Available online at www.sciencedirect.com
ScienceDirect
Journal of Empirical Finance 15 (2008) 265–286
Journal of EMPIRICAL FINANCE
www.elsevier.com/locate/jempfin

Finite sample accuracy and choice of sampling frequency in integrated volatility estimation ☆

Morten Ørregaard Nielsen [a,*], Per Frederiksen [b]

[a] Department of Economics, Cornell University, USA
[b] Equity Trading and Derivatives, Nordea Markets, Denmark

Accepted 15 December 2006
Available online 23 May 2007

Abstract

We consider the properties of three estimation methods for integrated volatility, i.e. realized volatility, Fourier, and wavelet estimation, when a typical sample of high-frequency data is observed. We employ several different generating mechanisms for the instantaneous volatility process, e.g. Ornstein–Uhlenbeck, long memory, and jump processes. The possibility of market microstructure contamination is also entertained using models with bid-ask bounce and price discreteness, in which case alternative estimators with theoretical justification under market microstructure noise are also examined. The estimation methods are compared in a simulation study which reveals a general robustness towards persistence or jumps in the latent stochastic volatility process. However, bid-ask bounce effects render realized volatility and especially the wavelet estimator less useful in practice, whereas the Fourier method remains useful and is superior to the other two estimators in that case. More strikingly, even compared to bias correction methods for microstructure noise, the Fourier method is superior with respect to RMSE while having only slightly higher bias. A brief empirical illustration with high-frequency GE data is also included.

JEL classification: C15; C22
Keywords: Bid-ask bounce; Finite sample bias; Integrated volatility; Long memory; Market microstructure; Monte Carlo simulation; Realized volatility; Wavelet

1. Introduction

The definition and analysis of realized volatility in financial return series has attracted considerable interest in the literature starting with French et al. (1987), see e.g. Andersen et al. (in press-a), Barndorff-Nielsen and Shephard (in press), and the references therein for reviews. Essentially, integrated instantaneous volatility is estimated consistently by realized volatility, i.e. by its sample analogue based on high-frequency return observations. This approach allows gathering much more detailed information on the properties of financial market volatility than previously.

☆ We thank Torben G. Andersen, Esben Høg, Asger Lunde, and an anonymous referee for valuable comments and suggestions, and Esben Høg and Asger Lunde for making their Gauss codes available to us. This paper was previously circulated under the title "Finite sample accuracy of integrated volatility estimators." The first author is grateful for financial support from the Danish Social Science Research Council (grant no. SSF 24–20–0181).
* Corresponding author. Mailing address: Department of Economics, Cornell University, 486 Uris Hall, Ithaca, NY 14853, USA. Tel.: +1 607 255 6338; fax: +1 607 255 2818.
E-mail address: mon2@cornell.edu (M.Ø. Nielsen).

0927-5398/$ - see front matter
doi:10.1016/j.jempfin.2006.12.005

High-frequency, or tick-by-tick, financial prices have recently become available for a host of different instruments and markets, and the analysis of the corresponding realized volatility, or realized variation, measures have already materialized in a number of new and important empirical insights related to the distributional properties and dynamic dependencies in financial market volatilities; see, e.g., Andersen and Bollerslev (1998), Andersen et al. (2001b, 2003), Barndorff-Nielsen and Shephard (2002a,b), and Andersen et al. (2006b). In particular, Andersen et al. (2006b) show how realized volatility measures may be used in the formulation of highly informative and directly testable distributional implications for discretely observed asset returns. As such, their tests may serve as a useful set of diagnostic tools in the specification of more accurate and empirically realistic continuous-time models.

Furthermore, high-frequency based quantities in general have been shown to be useful for the analysis of numerous problems. For example, several authors have applied filtering and smoothing techniques to realized volatility time series to obtain a time series of the underlying daily volatilities, e.g. Maheu and McCurdy (2002), Barndorff-Nielsen et al. (2004), Hansen and Lunde (2005), Koopman et al. (2005), and Owens and Steigerwald (2006). Other important and fruitful recent applications of high-frequency based quantities are the evaluation and comparison of volatility models, e.g. Andersen and Bollerslev (1998), Hansen et al. (2003), and Patton (2005), as well as reduced form volatility forecasting, see Andersen et al. (2003), Andersen et al. (2004), Andersen et al. (in press-b), and Ghysels et al. (2006), among others.

Two rivaling approaches for the estimation of integrated volatility have recently been introduced. Malliavin and Mancino (2002) suggest estimating integrated volatility by a Fourier transform based method and Høg and Lunde (2003) suggest employing a method based on the wavelet transform. The properties of the three estimation methods for integrated volatility, i.e. realized volatility, the Fourier estimator, and the wavelet estimator, when only a finite sample of the price process (albeit at a high frequency) is observed, have been examined only briefly in the literature.

Previously, Barucci and Reno (2002a,b) have compared the Fourier method to realized volatility in a Monte Carlo study using the Cox et al. (1985) model to generate the latent instantaneous volatility process, and their simulations show that the Fourier method compares favorably with realized volatility. However, Barucci and Reno (2002a,b) typically contrast a 5 min realized volatility estimator to a Fourier estimator using all observations (which are measured as often as every 14 s on average) creating a very uneven base for comparison, and furthermore they use interpolation between observations rather than the imputation scheme that the literature has settled upon. Within the same framework, i.e. using a Cox et al. (1985) model for the latent instantaneous volatility process, simulations by Høg and Lunde (2003) show that the wavelet method and the Fourier method are virtually indistinguishable with respect to bias and variance, although the wavelet method is computationally much faster than the Fourier method. However, a major drawback of both these studies is that they consider only one generating mechanism for the unobserved instantaneous volatility process, and in particular, their generating mechanism does not allow for any of the recently popularized features of integrated volatility, such as long memory, jumps, and market microstructure effects. In this paper, we examine all three estimators mentioned above and compare them within the same model setup. Unlike the previous studies, we try to even the playing field compared to Barucci and Reno (2002a,b) by using the same number of implied intra-daily returns for each estimator. Furthermore, we follow the literature and use imputation rather than interpolation, and we employ several different generating mechanisms for the instantaneous volatility process. In particular, we consider logarithmic Ornstein–Uhlenbeck processes, long memory processes (e.g. Comte and Renault (1996, 1998)), and jump processes (e.g. Andersen et al. (2002), Eraker et al. (2003), and Eraker (2004)). The possibility of market microstructure effects contaminating the data is also entertained in models that allow for a bid-ask bounce and possibly discrete tick sizes in the spirit of Roll (1984) and Hasbrouck (1999a,b). In the latter case we also consider alternative estimators by Barndorff-Nielsen and Shephard (2004, 2006) and Hansen and Lunde (2006) with theoretical justification under market microstructure noise.

The estimation methods are compared in a Monte Carlo study which reveals that the theoretical robustness of the estimators towards persistence or jumps in the stochastic process governing the latent volatility carries over to practice. On the other hand, irregularities such as bid-ask bounce effects, which the methods have no theoretical robustness against, in general render the wavelet estimator, and to a lesser degree realized volatility, less useful in practice. We find the Fourier method to be superior compared to the other two estimators in the case of market microstructure noise, and indeed this estimator remains very useful even in that case. More strikingly, even when compared to the bias correction methods designed specifically to handle market microstructure effects, the Fourier method is superior with respect to RMSE while having only slightly higher bias.

In a brief empirical illustration using high-frequency General Electric Co. (GE) mid-quote data, we examine volatility signature plots for the Fourier and wavelet estimators, generalizing the well known realized volatility signature plot of Andersen et al. (2000). The three volatility signature plots are compared, and in general may be used as a diagnostic tool to gauge the influence of market microstructure noise on the estimators.

The remainder of the paper is organized as follows. In the next section we present a typical model for asset returns and integrated volatility, which is the focus of the estimation. Section 3 presents the three estimation methods, realized volatility, the Fourier estimator, and the wavelet estimator. In Sections 4 and 5 we describe the setup of the Monte Carlo study for the models without market microstructure effects and the models including these effects, respectively. Section 5 also introduces three alternative estimators designed for the case with market microstructure contamination. Section 6 presents the simulation results in terms of the finite sample biases and RMSEs of the estimators in Sections 3 and 5, and 7 offers a brief empirical illustration with GE data. Section 8 concludes.

2. Integrated volatility and quadratic variation

Suppose the log-price of an asset, $p(t)$, follows a stochastic volatility model, where the basic Brownian motion is generalized to allow the volatility to vary over time, see e.g. Ghysels et al. (1996) or Barndorff-Nielsen and Shephard (2001) and the references therein for overviews of the vast literature on this topic. In particular, we assume $p(t)$ follows the stochastic differential equation model

$$\mathrm{d}p(t) = \mu(t)\mathrm{d}t + \sigma(t)\mathrm{d}w(t), \quad t \geq 0, \tag{1}$$

where the mean process $\mu(\cdot)$ and the instantaneous volatility process $\sigma(\cdot)>0$ are assumed to be independent of the Brownian motion $w(\cdot)$. Allowing the instantaneous volatility to be random and possibly exhibit serial correlation (which we shall do below, see the examples in Section 4), the model (1) will generate returns with unconditional distributions that are fat- tailed and have volatility clustering. This replicates more closely actually observed processes than constant volatility, and e.g. allows the model to overcome some of the shortcomings of the basic Black and Scholes (1973) option pricing model, see Hull and White (1987).

An important feature of the model (1) is that

$$p(t)| \int_0^t \mu(s)\mathrm{d}s, \sigma^{2*}(t) \sim N\left(\int_0^t \mu(s)\mathrm{d}s, \sigma^{2*}(t)\right), \tag{2}$$

where

$$\sigma^{2*}(t) = \int_0^t \sigma^2(s)\mathrm{d}s \tag{3}$$

is called the integrated volatility (or integrated variance) and is the object of interest. For pricing options, this is the relevant volatility measure, see Hull and White (1987), and for the econometrician this is the object to be estimated, see also Andersen and Bollerslev (1998) and the references in the introduction above. Thus, it is also the focal point of this paper.

Another important object is the quadratic variation of the process $p(t)$, denoted $[p](t)$, which is defined for any semimartingale (see e.g. Protter, 1990) by

$$[p](t) = p^2(t) - 2\int_0^t p(s-)\mathrm{d}p(s) \tag{4}$$

or equivalently

$$[p](t) = p\lim \sum_{j=1}^{M} (p(s_j) - p(s_j - 1))^2, \tag{5}$$

where $0=s_0<s_1<\ldots<s_M=t$ and the limit is taken for $\max_j|s_j-s_{j-1}|\to 0$ as $M\to\infty$.

Under some very general regularity conditions, which allow the instantaneous volatility process to exhibit many irregularities, e.g. jumps, long memory, or even nonstationarity, it was shown by Andersen and Bollerslev (1998) and Barndorff-Nielsen and Shephard (2001) that

$$[p](t) = \sigma^{2*}(t) \tag{6}$$

for the model (1). For the purpose of this paper we note that this implies that the object of interest, $\sigma^{2}*(t)$, can be estimated either directly via a parametric model or nonparametrically via the quadratic variation. The latter has been a very popular approach recently.

3. Estimation of integrated volatility

We next review briefly three different methods of estimation for the integrated volatility based on a sample of high-frequency return observations. First, we describe the very popular realized volatility approach which utilizes the connection to quadratic variation, and subsequently we describe the Fourier and wavelet estimators which estimate $\sigma^{2}*(t)$ via the Fourier and wavelet transforms, respectively.

To lighten the notation, we make some simplifying assumptions. We assume that only one day (or month or year) of observations is available and denote the intra-daily (or intra-monthly or intra-yearly) observations on the log-price of the asset by p_j. In principle, the time period could be any arbitrary period, but most often in empirical work either intra-daily or intra-monthly observations are considered in order to estimate integrated volatility on a daily or monthly basis. For instance, to obtain estimates of $\sigma^{2}*(t)$, where $t=1,\ldots, T$ denotes days, the econometrician employs intra-daily observations on the price process. For the purpose of this section we normalize by setting $T=1$ and consider the estimation of integrated volatility over one time period.

