

Bootstrap inference for pre-averaged realized volatility based on non-overlapping returns*

Sílvia Gonçalves[†] Ulrich Hounyo[‡] and Nour Meddahi[§]

Université de Montréal, University of Oxford, and Toulouse School of Economics

February 22, 2013

Abstract

The main contribution of this paper is to propose bootstrap methods for realized volatility-like estimators defined on pre-averaged returns. In particular, we focus on the pre-averaged realized volatility estimator proposed by Podolskij and Vetter (2009). This statistic can be written (up to a bias correction term) as the (scaled) sum of squared pre-averaged returns, where the pre-averaging is done over all possible non-overlapping blocks of consecutive observations. Pre-averaging reduces the influence of the noise and allows for realized volatility estimation on the pre-averaged returns. The non-overlapping nature of the pre-averaged returns implies that these are asymptotically independent, but possibly heteroskedastic. This motivates the application of the wild bootstrap in this context. We provide a proof of the first order asymptotic validity of this method for percentile and percentile-t intervals. Our Monte Carlo simulations show that the wild bootstrap can improve the finite sample properties of the existing first order asymptotic theory provided we choose the external random variable appropriately. We use empirical work to illustrate its use in practice.

Keywords: High frequency data, realized volatility, pre-averaging, market microstructure noise, wild bootstrap.

1 Introduction

The increasing availability of financial return series measured over higher and higher frequencies (e.g. every minute or every second) has revolutionized the field of financial econometrics over the last decade. Researchers and practitioners alike now routinely rely on high frequency data to estimate volatility (and functionals of it, such as regression and correlation coefficients).

One earlier popular estimator was realized volatility, computed as the sum of squared intraday returns. This is a consistent estimator of integrated volatility (a measure of the ex-post variation of

*We would like to thank Ilze Kalnina for many useful comments and discussions. This work was supported by grants FQRSC-ANR and SSHRC. The third author benefited from the financial support of the chair “Marché des risques et création de valeur” Fondation du risque/SCOR.

[†]Département de sciences économiques, CIREQ and CIRANO, Université de Montréal. Address: C.P.6128, succ. Centre-Ville, Montréal, QC, H3C 3J7, Canada. Tel: (514) 343 6556. Email: silvia.Goncalves@umontreal.ca.

[‡]Oxford-Man Institute, University of Oxford. Address: Eagle House, Walton Well Road, Oxford OX2 6ED, UK. Email: ulrich.hounyo@oxford-man.ox.ac.uk.

[§]Toulouse School of Economics, 21 allée de Brienne -Manufacture des Tabacs-31000, Toulouse, France. Email: nour.meddahi@tse-fr.eu.

asset prices over a given day) under quite general assumptions on the volatility process. However, one important assumption underlying the consistency of realized volatility is the assumption that markets are frictionless (so that asset prices are observed without any error). This assumption does not hold in practice. As the sampling frequency increases, market microstructure effects such as the existence of bid-ask bounds, rounding errors, discrete trading prices, etc, contribute to a discrepancy between the true efficient price process and the price observed by the econometrician (known as the market microstructure noise).

The negative impact of market microstructure effects on realized volatility is now an accepted fact in the econometrics literature of high frequency data. A number of alternative estimators have been proposed that take into account these effects (see e.g. Zhou (1996), Zhang et al. (2005), Hansen and Lunde (2006), Bandi and Russell (2008), Barndorff-Nielsen et al. (2008), Podolskij and Vetter (2009) and Jacod et al. (2009)). Although these estimators rely on a large number of high frequency returns, finite sample distortions associated with the first order normal approximation may persist even at large sample sizes, as shown by our simulations.

In this paper, we consider the bootstrap as an alternative method of inference. We focus on the pre-averaging approach of Podolskij and Vetter (2009), where we first “average” the observed noisy returns over given blocks of non-overlapping observations, and then apply the standard realized volatility estimator to the pre-averaged returns. By averaging returns, the impact of the market microstructure noise is lessened, thus justifying realized volatility-like estimation on the pre-averaged returns. The class of statistics that we consider can be written (up to a bias term) as the (scaled) sum of squared pre-averaged returns (using an appropriate weighting function) computed over non-overlapping intervals. Our proposal is to bootstrap the pre-averaged returns.

Jacod et al. (2009) propose a generalization of the pre-averaging approach of Podolskij and Vetter (2009) which entails the use of overlapping intervals and the use of a more general weighting function for the pre-averaging of returns over these intervals. In this paper, we consider the case of non-overlapping returns only. The main reason is that the structure of dependence of the pre-averaged returns is much simpler in this case as compared to the overlapping case, which simplifies inference significantly. In the non-overlapping case, the pre-averaged returns are independent asymptotically (as the number of blocks increases) but possibly heteroskedastic (due to stochastic volatility). Thus a wild bootstrap applied to the pre-averaged returns is asymptotically valid. In contrast, overlapping pre-averaged returns (as in Jacod et al. (2009)) are very strongly dependent because they rely on common returns. Therefore, the wild bootstrap is not appropriate and more sophisticated bootstrap methods are required. In particular, Hounyo, Gonçalves and Meddahi (2013) show that a combination of the wild bootstrap with the blocks of blocks bootstrap of Bühlmann and Künsch (1995) (see also Künsch (1989), Politis and Romano (1992)) is asymptotically valid when applied to the pre-averaging estimator of Jacod et al. (2009). Although more generally applicable, the wild blocks of blocks bootstrap has the disadvantage of requiring the choice of a block size (in addition to the choice of the external random variable). For

this reason, here we focus on the simpler non-overlapping case.

Our main contribution is to provide a proof of the validity of the wild bootstrap. Specifically, we follow the literature and model the observed price process as the sum of the true but latent price process (defined as a Brownian semimartingale process subject to stochastic volatility of a general nonparametric form) plus a noise term which captures the market microstructure noise. As in Podolskij and Vetter (2009), the noise is assumed i.i.d. Under these assumptions, the pre-averaged returns are asymptotically independent and play the role of the original returns in the realized volatility estimator when no market microstructure noise exists. Therefore, the proof of the validity of the wild bootstrap in the present context where market microstructure effects exist parallels the proof of the validity of the wild bootstrap in the context of Gonçalves and Meddahi (2009), where the wild bootstrap was proposed for realized volatility under no market microstructure effects. Nevertheless, an important difference between these two applications is the fact that the pre-averaging estimator of integrated volatility entails an analytical bias correction term. As it turns out, this bias correction is only important for the proper centering of the confidence intervals and does not impact the variance of the estimator. As a consequence, we show that no bias correction term is needed in the bootstrap world (because we can always center the bootstrap statistic at its own theoretical mean, without affecting the bootstrap variance). This simplifies the application of the bootstrap in this context and justifies an approach solely based on bootstrapping the pre-averaged returns (as the bias term typically depends on the highest available frequency returns, which we are not resampling in the proposed approach).

We first discuss conditions under which the wild bootstrap variance is a consistent estimator of the (conditional) variance of the pre-averaged realized volatility. Specifically, we show that a necessary condition for the consistency of the wild bootstrap variance is that $\mu_4^* - (\mu_2^*)^2 = \frac{2}{3}$, where $\mu_q^* \equiv E|v_j|^q$ and v_j denotes the external random variable used to generate the wild bootstrap pre-averaged returns $\bar{Y}_j^* = \bar{Y}_j \cdot v_j$, where \bar{Y}_j are the pre-averaged returns. Under this condition, the bootstrap distribution of the scaled difference between the bootstrap pre-averaged realized volatility and its conditional mean is consistent for the (conditional) distribution of the pre-averaged realized volatility estimator. This result justifies the asymptotic validity of bootstrap percentile intervals for integrated volatility. Although this type of intervals does not promise asymptotic refinements over the first-order asymptotic approximation, they are easier to implement as they do not require an explicit estimator of the variance¹. We then discuss the first-order asymptotic validity of bootstrap percentile- t intervals. In this case, we propose a consistent bootstrap variance estimator and show that the studentized bootstrap statistic based on this estimator is asymptotically normal for any choice of the external random variable, provided we center and scale the bootstrap statistic appropriately.

¹In the univariate context considered here, the estimator of the variance of the pre-averaged realized volatility estimator is rather simple (it is given by a (scaled) version of the realized quarticity of pre-averaged returns), but this is not necessarily the case for other applications. For instance, for realized regression and realized correlation coefficients defined by the pre-averaging approach, the variance estimator is obtained by the delta method (whose finite sample properties are often poor) and the bootstrap percentile method could be useful in that context.

