Asymptotic normality and finite-sample robustness of the Fourier spot volatility estimator in the presence of microstructure noise

Maria Elvira Mancino[*], Tommaso Mariotti[†], Giacomo Toscano[‡]

February 3, 2025

Abstract

We study the efficiency and robustness of the Fourier spot volatility estimator by Malliavin and Mancino [2002] when high-frequency prices are contaminated by microstructure noise. Firstly, we show that the estimator is consistent and asymptotically efficient in the presence of additive noise, establishing a Central Limit Theorem (CLT) with the optimal rate of convergence $n^{1/8}$. Additionally, we complete the asymptotic theory proved by Mancino and Recchioni [2015] in the absence of noise, obtaining a CLT with the optimal rate of convergence $n^{1/4}$. Feasible CLTs with the optimal convergence rate are also obtained, by proving the consistency of the Fourier estimators of the spot volatility of volatility and the quarticity in the presence of noise. Secondly, we introduce a feasible method for selecting the cutting frequencies of the estimator in the presence of noise, based on the optimization of the integrated asymptotic variance. Finally, we provide support to the accuracy and robustness of the method by means of a numerical study and an empirical exercise, which is conducted using tick-by-tick prices of three US stocks with different liquidity.

Keywords: non-parametric spot volatility estimation, Fourier estimator, stochastic volatility, high-frequency data.

[*]University of Florence, Dept. of Economics and Management, Via delle Pandette 9, 50127 Firenze, Italy: mariaelvira.mancino@unifi.it

[†]University of Turin, Dept. of Economics, Social Studies, Applied Mathematics and Statistics, Corso Unione Sovietica 218/bis, 10134 Torino, Italy: tommaso.mariotti@unito.it

[‡]University of Florence, Dept. of Economics and Management, Via delle Pandette 9, 50127 Firenze, Italy: giacomo.toscano@unifi.it

1 Introduction

Even though the non-parametric estimation of the time-varying volatility of financial assets has long been recognized as a relevant topic, see, e.g., Andersen et al. [2003] and the references therein, in the last couple of decades, the increasing availability of high-frequency market data has given an enormous impulse to the investigation of novel estimation methodologies and the exploration of their applications. Besides, the focus has shifted from the estimation of the volatility accumulated over a given time horizon, i.e., the integrated volatility, to the reconstruction of the trajectory of the volatility process, i.e., the instantaneous (or spot) volatility.

The benchmark for estimating the integrated volatility of an asset on a fixed time interval with high-frequency prices is given by the sum of the squared log-returns, i.e., the realized volatility. Under the Brownian semi-martingale assumption, a classical result (Jacod [1994]) gives that in the limit, as the time between two consecutive observations tends to zero, the realized volatility converges to the quadratic variation of the asset log-price process, i.e., the integrated volatility. Then, the estimation of the integrand function, namely, the spot volatility, goes through the differentiation of the sum of the squared log-returns by considering an asymptotically shrinking time window (as in Foster and Nelson [1996]). Nevertheless, two difficulties appear. Firstly, such empirical derivative required to get the spot volatility gives rise to numerical instabilities. Secondly, high-frequency prices are affected by the noise due to market microstructure, which causes a discrepancy between theoretical models, based on semimartingales, and the actual price observations, sampled at very fine intervals.

To reduce the numerical instabilities due to the differentiation of the quadratic variation, authors have extensively adopted kernel-smoothing procedures, see Kristensen [2010] for a study of the asymptotic properties of kernel-based spot estimators. Moreover, mod-

ifications of spot estimators have been proposed to make them robust to the presence of microstructure noise. Such modifications, which rely on the pre-processing of price observations, include the Two-Scale sub-sampling method by Zu and Boswijk [2014], which is also used by Mykland et al. [2019] for deriving the Smoothed Two-Scale estimator, and the Pre-Averaging method for kernel-based estimators by Aït-Sahalia and Jacod [2014], which is employed for obtaining the Pre-Averaging Kernel estimator by Figueroa-López and Wu [2022], the Pre-Averaging ReMeDI estimator by Li and Linton [2023] and the Local Method of Moments by Bibinger et al. [2019]. A different approach derives nonparametric spot volatility estimators from short-dated options, see Todorov [2019], Todorov and Zhang [2023].