The intra-daily data may be of the tick-by-tick type, where the price is observed at every trade or quote (tick), or of the fixed interval type, where the price is observed at fixed intervals, e.g. at 5 min intervals. Both these types of data, commonly denoted high-frequency data, are widely available on many types of assets. For applications to stock return data, see e.g. French et al. (1987), Andersen et al. (2001a), Eraker et al. (2003), or Eraker (2004), and for introductions to some of the very commonly used Olsen and Associates high-frequency exchange rate data sets, see e.g. Guillaume et al. (1997), Andersen and Bollerslev (1998), Andersen et al. (2001b), Dacorogna et al. (2001), or Barndorff-Nielsen and Shephard (2001, 2002a).

3.1. Realized volatility

Suppose n intra-daily log-price observations are available. It is often desirable to have observations that are evenly spaced in time. In particular, suppose M evenly spaced return observations are desired based on the n intra-daily and possibly irregularly spaced price observations. To avoid the problem of irregularly spaced data in high-frequency data sets, it is common to use the imputation scheme (in contrast to an interpolation scheme, see e.g. Dacorogna et al., 2001 or Barucci and Reno, 2002a), i.e. for each of the evenly spaced observations to use the last observed price. In this way a data set of $M+1$ evenly spaced intra-daily price observations can be constructed based on an irregularly spaced high-frequency data set (preferably with n much higher than M to avoid using the same observation more than once).

Using these $M+1$ evenly spaced price observations we denote the continuously compounded intra-daily returns by

$$r_j = p_j - p_{j-1}, \quad j = 1, \ldots, M. \tag{7}$$

Using Eq. (5), quadratic variation can be estimated by

$$\hat{\sigma}^{2*}_{\mathrm{RV},M} = \sum_{j=1}^{M} r_j^2, \tag{8}$$

which is denoted the realized volatility of the process $p(\cdot)$. If several days of intra-daily observations on $p(\cdot)$ (or $r(\cdot)$) were observed, a time series of daily observations on the realized volatility could be obtained, but that is not the issue here so we retain the assumption that only one day of observations is available. Some authors refer to the quantity (8) as the realized variance and reserve the name realized volatility for the square root of Eq. (8), but we shall use the more conventional name realized volatility.

Andersen and Bollerslev (1998) and Andersen et al. (2001b) noted that, using Eqs. (5) and (6), $\hat{\sigma}^{2*}_{\mathrm{RV},M}$ is by definition a consistent (in probability) estimator of integrated volatility in Eqs. (9) and (3). The consistency result does not require the observations to be evenly spaced, only that the maximum distance between observations goes to zero in the limit.

Barndorff-Nielsen and Shephard (2002a) strengthened the consistency result and showed that $\hat{\sigma}^{2*}_{\mathrm{RV},M}$ converges to σ^{2*} in probability at rate $\sqrt{M}$, and furthermore that $\hat{\sigma}^{2*}_{\mathrm{RV},M}$ satisfies

$$\frac{\hat{\sigma}^{2*}_{\mathrm{RV},M} - \sigma^{2*}}{\sqrt{\frac{2}{3}\sum_{j=1}^{M} r_j^4}} \rightarrow_d N(0,1) \quad \text{as } M \rightarrow \infty. \tag{9}$$

This is a mixed Gaussian asymptotic distribution theory since the denominator is itself random, and hence $\hat{\sigma}^{2*}_{\mathrm{RV},M}$ has fatter tails than the normal distribution. For some simulation evidence on the accuracy of the asymptotic distribution in Eq. (9), see Barndorff-Nielsen and Shephard (2002a) and Barndorff-Nielsen and Shephard (2005).

If there are many intra-daily observations available, the coarseness of the realized volatility estimator is governed by the choice of M. For example, if trading occurs 24 h per day as in the foreign exchange markets, choosing 5 min returns in Eqs. (7) and (8) corresponds to $M=288$, 15 min returns corresponds to $M=96$, and hourly returns corresponds to $M=24$. Choosing a higher number of intra-daily returns improves the precision of the estimator but at the same time makes it more sensitive towards microstructure effects in the market, e.g. measurement errors, price discreteness, bid-ask bounces, etc., see Section 5.

3.2. Fourier estimator

This estimator was suggested by Malliavin and Mancino (2002) and subsequently applied by Barucci and Reno (2002a,b). The Fourier method only requires that the quadratic variation in Eq. (4) or (5) is bounded. The method is based on the Fourier transform,

$$\int_0^{2\pi} \sigma^2(s)\mathrm{d}s = 2\pi a_0(\sigma^2), \tag{10}$$

where

$$a_0(\sigma^2) = \lim_{S\rightarrow\infty} \frac{\pi}{2S} \sum_{s=1}^{S} (a_s^2(\mathrm{d}p) + b_s^2(\mathrm{d}p)) \tag{11}$$

and the Fourier coefficients are given by

$$a_s(\mathrm{d}p) = \frac{1}{\pi}\int_0^{2\pi} \cos(st)\mathrm{d}p(t), \quad s\geq 1, \tag{12}$$

$$b_s(\mathrm{d}p) = \frac{1}{\pi}\int_0^{2\pi} \sin(st)\mathrm{d}p(t), \quad s\geq 1. \tag{13}$$

Hence, the n intra-daily and possibly irregularly spaced observations at times $t_1,\ldots,t_n$, i.e. on the interval $[t_1, t_n]$, need to be normalized into the interval $[0, 2\pi]$, and we denote the renormalized time points by $\tau_j=2\pi(t_j-t_1)/(t_n-t_1)$, $j=1,\ldots,n$. The formal justification for using Eqs. (10) and (11) was given by Malliavin and Mancino (2002). Barucci and Reno (2002b) derived the following approximations that we use to compute estimates of the Fourier coefficients $(a_s(\mathrm{d}p), b_s(\mathrm{d}p))$ in Eqs. (12) and (13),

$$\hat{a}_s(\mathrm{d}p) = \frac{p(\tau_n) - p(\tau_1)}{\pi} - \frac{1}{\pi}\sum_{j=2}^{n} p(\tau_{j-1})(\cos(s\tau_j) - \cos(s\tau_{j-1})), \quad s\geq 1, \tag{14}$$

$$\hat{b}_s(\mathrm{d}p) = \frac{1}{\pi}\sum_{j=2}^{n} p(\tau_{j-1})(\sin(s\tau_j) - \sin(s\tau_{j-1})), \quad s\geq 1. \tag{15}$$

To obtain the Fourier estimator of integrated volatility, we thus plug Eqs. (14) and (15) into Eq. (11) and use Eq. (10),

$$\hat{\sigma}^{2*}_{F,S} = \frac{\pi^2}{S}\sum_{s=1}^{S} (\hat{a}_s^2(\mathrm{d}p) + \hat{b}_s^2(\mathrm{d}p)). \tag{16}$$

The coarseness of the estimator is controlled by the user-chosen number S, i.e. the number of Fourier coefficients to include in the estimation, which is related to M for realized volatility by $S=M/2$. Higher S thus corresponds to choosing a finer grid (higher M) for realized volatility, e.g. choosing 5 min returns instead of 15 min returns. Note that by including only the lowest S frequencies in the Fourier estimator in Eq. (16), high-frequency noise or short-run noise is ignored by the estimator. Hence, by choosing a smaller number of (low) frequency ordinates to be used for estimation, i.e. by choosing S small, it is in principle possible to render the Fourier estimator invariant to short-run noise introduced by e.g. market microstructure effects.[1]

3.3. Wavelet estimator

Høg and Lunde (2003) suggest an alternative estimator which is based on the wavelet transform of $\mathrm{d}p(\cdot)$, instead of the Fourier transform based estimator above.

A function $y(j)\in\mathbb{L}^2$ with $j=0, 1,\ldots,2^K-1$, where $K\in\mathbb{N}$ and $\mathbb{L}^2$ is the space of square integrable functions, can be expanded into a wavelet series,

$$y(t) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} w_{j,k}\psi_{j,k}(t), \tag{17}$$

with wavelet coefficients

$$w_{j,k} = 2^{j/2}\int y(t)\psi_{j,k}(t)\mathrm{d}t, \tag{18}$$

where $\psi_{j,k}(t)\equiv 2^{-j/2}\psi(2^{-j}t-k)$ for $j, k\in\mathbb{N}$, is the collection of dilations (scales), j, and translations, k, of the wavelet function $\psi(t)$.

By design, the wavelets strength rests in its ability to simultaneously localize a process in time and scale. At high scales, the wavelet has a small centralized time support enabling it to focus in on short lived time phenomena like jumps. At low scales, the wavelet has a large time support allowing it to identify long periodic behavior. By moving from low to high scales, the wavelet zooms in on the behavior of a process at a particular point in time, identifying singularities, jumps, and cusps. Alternatively, the wavelet can zoom out to reveal the long, smooth features of a series. In our implementation we use the Haar wavelet as in Høg and Lunde (2003).

As was the case with the Fourier estimator, the time interval $[t_1, t_n]$ needs to be renormalized. For the wavelet estimator we renormalize into the interval [0, 1], and we denote the renormalized time points by $\tau_j=(t_j-t_1)/(t_n-t_1)$, $j=1,\ldots,n$. It is shown by Høg and Lunde (2003) that defining $K(n)=\mathrm{int}(\log_2(n))$ the wavelet estimator can be based on

$$\int_0^1 \sigma^2(s)\mathrm{d}s = \lim_{n\to\infty} 2^{-K(n)} \sum_{k=0}^{2^{K(n)}-1} w^2_{K(n),k}(\mathrm{d}p), \tag{19}$$

where $K(n)$ is the highest power of two which is less than or equal to n.

Given a time series of n possibly irregularly sampled observations $(t_j, p(t_j))$, $j=1,\ldots,n$, we construct the Haar wavelet coefficients

$$w_{j,k}(\mathrm{d}p) = 2^{j/2}[2p(2^{-j}(k+1/2)) - p(2^{-j}k) - p(2^{-j}(k+1))]. \tag{20}$$

To obtain the wavelet estimator we plug Eq. (20) into Eq. (19),

$$\hat{\sigma}^{2*}_{W,K} = \sum_{k=0}^{2^K-1}[2p(2^{-K}(k+1/2)) - p(2^{-K}k) - p(2^{-K}(k+1))]^2. \tag{21}$$

[1] Thus, the Fourier method is expected to work well when the high-frequency shocks are spurious noise and, essentially, the truncation of higher frequencies eliminates their effects. It is possible, though, to imagine situations where this elimination is not optimal, e.g. if the true latent volatility process also contains much short-run dynamics in addition to that stemming from the market microstructure noise, and the inability to distinguish the two might be considered a weakness of the Fourier method. However, in our simulations below we found no evidence in this direction for the wide array of volatility processes considered. In addition, the analysis of Hansen and Lunde (2006) shows that the noise from market microstructure effects constitutes the dominating part of the short-run dynamics.

The coarseness of the wavelet estimator is controlled by the user-chosen number K, i.e. the number of dilations (scales) to include in the estimation, which is related to M by $2^K = M$ or $K = \log_2(M)$. Choosing K high is similar to choosing a high S for the Fourier estimator and a high M for realized volatility.

4. Monte Carlo setup

In this section we describe the simulation setup used to investigate the biases and root mean squared errors (RMSEs) of the three estimation methods described above. The objective of the simulation exercise is to shed light on which method most accurately estimates the integrated volatility in practical applications with realistic sample sizes.