We provide a set of Monte Carlo experiments that compare the finite sample performance of the bootstrap with the existing mixed normal approximation. Our results show that the choice of the external random variable is rather important in finite samples. In particular, percentile intervals that do not satisfy the moment condition $\mu_4^* - (\mu_2^*)^2 = \frac{2}{3}$ behave quite poorly in finite samples, confirming our theoretical result. In contrast, asymptotically valid percentile intervals behave similarly to the asymptotic theory-based intervals and both are dominated by percentile- t bootstrap intervals. Although percentile- t intervals are asymptotically valid for any choice of the external random variable, their finite sample performance is also influenced by this choice. Our results show that matching the first four cumulants (including the variance but also the mean, the skewness and the kurtosis) of the studentized statistic is important for good coverage properties. The optimal choice proposed by Gonçalves and Meddahi (2009) fails to do so when the sample size is small and therefore does not work well in the simulations. This suggests that a different choice may be optimal in the present context. Deriving such a choice would require the development of an Edgeworth expansion for the studentized statistic based on the pre-averaged realized volatility estimator and is outside the scope of this paper. This is a non-trivial exercise given that the presence of the bias correction in the pre-averaged realized volatility estimator has an impact on the higher order cumulants, as our simulations shows. Instead, we show by simulation that a specific choice of the external random variable that does well in mimicking the first four cumulants of the statistic of interest has good finite sample coverage properties in the context of our Monte Carlo design.

The remainder of this paper is organized as follows. In Section 2, we introduce the basic model and the main assumptions. Furthermore, we review the existing first-order asymptotic theory. We also introduce the Monte Carlo design underlying all simulations in the paper and discuss the coverage probability results for the first-order asymptotic approach for nominal 95% two-sided symmetric intervals. In Section 3, we introduce our resampling method and prove its first-order asymptotic validity. In Section 4 we discuss the Monte Carlo results for bootstrap two-sided intervals. Section 5 contains an empirical application and Section 6 concludes. In the Appendix we give some technical results and present tables that illustrate the finite sample properties of the proposed procedures.

2 Setup, assumptions and review of existing results

2.1 Setup and assumptions

Let X denote the unobservable efficient log-price process defined on a probability space (Ω, \mathcal{F}, P) equipped with a filtration $(\mathcal{F}_t)_{t \geq 0}$. We model X as a Brownian semimartingale process defined by the equation

$$X_t = X_0 + \int_0^t \mu_s ds + \int_0^t \sigma_s dW_s, \quad t \geq 0, \quad (1)$$

where μ is a predictable locally bounded drift term, σ is an adapted càdlàg spot volatility process and W a standard Brownian motion. The object of interest is the quadratic variation of X given by

$$\int_0^T \sigma_s^2 ds,$$

also known as the integrated volatility. Without loss of generality, we let $T = 1$ and define $IV \equiv \int_0^1 \sigma_s^2 ds$ as the integrated volatility of X over a given time interval $[0, 1]$, which we think of as a given day.

The presence of market frictions such as price discreteness, rounding errors, bid-ask spreads, gradual response of prices to block trades, etc, prevent us from observing the true efficient price process X . Instead, we observe a noisy price process Y , given by

$$Y_t = X_t + \epsilon_t,$$

where ϵ_t represents the noise term that collects all the market microstructure effects. We assume that ϵ_t is i.i.d. and that ϵ_t is independent of X_t . Assumption 1 below collects these assumptions.

Assumption 1

- (i) The noise component ϵ_t is i.i.d. $(0, \omega^2)$ with $E|\epsilon_t|^{8+\varepsilon} < \infty$ for some $\varepsilon > 0$.
- (ii) ϵ_t is independent from the latent log-price X_t .

Assumption 1 is standard in the literature on market microstructure noise robust estimators of integrated volatility (see, among others, Zhang et al. (2005), Barndorff-Nielsen et al. (2008), Podolskij and Vetter (2009)). Nevertheless, empirically the i.i.d. assumption on ϵ and the independence between X and ϵ may be too strong a set of assumptions, especially at the highest frequencies. See e.g. Hansen and Lunde (2006), Zhang et al. (2011b), Diebold and Strasser (2012) for more on this issue. For simplicity, we will maintain these assumptions throughout.

Although for consistency of the pre-averaging estimator, $4 + \varepsilon$ moments of ϵ_t suffice (see, in particular, Theorem 1 of Podolskij and Vetter (2009) with $r = 2$ and $l = 0$), here we impose a stronger moment condition that requires the existence of $8 + \varepsilon$ moments. This is because we are interested in approximating the entire distribution of the studentized statistic based on the pre-averaging realized volatility estimator and we need a consistent estimator of its conditional variance. Consistency of the variance estimator requires this strengthening of the moment condition (see again Theorem 1 of Podolskij and Vetter (2009) with $r = 4$ and $l = 0$). Note that in contrast to Podolskij and Vetter (2009), we do not need to impose a Gaussianity assumption on ϵ , nor do we need to restrict the volatility process σ to be a Brownian semi-martingale. These assumptions are needed when studying the asymptotic properties of bipower or multipower pre-averaging statistics but can be dispensed with in the case of squared averaged returns (see Vetter (2008), p.49, for more details on this).

2.2 The pre-averaging approach

Suppose we observe Y at regular time points $\frac{i}{n}$, for $i = 0, \dots, n$, from which we compute n intraday returns at frequency $\frac{1}{n}$,

$$r_i \equiv Y_{\frac{i}{n}} - Y_{\frac{i-1}{n}}, \quad i = 1, \dots, n.$$

Given that $Y = X + \epsilon$, we can write

$$r_i = \left(X_{\frac{i}{n}} - X_{\frac{i-1}{n}} \right) + \left(\epsilon_{\frac{i}{n}} - \epsilon_{\frac{i-1}{n}} \right) \equiv r_i^e + \Delta\epsilon_i,$$

where $r_i^e = X_{\frac{i}{n}} - X_{\frac{i-1}{n}}$ denotes the $\frac{1}{n}$ -frequency return on the efficient price process.

We can show that

$$r_i = r_i^e + \Delta\epsilon_i = O_P\left(\frac{1}{\sqrt{n}}\right) + O_P(1). \quad (2)$$

Since X follows a stochastic volatility model given by (1), r_i^e is (conditionally on the path of σ and μ) independent and heteroskedastic with (conditional) variance given by $\int_{(i-1)/n}^{i/n} \sigma_s^2 ds$. The order of magnitude of r_i^e is thus $O_P\left(\frac{1}{\sqrt{n}}\right)$. In contrast, under Assumption 1, the difference $\Delta\epsilon_i \equiv \epsilon_{\frac{i}{n}} - \epsilon_{\frac{i-1}{n}}$ is an $MA(1)$ process whose order of magnitude is $O_P(1)$.

The decomposition in (2) shows that the noise completely dominates the observed return process as $n \rightarrow \infty$. This in turn implies that the usual realized volatility estimator is biased and inconsistent.

Moreover, even though the efficient returns r_i^e are conditionally independent, this is no longer the case for the observed returns. More specifically, the i.i.d. assumption on ϵ_t implies that the observed returns r_i are (conditionally) one-dependent due to the $MA(1)$ structure induced by the i.i.d. noise process.

Several approaches have been considered in the literature. Zhang et al. (2005) proposed a subsampling approach and derived the two times scale realized volatility estimator. This estimator amounts to using a linear combination of realized volatility estimators computed on subsamples (the slow scale) and an analytical bias correction term that relies on a realized volatility computed on a fast scale. Barndorff-Nielsen et al. (2008) proposed the realized kernel estimators, where linear combinations of autocovariances are considered. More recently, Podolskij and Vetter (2009) introduced the pre-averaging approach based on non-overlapping blocks. This was further generalized in Jacod et al. (2009) to allow for overlapping blocks.

In this paper we focus on bootstrapping the pre-averaged realized volatility estimator of Podolskij and Vetter (2009). As we mentioned before, our proposal is to bootstrap the pre-averaged returns. By focusing only on non-overlapping intervals, we can apply the wild bootstrap method to the pre-averaged returns. The dependence structure of the pre-averaged returns becomes much stronger under overlapping intervals and invalidates the use of the wild bootstrap. See Hounyo, Gonçalves and Meddahi (2013) for a bootstrap method that is valid in this context and which combines the wild bootstrap with a blocks of blocks bootstrap.

Next we describe the pre-averaging approach of Podolskij and Vetter (2009). This approach depends on two tuning parameters K and L , which denote two different block sizes. Specifically, let K denote the size of a block of K consecutive $\frac{1}{n}$ -horizon returns. Within each non-overlapping block of size K , we consider the set of all overlapping blocks of size L , where L is a fraction of K . For a given (non-overlapping) block of size K , there will be such $K - L + 1$ blocks of size L .