This paper studies asymptotic error normality and finite-sample robustness, in the presence of microstructure noise, of an alternative non-parametric estimator of the spot volatility, namely, the Fourier estimator introduced by Malliavin and Mancino [2002]. The Fourier method is a global method, in the sense that it aims to estimate the volatility path over the interval of interest. It reconstructs the spot volatility as a series expansion. The coefficients of this series are computed through a convolution formula from the estimated Fourier coefficients of the log-price increment. Specifically, for a given sample size n, the efficient implementation of the Fourier spot volatility estimator requires the balancing of only two parameters: the cutting frequency N in the convolution formula (2) that reconstructs the Fourier coefficients of the volatility, and the number M of the volatility Fourier coefficients to be used in the Fourier inversion formula (1), which yields the estimated volatility path. In the presence of noise, the definition of the Fourier estimator is unchanged with respect to the no-noise case, and the efficiency of the estimator is achieved by carefully reducing N, that is, by cutting out the highest frequencies in the construction of the Fourier coefficients of the volatility. This way, the Fourier method allows obtaining a noise-robust estimator of the spot volatility without the need for any manipulation of the

original observations, such as, e.g., sparse sampling, or any bias correction.

The convolution formula (2) exploits the contribution of weighted realized auto-covariances to deal with the presence of noise, with N determining the weight. The role of realized auto-covariances has been studied in the early work of Zhou [1996] and, later, in Barndorff-Nielsen et al. [2008], for integrated volatility estimators. In this spirit, the Fourier spot volatility estimator can be related to a local version of the infinite-lag realized kernel estimator, see Remark 2.1 in Section 2 and [Barndorff-Nielsen et al., 2008, Section 4.5]. Removing the cross products of the log-returns, as proposed in [Mancini et al., 2015, Proposition 4.2], yields an estimator that is substantially different from the Fourier spot volatility estimator. In fact, it loses the robustness to microstructure noise, as shown in Mancino and Recchioni [2015].

The asymptotic properties of the Fourier volatility estimator in the absence of noise have been investigated by several papers. The weak convergence of the asymptotic error distribution of the Fourier spot covariance estimator was first derived in Malliavin and Mancino [2009], with a sub-optimal rate. Clement and Gloter [2011] then proved the asymptotic normality of the Fourier integrated covariance estimator with the optimal convergence rate and variance of the asymptotic distribution, under the assumption of irregular and non-synchronous observations. A CLT for the Fourier spot volatility estimator was proved in Mancino and Recchioni [2015], with a slightly sub-optimal rate of convergence. Besides, Cuchiero and Teichmann [2015] proposed a spot covariance estimator based on a modification of the classical Fourier method and studied its asymptotic normality in the presence of jumps. Finally, a different but related study is present in Park et al. [2016], where the authors prove the consistency and asymptotic normality for a class of Fourier-type estimators of the Fourier coefficients of the covariance, named Fourier Realized Kernels. Specifically, in Park et al. [2016], a CLT for any Fourier coefficient is derived under some general conditions that allow for microstructure noise effects and asynchronicity. The rate

of convergence depends on the relative liquidity between assets and reaches $n^{1/5}$ if the assets are synchronous. The convergence rate of the spot Fourier Realized Kernel estimator is not obtained.

In the presence of microstructure noise, the finite-sample performance of the Fourier method has been extensively analyzed: for what concerns the Fourier estimator of the integrated variance, see, e.g., Barucci and Renò [2002], Hansen and Lunde [2006], Nielsen and Frederiksen [2008], Mancino and Sanfelici [2008]; for the Fourier spot volatility estimator, Mancino and Recchioni [2015] performed a thorough simulation study that included auto-correlated and price-dependent noise patterns. More recently, Mariotti et al. [2023] have evaluated the finite-sample efficiency of the different spot volatility estimators, including the one based on the Fourier method, with high-frequency samples generated via the limit-order-book simulator by Huang et al. [2015], which is capable of reproducing several microstructural characteristics of financial markets.