For the Monte Carlo study we let the log-price $p(t)$ be generated by

$$\mathrm{d}p(t) = \sigma(t)\mathrm{d}W_1(t), \tag{22}$$

and assume the instantaneous volatility process $\sigma(t)$ follows

$$\text{Model A}: \quad \mathrm{d}\ln\sigma^2(t) = \alpha(\beta - \ln\sigma^2(t))\mathrm{d}t + v\mathrm{d}W_d(t) \tag{23}$$

or

$$\text{Model B}: \quad \mathrm{d}\sigma^2(t) = \alpha(\beta - \sigma^2(t))\mathrm{d}t + \sigma(t)v\mathrm{d}W_2(t) + \kappa(t)\mathrm{d}q(t), \tag{24}$$

where $W_1(\cdot)$ and $W_2(\cdot)$ are independent standard Brownian motions. Both volatility models ensure, under regularity conditions, that volatility cannot become negative.

In Model A, $W_d(\cdot)$ is a fractional Brownian motion of order d independent of $W_1(\cdot)$. It may be represented by the Holmgren–Riemann–Liouville fractional integral

$$W_d(t) = \int_0^t \frac{(t-s)^2}{\Gamma(d+1)}\mathrm{d}W(s), \quad t > 0, \tag{25}$$

where $W(\cdot)$ is a standard Brownian motion, see e.g. Comte and Renault (1996) and Marinucci and Robinson (1999). Sometimes Eq. (25) is called a type II fractional Brownian motion, see Marinucci and Robinson (1999) for details on the generation and simulation of this process. The generating mechanism for the instantaneous volatility process in Model A, i.e. Eq. (23), is a logarithmic Ornstein–Uhlenbeck process, driven by a possibly fractional Brownian motion. Thus, Model A allows the stochastic process driving volatility to exhibit long memory, which is also empirically well founded, see e.g. Andersen et al. (2001a), Andersen et al. (2001b, 2003), Andersen et al. (2004), and the references therein. Note that when $d = 0$, Eq. (23) is a standard (logarithmic) Ornstein–Uhlenbeck process. For a detailed discussion of the implications of using $d \neq 0$ in Model A, see Comte and Renault (1998).

In Model B the allowance for long memory has been substituted with allowance for jumps, represented by the jump process $\mathrm{d}q(t)$. In particular, the process in Eq. (24) is a Cox et al. (1985) (or CIR), square-root model for the volatility process with the addition of a (positive) jump process. The arrival of jumps is assumed to follow a Poisson process with intensity given by $\lambda_0 + \lambda_1\sigma^2(t)$, i.e. the arrival of jumps is assumed to depend on the volatility through the parameter λ_1. The magnitude of the jumps, $\kappa(t)$, is assumed to be exponentially distributed with mean μ. This setup of the jump process follows that in Andersen et al. (2002) (although they have the jump process in the mean and assume $\kappa(t)$ to be log-normally distributed), Eraker et al. (2003), and Eraker (2004). Note that when $\kappa(t) = 0$, i.e. in the absence of jumps, Eq. (24) is a standard CIR model.

For each Monte Carlo DGP we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p(t), \sigma(t))$, $t = 1, \ldots, T$, assuming 24 h trading, i.e. a total of $T = 86,400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean τ. We choose the value $\tau = 14$ (also used by Barucci and Reno, 2002a) corresponding to approximately 6171 observations per day and to the mean time between observations in the DEM-USD exchange rate time series analyzed by e.g. Andersen and Bollerslev (1998). This procedure is repeated for $L=10,000$ days (replications), which then each have a different number (n) of irregularly spaced observations.

To evaluate the performance of the estimation methods, we calculate, for each method and for each day l, the relative error statistic

$$\pi_l = \frac{\hat{\sigma}_l^{2*} - \sum_{t=1}^{T}(p_l(t) - p_{l-1}(t))^2}{\sum_{t=1}^{T}(p_l(t) - p_{l-1}(t))^2}, \quad l = 1, \ldots, L, \tag{26}$$

where $\hat{\sigma}_l^{2*}$ denotes any of the estimators above and $\sum_{t=1}^{T}(p_l(t)-p_{l-1}(t))^2$ is the "true" integrated volatility on day l. In the following we focus on the average and RMSE of π_l, which can be given natural interpretations as (relative) bias and RMSE of the estimators, respectively.

We define them in the usual way as

$$\text{Bias} = \bar{\pi} = \frac{1}{L}\sum_{l=1}^{L} \pi_l \tag{27}$$

and

$$\text{RMSE} = \left(\frac{1}{L}\sum_{l=1}^{L} \pi_l^2\right)^{1/2} = \sqrt{\bar{\pi}^2 + s_\pi^2}, \tag{28}$$

where $s_\pi^2 = L^{-1}\sum_{l=1}^{L}(\pi_l - \bar{\pi})^2$ is the sample variance of π_l. Note that, since we divide by the "true" integrated volatility in Eq. (26), the bias $\bar{\pi}$ in Eq. (27) is a relative bias and not an absolute bias.

5. Microstructure effects in volatility

To expand further on our Monte Carlo study and introduce more empirical realism, this section applies the estimation methods on data contaminated by market microstructure effects. Inspired by Roll (1984), we introduce a bid-ask bounce effect, see also Campbell et al. (1997, pp. 100–101). Intuitively, as random buys and sells arrive at the market, prices can bounce back and forth between the ask and the bid prices, thus creating spurious volatility and serial correlation in returns, even if the fundamental value of the asset remains unchanged.

As before, let $p(t)$ denote the simulated log-price at time t when no microstructure effects are present and introduce the order-driven indicator variable $I(t)$, indicating whether, at time t, the observed price is an ask (buyer-initiated, $I(t)=1$) or a bid (seller-initiated, $I(t)=-1$) price. The new contaminated price $p^*(t)$ becomes

$$p^*(t) = p(t) + \frac{\xi}{2}I(t) \tag{29}$$

so that

$$\mathrm{d}p^*(t) = \mathrm{d}p(t) + \frac{\xi}{2}(I(t) - I(t-1)) = \mathrm{d}p(t) + \frac{\xi}{2}\mathrm{d}I(t), \tag{30}$$

where ξ is the percentage spread and the $I(t)$ are independently (across t and from p) and identically distributed with $\Pr(I(t)=1)=\Pr(I(t)=-1)=1/2$.

Since d$p(t)$ and d$I(t)$ are independent, at least theoretically, the (instantaneous) variance, covariance, and first-order autocorrelation (the setup does not introduce any higher-order serial correlation) of the contaminated asset return can be easily calculated from Eq. (30) and are given as

$$\mathrm{Var}(\mathrm{d}p^*(t)) = \sigma^2(t) + \frac{\xi^2}{2}, \tag{31}$$

$$\mathrm{Cov}(\mathrm{d}p^*(t), \mathrm{d}p^*(t-1)) = -\frac{\xi^2}{4}, \tag{32}$$

$$\mathrm{Corr}(\mathrm{d}p^*(t), \mathrm{d}p^*(t-1)) = -\frac{\xi^2/4}{\sigma^2(t) + \xi^2/2} \leq 0. \tag{33}$$

Hence, we notice that d$p^*(t)$ exhibits spurious volatility and negative serial correlation as a result of the bid-ask bounce, and a larger spread, ξ, implies a higher spurious volatility.

Roll (1984) estimates the effective spreads of NYSE and AMEX stocks using returns data from 1963 to 1982, and finds the average effective spread to be 0.298% using daily returns and 1.74% using weekly returns. This corresponds roughly to $\xi=0.003$ and $\xi=0.017$ in Eq. (29). To ensure that these somewhat outdated estimates are still relevant, we have estimated the effective spreads of six highly traded stocks and three common stock indices using daily returns following the estimation method of Roll (1984). The results are presented in Table 1 for the two periods 1.1.1995–6.30.2004 and 1.1.2001–6.30.2004. The former period is typical for recent empirical applications, where roughly a decade of high-frequency data is employed, and the latter period is after the so-called "decimalization" in 2000 which would presumably have lowered the effective spread. However, it is clear from Table 1 that effective bid-ask spreads are only slightly lower than the earlier results by Roll (1984), and in particular our estimated values correspond to $\xi \in (0.0019, 0.0108)$.

It is well known that sampling at the highest possible frequency induces bias due to market microstructure effects, see e.g. Andreou and Ghysels (2002) and Oomen (2002), because intra-day returns become (spuriously) autocorrelated. In applications it is common to use lower-frequency intra-day returns to alleviate the problem, see e.g. Andersen and Bollerslev (1997) and Andersen et al. (2000), since the (spurious) autocorrelation only lives for a short period of time (Hansen and Lunde, 2006). However, it is only recently that a formal justification for this approach has been established by Bandi and Russell (2005). Unfortunately, the use of lower-frequency observations inevitably implies a loss of information in the sense that fewer data points are applied, resulting in an inefficient measure of volatility. Several bias reduction methods have been suggested in the literature, e.g. the (moving average or autoregressive) filtering techniques by Andersen et al. (2001a) and Bollen and Inder (2002), the subsampling techniques by Zhang et al. (2005), and the autocorrelation correction methods by Hansen and Lunde (2006), see also Barndorff-Nielsen et al. (2006).

In the present setup with market microstructure contamination, we extend the investigation of the three above-mentioned estimators to include also the Newey–West correction applied by e.g. French et al. (1987) and Zhou (1996) and the related bias correction procedure suggested by Hansen and Lunde (2006).[2] Furthermore, since recent literature concerning jump detection in financial markets has focused on bipower variation, see e.g. Andersen et al. (in press-b), Barndorff-Nielsen and Shephard (2004, 2006), and Huang and Tauchen (2005), we also briefly explore the performance of (2nd lag) realized bipower variation under market microstructure contamination.

Under the process in Eq. (29), realized volatility is no longer consistent for integrated volatility, but instead

$$\hat{\sigma}^{2*}_{\mathrm{RV},M}(t) \rightarrow_p [p^*](t) = \sigma^{2*}(t) + \frac{\xi^2}{4}\sum_{j=1}^{M}(\mathrm{d}I(j))^2, \tag{34}$$

[2] Note that if the price observations are not contaminated by noise, applying the bias correction would not be efficient relative to standard (uncorrected) realized volatility.

Table 1
Estimates of the bid-ask spread

Sample period	AA (%)	GE (%)	IBM (%)	IP (%)	JNJ (%)	JPM (%)	DJIA (%)	S&P 100 (%)	S&P 500 (%)
1.1.1995–6.30.2004	0.62	0.66	1.08	0.56	0.68	0.53	0.19	0.50	0.27
1.1.2001–6.30.2004	0.94	0.38	0.36	0.41	0.50	0.80	0.47	0.49	0.41

Note: Following Roll (1984) this table presents estimated percentage bid-ask spreads using the estimator $2\sqrt{-\widehat{\mathrm{Cov}}(\mathrm{d}p^*(t),\mathrm{d}p^*(t-1))}$ applied to daily data. The data sample periods are given in the first column.

which diverges as $M\rightarrow\infty$. Thus, $\hat{\sigma}^{2*}_{\mathrm{RV},M}(t)$ estimates the sum of the "true" latent integrated volatility and the summing of an infinite number of (squared) microstructure contributions, i.e. the "true" latent integrated volatility is stochastically dominated by the microstructure component.