Assume that n/K is an integer so that the number of non-overlapping blocks of size K is n/K . For $j = 1, \dots, n/K$, the pre-averaged return \bar{Y}_j is obtained as follows:

$$\bar{Y}_j = \frac{1}{K - L + 1} \sum_{i=(j-1)K}^{jK-L} \left(\sum_{l=1}^L r_{l+i} \right).$$

This amounts to computing the sum of $\frac{1}{n}$ -horizon returns over each block of size L and then averaging the result over all possible such overlapping blocks. An alternative expression for \bar{Y}_j is as follows:

$$\bar{Y}_j = \sum_{i=1}^K g(i, K, L) r_{i+(j-1)K},$$

where for every $i = 1, \dots, K$, the weighting function $g(i, K, L)$ is defined as

$$g(i, K, L) = \begin{cases} \frac{i}{K-L+1}, & \text{if } i \in \{1, \dots, L\} \\ \frac{L}{K-L+1}, & \text{if } i \in \{L+1, \dots, K-L\} \\ \frac{K-i+1}{K-L+1}, & \text{if } i \in \{K-L+1, \dots, K\} \end{cases},$$

and where we can show that $\sum_{i=1}^K g(i, K, L) = L$.

The effect of pre-averaging is to reduce the impact of the noise in the pre-averaged return. Specifically, we can show that by pre-averaging returns over blocks of size K in this particular manner, we reduce the variance by a factor of about $\frac{1}{K}$. To be more precise, Podolskij and Vetter (2009) show that

$$\bar{Y}_j = \bar{r}_j^e + \Delta \bar{\epsilon}_j = O_P \left(\sqrt{\frac{L}{n}} \right) + O_P \left(\frac{1}{\sqrt{K-L}} \right), \quad (3)$$

where \bar{r}_j^e and $\Delta \bar{\epsilon}_j$ denote the pre-averaged versions of the efficient returns and the difference of the noise process, respectively. Thus, comparing (2) with (3), we see that pre-averaging manages to reduce the impact of the noise from $O_P(1)$ to $O_P\left(\frac{1}{\sqrt{K-L}}\right)$. Since L is a fraction of K , i.e. $L \sim \frac{1}{c_2}K$, for some $c_2 > 1$, the order of magnitude of the noise in (3) is $O_P\left(\frac{1}{\sqrt{K}}\right)$. The overall implication is that we can compute a realized volatility-like estimator on the pre-averaged returns \bar{Y}_j . This is the essence of the pre-averaging approach.

To give the explicit formula of the pre-averaging realized volatility estimator of Podolskij and Vetter (2009), we need to introduce some additional notation. In particular, we let

$$L = \frac{1}{c_2}K, \quad (4)$$

with $c_2 > 1$, and

$$K = c_1 c_2 \sqrt{n}, \quad (5)$$

where $c_1 > 0$, and c_1 and c_2 are two tuning parameters that need to be chosen. These choices of K and L imply that the two terms in (3) are balanced and equal to $O_P(n^{-1/4})$.

Under Assumption 1, and assuming that K and L satisfy the conditions (4) and (5), respectively, Podolskij and Vetter (2009) [cf. Theorem 1] show that

$$p \lim_{n \rightarrow \infty} \left(\sum_{j=1}^{n/K} \bar{Y}_j^2 \right) = \frac{\nu_1}{c_1 c_2} \int_0^1 \sigma_s^2 ds + \frac{\nu_2}{c_1 c_2} \omega^2,$$

where $\omega^2 = \text{Var}(\epsilon_i)$ and where

$$\nu_1 = \frac{c_1 \left(3c_2 - 4 + \max \left((2 - c_2)^3, 0 \right) \right)}{3(c_2 - 1)^2}, \quad \nu_2 = \frac{2(\min((c_2 - 1), 1))}{c_1(c_2 - 1)^2}.$$

Two implications can be obtained from this result. First, the particular weighting scheme induced by the pre-averaging approach introduces a scaling factor given by $\frac{\nu_1}{c_1 c_2}$ when estimating $\int_0^1 \sigma_s^2 ds$. This implies that we need to scale $\sum_{j=1}^{n/K} \bar{Y}_j^2$ by $\frac{c_1 c_2}{\nu_1}$. Second, although the pre-averaging approach reduces the order of magnitude of the noise, it does not completely eliminate its influence. In particular,

$$p \lim_{n \rightarrow \infty} \left(\frac{c_1 c_2}{\nu_1} \sum_{j=1}^{n/K} \bar{Y}_j^2 \right) = \int_0^1 \sigma_s^2 ds + \underbrace{\frac{\nu_2}{\nu_1} \omega^2}_{\text{Bias term}},$$

where the bias term is proportional to the variance of the noise ω^2 . A consistent estimator of ω^2 is given by the realized volatility estimator computed on the n highest frequency returns r_i , divided by $2n$, i.e.

$$\hat{\omega}^2 = \frac{\sum_{i=1}^n r_i^2}{2n} \xrightarrow{P} \omega^2.$$

This suggests the following consistent estimator of integrated volatility:

$$PRV_n = \underbrace{\frac{c_1 c_2}{\nu_1} \sum_{j=1}^{n/K} \bar{Y}_j^2}_{\text{RV-like estimator}} - \underbrace{\frac{\nu_2}{\nu_1} \hat{\omega}^2}_{\text{bias correction term}}.$$

2.3 First-order asymptotic distribution theory

Under Assumption 1, and assuming that K and L are chosen according to (4) and (5), Podolskij and Vetter (2009) (cf. Corollary 1) show that

$$\frac{n^{1/4} \left(PRV_n - \int_0^1 \sigma_s^2 ds \right)}{\sqrt{V}} \xrightarrow{st} N(0, 1), \quad (6)$$

where \xrightarrow{st} denotes stable convergence (see Christensen and al. (2009), p. 119 for a definition of stable convergence), and

$$V = \frac{2c_1^2 c_2^2}{\nu_1^2} \int_0^1 (\nu_1 \sigma_s^2 + \nu_2 \omega^2)^2 ds$$

is the conditional variance of PRV_n .

By Theorem 1 of Podolskij and Vetter (2009), a consistent estimator of V is given by

$$\hat{V}_n = \frac{2c_1^2 c_2^2}{3\nu_1^2} \sqrt{n} \sum_{j=1}^{n/K} |\bar{Y}_j|^4.$$

This estimator has the form of a realized quarticity estimator applied to the pre-averaged returns \bar{Y}_j . Together with the CLT result (6), it implies that (cf. equation (3.19) in Podolskij and Vetter (2009))

$$T_n \equiv \frac{n^{1/4} \left(PRV_n - \int_0^1 \sigma_s^2 ds \right)}{\sqrt{\hat{V}_n}} \rightarrow^d N(0, 1).$$

We can use this feasible asymptotic distribution result to build confidence intervals for integrated volatility. In particular, a two-sided feasible $100(1 - \alpha)\%$ level interval for $\int_0^1 \sigma_s^2 ds$ is given by:

$$IC_{Feas, 1-\alpha} = \left(PRV_n - z_{1-\alpha/2} n^{-1/4} \sqrt{\hat{V}_n}, PRV_n + z_{1-\alpha/2} n^{-1/4} \sqrt{\hat{V}_n} \right),$$

where $z_{1-\alpha/2}$ is such that $\Phi(z_{1-\alpha/2}) = 1 - \alpha/2$, and $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. For instance, $z_{0.975} = 1.96$ when $\alpha = 0.05$.

2.4 Finite sample properties of the feasible asymptotic approach

In this section we assess by Monte Carlo simulation the accuracy of the feasible asymptotic theory of the pre-averaging approach of Podolskij and Vetter (2009). We find that this approach leads to important coverage probability distortions when returns are not sampled too frequently. This motivates the bootstrap as an alternative method of inference in this context.

We consider two data generating processes in our simulations. First, following Zhang et al. (2005), we use the one-factor stochastic volatility (SV1F) model of Heston (1993) as our data-generating process, i.e.

$$dX_t = (\mu - \nu_t/2) dt + \sigma_t dB_t,$$

and

$$d\nu_t = \kappa(\alpha - \nu_t) dt + \gamma(\nu_t)^{1/2} dW_t,$$

where $\nu_t = \sigma_t^2$, B and W are two Brownian motions, and we assume $Corr(B, W) = \rho$. The parameter values are all annualized. In particular, we let $\mu = 0.05/252$, $\kappa = 5/252$, $\alpha = 0.04/252$, $\gamma = 0.05/252$, $\rho = -0.5$. The size of the market microstructure noise is an important parameter. We follow Barndorff-Nielsen et al. (2009) and model the noise magnitude as $\xi^2 = \omega^2 / \sqrt{\int_0^1 \sigma_s^4 ds}$. We fix ξ^2 be equal to 0.0001, 0.001 and 0.01, and let $\omega^2 = \xi^2 \sqrt{\int_0^1 \sigma_s^4 ds}$. These values are motivated by the empirical study of Hansen and Lunde (2006), who investigate 30 stocks of Dow Jones Industrial Average.