The previous studies have explored the finite-sample efficiency of the Fourier volatility estimator of the volatility in the presence of noise, without investigating the asymptotic properties of the estimator, namely the rate of convergence and the asymptotic variance. The present paper fills this gap providing novel results on the (pointwise) asymptotic normality of the Fourier estimator, which show that the latter reaches the optimal rate of convergence in the presence of additive microstructure noise. Specifically, in this paper we first fine-tune the result of Mancino and Recchioni [2015], that is, under the assumption that the volatility process is also a continuous Itô semimartingale, we prove that the Fourier spot volatility estimator attains the optimal rate of convergence in the absence of noise, that is, $n^{1/4}$. Then, we prove that the Fourier estimator attains the optimal rate of convergence also in the presence of additive microstructure noise, that is, $n^{1/8}$. This is the convergence rate of the spot volatility estimators in [Aït-Sahalia and Jacod, 2014, Section 8.7] and [Figueroa-López and Wu, 2022, Theorem 2.2], while the Two-scale estimator in [Zu and

Boswijk, 2014, Theorem 2] attains the convergence rate $n^{1/12}$. Furthermore, we prove the consistency of the Fourier estimators of the spot volatility of volatility and quarticity in the presence of noise. This novel result allows us to establish the feasible CLT with optimal convergence rate not only in the no-noise case but also in the noisy case. As mentioned above, the Fourier estimator does not need any modification to be efficient in the presence of noise, but only requires a proper reduction of the cutting frequency N; precisely, in the absence of noise, N is chosen equal to the Nyquist frequency $n/2$ (see Theorem 3.1); instead, in the presence of noise, it has order $n^{1/2}$ (see Theorem 3.6).

For the selection of the cutting frequencies N and M in the presence of noise, we introduce a numerical feasible adaptive method that minimizes the asymptotic variance of the estimation error of the Fourier estimator. Specifically, this method relies on the gradient-descent algorithm and optimizes the conditional mean integrated squared error, which is approximated, at high-frequencies, by the integrated asymptotic variance (see [Zu and Boswijk, 2014, Section 3.4]). The introduction of a feasible estimation strategy, rooted in asymptotic theory, represents the main innovation of the finite-sample study conducted in this paper, compared to the one provided in [Mancino and Recchioni, 2015, Section 4].

We evaluate the robustness of the feasible adaptive method in different simulated scenarios, where we also compare the finite-sample performance of the Fourier estimator with that of alternative noise-robust spot volatility estimators. The results of the simulation study provide support to the robustness of the feasible method and the relative finite-sample efficiency of the Fourier estimator. Finally, we test the feasible method for selecting N and M in an empirical exercise, conducted with tick-by-tick prices of three US stocks with different liquidity, sampled over the period December 1, 2022 - December 30, 2022. The results of the test suggest that the feasible method is satisfactorily accurate with real market prices, for the different liquidity levels considered.

The paper is organized as follows. Section 2 contains the definition of the Fourier

estimator of the instantaneous volatility. Asymptotic results are detailed in Section 3. Section 4 illustrates the simulation study, while Section 5 contains the empirical exercise. Section 6 concludes. The proofs are provided in the supplementary material, along with additional numerical results.

2 The Fourier estimator of spot volatility

The Fourier estimator of spot volatility was originally proposed in Malliavin and Mancino [2002]. Its implementation consists of three steps. In the first step, one computes the discrete Fourier coefficients of the log-returns. The second step then entails computing the convolution of these coefficients, which yields the Fourier coefficients of the volatility. Finally, the last step involves the reconstruction of the volatility path via the Fourier-Fejér inversion formula.

Suppose that the asset log-price, denoted by p, satisfies the following assumptions:

(A.I) the process p is a continuous Itô semimartingale satisfying

$$dp(t) = \sigma(t)\, dW_t + b(t)\, dt,$$

where W is a standard Brownian motion on a filtered probability space $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, P)$, which satisfies the usual conditions.