As emphasized by Barndorff-Nielsen and Shephard (2004, 2006), who consider both jump processes and market microstructure noise processes, if the noise (or jump) component is of finite activity, separate nonparametric identification of the two components in Eq. (34) is possible using bipower variation measures. In our setup with first order serial correlation, the 2nd lag realized bipower variation measure,

$$\hat{\sigma}^{2*}_{\mathrm{BV},M}(t)=\frac{\pi}{2}\frac{M}{M-2}\sum_{j=3}^{M}|r_{t,j}||r_{t,j-2}|, \tag{35}$$

would, in the finite activity case, converge in probability (as $M\rightarrow\infty$) to $\sigma^{2*}(t)$. Hence, the microstructure noise component (or jump component) could be identified as $\hat{\sigma}^{2*}_{\mathrm{RV},M}-\hat{\sigma}^{2*}_{\mathrm{BV},M}$, see Barndorff-Nielsen and Shephard (2004, 2006) and Huang and Tauchen (2005), but the identification of jump or noise components is not the focus of the present study. Instead, we wish to examine the impact of the market microstructure effects on the performance of $\hat{\sigma}^{2*}_{\mathrm{BV},M}$, since Andersen et al. (in press-b, p. 16) conjecture that "the staggering should help alleviate the confounding influences of the market microstructure noise, resulting in empirically more accurate finite sample approximations." To what extent this is true under our particular noise process (Eq. (29)) will be examined in the simulations below.

Hansen and Lunde (2006) consider the estimator

$$\hat{\sigma}^{2*}_{\mathrm{HL},M}(t)=\sum_{j=1}^{M}r_{t,j}^{2}+2\frac{M}{M-1}\sum_{j=2}^{M}r_{t,j}r_{t,j-1}, \tag{36}$$

which is robust to first order serial correlation. More general estimators robust to higher order serial correlation are also considered by Hansen and Lunde (2006) and Barndorff-Nielsen et al. (2006), but Hansen and Lunde (2006) argue that the first order correction in Eq. (36) is sufficient, at least for the sampling frequencies and type of noise considered here. Furthermore, Hansen and Lunde (2006) show that $\hat{\sigma}^{2*}_{\mathrm{HL},M}(t)$ is an unbiased estimator of the "true" integrated volatility $\sigma^{2*}(t)$.

A similar bias correction method is based on the Bartlett kernel, well known from the Newey and West (1987) covariance estimator. The implementation by French et al. (1987) and Zhou (1996) in the context of integrated volatility estimation is

$$\hat{\sigma}^{2*}_{\mathrm{NW},M}(t)=\sum_{j=1}^{M}r_{t,j}^{2}+2\sum_{h=1}^{q_M}\left(1-\frac{h}{q_M+1}\right)\sum_{j=h+1}^{M}r_{t,j}r_{t,j-h} \tag{37}$$

using $q_M=1$. In our implementation we use $q_M=\mathrm{int}\,(4(M/100)^{2/9})$, which is typical in the time series literature, e.g. Andrews and Monahan (1992). As mentioned in Hansen and Lunde (2006), this estimator may not be unbiased like $\hat{\sigma}^{2*}_{\mathrm{HL},M}$ but it may have a smaller asymptotic MSE. Of course, given the nature of the noise that we assume, the Hansen and Lunde (2006) estimator is presumably superior to the Newey and West (1987) correction, which is robust to a more general noise structure involving higher order autoregressive or moving average terms — that may or may not be relevant in practice depending on the specific market.

Finally, to add even more realism to the microstructure effects, we consider also the following model in the spirit of Hasbrouck (1999a,b),

$$p^*(t) = \begin{cases} \text{ceil}(p(t) + \xi/2, \text{tick}), & \text{if } I(t) = 1, \\ \text{floor}(p(t) - \xi/2, \text{tick}), & \text{if } I(t) = -1, \end{cases} \tag{38}$$

where, as before, $p(t)$ is the latent price with no microstructure effects, ξ is the bid-ask spread, and $I(t)$ is the indicator variable denoting whether the observed price is a bid or an ask. Here, ceil(x, y) is the ceiling function which returns the smallest multiple of y that is greater than or equal to x, floor(x, y) is the floor function returning the highest multiple of y that is less than or equal to x, and tick denotes the tick size.

Thus, the model (38) includes both a bid-ask spread and price discreteness. In the simulations below, we simulate from Eq. (38) applying the same values of ξ as in the pure bid-ask model (29) and tick size equal to 1/8, 1/16, and also 1/100 to reflect the recent "decimalization" in 2000.

6. Simulation results

To introduce as much empirical realism as possible, the parameter values chosen for the simulations of the volatility processes are based on the estimation results of Roll (1984), Andersen et al. (2002), and Eraker (2004).

The parameter values used for the simulations of Model A in Eq. (23) are inspired by the analysis of the S&P 500 stock index in Andersen et al. (2002, Table III, p. 1256), who found the estimated parameter values $(\hat{\alpha}, \hat{\beta}, \hat{v})=(0.0062, -1, 0.0374)$ (in our notation) for their model without jumps. In particular, we use the parameter values $\alpha \in \{0, 0.0062, 0.0124\}$, $\beta=-1$, and $v \in \{0.0374, 0.1122\}$ in our simulations. Furthermore, since it is empirically well founded that financial volatility time series exhibit long memory (see e.g. Andersen et al. (2001a), Andersen et al. (2001b, 2003), Andersen et al. (in press-a), and the references therein), we let the process (23) be governed by a stationary fractional Brownian motion with long memory parameter $d \in \{0, 0.15, 0.30, 0.45\}$. The studies by Andersen et al. (2001a), Andersen et al. (2001b, 2003), and Andersen et al. (in press-a) find long memory in volatility with a parameter around 0.3–0.45. The starting values for Model A were chosen as $p(0)=\ln(100)$ and $\sigma(0)=e^{\beta}$. This makes for easy comparisons across days.

In Model B, the square-root jump model (24), we follow the results on jumps in the volatility of the S&P 500 stock index by Eraker (2004, Table III, p. 49), who found the parameter estimates $(\hat{\alpha}, \hat{\beta}, \hat{v})=(0.023, 0.943, 0.137)$ which we also employ in our simulations of Eq. (24). For the jump process, Eraker (2004) found the estimates $(\hat{\mu}, \hat{\lambda}_0, \hat{\lambda}_1)=(1.530, 0.002, 1.298)$, so we simulate our process using the parameter values $\mu \in \{0.7515, 1.530\}$, $\lambda_0 \in \{0.002, 0.01\}$, and $\lambda_1 \in \{0, 1.298, 2.596\}$. The starting values for Model B were chosen as $p(0)=\ln(100)$ and $\sigma(0)=\beta$.

For the model in Eq. (29) including market microstructure distortions, the parameter choice is inspired by Roll (1984, Table I, pp. 1132–1133) and our own estimates of the more recent effective spreads of some highly traded stocks and stock indices in Table 1. Roll (1984) finds the average effective spread of NYSE and AMEX stocks, using returns data from 1963 to 1982, to be 0.298% and 1,74% for daily and weekly returns, respectively. As shown in Table 1 we estimate effective spreads ranging from 0.19% to 1.08% for six highly traded stocks and three common stock indices using daily returns data and the two sample periods 1.1.1995–6.30.2004 and 1.1.2001–6.30.2004. The former period is typical for recent applications where roughly a decade of high-frequency data is employed and the latter period is after the so-called "decimalization" in 2000 which would presumably have lowered the effective spread. Thus, according to our results in Table 1, the effective bid-ask spreads are slightly lower than (but still comparable to) the older results by Roll (1984), and in particular we use $\xi \in \{0.003, 0.005, 0.01\}$ in simulating Eq. (29). Hence, our simulation is a little conservative compared to Roll (1984), and probably more in line with the current average microstructure frictions such as bid-ask spreads, especially since the so-called "decimalization" in 2000. We simulate with $\alpha=0.0124$, $\beta=-1$, and $d=0$ in model (23), and to measure the effect of the bid-ask bounce for different values of the volatility of volatility, we simulate using $v \in \{0.05, 0.5\}$. The starting values were chosen as in Model A. For the model (38), including discrete tick sizes, we simulate using $(\alpha, \beta, v, d)=(0.0124, -1, 0.0374, 0)$, spread $\xi \in \{0.003, 0.01\}$, and tick $\in \{1/8, 1/16, 1/100\}$, and again the starting values were chosen as in Model A.

We implement realized volatility using 1 min returns ($M=1440$), 5 min returns ($M=288$), and 15 min returns ($M=96$). The first two values are close to the optimal sampling frequency, in the presence of market microstructure noise, of 1.5 min found by Bandi and Russell (2005), which is shorter than the 5 min frequency used in most empirical work on the subject, see e.g. Andersen et al. (2001a) and the above references. Previously, Andersen et al. (2000) conjectured that optimal sampling should be based on 15–20 min returns, corresponding to our third implementation of realized volatility.

Note that the results by Bandi and Russell (2005) depend crucially on the very high liquidity of the simulated stocks. Consequently, Bandi and Russell (2006) find that less liquid stocks require lower sampling frequencies leading to optimal sampling schemes that are close to the standard 5 min frequency. For related (theoretical) results on optimal sampling schemes for maximum likelihood estimation of diffusions in the presence of market microstructure noise, see Ait-Sahalia et al. (2005).

The Fourier method is implemented using $S=720$ corresponding to 1 min returns, $S=144$ corresponding to 5 min returns, and $S=48$ corresponding to 15 min returns. Finally, the wavelet estimator is implemented using slightly different values since it is limited to powers of 2. In particular, we use $K=10$ corresponding to $M=2^{10}=1024$, which is a little less than for realized volatility and the Fourier estimator, $K=8$ corresponding to $M=2^8=256$, which is again a little less, and $K=7$ corresponding to $M=2^7=128$, which is a little more than for realized volatility and the Fourier estimator. That is, our implementations of realized volatility and of the Fourier estimator match with respect to the coarseness of the estimators, whereas the implementation of the wavelet estimator is a little off compared to the other two. The reason for this is two-fold. First, the wavelet estimator must be implemented using a power of 2 which limits the possible choices of coarseness. Second, realized volatility is by far the most applied estimator in empirical studies, the 5 min estimator being particularly popular, and we thus wish to include that particular implementation here. Hence, when reading our results below, one should take into account this slight discrepancy in the implementation of the wavelet estimator.

In particular, we thus even the playing field between the realized volatility estimator and the Fourier estimator by matching the implied number of intra-daily returns used in the estimation. The Monte Carlo analysis in Barucci and Reno (2002a,b) typically compare a 5 min realized volatility estimator to a Fourier estimator using all observations. Also in contrast to Barucci and Reno (2002a,b) we use the imputation rather than the interpolation technique, see also Dacorogna et al. (2001), as this is what the literature has settled upon.

In Tables 2–11 the results of our Monte Carlo study are presented in terms of the relative biases in Eq. (27) (multiplied by 1000) and the RMSEs in Eq. (28) of the estimators. All calculations were made using the computer language Gauss v5.0.