We also consider the two-factor stochastic volatility (SV2F) model analyzed by Barndorff-Nielsen

et al. (2009), where ²

$$\begin{aligned}
dX_t &= \mu dt + \sigma_t dB_t, \\
\sigma_t &= s\text{-exp}(\beta_0 + \beta_1 \tau_{1t} + \beta_2 \tau_{2t}), \\
d\tau_{1t} &= \alpha_1 \tau_{1t} dt + dB_{1t}, \\
d\tau_{2t} &= \alpha_2 \tau_{2t} dt + (1 + \phi \tau_{2t}) dB_{2t}, \\
\text{corr}(dW_t, dB_{1t}) &= \varphi_1, \text{corr}(dW_t, dB_{2t}) = \varphi_2.
\end{aligned}$$

We follow Huang and Tauchen (2005) and set $\mu = 0.03$, $\beta_0 = -1.2$, $\beta_1 = 0.04$, $\beta_2 = 1.5$, $\alpha_1 = -0.00137$, $\alpha_2 = -1.386$, $\phi = 0.25$, $\varphi_1 = \varphi_2 = -0.3$. We initialize the two factors at the start of each interval by drawing the persistent factor from its unconditional distribution, $\tau_{10} \sim N\left(0, \frac{-1}{2\alpha_1}\right)$, by starting the strongly mean-reverting factor at zero.

We simulate data for the unit interval $[0, 1]$ and normalize one second to be $1/23400$, so that $[0, 1]$ is thought to span 6.5 hours. The observed Y process is generated using an Euler scheme. We then construct the $\frac{1}{n}$ -horizon returns $r_i \equiv Y_{i/n} - Y_{(i-1)/n}$ based on samples of size n .

The pre-averaging approach requires the choice of the tuning parameters c_1 and c_2 . Podolskij and Vetter (2009) give the optimal values of c_1 and c_2 that minimize the conditional variance V of the PRV_n estimator when the volatility process is constant. In our simulations, we followed Podolskij and Vetter (2009) and let $c_2 = 1.6$ and $c_1 = 1$. These choices may not be optimal under stochastic volatility, but since we will compute the bootstrap statistics using these same values, they allow for a meaningful comparison of the different intervals for integrated volatility (asymptotic theory-based and bootstrap intervals).

Table 1 gives the actual rates of 95% confidence intervals of integrated volatility for the SV1F and the SV2F models, respectively, computed over 10,000 replications. Results are presented for eight different samples sizes: $n = 23400, 11700, 7800, 4680, 1560, 780, 390$ and 195, corresponding to “1-second”, “2-second”, “3-second”, “5-second”, “15-second”, “30-second”, “1-minute” and “2-minute” frequencies (this table also includes results for the bootstrap methods but those results will be discussed later in Section 3.)

For the two models, all intervals tend to undercover. The degree of undercoverage is especially large for smaller values of n , when sampling is not too frequent. The SV2F model exhibits overall larger coverage distortions than the SV1F model, for all sample sizes. Results are not very sensitive to the noise magnitude.

²The function $s\text{-exp}$ is the usual exponential function with a linear growth function splined in at high values of its argument: $s\text{-exp}(x) = \exp(x)$ if $x \leq x_0$ and $s\text{-exp}(x) = \frac{\exp(x_0)}{\sqrt{x_0 - x_0^2 + x^2}}$ if $x > x_0$, with $x_0 = \log(1.5)$.

3 The bootstrap

In this section we provide a bootstrap method for inference on integrated volatility based on the pre-averaging approach of Podolskij and Vetter (2009). Our proposal is to bootstrap the pre-averaged returns \bar{Y}_j , $j = 1, \dots, n/K$. Because non-overlapping intervals are used, the pre-averaged returns \bar{Y}_j are asymptotically independent, as $n \rightarrow \infty$. In fact, we can show that they are one-dependent, i.e. \bar{Y}_j is independent of \bar{Y}_m whenever $|m - j| > 1$. Moreover, the amount of dependence between two consecutive squared pre-averaged returns is very small and it is only due to edge effects. Specifically, $Cov(\bar{Y}_j^2, \bar{Y}_{j+1}^2) = O(\frac{1}{n^2}) = o(1)$ as $n \rightarrow \infty$.

Since pre-averaged returns are asymptotically independent but possibly heteroskedastic (due to the fact that volatility is time-varying) a wild bootstrap approach is appropriate. The wild bootstrap method was introduced by Wu (1986), and further studied by Liu (1988) and Mammen (1993), in the context of cross-section linear regression models subject to unconditional heteroskedasticity in the error term. Gonçalves and Meddahi (2009) applied the wild bootstrap method in the context of realized volatility under no market microstructure noise. Our approach here follows Gonçalves and Meddahi (2009), but instead of bootstrapping the $\frac{1}{n}$ -horizon raw returns r_i , we propose to bootstrap the pre-averaged returns \bar{Y}_j .

The bootstrap pseudo-data is given by

$$\bar{Y}_j^* = \bar{Y}_j \cdot v_j, \quad j = 1, \dots, n/K,$$

where the external random variable v_j is an i.i.d. random variable independent of the data and whose moments are given by $\mu_q^* \equiv E^* |v_j|^q$. As usual in the bootstrap literature, P^* (E^* and Var^*) denotes the probability measure (expected value and variance) induced by the bootstrap resampling, conditional on a realization of the original time series. In addition, for a sequence of bootstrap statistics Z_n^* , we write $Z_n^* = o_{P^*}(1)$ in probability, or $Z_n^* \rightarrow^{P^*} 0$, as $n \rightarrow \infty$, in probability, if for any $\varepsilon > 0$, $\delta > 0$, $\lim_{n \rightarrow \infty} P[P^*(|Z_n^*| > \delta) > \varepsilon] = 0$. Similarly, we write $Z_n^* = O_{P^*}(1)$ as $n \rightarrow \infty$, in probability if for all $\varepsilon > 0$ there exists a $M_\varepsilon < \infty$ such that $\lim_{n \rightarrow \infty} P[P^*(|Z_n^*| > M_\varepsilon) > \varepsilon] = 0$. Finally, we write $Z_n^* \rightarrow^{d^*} Z$ as $n \rightarrow \infty$, in probability, if conditional on the sample, Z_n^* weakly converges to Z under P^* , for all samples contained in a set with probability P converging to one.

The bootstrap pre-averaged realized volatility estimator is given by

$$PRV_n^* = \frac{c_1 c_2}{\nu_1} \sum_{j=1}^{n/K} \bar{Y}_j^{*2}.$$

Although the pre-averaged realized volatility estimator PRV_n contains a bias correction term, we do not consider bias correction in the bootstrap world. The reason is twofold. First, our goal is not to estimate consistently the integrated volatility using the bootstrap. Instead, our goal is to use the bootstrap to approximate the distribution of statistics based on PRV_n , for instance we would like to approximate the distribution of the t -statistic T_n defined in the previous section. We can easily show

that

$$E^*(PRV_n^*) = \mu_2^* \frac{c_1 c_2}{\nu_1} \sum_{j=1}^{n/K} \bar{Y}_j^2.$$

This is a biased estimator of integrated volatility, but we can correctly center our bootstrap statistics using this theoretical bootstrap mean. Since the bias correction term does not affect the variance of the pre-averaging estimator, as long as the bootstrap method is able to consistently estimate this variance, no bias correction is needed in the bootstrap world. The second reason why we do not consider bootstrap bias correction is that the bootstrap bias correction term would involve the bootstrap highest frequency returns r_i^* , which are not available under our proposed method.

We can show that

$$Var^*(n^{1/4}PRV_n^*) = (\mu_4^* - \mu_2^{*2}) \underbrace{\frac{c_1^2 c_2^2}{\nu_1^2} \sqrt{n} \sum_{j=1}^{n/K} |\bar{Y}_j|^4}_{\equiv \frac{3}{2} \hat{V}_n}.$$

It follows then that a sufficient condition for the bootstrap to provide a consistent estimator of the conditional variance of $n^{1/4}PRV_n$ is that $\mu_4^* - \mu_2^{*2} = \frac{2}{3}$. Under this condition, the bootstrap can be used to approximate the quantiles of the distribution of the root

$$n^{1/4} \left(PRV_n - \int_0^1 \sigma_s^2 ds \right),$$

thus justifying the construction of bootstrap percentile confidence intervals.