(A.II) The spot volatility[1] process σ^2 is an Itô semimartingale satisfying

$$d\sigma^2(t) = \gamma(t)dZ_t + b_v(t)dt,$$

where Z is a standard Brownian motion adapted to the filtration $\mathcal{F}$ and such that $d\langle W, Z\rangle_t = \phi\, dt$, with constant ϕ.

[1]We follow the relevant econometric literature and use the term volatility as a synonym of variance, thus we refer to $\sigma^2(t)$ as the volatility process.

(A.III) The processes σ, b, γ and b_v are continuous and adapted stochastic processes defined on the same probability space $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, P)$ and such that

$$E\left[\int_0^T \sigma^8(t)dt\right] < \infty\,,\; E\left[\int_0^T b^4(t)dt\right] < \infty\,,\; E\left[\int_0^T \gamma^4(t)dt\right] < \infty\,,\; E\left[\int_0^T b_v^2(t)dt\right] < \infty.$$

The processes are specified in such a way that σ and γ are a.s. positive.

Suppose that p is observed at discrete, irregularly-spaced points in time: $\{0 = t_{0,n} \leq \ldots t_{i,n} \ldots \leq t_{[nT],n} = T\}$. For brevity, we omit the second index n. Let $\rho(n) := \max_{0\leq h\leq [nT]-1} |t_{h+1} - t_h|$ and suppose that $\rho(n) \to 0$ as $n \to \infty$. Consider the following interpolation formula

$$p_n(t) := \sum_{i=0}^{[nT]-1} p_{t_i} I_{[t_i, t_{i+1}[}(t)$$

and define the discrete Fourier coefficients of the log-price increment as

$$c_k(dp_n) := \frac{1}{T}\sum_{i=0}^{n-1} e^{-\mathrm{i}\frac{2\pi}{T}t_i k}\delta_i(p),$$

with $\delta_i(p) := p_{t_{i+1}} - p_{t_i}$. According to Malliavin and Mancino [2009], for any $t \in (0,T)$, the spot volatility estimator is defined as

$$\widehat{\sigma}^2_{nNM}(t) := \sum_{|k|\leq M}\left(1 - \frac{|k|}{M+1}\right) e^{\mathrm{i}\frac{2\pi}{T}tk} c_k(\sigma^2_{nN}), \tag{1}$$

where $c_k(\sigma^2_{nN})$ is an unbiased estimator of the k-th Fourier coefficient of the volatility process, obtained through the following convolution formula:

$$c_k(\sigma^2_{nN}) := \frac{T}{2N+1}\sum_{|h|\leq N} c_h(dp_n)c_{k-h}(dp_n). \tag{2}$$

In the absence of noise, the consistency and asymptotic normality with sub-optimal convergence rate of the spot volatility estimator (1) for fixed $t \in (0,T)$ are proved by Mancino and Recchioni [2015], where, instead of Assumption (A.II), the authors assume that the volatility process is a.s. Hölder continuous with parameter $\nu \in (0, 1/2)$.

By means of the rescaled Dirichlet kernel $D_N(x) := \frac{1}{2N+1}\sum_{|k|\leq N} e^{\mathrm{i}k\frac{2\pi}{T}x}$, it is possible to express (2) as

$$c_k(\sigma^2_{nN}) = \frac{1}{T}\sum_{i=0}^{n-1}\sum_{j=0}^{n-1} D_N(t_j - t_i)e^{-\mathrm{i}k\frac{2\pi}{T}t_j}\delta_i(p)\delta_j(p).$$

Therefore, the Fourier spot volatility estimator (1) can be rewritten, using two kernels, as

$$\widehat{\sigma}^2_{nNM}(t) = \frac{1}{T}\sum_{i=0}^{n-1}\sum_{j=0}^{n-1} F_M(t-t_j)D_N(t_j-t_i)\delta_i(p)\delta_j(p), \tag{3}$$

where $F_M(x) := \sum_{|k|\leq M}\left(1-\frac{|k|}{M+1}\right)e^{\mathrm{i}\frac{2\pi}{T}kx}$ is the Fejér kernel.