Tables 2–5 display the results for the estimators when instantaneous volatility is governed by a logarithmic fractional Ornstein–Uhlenbeck process, Model A. Tables 6 and 7 display the results for the square-root jump process, Model B. Next, Tables 8–11 display the results when bid-ask bounce effects are introduced into the price process and

Table 2
Simulation results for Model A I: relative bias (×1000)

α	d	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.0000	0.00	−0.7055	−0.5915	−0.4304	−1.0062	−0.8192	−0.6201	0.4496	−1.7567	−0.4530
	0.15	0.4583	0.0927	0.2498	1.4775	−0.1334	−0.5432	−0.0119	−0.9714	−0.1007
	0.30	0.4863	−1.3222	−0.8911	−0.6410	−0.9654	−0.7419	−0.1815	−1.4819	−0.7038
	0.45	−2.7630	−3.0976	−1.4876	−3.6948	−2.3962	−1.0222	−2.3314	−2.7952	−1.0631
0.0062	0.00	0.0755	0.5924	−0.1730	−1.6686	−0.8413	−0.2669	−0.6830	−0.6474	−0.5130
	0.15	2.0025	0.4435	0.1936	1.4967	0.1044	−0.2624	1.5461	−1.5653	0.0488
	0.30	0.4254	1.3451	−0.1448	2.1558	0.7306	−0.3160	−1.9598	−0.9732	−0.1840
	0.45	−5.3383	−3.3919	−2.0422	−4.7310	−2.9846	−1.8246	−1.2312	−1.9397	−2.7308
0.0124	0.00	−1.7707	−1.4212	−0.2095	−2.2512	−1.1830	−0.4279	−0.1238	−0.2991	−0.4353
	0.15	−0.8734	−0.5482	−0.9901	−0.8381	−1.0543	−1.0814	1.4513	−2.6114	−0.8863
	0.30	−0.6384	−0.0582	−0.3482	0.7188	0.1207	−0.2306	−0.3850	0.1175	−0.2032
	0.45	0.1016	−0.6692	−0.3883	−1.6604	−0.8492	−0.3214	−1.6346	−0.1256	−1.6033

Note: The table reports the average of the relative error statistic in Eq. (26), i.e. the relative bias, multiplied by 1000 from $L=10{,}000$ simulated replications of Model A in Eq. (23) with $\beta=-1$ and $\nu=0.0374$. For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p(t), \sigma(t))$, $t=1,\ldots,T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

Table 3
Simulation results for Model A I: RMSE

α	d	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.0000	0.00	0.1435	0.0829	0.0394	0.1436	0.0829	0.0395	0.1261	0.0876	0.0454
	0.15	0.1449	0.0835	0.0395	0.1451	0.0841	0.0394	0.1266	0.0883	0.0454
	0.30	0.1522	0.0875	0.0417	0.1511	0.0881	0.0414	0.1316	0.0930	0.0481
	0.45	0.2366	0.1344	0.0635	0.2299	0.1328	0.0630	0.2004	0.1420	0.0740
0.0062	0.00	0.1443	0.0831	0.0394	0.1453	0.0825	0.0390	0.1259	0.0880	0.0455
	0.15	0.1451	0.0834	0.0394	0.1447	0.0841	0.0392	0.1267	0.0889	0.0457
	0.30	0.1540	0.0880	0.0407	0.1539	0.0885	0.0405	0.1323	0.0939	0.0471
	0.45	0.2324	0.1348	0.0625	0.2283	0.1332	0.0619	0.2010	0.1430	0.0733
0.0124	0.00	0.1444	0.0831	0.0391	0.1444	0.0821	0.0389	0.1253	0.0878	0.0447
	0.15	0.1454	0.0831	0.0392	0.1458	0.0833	0.0391	0.1250	0.0882	0.0453
	0.30	0.1524	0.0880	0.0417	0.1521	0.0880	0.0413	0.1329	0.0928	0.0482
	0.45	0.2332	0.1368	0.0638	0.2284	0.1339	0.0625	0.2040	0.1430	0.0747

Note: The table reports the root mean squared error of the relative error statistic in Eq. (26) from $L=10{,}000$ simulated replications of Model A in Eq. (23) with $\beta=-1$ and $\nu=0.0374$. For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

the true latent instantaneous volatility is governed by a simple logarithmic Ornstein–Uhlenbeck process, i.e. by Model A with $d=0$, and lastly Tables 12–15 present the results with both bid-ask bounce and discrete tick sizes.

We consider first the Monte Carlo results for the three estimation methods under Model A, where the underlying instantaneous volatility follows a logarithmic Ornstein–Uhlenbeck process possibly governed by a fractional Brownian motion.

The relative biases are generally very low (less than 1%) as evident from Tables 2 and 4, and seem to be lower for higher sampling frequencies which was to be expected. Intuitively, one might expect that the performance of the estimators would be poor when the simulated series exhibit a high degree of persistence in the form of long memory or even nonstationarity, even though all three estimation methods are robust to such persistence in theory. However, the tables reveal that the methods are very robust towards such persistence even in finite samples. When α and d are zero

Table 4
Simulation results for Model A II: Relative bias (×1000)

α	d	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.0000	0.00	1.2748	−0.3355	−0.7811	0.7110	−0.2561	−0.9026	0.0917	−2.0880	−0.7733
	0.15	−1.4604	−1.1342	−1.0978	−1.5769	−0.9251	−0.6643	0.8388	−1.9007	−0.3863
	0.30	−0.5957	−1.6722	−0.3615	−1.5510	−1.3847	−0.7577	−0.7989	−1.1150	−0.6636
	0.45	3.6261	−2.4180	−2.0503	1.3231	−0.9371	−1.1779	−2.8642	−2.8783	−1.8356
0.0062	0.00	0.2415	0.3961	−0.9228	−1.3742	−1.1969	−0.8255	−1.8643	−1.3787	−0.4259
	0.15	−0.7589	−0.4210	0.5303	−2.2068	−0.3985	0.2660	−0.7397	0.1676	0.3289
	0.30	−2.2294	−2.1213	−0.7497	−1.0979	−0.8559	−0.7321	−0.7992	−2.6153	−1.1359
	0.45	4.0332	−2.5570	−2.8106	0.5308	−2.0275	−0.7129	−8.6316	4.1822	−2.0436
0.0124	0.00	0.8175	1.0354	−0.3810	1.5470	−0.2110	−0.6553	−0.4869	−1.4444	−0.4940
	0.15	−0.1336	−1.0675	−0.9770	−1.4175	−1.4680	−0.9212	−2.2219	−1.4821	−1.0009
	0.30	−0.6238	−2.9059	−1.3050	−3.4262	−2.0246	−1.2529	−0.3346	−0.1858	−2.0064
	0.45	0.0824	−2.5836	−1.8345	0.8803	−1.9273	−0.9515	−3.0580	1.5135	−3.1962

Note: The table reports the average of the relative error statistic in Eq. (26), i.e. the relative bias, multiplied by 1000 from $L=10{,}000$ simulated replications of Model A in Eq. (23) with $\beta=-1$ and $\nu=0.1122$. For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

Table 5
Simulation results for Model A II: RMSE

α	d	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.0000	0.00	0.1444	0.0834	0.0392	0.1441	0.0843	0.0391	0.1252	0.0878	0.0454
	0.15	0.1475	0.0853	0.0399	0.1470	0.0844	0.0399	0.1270	0.0901	0.0464
	0.30	0.1928	0.1126	0.0529	0.1876	0.1108	0.0523	0.1687	0.1193	0.0604
	0.45	0.3972	0.2241	0.1080	0.3802	0.2201	0.1060	0.3331	0.2425	0.1244
0.0062	0.00	0.1427	0.0830	0.0392	0.1444	0.0823	0.0391	0.1260	0.0889	0.0452
	0.15	0.1473	0.0850	0.0403	0.1476	0.0847	0.0400	0.1289	0.0904	0.0462
	0.30	0.1942	0.1106	0.0518	0.1906	0.1112	0.0515	0.1666	0.1187	0.0611
	0.45	0.4040	0.2298	0.1073	0.3801	0.2268	0.1059	0.3363	0.2437	0.1229
0.0124	0.00	0.1447	0.0823	0.0392	0.1453	0.0832	0.0388	0.1244	0.0865	0.0454
	0.15	0.1463	0.0849	0.0401	0.1474	0.0847	0.0404	0.1266	0.0893	0.0460
	0.30	0.1929	0.1106	0.0525	0.1910	0.1103	0.0519	0.1679	0.1199	0.0607
	0.45	0.3852	0.2204	0.1065	0.3707	0.2172	0.1055	0.3291	0.2399	0.1237

Note: The table reports the root mean squared error of the relative error statistic in Eq. (26) from $L=10{,}000$ simulated replications of Model A in Eq. (23) with $\beta=-1$ and $\nu=0.0374$. For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

the simulated process is very persistent since the instantaneous volatility is then governed by a continuous-time random walk. This persistence becomes even more pronounced when d is greater than zero (and $\alpha=0$), as the order of integration then becomes $1+d$, but we still do not find any indication of contradiction of the theoretical robustness.

However, we do find the highest relative biases when the series exhibit slow mean reversion ($\alpha \leq 0.062$) and long memory ($d=0.45$). This is interesting since many empirical studies have concluded that integrated volatility time series exhibit long memory of roughly this magnitude. Furthermore, Table 4 indicates that when increasing the volatility of volatility ($\nu=0.1122$) we generally encounter slightly higher relative biases.

Increasing the volatility of volatility has a very interesting but expected impact on the variability of the estimators. Tables 3 and 5 show a clear pattern when the series exhibit relatively strong long memory ($d \geq 0.3$), in which case the RMSE (mainly variance since biases are very small) of all the estimators increases noticeably, in accordance with the

Table 6
Simulation results for Model B: Relative bias (×1000)

μ	λ_0	λ_1	Realized volatility			Fourier			Wavelet		
			15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.7515	0.0002	0.000	−0.4096	−0.1961	−0.0101	−0.1788	−1.0763	−0.5357	−0.7135	−1.7050	−0.5489
		1.298	0.7465	0.1119	−0.5220	−0.1551	0.2411	−0.3065	0.3924	−0.6000	−0.1430
		2.596	0.0689	0.1546	−1.1297	0.7316	−0.0242	−0.9588	−0.0525	−2.1016	−1.4776
	0.0100	0.000	0.7938	−0.6632	−0.0514	−0.3697	−1.1959	−0.5438	−2.0806	−0.6514	−1.1300
		1.298	−1.2241	0.5597	0.1855	−1.5922	−0.0911	0.0822	−0.2367	0.7899	−0.0018
		2.596	−0.8487	−0.6688	−0.0896	0.0389	−0.1353	−0.2071	0.4876	−0.8325	0.2767
1.5300	0.0002	0.000	−0.8915	0.1760	0.2460	−0.1027	−0.2596	−0.1420	−1.8248	0.9044	−0.2842
		1.298	−0.1182	−1.6774	−1.6449	−0.8492	−1.3033	−2.0398	−2.6302	−3.7401	−2.2957
		2.596	−2.8403	−2.2980	−1.0465	−0.5911	−2.1993	−0.6012	−2.9168	−2.5319	−0.9122
	0.0100	0.000	−0.3969	−1.0942	−0.0919	−0.5675	−0.7069	−0.2713	0.6571	−1.5714	−0.1078
		1.298	−1.3928	−1.3416	−0.7572	−1.1821	−0.5097	−0.7751	−3.0634	−1.8643	−0.3814
		2.596	−3.3209	−2.2607	−1.1896	−2.2697	−1.2317	−1.2272	−3.1931	−2.2495	−1.2734

Note: The table reports the average of the relative error statistic in Eq. (26), i.e. the relative bias, multiplied by 1000 from $L=10{,}000$ simulated replications of Model B in Eq. (24) with $\alpha=0.023$, $\beta=0.943$, and $\nu=0.137$. For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

Table 7
Simulation results for Model B: RMSE

μ	λ_0	λ_1	Realized volatility			Fourier			Wavelet		
			15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.7515	0.0002	0.000	0.1451	0.0838	0.0395	0.1456	0.0843	0.0396	0.1245	0.0891	0.0462
		1.298	0.1546	0.0888	0.0410	0.1509	0.0876	0.0413	0.1312	0.0941	0.0479
		2.596	0.1624	0.0933	0.0445	0.1597	0.0931	0.0439	0.1410	0.0995	0.0511
	0.0100	0.000	0.1453	0.0832	0.0392	0.1450	0.0839	0.0394	0.1255	0.0893	0.0456
		1.298	0.1531	0.0890	0.0412	0.1527	0.0887	0.0413	0.1315	0.0936	0.0480
		2.596	0.1628	0.0944	0.0442	0.1620	0.0942	0.0438	0.1415	0.0996	0.0505
1.5300	0.0002	0.000	0.1449	0.0826	0.0394	0.1450	0.0830	0.0392	0.1255	0.0873	0.0457
		1.298	0.1609	0.0936	0.0440	0.1606	0.0947	0.0440	0.1409	0.0976	0.0512
		2.596	0.1893	0.1105	0.0518	0.1904	0.1096	0.0514	0.1661	0.1169	0.0600
	0.0100	0.000	0.1460	0.0836	0.0396	0.1453	0.0841	0.0392	0.1275	0.0899	0.0455
		1.298	0.1627	0.0931	0.0437	0.1610	0.0937	0.0434	0.1405	0.0985	0.0509
		2.596	0.1885	0.1096	0.0512	0.1877	0.1087	0.0509	0.1631	0.1164	0.0600

Note: The table reports the root mean squared error of the relative error statistic in Eq. (26) from $L=10,000$ simulated replications of Model B in Eq. (24) with $\alpha=0.023$, $\beta=0.943$, and $v=0.137$. For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

distribution theory in Eq. (9) for realized volatility. On the other hand, when the integrated volatility series are simulated with relatively weak long memory ($d \leq 0.15$), the higher variability of the volatility does not imply less accuracy (higher RMSE), i.e. the RMSEs in Table 5 are nearly identical to those in Table 3 when d is small. Overall, the RMSEs are higher if lower sampling frequencies are used. Hence, if possible, one should in empirical investigations use a relatively high sampling frequency even if the latent integrated volatility series exhibits long memory or nonstationarity.