These results are summarized in the following theorem.

Theorem 3.1. *Suppose Assumption 1 holds and let K and L satisfy the conditions (4) and (5), respectively. Suppose that $\{\bar{Y}_j^* = \bar{Y}_j \cdot v_j : j = 1, \dots, n/K\}$, where $v_j \sim i.i.d.$ such that for any $\delta > 0$, $\mu_{2(2+\delta)}^* = E^* |v_j|^{2(2+\delta)} < \infty$. If $\mu_4^* - (\mu_2^*)^2 = \frac{2}{3}$, then as $n \rightarrow \infty$,*

- (1) $V_n^* \equiv Var^*(n^{1/4}PRV_n^*) \xrightarrow{P} V \equiv \frac{2c_1^2 c_2^2}{\nu_1^2} \int_0^1 (\nu_1 \sigma_s^2 + \nu_2 \omega^2)^2 ds.$
- (2) $\sup_{x \in \mathbb{R}} \left| P^*(n^{1/4}(PRV_n^* - E^*(PRV_n^*)) \leq x) - P\left(n^{1/4}\left(PR V_n - \int_0^1 \sigma_s^2 ds\right)\right) \right| \xrightarrow{P} 0.$

An example of a random variable that satisfies the condition $\mu_4^* - \mu_2^{*2} = \frac{2}{3}$ is

$$v_j \sim i.i.d. N\left(0, \sqrt{3}/3\right).$$

Theorem 3.1 justifies using the wild bootstrap to construct bootstrap percentile intervals for integrated volatility. Specifically, a $100(1 - \alpha)\%$ symmetric bootstrap percentile interval for integrated volatility based on the bootstrap is given by

$$IC_{perc, 1-\alpha}^* = \left(PRV_n - n^{-1/4} p_{1-\alpha}^*, PRV_n + n^{-1/4} p_{1-\alpha}^* \right), \quad (7)$$

where $p_{1-\alpha}^*$ is the $1 - \alpha$ quantile of the bootstrap distribution of $|n^{1/4}(PRV_n^* - E^*(PRV_n^*))|$.

Bootstrap percentile intervals do not promise asymptotic refinements. Next, we propose a consistent bootstrap variance estimator that allows us to form bootstrap percentile- t intervals. More specifically, we can show that the following bootstrap variance estimator consistently estimates V_n^* for any choice of the external random variable v_j :

$$\hat{V}_n^* = \frac{\mu_4^* - \mu_2^{*2}}{\mu_4^*} \frac{c_1^2 c_2^2}{\nu_1^2} n^{1/2} \sum_{j=1}^{n/K} \bar{Y}_j^{*4}.$$

Our proposal is to use this estimator to construct a bootstrap studentized statistic,

$$T_n^* \equiv \frac{n^{1/4} (PRV_n^* - E^*(PRV_n^*))}{\sqrt{\hat{V}_n^*}},$$

the bootstrap analogue of T_n .

Theorem 3.2. *Suppose Assumption 1 holds such that for any $\delta > 0$, $E|\epsilon_t|^{2(8+\delta)} < \infty$, and let K and L satisfy the conditions (4) and (5), respectively. Suppose that $\{\bar{Y}_j^* = \bar{Y}_j \cdot v_j : j = 1, \dots, n/K\}$, where $v_j \sim i.i.d.$ such that $\mu_8^* = E^*|v_j|^8 < \infty$. It follows that as $n \rightarrow \infty$, $\sup_{x \in \mathbb{R}} |P^*(T_n^* \leq x) - P(T_n \leq x)| \xrightarrow{P} 0$.*

Theorem 3.2 justifies constructing bootstrap percentile- t intervals. In particular, a $100(1 - \alpha)\%$ symmetric bootstrap percentile- t interval for integrated volatility is given by

$$IC_{perc-t, 1-\alpha}^* = \left(PRV_n - q_{1-\alpha}^* n^{-1/4} \sqrt{\hat{V}_n}, PRV_n + q_{1-\alpha}^* n^{-1/4} \sqrt{\hat{V}_n} \right), \quad (8)$$

where $q_{1-\alpha}^*$ is the $(1 - \alpha)$ -quantile of the bootstrap distribution of $|T_n^*|$. The first order asymptotic validity of the bootstrap requires a strengthening of the moment condition on ϵ_t when applied to the feasible statistic T_n .

4 Monte Carlo results for the bootstrap

In this section, we compare the finite sample performance of the bootstrap with the first order asymptotic theory for confidence intervals of integrated volatility. In our simulations, bootstrap intervals use 999 bootstrap replications for each of the 10,000 Monte Carlo replications.

We consider bootstrap percentile and bootstrap percentile- t intervals, computed at the 95% level using (7) and (8), respectively.

To generate the bootstrap data we use three different external random variables.

WB1 $v_j \sim i.i.d. N(0, \sqrt{3}/3)$, implying that $\mu_2^* = \sqrt{3}/3$ and $\mu_4^* = 1$.

WB2 A two point distribution $v_j \sim i.i.d.$ such that:

$$v_j = \begin{cases} \left(\frac{2}{3}\right)^{1/4} \frac{-1+\sqrt{5}}{2}, & \text{with prob } p = \frac{\sqrt{5}-1}{2\sqrt{5}} \\ \left(\frac{2}{3}\right)^{1/4} \frac{-1-\sqrt{5}}{2}, & \text{with prob } 1-p = \frac{\sqrt{5}+1}{2\sqrt{5}} \end{cases},$$

for which $\mu_2^* = 2\sqrt{2/3}$ and $\mu_4^* = 10/3$.

WB3 The two point distribution proposed by Gonçalves and Meddahi (2009), where $v_j \sim$ i.i.d. such that:

$$v_j = \begin{cases} \frac{1}{5}\sqrt{31 + \sqrt{186}}, & \text{with prob } p = \frac{1}{2} - \frac{3}{\sqrt{186}} \\ -\frac{1}{5}\sqrt{31 - \sqrt{186}}, & \text{with prob } 1 - p \end{cases},$$

for which we have $\mu_2^* = 1$ and $\mu_4^* = 31/25$.

The condition $\mu_4^* - (\mu_2^*)^2 = \frac{2}{3}$ is satisfied for the first two choices (WB1 and WB2) but not for WB3. The implication is that WB1 and WB2 are valid for percentile intervals but not WB3. Note however that all three choices of v_j are asymptotically valid when used to construct bootstrap percentile- t intervals.

Table 1 shows the actual coverage probability rates of nominal 95% symmetric bootstrap intervals for integrated volatility based on WB1, WB2 and WB3 for each of the two models (SV1F and SV2F). Both percentile and percentile- t intervals are considered. Results based on the asymptotic normal distribution are also included (under the label CLT). As already discussed in Section 2.4, results are not very sensitive to the choice of ξ^2 and distortions are larger (both based on asymptotic theory and on the bootstrap) for the SV2F than for the SV1F model. These trends are also present for the bootstrap.

Starting with the bootstrap percentile intervals, we see that these are close to the CLT-based intervals for WB1 and WB2 (when the condition $\mu_4^* - (\mu_2^*)^2 = \frac{2}{3}$ is satisfied) whereas coverage rates for percentile intervals based on WB3 are systematically much lower than 95% even for the largest sample sizes. This confirms the theoretical prediction of asymptotic invalidity for these intervals. The results also confirm that the bootstrap percentile intervals do not outperform the asymptotic theory-based intervals. Nevertheless, choosing v_j to match the variance of the pre-averaging estimator may result in better percentile- t intervals, as a comparison the different bootstrap methods shows for this type of intervals. Specifically, although WB2 and WB3 both undercover for smaller sample sizes, WB2 outperforms WB3 significantly for the smaller samples sizes. For instance, for SV1F, WB3 covers IV 81.41% of the time when $n = 195$ whereas WB2 does so 91.05%. These rates decrease to 71.89% and 86.78% for the SV2F model, respectively. In contrast, the WB1 method covers IV with a rate equal to 97.91% for SV1F and 94.72% for SV2F, when $n = 195$. In general, the results show that percentile- t intervals based on WB1 are too conservative, yielding coverage rates larger than 95%, especially for the SV1F model. WB2 intervals tend to be closer to the desired nominal level than the WB3 method, without being conservative. Overall, the results suggest that the choice of v_j is important in finite samples.

In order to gain further insight into why the different choices of v_j matter in finite samples, we computed the first four cumulants of T_n and of its bootstrap analogue T_n^* . The results are presented in Table 2, which also reports the coverage rates of symmetric intervals based on these studentized

statistics. Results are only given for $\xi^2 = 0.01$. For T_n , we report the mean, the standard error, the excess skewness and the excess kurtosis across the 10,000 simulations. For T_n^* , the numbers correspond to the average value (across the 10,000 simulations) of the bootstrap mean, standard error, excess skewness and excess kurtosis computed for each simulation across the 999 bootstrap replications.