Remark 2.1 *The form (3) for the Fourier spot volatility estimator highlights the presence of the cross-products arising from the convolution formula (2). In the case of the integrated volatility, the contribution of auto-covariances has early been considered by Zhou [1996] and, later, by Barndorff-Nielsen et al. [2008], who propose the* realised kernel *to correct for the bias of realized-variance-type estimators in the presence of microstructure noise. The definition of the realised kernel assigns zero weight to all the autocovariances at lags higher than a fixed order. However, the authors also consider realised kernels with infinite lags (see [Barndorff-Nielsen et al., 2008, Sect. 4.5]) and different kernels, including the Dirichlet kernel, which produces an estimator closely related to the Fourier estimator of the integrated volatility. Confirming their findings, in Section 3.2 we will prove that the presence of the cross-products produced by the convolution formula (2) is crucial to ensure the robustness of the Fourier estimator in the presence of microstructure noise. Specifically, in the presence of noise, the Fourier estimator remains efficient with an appropriate choice of the frequency N, which controls the convolution estimating the volatility Fourier coefficients. On the contrary, if the auto-covariances in eq. (2) are ignored, the Fourier estimator in (3) can be written as a kernel-type spot volatility estimator, namely*

$$\widehat{\sigma}^2_{nM}(t) = \frac{1}{T}\sum_{j=0}^{n-1} F_M(t-t_j)(\delta_j(p))^2. \tag{4}$$

Kernel-type spot volatility estimators as (4), implemented with several kernels, are studied - in the absence of noise - in Kristensen [2010], Mancini et al. [2015], Figueroa-López and Li [2020], among others. However, these kernel-type estimators are not robust in the presence of noise if applied to the original sample data, as shown in Mancino and Recchioni

[2015]: therefore, the authors need to employ the pre-averaging technique and introduce a bias correction, see Zu and Boswijk [2014], Figueroa-López and Wu [2022].

Remark 2.2 *The estimation method introduced in this Section applies to continuous Itô semimartingale models. For what concerns semimartingale models allowing for jumps both in the price and the volatility, in the absence of microstructure noise, the Fourier methodology has been extended by Cuchiero and Teichmann [2015]. Their approach to estimating the spot volatility process (and the spot covariance process) in the presence of price jumps relies on a modification of classical jump-robust estimators of the integrated volatility, which are employed to estimate the Fourier coefficients of the volatility trajectory (which are integrated functions of the volatility, indeed). Specifically, their approach requires two steps. In the first step, several jump-robust estimators, such as the power-variation estimators considered by Barndorff-Nielsen et al. [2001] and the estimator based on the realized Laplace transform of the volatility introduced by Todorov and Tauchen [2012], can be used for obtaining the coefficients of the volatility. Then, in the second step, the path of the spot volatility is reconstructed using the Fourier-Fejér inversion formula and, according to the theory, it converges to* $\frac{\sigma^2(t-)+\sigma^2(t)}{2}$ *for any fixed time of discontinuity and* $\sigma^2(t)$ *for any continuity time. The obtained instantaneous volatility estimator is consistent and a CLT holds with a rate of convergence equal to* $(1-\nu)/2$*, where* ν *is the Hölder continuity (between two jumps) of the volatility path. This jump-robust Fourier estimator shares the main feature of the Fourier method but loses the robustness in the presence of noise, which is the main interest of the present paper, see also [Cuchiero and Teichmann, 2015, Remark 3.2]. Combining the two issues represents an interesting challenge to be addressed by future research.*

Remark 2.3 *The positivity of the Fourier estimator (1) has been proved in Malliavin and Mancino [2009], Remark 2.3, and follows from the positivity of the Fourier transform of the Fejér kernel together with the convolution formula (2). Furthermore, it is worth notic-*

ing that the definition of the Fourier spot volatility estimator does not require considering equidistant observations. However, for simplicity, in Section 3 we will refer to the latter sampling scheme to derive the asymptotic theory. The general case with irregular observations can be obtained as in Malliavin and Mancino [2009].