However, an important caveat here is the risk of contamination from market microstructure effects which would lead to the selection of a lower sampling frequency. This is discussed in detail below, where indeed our simulations support this important point.

Focusing more explicitly on the three individual estimators, we do not find any noticeable differences between them. However, regardless of the size of the mean reversion, α, and the size of the volatility of the volatility, v, it seems that the Fourier method provides the lowest relative bias and RMSE when looking at the empirically interesting scenario of high d and high sampling frequency, although only marginally so.

Table 8
Simulation results for Model A with bid-ask Bounce I: Relative bias ($\times 1000$)

ξ	v	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.003	0.05	0.3586	2.0200	16.7259	−0.5529	−0.8312	9.7284	2.8724	9.7648	34.9593
	0.50	−1.2278	3.6067	17.5627	−0.7294	0.3012	10.7018	7.5200	9.8177	35.4649
0.005	0.05	3.8316	10.1418	48.0333	0.6336	1.6777	28.3677	12.7942	25.2158	97.3825
	0.50	1.7986	8.5590	48.8193	−1.3183	0.0210	28.9011	13.5345	25.8629	99.6641
0.01	0.05	12.2471	40.2888	192.4841	0.4416	6.6724	114.7427	52.6355	103.6826	390.1397
	0.50	14.0446	39.6586	195.3937	1.1314	6.1340	116.5520	52.9863	104.7563	396.4723

Note: The table reports the average of the relative error statistic in Eq. (26), i.e. the relative bias, multiplied by 1000 from $L=10{,}000$ simulated replications of Model A in Eq. (24) with $\alpha=0.0124$, $\beta=-1$, $d=0$, and a bid-ask bounce effect as in Eq. (30). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p^*(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

Table 9
Simulation results for Model A with bid-ask Bounce I: RMSE

ξ	v	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.003	0.05	0.1438	0.0842	0.0434	0.1435	0.0837	0.0409	0.1272	0.0904	0.0586
	0.50	0.1462	0.0854	0.0444	0.1467	0.0851	0.0419	0.1292	0.0902	0.0599
0.005	0.05	0.1458	0.0850	0.0629	0.1454	0.0837	0.0491	0.1272	0.0941	0.1094
	0.50	0.1466	0.0862	0.0659	0.1465	0.0853	0.0507	0.1305	0.0960	0.1153
0.01	0.05	0.1458	0.0962	0.1981	0.1445	0.0841	0.1229	0.1413	0.1421	0.3951
	0.50	0.1492	0.0974	0.2084	0.1470	0.0852	0.1290	0.1448	0.1479	0.4173

Note: The table reports the root mean squared error of the relative error statistic in Eq. (26) from $L=10{,}000$ simulated replications of Model A in Eq. (24) with $\alpha=0.0124$, $\beta=-1$, $d=0$, and a bid-ask bounce effect as in Eq. (30). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p^*(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

An interesting irregularity that has turned up in empirical studies, e.g. Eraker et al. (2003) and Eraker (2004), is the possible presence of jumps in the instantaneous volatility process. Tables 6 and 7 present the results for the square-root jump process, Model B. Since the biases are very low, we again find support of the theoretical robustness mentioned in Section 2. Nonetheless, there is a tendency that a larger magnitude of the jump, μ, weakens the performance of the estimators as the (absolute sizes of the) relative biases and RMSEs increase when μ increases. This is generally also the case when increasing the intensity of the arrival of jumps, but only for the part of the intensity that is directly linked to the volatility. That is, there is no clear relationship between the size of the minimum intensity, λ_0, and the performance of the estimators. On the contrary, increasing λ_1 such that the arrival intensity is more sensitive towards the size of the volatility generally implies an increase in the relative biases and as expected an even more pronounced impact is found on the variability of the estimators (the RMSEs). Here we find a noticeable increase when intensity increases.

As under the persistence scenario in Tables 2–5, it is recommended in Model B to use a higher sampling frequency as the relative biases and especially the RMSEs become smaller. Increasing the sampling frequency also seems to imply that the RMSEs of the estimators become relatively insensitive towards the magnitude of the jumps as the RMSEs are almost identical for high frequencies (across μ and λ_0). Again we find that the Fourier method seems to provide better relative bias and RMSE compared to the other estimators, although only marginally so.

Ultimately, the conclusions of the above scenarios are that the estimators, in general, are robust towards irregularities in the instantaneous volatility process such as long memory, nonstationarity, and jumps. The performance of the estimators depends on the characteristics of the irregularity and especially on the sampling frequency.

Table 10
Simulation results for Model A with bid-ask Bounce II: Relative bias (×1000)

ξ	v	2nd lag realized bipower variation			Realized volatility HL			Realized volatility NW		
		15 min	5 min	1 min	15 min	5 min	1 min	15 min	5 min	1 min
0.003	0.05	0.6040	1.4069	−21.1331	−3.8274	−1.3772	−1.0926	−4.0726	−2.0171	0.7270
	0.50	−2.0761	2.6984	−20.5927	−2.9800	−1.6996	−0.3671	−3.0280	−1.5966	1.2103
0.005	0.05	3.9707	9.5815	11.1772	−0.3028	1.5037	0.3467	−1.1657	1.3116	6.7039
	0.50	0.0750	7.6059	11.7717	−3.6051	−1.8898	−0.4615	−1.6500	−0.7138	4.7324
0.01	0.05	12.9514	40.3629	159.7021	−1.5864	−0.0097	2.9440	1.4349	5.7893	24.9240
	0.50	14.6658	38.6138	162.7062	−0.7714	1.3298	2.3916	4.0392	7.3932	25.1500

Note: The table reports the average of the relative error statistic in Eq. (26), i.e. the relative bias, multiplied by 1000 from $L=10{,}000$ simulated replications of Model A in Eq. (24) with $\alpha=0.0124$, $\beta=-1$, $d=0$, and a bid-ask bounce effect as in Eq. (30). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p^*(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the 2nd lag realized bipower variation, Hansen–Lunde (HL) realized volatility, and Newey–West realized volatility estimators, respectively, each applying three different sampling frequencies.

Table 11
Simulation results for Model A with bid-ask Bounce II: RMSE

ξ	ν	2nd lag realized bipower variation			Realized volatility HL			Realized volatility NW		
		15 min	5 min	1 min	15 min	5 min	1 min	15 min	5 min	1 min
0.003	0.05	0.1653	0.0971	0.0480	0.2466	0.1432	0.0657	0.2357	0.1665	0.0859
	0.50	0.1685	0.0979	0.0488	0.2488	0.1474	0.0656	0.2381	0.1688	0.0882
0.005	0.05	0.1662	0.0967	0.0453	0.2496	0.1452	0.0660	0.2389	0.1686	0.0871
	0.50	0.1679	0.0983	0.0495	0.2521	0.1472	0.0668	0.2417	0.1702	0.0884
0.01	0.05	0.1675	0.1070	0.1677	0.2490	0.1460	0.0694	0.2381	0.1679	0.0904
	0.50	0.1711	0.1077	0.1803	0.2527	0.1489	0.0703	0.2420	0.1699	0.0920

Note: The table reports the root mean squared error of the relative error statistic in Eq. (26) from $L=10{,}000$ simulated replications of Model A in Eq. (30) with $\alpha=0.0124$, $\beta=-1$, $d=0$, and a bid-ask bounce effect as in Eq. (30). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p^*(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86{,}400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the 2nd lag realized bipower variation, Hansen–Lunde (HL) realized volatility, and Newey–West realized volatility estimators, respectively, each applying three different sampling frequencies.

We next consider the model (29) in Section 5, i.e. Model A with a bid-ask bounce effect. When introducing this market microstructure effect, and thereby introducing spurious volatility and serial correlation, we do not expect the estimators to remain as well behaved as above. This fact is reflected in Tables 8 and 9. Almost all biases are positive, revealing that the methods overestimate the true latent integrated volatility. That is, the bid-ask effect introduces spurious volatility and negative serial correlation which causes the estimators to overestimate the underlying integrated volatility. For example, the relative bias of the wavelet estimator (with $K=10$) is approximately 0.4 with a 1% spread, see Table 8, meaning that it overestimates the true integrated volatility by as much as 40% in this case.

As expected, we generally find that the performance of the estimators declines rapidly as the sampling frequency increases. This is in agreement with the literature where it is widely accepted that increasing the sampling frequency worsens the impact from market microstructure effects. This tendency becomes even more distinct when the spread (in percentage of the underlying asset price) increases. Hence, in general, when conducting an empirical analysis of financial time series contaminated by bid-ask bounce effects (and in general by market microstructure effects), where the spread is known to be significant (even if it is as small as 0.3%), it is extremely important not to use too high sampling frequency. Indeed, the lowest biases in our simulations in Table 8 are found for the very lowest sampling frequencies. Furthermore, the RMSEs in Table 9 indicate the usual trade-off between bias and variance in this model.

In the two scenarios of persistence and jumps in the integrated volatility process, i.e. Models A and B in Tables 2–7, our Monte Carlo study did not reveal any apparent differences between the three estimators. Now, when introducing an

Table 12
Simulation results for Model A with bid-ask bounce and discreteness I: Relative bias (×1000)

Tick size	ξ	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	$S=48$	$S=144$	$S=720$	$K=7$	$K=8$	$K=10$
0.1250	0.003	4.3283	8.5126	38.7365	0.2191	1.3101	23.1678	12.2703	21.6713	78.2102
	0.010	17.4055	50.3126	251.5964	1.0044	7.8054	150.2341	69.5755	135.2750	509.1845
0.0625	0.003	1.7312	6.5205	26.5182	0.3841	1.4168	15.7141	7.4327	14.3163	53.9762
	0.010	13.9316	44.5144	220.3049	0.4571	6.5283	131.5391	57.6773	119.9083	445.5253
0.010	0.003	1.1767	3.9023	18.3871	0.1993	1.2582	10.9835	6.1237	10.7230	37.0455
	0.010	10.9380	37.7042	195.8516	−0.1544	4.6129	117.1674	50.0171	105.8323	398.7507

Note: The table reports the average of the relative error statistic in Eq. (26), i.e. the relative bias, multiplied by 1000 from $L=10,\ 000$ simulated replications of Model A in Eq. (24) with $\alpha=0.0124$, $\beta=-1$, d=0, $\nu=0.0374$, a bid-ask bounce effect as in Eq. (30), and price discreteness as in Eq. (38). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p^*(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of $T=86,\ 400$ s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean $\tau=14$ s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

Table 13
Simulation results for Model A with bid-ask bounce and discreteness I: RMSE

Tick size	ξ	Realized volatility			Fourier			Wavelet		
		15 min	5 min	1 min	S=48	S=144	S=720	K=7	K=8	K=10
0.1250	0.003	0.1437	0.0845	0.0575	0.1436	0.0832	0.0467	0.1274	0.0924	0.0951
	0.010	0.1468	0.1011	0.2575	0.1445	0.0848	0.1576	0.1515	0.1690	0.5162
0.0625	0.003	0.1433	0.0840	0.0484	0.1449	0.0833	0.0429	0.1252	0.0909	0.0727
	0.010	0.1468	0.0979	0.2256	0.1436	0.0839	0.1391	0.1441	0.1551	0.4508
0.010	0.003	0.1444	0.0839	0.0441	0.1449	0.0832	0.0408	0.1260	0.0898	0.0601
	0.010	0.1448	0.0945	0.2012	0.1446	0.0842	0.1250	0.1392	0.1430	0.4037

Note: The table reports the root mean squared error of the relative error statistic in Eq. (26) from L=10, 000 simulated replications of Model A in Eq. (24) with α=0.0124, β=−1, d=0, ν=0.0374, a bid-ask bounce effect as in Eq. (30), and price discreteness as in Eq. (38). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for (p^* (t), σ (t)), t=1, ..., T, assuming 24 h trading, i.e. a total of T=86, 400 s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean τ=14 s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the realized volatility, Fourier, and Wavelet estimators, respectively, each applying three different sampling frequencies.

irregularity that the methods are not theoretically equipped to handle, we can really unveil the strengths of the three estimators. The wavelet estimator of integrated volatility generally provides the highest biases and RMSEs for all sizes of the spread and for any frequency, thus rendering the wavelet estimator less useful in case of bid-ask bounce effects. The realized volatility estimator is also heavily biased in the presence of the bid-ask bounce, but not quite as badly as the wavelet estimator.