Starting with T_n , the results show that this statistic is centered at a negative value across the different sample sizes. The negative bias decreases as n increases, but it can be quite large when n is small. Since the asymptotic normal distribution is centered at zero, it completely misses this downward bias. We can also see that the finite sample distribution of T_n is more dispersed than the $N(0, 1)$ distribution (its standard error is larger than 1), and that it is strongly negatively skewed (the excess skewness is very negative) and fat-tailed (the excess kurtosis is positive). All these features explain the undercoverage of the CLT approach. In contrast, the bootstrap cumulants of T_n^* replicate to a better degree the finite sample patterns of the four cumulants of T_n depending on the choice of v_j . Specifically, we can see that the three choices of v_j typically induce a negative bias as well as negative excess skewness and positive excess kurtosis (an exception is WB3 for the smaller sample sizes). Nevertheless, WB1 implies too strong a correction. For instance, the bias of T_n^* is more negative than it should be on average as well as its excess skewness. This means that the bootstrap distribution of T_n^* is on average to the left of the finite sample distribution of T_n , resulting in too large a critical value, which explains the overcoverage problem noted in Table 1. In contrast, for the smaller sample sizes, WB2 and WB3 imply too little a correction in terms of the bias, which implies that these bootstrap distributions are on average centered to the right of the true distribution of T_n . This contributes to too small a critical value and to some undercoverage.

Overall, the results suggest that WB3 does a poorer job at capturing the first four cumulants than WB2, especially for the smaller sample sizes. This suggests that the optimal choice of v_j proposed by Gonçalves and Meddahi (2009) in the context of realized volatility without market microstructure noise is no longer optimal in the context of pre-averaging realized volatility. The presence of the bias correction term in the definition of PRV_n implies that the Edgeworth expansions derived in Gonçalves and Meddahi (2009) do not apply in the pre-averaging approach considered here. Thus, although bias correction does not have an impact to first order on the asymptotic variance of PRV_n , it likely has an impact on the higher order cumulants, as our Monte Carlo simulation results suggest. Deriving the optimal choice of the external random variable in this context is an interesting research question which we will consider elsewhere.

5 Empirical results

As a brief illustration, in this section we implement the proposed wild bootstrap method to real high frequency data, and compare it to the existing feasible asymptotic procedure of Podolskij and Vetter (2009). The data consists of transaction log prices of General Electric (GE) shares carried out on the

New York Stock Exchange (NYSE) in December 2011. For each day, we consider data from the regular exchange opening hours from time stamped between 9:30 a.m. until 4 p.m. Our procedure for cleaning the data is exactly identical to that used by Barndorff-Nielsen et al. (2008).

We implement the pre-averaged realized volatility estimator of Jacod et al. (2009) on returns recorded every S transactions, where S is selected each day so that there are approximately 1493 observations a day. This means that on average these returns are recorded roughly every 15 seconds. Table 3 in the Appendix provides the number of transactions per day (8 on average) and the sample size for the pre-averaged returns. The pre-averaged realized volatility estimator is implemented with $c_2 = 1.6$ and $c_1 = 1$.

We consider bootstrap percentile- t intervals, computed at the 95% level using (8), where v_j is generated using WB2 (our best choice according to the Monte Carlo simulations). The results are displayed in Figure 1 in terms of daily 95% confidence intervals (CIs) for integrated volatility. Two types of intervals are presented: our proposed wild bootstrap method and the existing feasible asymptotic procedure Podolskij and Vetter (2009). The pre-averaged realized volatility estimate is in the center of both confidence intervals by construction.

The confidence intervals for IV based on the bootstrap method are usually wider than the confidence intervals using the feasible asymptotic theory. Nevertheless, as our Monte Carlo simulations showed, the latter typically have undercoverage problems whereas the bootstrap intervals have coverage rates closer to the desired level. Therefore if the goal is to control the coverage probability, shorter intervals are not necessarily better. The figures also show a lot of variability in the daily estimate of integrated volatility.

6 Conclusion

In this paper we propose the wild bootstrap as a method of inference for integrated volatility in the context of the pre-averaged realized volatility estimator proposed by Podolskij and Vetter (2009). The wild bootstrap is motivated by the fact that non-overlapped pre-averaged returns are asymptotically independent but possibly heteroskedastic (in the context of stochastic volatility models). We provide a set of conditions under which this method is asymptotically valid to first order. Both percentile and percentile- t bootstrap intervals are considered. Our Monte Carlo simulations show that the bootstrap can improve upon the mixed Gaussian inference derived by Podolskij and Vetter (2009) provided we choose the external random variable appropriately.

An important question for future research is the optimal choice of the external random variable in this context. This is not an easy question because it requires developing Edgeworth expansions for the statistics of interest in the original sample and the bootstrap samples. Since the pre-averaged realized volatility estimator depends on a bias correction term, its Edgeworth expansion will reflect the contribution of this term at higher orders and render the analysis rather complex. We plan on

investigating this issue in future work.

Appendix A

Table 1. Coverage rate of Nominal 95 % intervals

n	SV1F								SV2F							
	CLT	WB1		WB2		WB3		CLT	WB1		WB2		WB3			
		Perc	Perc-t	Perc	Perc-t	Perc	Perc-t		Perc	Perc-t	Perc	Perc-t	Perc	Perc-t		
$\xi^2 = 0.0001$																
195	77.54	77.49	97.91	76.42	91.05	61.11	81.41	69.49	69.38	94.72	68.51	86.78	55.51	71.89		
390	84.85	84.47	98.42	83.51	93.71	66.76	90.20	77.97	77.64	96.17	76.89	89.88	62.87	82.42		
780	86.82	86.11	98.43	85.41	93.94	67.73	92.91	80.61	80.19	96.24	79.09	90.17	63.36	85.87		
1560	88.89	88.13	98.36	87.74	93.93	69.58	94.36	83.36	82.89	96.63	82.03	90.87	65.16	89.07		
4680	91.49	90.65	98.63	90.78	94.69	72.56	96.59	86.17	85.59	96.76	85.41	91.74	67.71	91.92		
7800	92.78	92.04	98.56	92.34	95.12	73.24	96.97	89.50	88.66	97.46	88.59	93.33	70.11	94.21		
11700	93.01	92.41	98.35	92.63	95.11	73.40	97.16	89.05	88.34	97.09	88.27	93.15	70.23	94.15		
23400	93.48	92.85	98.06	93.09	94.89	74.26	97.56	89.89	89.06	96.86	89.33	92.81	71.13	94.67		
$\xi^2 = 0.001$																
195	77.63	77.46	97.90	76.56	90.92	61.18	81.54	69.72	69.81	94.78	68.83	86.59	55.86	71.86		
390	85.02	84.48	98.50	83.66	93.75	66.71	90.38	77.95	77.73	96.14	76.97	89.93	63.17	82.57		
780	86.81	86.11	98.43	85.22	93.86	67.91	92.76	80.55	80.23	96.14	79.36	90.25	63.68	85.94		
1560	88.91	88.13	98.48	87.74	93.94	69.51	94.46	83.26	82.70	96.66	82.10	90.91	65.16	89.12		
4680	91.47	90.76	98.67	90.78	94.78	72.55	96.59	86.33	85.66	96.68	85.37	91.94	67.86	91.93		
7800	92.86	91.91	98.56	92.37	95.06	73.47	97.00	89.56	88.73	97.54	88.53	93.36	70.09	94.25		
11700	92.98	92.25	98.32	92.56	95.14	73.57	97.10	88.95	88.19	97.00	88.15	93.04	70.16	94.23		
23400	93.45	92.88	98.12	93.12	94.92	74.18	97.51	90.01	89.16	96.80	89.40	92.83	71.22	94.80		
$\xi^2 = 0.01$																
195	77.93	77.67	97.67	76.86	91.23	61.69	81.35	70.17	70.12	94.80	69.28	86.74	56.07	72.73		
390	85.09	84.57	98.35	83.61	93.59	66.85	90.28	78.59	78.46	96.42	77.43	90.00	62.89	83.60		
780	86.75	86.29	98.38	85.31	93.47	67.96	92.75	81.29	80.90	96.33	79.92	90.36	63.86	86.54		
1560	89.03	88.12	98.41	87.74	94.02	69.16	94.54	83.45	82.68	96.51	82.20	91.06	65.40	89.59		
4680	91.42	90.54	98.78	90.66	94.62	72.39	96.64	86.78	86.04	96.57	85.67	91.97	68.07	92.17		
7800	92.61	91.77	98.63	92.24	94.90	73.48	97.03	89.41	88.67	97.50	88.65	93.26	70.26	94.31		
11700	93.22	92.36	98.43	92.72	94.92	73.63	97.17	89.09	88.4	96.97	88.42	92.93	70.17	94.30		
23400	93.40	92.89	98.09	93.10	94.75	74.20	97.58	90.13	89.33	96.79	89.41	92.96	71.17	94.71		

Notes: CLT-intervals based on the Normal; WB1 wild bootstrap intervals based on the external random variable WB1; WB2 wild bootstrap intervals based on the external random variable WB2; WB3 wild bootstrap intervals based on the external random variable WB3. 10,000 Monte Carlo trials with 999 bootstrap replications each.

