997	0.994	0.991	0.940	0.994		0.991	1.079	1.272	1.393	2.306	1.689	1.394
	Model 3: BM + Sparsity	1.000	1.000	0.999	0.999	0.999		0.999	1.077	1.277	1.396	1.300	1.698	1.399
	Model 4: BM + 1 Jump	1.250	1.039	1.023	1.007	1.006		1.007	75.432	2.521	1.812	1.349	1.750	1.459
	Model 5: BM + 4 Jumps	1.250	1.073	1.051	1.025	1.024		1.027	50.138	4.509	2.876	1.627	2.243	2.059
	Model 6: BM + Noise	1.003	1.002	1.002	1.002	1.002		1.002	1.086	1.295	1.418	1.301	1.714	1.417
	Model 7: BM + 1 Bounceback	1.001	1.001	1.001	1.001	1.001		1.001	1.085	1.283	1.417	1.312	1.691	1.410
300-sec window	Model 1: BM	1.001	1.001	1.001	1.001	1.000		1.000	5.495	6.674	7.400	7.046	8.882	7.315
	Model 2: SV-U	0.988	0.976	0.965	0.881	0.976		0.966	5.353	6.212	6.763	9.621	8.162	6.745
	Model 3: BM + Sparsity	0.999	0.999	0.999	0.999	0.999		0.999	5.458	6.570	7.250	7.035	8.725	7.233
	Model 4: BM + 1 Jump	1.251	1.077	1.055	1.034	1.025		1.027	83.174	12.268	10.069	8.159	9.383	7.985
	Model 5: BM + 4 Jumps	1.246	1.132	1.108	1.092	1.077		1.086	55.577	19.167	15.380	12.986	13.976	13.369
	Model 6: BM + Noise	1.000	1.000	1.000	0.999	0.999		0.999	5.478	6.556	7.245	6.969	8.685	7.211
	Model 7: BM + 1 Bounceback	1.000	1.000	1.000	1.001	1.001		1.001	5.546	6.620	7.299	7.191	8.715	7.241