Probably the most striking finding is that the Fourier method is vastly superior in this scenario with market microstructure effects. The Fourier estimator is practically unaffected with respect to both bias and RMSE except in the case with the highest sampling frequency. Even then the bias and RMSE are vastly superior to those of the other methods. The superiority of the Fourier estimator in the presence of market microstructure contamination can be attributed to the decomposition of the integrated variance into components of varying frequencies by the Fourier transform. That is, including only the lowest S frequencies in the Fourier estimator in Eq. (16) implies that high-frequency noise or short-run noise is ignored by the estimator. Hence, by choosing a smaller number of (low) frequency ordinates to be used for estimation, i.e. by choosing S small, it is in principle possible to render the Fourier estimator invariant to short-run noise introduced by market microstructure effects.

In Tables 10 and 11 we present the relative biases and RMSEs of the three alternative integrated volatility estimators from Section 5, i.e. the 2nd lag realized bipower variation in Eq. (35), the Hansen and Lunde (2006) estimator in Eq. (36), and the Newey and West (1987) estimator in Eq. (37). The tables show that the 2nd lag realized bipower

Table 14
Simulation results for Model A with bid-ask bounce and discreteness II: Relative bias (×1000)

Tick size	ξ	2nd lag realized bipower variation			Realized volatility HL			Realized volatility NW		
		15 min	5 min	1 min	15 min	5 min	1 min	15 min	5 min	1 min
0.1250	0.003	4.0349	7.6295	−0.3302	0.6790	0.4415	1.0380	0.6549	1.2883	5.4225
	0.010	17.0115	49.1670	219.2320	0.8421	1.2727	3.3692	6.0186	9.4780	31.9380
0.0625	0.003	2.1973	6.2719	−11.1779	−1.8739	−0.1687	0.5906	−1.2866	0.1961	3.8183
	0.010	14.6494	44.1842	188.2374	−0.9695	−0.8188	3.0739	1.9532	6.4365	27.3608
0.010	0.003	1.1387	3.0761	−19.1488	−0.2764	0.1794	0.6930	0.5616	0.9375	2.6133
	0.010	11.3343	37.1567	163.5724	−1.3150	−1.1456	1.5831	2.2341	4.7169	23.4356

Note: The table reports the average of the relative error statistic in Eq. (26), i.e. the relative bias, multiplied by 1000 from L=10,000 simulated replications of Model A in Eq. (24) with α=0.0124, β=−1, d=0, ν=0.0374, a bid-ask bounce effect as in Eq. (30), and price discreteness as in Eq. (38). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for (p^* (t), σ (t)), t=1, ..., T, assuming 24 h trading, i.e. a total of T=86,400 s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean τ=14 s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the 2nd lag realized bipower variation, Hansen–Lunde (HL) realized volatility, and Newey–West realized volatility estimators, respectively, each applying three different sampling frequencies.

Table 15
Simulation results for Model A with bid-ask bounce and discreteness II: RMSE

Tick size	ξ	2nd lag realized bipower variation			Realized volatility HL			Realized volatility NW		
		15 min	5 min	1 min	15 min	5 min	1 min	15 min	5 min	1 min
0.1250	0.003	0.1651	0.0960	0.0454	0.2473	0.1439	0.0657	0.2376	0.1667	0.0865
	0.010	0.1684	0.1114	0.2271	0.2477	0.1467	0.0699	0.2375	0.1674	0.0929
0.0625	0.003	0.1646	0.0956	0.0453	0.2483	0.1441	0.0652	0.2385	0.1678	0.0864
	0.010	0.1697	0.1097	0.1956	0.2480	0.1466	0.0697	0.2348	0.1668	0.0911
0.010	0.003	0.1652	0.0956	0.0474	0.2477	0.1442	0.0644	0.2379	0.1676	0.0864
	0.010	0.1671	0.1054	0.1714	0.2496	0.1446	0.0691	0.2375	0.1676	0.0897

Note: The table reports the root mean squared error of the relative error statistic in Eq. (26) from L=10, 000 simulated replications of Model A in Eq. (24) with α=0.0124, β=−1, d=0, v=0.0374, a bid-ask bounce effect as in Eq. (30), and price discreteness as in Eq. (38). For each replication we simulate (through simple Euler discretization) artificial time series of second-by-second intra-daily data points for $(p^*(t), \sigma(t))$, $t=1, \ldots, T$, assuming 24 h trading, i.e. a total of T=86,400 s. Then we sample from this log-price "process" by assuming that the time difference between successive observations is exponentially distributed with mean τ=14 s, such that each simulated replication has a different number (n) of irregularly spaced observations. The results are presented for the 2nd lag realized bipower variation, Hansen–Lunde (HL) realized volatility, and Newey–West realized volatility estimators, respectively, each applying three different sampling frequencies.

variation estimator is greatly affected by the spurious volatility and the negative serial correlation introduced by the bid-ask bounce effect. For small ξ and high M the negative serial correlation causes the 2nd lag realized bipower variation estimator to underestimate the true latent integrated volatility, whereas for higher ξ the 2nd lag realized bipower variation is nearly as biased as realized volatility. Thus, the 2nd lag realized bipower variation seems to be slightly more robust to market microstructure effects, as conjectured by Andersen et al. (in press-b, p. 16), only in the case of a large spread and high sampling frequency.

Tables 10 and 11 also demonstrate that the bias correction suggested by Hansen and Lunde (2006) works very well, and in particular outperforms the Newey and West (1987) correction since (at least for higher frequencies) $\hat{\sigma}^{2*}_{\mathrm{HL},M}$ compares favorably to $\hat{\sigma}^{2*}_{\mathrm{NW},M}$ both in terms of bias and RMSE. Of course, given the nature of the noise in model (29), the Hansen and Lunde (2006) estimator is presumably superior to the Newey and West (1987) correction, which is robust to a more general noise structure involving higher order autoregressive or moving average terms — that may or may not be relevant in practice depending on the specific market.

More interestingly, the autocorrelation bias correction methods are only noticeably superior to the Fourier method for the very highest sampling frequency (M=1440 and S=720). In fact, the RMSEs for $\hat{\sigma}^{2*}_{\mathrm{HL},M}$ are generally much higher than the RMSEs for the Fourier method except for the combination of highest sampling frequency and highest spread. This implies that the Fourier method remains a very attractive estimator in the presence of market microstructure effects even in comparison with methods specifically designed to handle such contamination. However, if one wishes to incorporate the full information available in ultra high-frequency data the need for bias correction methods remains.

To add further empirical realism to the simulation study, Tables 12–15 show the results of introducing both bid-ask bounce effects and price discreteness. The discreteness follows the latest evolution in tick size, i.e. we consider tick sizes 1/8, 1/16, and 1/100. For easy comparison to the results in Tables 8–11, the bid-ask spread is either 0.3% or 1%.

As seen from the tables, limiting the tick size to 1/100 causes no extra bias or RMSE compared to the results in Tables 8–11 for bid-ask bounce effect only. Since the U.S. stock markets are now fully decimalized, this implies that price discreteness no longer appears to be an issue in terms of volatility estimation and modeling, at least compared to other market microstructure effects. As expected, we see that the biases increase with tick size although the effect is moderate compared to changing the bid-ask spread. This implies that introducing price discreteness, and thereby increasing the complexity of the microstructure effects, does not change the general results from Tables 8–11, but only acts to support the conclusions. For instance, we find the wavelet estimator to be seriously biased when the bid-ask spread is 1% with a tick size of 1/8, and as before, that the Fourier estimator performs superiorly.

Again we see that the Hansen and Lunde (2006) estimator works very well, the Newey–West estimator works nearly as well, whereas the 2nd lag realized bipower variation is just as biased as the realized volatility. More interestingly, the bias correction methods are still only noticeably superior to the Fourier method in terms of bias for the highest sampling frequency ($M=1440$ and $S=720$), but the bias correction methods have (much) larger RMSEs. Overall, this reinforces the conclusion that the Fourier method is an attractive estimator in the presence of market microstructure noise, even when compared with methods specifically designed to handle such contamination.

7. Empirical illustration

To illustrate empirically the three integrated volatility estimators we consider so-called volatility signature plots for General Electric Co. (GE) mid-quotes for realized volatility as well as the Fourier and wavelet methods. The volatility signature plot advocated by Andersen et al. (2000) graphs the average sample values measured over a long time span – preferably several years – of RV for different high-frequency return sampling frequencies, i.e. M. These different realized volatility measures should, in the absence of microstructure frictions, provide reasonable estimates of the same average return variability irrespective of the underlying sampling frequency. Meanwhile, if there are significant microstructure biases present for the measures based on the highest sampling frequencies, this should be reflected in systematic deviations of the corresponding average variation measures relative to those computed from returns sampled at lower frequencies.

As explained by Hansen and Lunde (2006), this diagnostic tool is quite useful for gauging the impact of noise in the various volatility estimators. The use of mid-quotes (as opposed to the actual transaction prices used by Andersen et al., 2000) reduces the influence of microstructure effects, e.g. Hansen and Lunde (2006). Various other generalized volatility signature plots for bipower variation, etc., have been introduced by Andersen et al. (2006a,b).

Fig. 1 displays the volatility signature plots for GE mid-quotes using daily averages over the sample period January 1, 2001 to June 30, 2004, for realized volatility as well as the Fourier and wavelet methods. To make the comparison simple, the log-scale horizontal axis of the figure measures the implied value of M for each method, i.e. M for realized volatility, $M=2S$ for the Fourier method, and $M=2^K$ for the wavelet method, and each method is plotted for a range of values of the implied M. Together, the three plots in Fig. 1 give a simple indication of the influence of market microstructure effects, and generally the volatility signature plots may be used as simple diagnostic tools to gauge the effect of market microstructure noise on the different volatility estimators. Moreover, since our simulations indicate that realized volatility and the wavelet estimator are more biased in the presence of market microstructure noise, the fact that the Fourier estimator is above realized volatility and the wavelet estimator for empirically realistic sampling frequencies (90 s to 30 min or M between 13 and 260), indicates that the actual volatility might be higher on average than predicted by the much used realized volatility.