Table 2. Summary results for the studentized statistic T_n and its bootstrap analogue T_n^*

$\xi^2 = 0.01$	SV1F				SV2F			
	T_n	T_n^{*WB1}	T_n^{*WB2}	T_n^{*WB3}	T_n	T_n^{*WB1}	T_n^{*WB2}	T_n^{*WB3}
<hr/>								
$n = 195$								
Mean	-1.109	-1.802	-0.552	-0.413	-1.798	-2.319	-0.752	-0.440
Standard error	2.356	4.526	2.135	1.298	3.921	5.867	2.466	1.264
Excess Skewness	-3.717	-5.751	-3.024	-0.009	-6.000	-6.041	-2.724	0.158
Excess Kurtosis	29.279	65.763	14.765	-1.206	72.426	71.078	11.29	-1.255
Cov two-sided	77.93	97.67	91.23	81.35	70.17	94.8	86.74	72.73
<hr/>								
$n = 390$								
Mean	-0.594	-1.275	-0.393	-0.373	-0.997	-1.676	-0.590	-0.410
Standard error	1.661	2.969	1.755	1.317	2.325	3.776	2.119	1.283
Excess Skewness	-2.454	-3.940	-2.685	-0.169	-3.235	-4.218	-2.591	0.022
Excess Kurtosis	11.775	31.511	12.71	-0.991	18.615	35.287	10.775	-1.13
Cov two-sided	85.09	98.35	93.59	90.28	78.59	96.42	90.00	83.60
<hr/>								
$n = 780$								
Mean	-0.519	-0.991	-0.297	-0.338	-0.850	-1.342	-0.482	-0.384
Standard error	1.464	2.304	1.517	1.323	1.982	2.924	1.884	1.293
Excess Skewness	-1.973	-2.974	-2.235	-0.279	-2.620	-3.306	-2.391	-0.077
Excess Kurtosis	8.798	17.221	9.053	-0.767	12.026	21.371	9.370	-1.005
Cov two-sided	86.75	98.38	93.47	92.75	81.29	96.33	90.36	86.54
<hr/>								
$n = 1560$								
Mean	-0.409	-0.788	-0.228	-0.299	-0.669	-1.094	-0.395	-0.357
Standard error	1.297	1.909	1.353	1.323	1.704	2.389	1.690	1.300
Excess Skewness	-1.369	-2.381	-1.829	-0.366	-2.203	-2.672	-2.144	-0.166
Excess Kurtosis	3.714	10.773	6.193	-0.519	9.133	13.287	7.579	-0.858
Cov two-sided	89.03	98.41	94.02	94.54	83.45	96.51	91.06	89.59
<hr/>								
$n = 23400$								
Mean	-0.191	-0.375	-0.095	-0.170	-0.334	-0.560	-0.183	-0.247
Standard error	1.064	1.277	1.085	1.281	1.225	1.490	1.247	1.295
Excess Skewness	-0.592	-1.226	-0.859	-0.419	-1.117	-1.504	-1.331	-0.364
Excess Kurtosis	0.599	2.611	1.357	0.066	2.518	3.794	3.053	-0.298
Cov two-sided	93.40	98.09	94.75	97.58	90.13	96.79	92.96	94.71

Notes: T_n studentized statistic; T_n^{*WB1} studentized wild bootstrap statistic based on WB1; T_n^{*WB2} studentized wild bootstrap statistic based on WB2; T_n^{*WB3} studentized wild bootstrap statistic based on WB3. 10,000 Monte Carlo trials with 999 bootstrap replications each.

Table 3: Summary statistics

Days	Trans	n	S
1 Dec	11924	1491	8
2 Dec	11681	1461	8
5 Dec	10538	1506	7
6 Dec	12959	1440	9
7 Dec	11360	1420	8
8 Dec	10064	1438	7
9 Dec	12120	1515	8
12 Dec	12082	1511	8
13 Dec	10379	1483	7
14 Dec	12616	1577	8
15 Dec	10869	1553	7
16 Dec	12265	1534	8
19 Dec	11119	1589	7
20 Dec	12623	1578	8
21 Dec	13270	1475	9
22 Dec	14765	1476	10
23 Dec	10970	1568	7
27 Dec	10206	1458	7
28 Dec	9580	1597	6
29 Dec	10876	1554	7
30 Dec	9839	1406	7

“Trans” denotes the number of transactions, n is the sample size used to calculate the pre-averaged realized volatility, we have sampled every S th transaction price, so the period over which returns are calculated is roughly 15 seconds.




Figure 1: 95% Confidence Intervals (CI's) for the daily IV, for each regular exchange opening days in December 2011, calculated using the asymptotic theory of Podolskij and Vetter (2009) (CI's with bars), and the wild bootstrap method using WB2 as external random variable (CI's with lines). The pre-averaging realized volatility estimator is the middle of all CI's by construction. Days on the x -axis.

Appendix B

Proof of Theorem 3.1. For part (1), given that $\bar{Y}_j^* = \bar{Y}_j v_j$, where v_j are i.i.d. with $\mu_q^* = E^* |v_j|^q$, for any $q > 0$, we have that

$$\begin{aligned} V_n^* &= \text{Var}^* \left(n^{1/4} \frac{c_1 c_2}{\nu_1} \sum_{j=1}^{n/K} \bar{Y}_j^{*2} \right) = \text{Var}^* \left(n^{1/4} \frac{c_1 c_2}{\nu_1} \sum_{j=1}^{n/K} \bar{Y}_j^2 v_j^2 \right) \\ &= \left(\mu_4^* - (\mu_2^*)^2 \right) \frac{c_1^2 c_2^2}{\nu_1^2} n^{1/2} \sum_{j=1}^{n/K} \bar{Y}_j^4 = \hat{V}_n \end{aligned}$$

under the condition that $\mu_4^* - (\mu_2^*)^2 = \frac{2}{3}$. Thus,

$$V_n^* \xrightarrow{P} V = \frac{2c_1^2 c_2^2}{\nu_1^2} \int_0^1 (\nu_1 \sigma_u^2 + \nu_2 \omega^2)^2 du,$$

by an application of Theorem 1 of Podolskij and Vetter (2009) (where we set $r = 4$ and $l = 0$).

For part (2), let $S_n^* = \sum_{j=1}^{n/K} z_j^*$, where $z_j^* = \frac{c_1 c_2}{\nu_1} n^{1/4} (\bar{Y}_j^{*2} - E^*(\bar{Y}_j^{*2}))$. Note that $E^*(z_j^*) = 0$ and that

$$\text{Var}^* \left(\sum_{j=1}^{n/K} z_j^* \right) = V_n^* \xrightarrow{P} V,$$

by part (1). Moreover, since $z_1^*, \dots, z_{n/K}^*$ are conditionally independent, by the Berry-Esseen bound, for some small $\delta > 0$ and for some constant $C > 0$,

$$\sup_{x \in \mathbb{R}} \left| P^*(S_n^* \leq x) - \Phi \left(x / \sqrt{V} \right) \right| \leq C \sum_{j=1}^{n/K} E^* |z_j^*|^{2+\delta},$$

which converges to zero in probability as $n \rightarrow \infty$. Indeed, we have that

$$\begin{aligned} \sum_{j=1}^{n/K} E^* |z_j^*|^{2+\delta} &= \left| \frac{c_1 c_2}{\nu_1} \right|^{2+\delta} \sum_{j=1}^{n/K} E^* \left| n^{1/4} (\bar{Y}_j^{*2} - E^*(\bar{Y}_j^{*2})) \right|^{2+\delta} \\ &\leq 2 \left| \frac{c_1 c_2}{\nu_1} \right|^{2+\delta} n^{\frac{(2+\delta)}{4}} \sum_{j=1}^{n/K} E^* |\bar{Y}_j^{*2}|^{2+\delta} \\ &\leq 2 \left| \frac{c_1 c_2}{\nu_1} \right|^{2+\delta} E^* |v_1|^{2(2+\delta)} n^{-\frac{\delta}{4}} \left(n^{\frac{(1+\delta)}{2}} \sum_{j=1}^{n/K} |\bar{Y}_j^2|^{2+\delta} \right) \\ &= O_p \left(n^{-\frac{\delta}{4}} \right) = o_p(1), \end{aligned}$$

since $E |v_1|^{2(2+\delta)} \leq \Delta < \infty$ by assumption, and given that by Theorem 1 of Podolskij and Vetter (2009)

$$n^{\frac{(1+\delta)}{2}} \sum_{j=1}^{n/K} |\bar{Y}_j^2|^{2+\delta} \xrightarrow{P} \frac{\mu_{2(2+\delta)}}{c_1 c_2} \int_0^1 (\nu_1 \sigma_u^2 + \nu_2 \omega^2)^{2+\delta} du,$$

which is bounded since σ is an adapted càdlàg spot volatility process and locally bounded away from

zero, and $E(\epsilon_t^2) = \omega^2 < \Delta < \infty$.