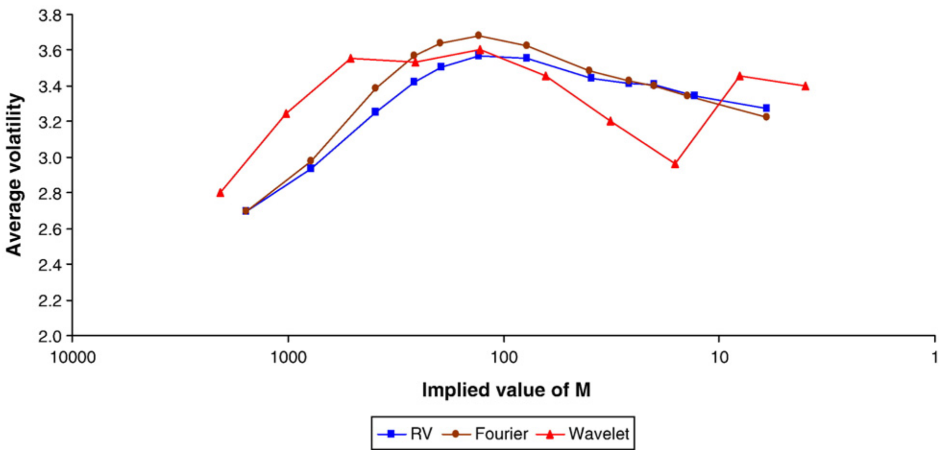

Fig. 1. Volatility signature plots for GE based on realized volatility, fourier, and wavelet methods. Note: The figure displays volatility signature plots for General Electric (GE) mid-quotes using daily averages over the sample period January 1, 2001 to June 30, 2004, for realized volatility as well as the Fourier and wavelet methods. Volatility signature plots graph the average sample values (measured over a long time span) of volatility estimators for different high-frequency return sampling frequencies. To make the comparison simple, the log-scale horizontal axis of the figure measures the implied value of M for each method, i.e. M for realized volatility, $M=2S$ for the Fourier method, and $M=2^K$ for the wavelet method, and each method is plotted for a range of values of the implied M.

8. Concluding remarks

We have considered the properties of three estimation methods for integrated volatility, i.e. realized volatility, the Fourier estimator, and the wavelet estimator, when only a finite (high-frequency) sample of the price process is observed. We considered several different generating mechanisms for the instantaneous volatility: Ornstein–Uhlenbeck, long memory, and jump processes, and the possibility of market microstructure effects contaminating the data was also entertained in models that allow for a bid-ask bounce effect and price discreteness. In the latter case we also considered alternative estimators with theoretical justification under market microstructure noise.

Our simulation study reveals that the theoretical robustness of the estimators towards persistence or jumps in the stochastic process governing the latent volatility carries over to practice. On the other hand, irregularities such as bid-ask bounce effects, which the methods have no theoretical robustness against, in general render the wavelet estimator, and to a lesser degree realized volatility, less useful in practice. However, we find the Fourier method to be superior compared to the other two estimators in the case of market microstructure noise, and indeed this estimator remains very useful in that case. Even more strikingly, when compared to the bias correction methods designed specifically to handle market microstructure effects, the Fourier method is superior with respect to RMSE while having only slightly higher bias.

In a brief empirical illustration using high-frequency GE mid-quote data, we examined volatility signature plots for the Fourier and wavelet estimators, generalizing the well known realized volatility signature plot. The three volatility signature plots were compared, and in general may be used as a diagnostic tool to gauge the influence of market microstructure noise on the estimators.

References

Ait-Sahalia, Y., Mykland, P.A., Zhang, L., 2005. How often to sample a continuous-time process in the presence of market microstructure noise. Review of Financial Studies 18, 351–416.

Andersen, T.G., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance 4, 115–158.

Andersen, T.G., Bollerslev, T., 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. International Economic Review 39, 885–905.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2000. Great realizations. Risk 13, 105–108.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001a. The distribution of realized stock return volatility. Journal of Financial Economics 61, 43–76.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2001b. The distribution of exchange rate volatility. Journal of the American Statistical Association 96, 42–55.

Andersen, T.G., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time equity return models. Journal of Finance 57, 1239–1284.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modelling and forecasting realized volatility. Econometrica 71, 579–625.

Andersen, T.G., Bollerslev, T., Meddahi, N., 2004. Analytical evaluation of volatility forecasts. International Economic Review 45, 1079–1110.

Andersen, T.G., Bollerslev, T., Diebold, F.X., in press-a. Parametric and nonparametric volatility measurement. In: L.P. Hansen, Y. Ait-Sahalia (Eds), 'Handbook of Financial Econometrics', North-Holland, Amsterdam.

Andersen, T.G., Bollerslev, T., Diebold, F.X., in press-b. Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. Review of Economics and Statistics.

Andersen, T.G., Bollerslev, T., Frederiksen, P.H., Nielsen, M.Ø., 2006a. Comment on Peter R. Hansen and Asger Lunde: realized variance and market microstructure noise. Journal of Business and Economic Statistics 24, 173–179.

Andersen, T.G., Bollerslev, T., Frederiksen, P.H., Nielsen, M.Ø., 2006b. Continuous-time Models, Realized Volatilities, and Testable Distributional Implications for Daily Stock Returns. Working Paper. Northwestern University.

Andreou, E., Ghysels, E., 2002. Rolling-sample volatility estimators: some new theoretical, simulation, and empirical results. Journal of Business and Economic Statistics 20, 363–376.

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

Bandi, F.M., Russell, J.R., 2005. Microstructure Noise, Realized Variance, and Optimal Sampling. Working Paper. University of Chicago Graduate School of Business.

Bandi, F.M., Russell, J.R., 2006. Separating microstructure noise from volatility. Journal of Financial Economics 79, 655–692.

Barndorff-Nielsen, O.E., Shephard, N., 2001. Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics (with discussion). Journal of the Royal Statistical Society. Series B 63, 167–241.

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

Barndorff-Nielsen, O.E., Shephard, N., 2002b. Estimating quadratic variation using realized variance. Journal of Applied Econometrics 17, 457–477.

Barndorff-Nielsen, O.E., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps (with discussion). Journal of Financial Econometrics 2, 1–48.

Barndorff-Nielsen, O.E., Shephard, N., 2005. How accurate is the asymptotic approximation to the distribution of realised variance? In: Andrews, D.W.K., Stock, J.H. (Eds.), Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg. Cambridge University Press, Cambridge, pp. 306–331.

Barndorff-Nielsen, O.E., Shephard, N., 2006. Econometrics of testing for jumps in financial economics using bipower variation. Journal of Financial Econometrics 4, 1–30.

Barndorff-Nielsen, O.E., Shephard, N., in press. Variation, jumps, market frictions, and high frequency data in financial econometrics. In: R. Blundell, T. Persson, W.K. Newey, (Eds), Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress. Cambridge University Press, Cambridge, UK.

Barndorff-Nielsen, O.E., Nielsen, B., Shephard, N., Ysusi, C., 2004. Measuring and forecasting financial variability using realised variance with and without a model. In: Harvey, A.C., Koopman, S.J., Shephard, N. (Eds.), State Space and Unobserved Component Models: Theory and Applications. Proceedings of a Conference in Honour of James Durbin. Cambridge University Press, Cambridge, pp. 205–235.

Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., Shephard, N., 2006. Designing Realised Kernels to Measure the Ex-post Variation of Equity Prices in the Presence of Noise. Working Paper. Nuffield College, Oxford University.

Barucci, E., Reno, R., 2002a. On measuring volatility and the GARCH forecasting performance. Journal of International Financial Markets, Institutions and Money 12, 183–200.

Barucci, E., Reno, R., 2002b. On measuring volatility of diffusion processes with high frequency data. Economics Letters 74, 371–378.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

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

Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton.

Comte, F., Renault, E., 1996. Long-memory continuous-time models. Journal of Econometrics 73, 101–149.

Comte, F., Renault, E., 1998. Long-memory in continuous-time stochastic volatility models. Mathematical Finance 8, 291–323.

Cox, J.C., Ingersoll, J.E., Ross, S.A., 1985. A theory of the term structure of interest rates. Econometrica 53, 385–408.

Dacorogna, M.M., Gencay, R., Müller, U.A., Pictet, O.V., Olsen, R.B., 2001. An Introduction to High-Frequency Finance. Academic Press, San Diego.

Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. Journal of Finance 59, 1367–1403.

Eraker, B., Johannes, M., Polson, N., 2003. The impact of jumps in volatility and returns. Journal of Finance 58, 1269–1300.

French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–30.

Ghysels, E., Harvey, A.C., Renault, E., 1996. Stochastic volatility. In: Rao, C.R., Maddala, G.S. (Eds.), Statistical Methods in Finance. North-Holland, Amsterdam, pp. 119–191.

Ghysels, E., Santa-Clara, P., Valkanov, R., 2006. Predicting volatility: getting the most out of return data sampled at different frequencies. Journal of Econometrics 131, 59–95.

Guillaume, D.M., Dacorogna, M.M., Davé, R.R., Müller, U.A., Olsen, R.B., Pictet, O.V., 1997. From the bird's eye to the microscope: a survey of new stylized facts of the intra-daily foreign exchange markets. Finance and Stochastics 1, 95–129.

Hansen, P.R., Lunde, A., 2005. A realized variance for the whole day based on intermittent high-frequency data. Journal of Financial Econometrics 3, 525–554.

Hansen, P.R., Lunde, A., 2006. Realized variance and market microstructure noise (with discussion). Journal of Business and Economic Statistics 24, 127–218.

Hansen, P.R., Lunde, A., Nason, J.M., 2003. Choosing the best volatility models: the model confidence set approach. Oxford Bulletin of Economics and Statistics 65, 839–861.

Hasbrouck, J., 1999a. The dynamics of discrete bid and ask quotes. Journal of Finance 54, 2109–2142.

Hasbrouck, J., 1999b. Security bid/ask dynamics with discreteness and clustering: simple strategies for modeling and estimation. Journal of Financial Markets 2, 1–28.

Høg, E., Lunde, A., 2003. Wavelet Estimation of Integrated Volatility. Working Paper. Aarhus School of Business.

Huang, X., Tauchen, G., 2005. The relative contribution of jumps to total price variance. Journal of Financial Econometrics 3, 456–499.

Hull, J.C., White, A., 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance 42, 281–300.

Koopman, S.J., Jungbacker, B., Hol, E., 2005. Forecasting daily variability of the S&P 100 stock index using historical, realised and implied volatility measurements. Journal of Empirical Finance 12, 445–475.

Maheu, J.M., McCurdy, T.H., 2002. Nonlinear features of realized FX volatility. Review of Economics and Statistics 84, 668–681.

Malliavin, P., Mancino, M.E., 2002. Fourier series method for measurement of multivariate volatilities. Finance and Stochastics 6, 49–61.

Marinucci, D., Robinson, P.M., 1999. Alternative forms of fractional Brownian motion. Journal of Statistical Planning and Inference 80, 111–122.

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

Oomen, R.C.A., 2002. Modelling Realized Variance When Returns are Serially Correlated. Working Paper. Warwick Business School.

Owens, J.P., Steigerwald, D.G., 2006. Noise reduced realized volatility: a Kalman filter approach. In: Fomby, T.B., Terrell, D. (Eds.), Advances in Econometrics Volume 20: Econometric Analysis of Financial and Economic Time Series — Part A. Elsevier Science Ltd., Oxford, pp. 211–229.

Patton, A., 2005. Volatility Forecast Evaluation and Comparison Using Imperfect Volatility Proxies. Working Paper. London School of Economics.

Protter, P., 1990. Stochastic Integration and Differential Equations: A New Approach. Springer-Verlag, New York.

Roll, R., 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. Journal of Finance 39, 1127–1139.

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, 1394–1411.

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