Proof of Theorem 3.2 Given that $T_n \xrightarrow{d} N(0, 1)$ (cf. Corollary 1 of Podolskij and Vetter (2009)), it suffices to show that $T_n^* \xrightarrow{d^*} N(0, 1)$ in probability. Let

$$H_n^* = \frac{n^{1/4} (PRV_n^* - E^*(PRV_n^*))}{\sqrt{V_n^*}},$$

and note that

$$T_n^* = H_n^* \sqrt{\frac{V_n^*}{\hat{V}_n^*}},$$

where \hat{V}_n^* is defined in the main text. Theorem 3.1 proved that $H_n^* \xrightarrow{d^*} N(0, 1)$ in probability. Thus, it suffices to show that $\hat{V}_n^* - V_n^* \xrightarrow{P^*} 0$ in probability. In particular, we show that (1) $Bias^*(\hat{V}_n^*) = 0$, and (2) $Var^*(\hat{V}_n^*) \xrightarrow{P^*} 0$. It is easy to verify that (1) holds by the definition of \hat{V}_n^* and V_n^* . To prove (2), note that

$$\begin{aligned} Var^*(\hat{V}_n^*) &= E^*(\hat{V}_n^* - V_n^*)^2 = \left(\frac{\mu_4^* - (\mu_2^*)^2}{\mu_4^*} \right)^2 n \frac{c_1^2 c_2^2}{\nu_1^2} E^* \left(\sum_{j=1}^{n/K} (\bar{Y}_j^4 v_j^4 - \mu_4^* \bar{Y}_j^4) \right)^2 \\ &= \left(\frac{\mu_4^* - (\mu_2^*)^2}{\mu_4^*} \right)^2 n \frac{c_1^2 c_2^2}{\nu_1^2} \sum_{j=1}^{n/K} \bar{Y}_j^8 E^*(v_j^4 - \mu_4^*)^2 \\ &= \left(\frac{\mu_4^* - (\mu_2^*)^2}{\mu_4^*} \right)^2 (\mu_8^* - \mu_4^{*2}) \frac{c_1^2 c_2^2}{\nu_1^2} n^{-\frac{1}{2}} n^{\frac{3}{2}} \sum_{j=1}^{n/K} \bar{Y}_j^8 \\ &= O_P\left(n^{-\frac{1}{2}}\right) = o_P(1), \end{aligned}$$

where we have used the independence of v_j over j to justify the third equality and Theorem 1 of Podolskij and Vetter (2009) (with $r = 8$ and $l = 0$) to justify the fact that $n^{\frac{3}{2}} \sum_{j=1}^{n/K} \bar{Y}_j^8 = O_P(1)$. This requires strengthening the moment condition on ϵ by assuming that $E|\epsilon|^{2(8+\epsilon)} < \infty$.

References

- [1] Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2003. "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 529-626.
- [2] Bandi, F., and J. Russell (2008). "Microstructure noise, realized variance, and optimal sampling," *Review of Economic Studies*, 75(2), 339-369.
- [3] Barndorff-Nielsen, O., S.E. Graversen, J. Jacod, and N. Shephard (2006). "Limit theorems for bipower variation in financial econometrics," *Econometric Theory*, 22, 677-719.
- [4] Barndorff-Nielsen, O., S. E. Graversen, J. Jacod, M. Podolskij, and N. Shephard (2006). "A central limit theorem for realised power and bipower variations of continuous semimartingales." In Y. Kabanov, R. Lipster, and J. Stoyanov (Eds.), *From Stochastic Analysis to Mathematical Finance*, Festschrift for Albert Shiryaev, 33 - 68. Springer.
- [5] Barndorff-Nielsen, O. and N. Shephard (2002). "Econometric analysis of realized volatility and its use in estimating stochastic volatility models," *Journal of the Royal Statistical Society, Series B*, 64, 253-280.

- [6] Barndorff-Nielsen, O., P. Hansen, A. Lunde, and N. Shephard (2008). “Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise,” *Econometrica*, 76, 1481-1536.
- [7] Barndorff-Nielsen, O., P. Hansen, A. Lunde, and N. Shephard (2011). “Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading,” *Journal of Econometrics*, 162, 149-169.
- [8] Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., Shephard, N. (2009). “Realised kernels in practice: trades and quotes,” *Econometrics Journal* 12, C1-C32.
- [9] Christensen, K., S. Kinnebrock, and M. Podolskij (2010). “Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data,” *Journal of Econometrics*, 159, 116-133.
- [10] Diebold, F.X. and Strasser, G.H. (2012). “On the correlation structure of microstructure noise: a financial economic approach”, *Review of Economics Studies*, forthcoming.
- [11] Gonçalves, S. and N. Meddahi (2009). “Bootstrapping realized volatility,” *Econometrica*, 77(1), 283-306.
- [12] Heston, S. (1993). “Closed-form solution for options with stochastic volatility with applications to bonds and currency options,” *Review of Financial Studies*, 6, 327-343.
- [13] Hansen, P.R. and A. Lunde (2006). “Realized variance and market microstructure noise,” *Journal of Business and Economics Statistics*, 24, 127-161.
- [14] Hounyo, U. , Gonçalves, S., and N. Meddahi (2013). “Bootstrapping pre-averaged realized volatility under market microstructure noise,” manuscript.
- [15] Jacod, J., Y. Li, P. Mykland, M. Podolskij, and M. Vetter (2009). “Microstructure noise in the continuous case: the pre-averaging approach,” *Stochastic Processes and Their Applications*, 119, 2249-2276.
- [16] Kalnina, I., and O. Linton (2008). “ Estimating quadratic variation consistently in the presence of endogenous and diurnal measurement error,” *Journal of Econometrics*, 147(1), 47-59.
- [17] Katz, M.L. (1963). “Note on the Berry-Esseen theorem,” *Annals of Mathematical Statistics* 34, 1107-1108.
- [18] Künsch, H.R. (1989). “The jackknife and the bootstrap for general stationary observations,” *Annals of Statistics* 17, 1217-1241.
- [19] Liu, R.Y. (1988). “Bootstrap procedure under some non-i.i.d. models,” *Annals of Statistics* 16, 1696-1708.
- [20] Mammen, E. (1993). “Bootstrap and wild bootstrap for high dimensional linear models,” *Annals of Statistics* 21, 255-285.
- [21] Podolskij, M., and M. Vetter (2009). “Estimation of volatility functionals in the simultaneous presence of microstructure noise and jumps,” *Bernoulli*, 15(3), 634-658.
- [22] Politis, D. N. and Romano, J. P. (1992). “A general resampling scheme for triangular arrays of α -mixing random variables,” *Annals of Statistics* 20, 1985-2007.
- [23] Politis, D.N., Romano, J.P., Wolf, M. (1999). “Subsampling,” Springer-Verlag, New York.
- [24] Vetter, M. (2008). “Estimation methods in noisy diffusion models,” *Ph.D. Thesis. Bochum University*.
- [25] Zhang, L. (2011). “Estimating covariation: Epps effect, microstructure noise,” *Journal of Econometrics*, 160, 33-47.
- [26] Zhang, L, P.A. Mykland, and Y. Aït-Sahalia (2005). “A tale of two time-scales: determining integrated volatility with noisy high frequency data,” *Journal of the American Statistical Association*, 100, 1394-1411.
- [27] Zhang, L., Mykland, P. and Y. Aït-Sahalia (2011a). “Edgeworth expansions for realized volatility and related estimators,” *Journal of Econometrics*, 160, 190-203.
- [28] Zhang, L., Mykland, P. and Y. Aït-Sahalia, (2011b). “Ultra high frequency volatility estimation with dependent microstructure noise,” *Journal of Econometrics*, 160, 160-165.