3 Asymptotic theory

3.1 Asymptotic Normality in the absence of microstructure noise

In Mancino and Recchioni [2015], the authors show that the estimator defined in (1) achieves a slightly sub-optimal rate of convergence, under the hypothesis that the volatility process is a random function with Hölder-continuous paths. In this Section, we fine-tune this result by proving that the estimator can reach the optimal rate of convergence, $n^{1/4}$, as it is for the one-sided uniform kernel estimator of [Aït-Sahalia and Jacod, 2014, Th. 8.8], and the more general kernel estimator of [Figueroa-López and Li, 2020, Remark 2.2], under the assumption that the volatility process is a Brownian semimartingale. Moreover, we show that the estimator is efficient in terms of smaller variance of the asymptotic distribution in the sense of Jacod and Rosenbaum [2013], see also Remark 3.3.

Theorem 3.1 *Let assumptions $(A.I)$, $(A.II)$ and $(A.III)$ hold. Moreover, assume that $\lim_{n,N\to\infty} N/n = c > 0$. Then, for any fixed $t \in (0,T)$, the following stable convergence in law holds, as $n, N, M \to \infty$:*

- if $\lim_{n,M\to\infty} Mn^{-1/\tau} = a > 0$, for $1 < \tau < 2$,

$$n^{1/2}M^{-1/2}\left(\widehat{\sigma}^2_{nNM}(t) - \sigma^2(t)\right) \to \mathcal{N}\left(0, \frac{4}{3}(1+2K(2c))\,\sigma^4(t)\right); \tag{5}$$

- if $\lim_{n,M\to\infty} Mn^{-1/2} = a > 0$,

$$n^{1/2}M^{-1/2}\left(\widehat{\sigma}^2_{nNM}(t) - \sigma^2(t)\right) \to \mathcal{N}\left(0, \frac{4}{3}(1+2K(2c))\sigma^4(t) + \frac{T}{3a^2}\,\gamma^2(t)\right), \tag{6}$$

with

$$K(c) := \frac{1}{2c^2} r(c)(1 - r(c)), \tag{7}$$

where $r(x) = x - [x]$ *and* $[x]$ *is the integer part of* x.

Remark 3.2 *(Convergence Rate)*

The asymptotic result in (5) is in line with the one obtained in Mancino and Recchioni [2015], where the rate is $n^{1/4}(\log n)^{-1/2}$*, which is optimal up to a logarithmic correction. The authors assume that* $N/n \sim c > 0$ *and* $Mn^{-1/\tau} \sim a > 0$*, for* $1 < \tau < 2$*. Note that in this case* $Mn^{-1/2} \to \infty$*. It is worth noting that the CLT in Mancino and Recchioni [2015] holds under the assumption that* σ^2 *is a.s. Hölder continuous in* $[0, T]$ *with parameter* $\nu \in (0, 1/2)$*. Therefore, rough volatility models are also included, at the cost of a slower convergence rate.*

Letting $\beta = \sqrt{T}/a$*, the asymptotic result in (5) corresponds to the case* $\beta = 0$ *in [Aït-Sahalia and Jacod, 2014, Th. 8.6], where the rate is sub-optimal, while the asymptotic result in (6) corresponds to the case* $\beta \in (0, \infty)$*, which instead attains the optimal rate* $n^{1/4}$*. An analogous result is obtained in [Figueroa-López and Wu, 2022, Th. 2.1], where alternative localization kernels are considered.*

Remark 3.3 *(Asymptotic Variance)*

The function $K(c)$ *in (7) is obtained in Clement and Gloter [2011] and arises from the behavior of the discretized Dirichlet kernel, in the case when* N *has the same order of the sample size* n *(see Lemma 2.4 in the supplementary material). Note that* $K(2c)$ *is nonnegative for any positive* c *and equal to zero if* $c = (1/2)k$*,* $k = 1, 2, \ldots$*. If* $k = 1$*, i.e.,* $c = 1/2$*, then* N *is equal to the Nyquist frequency, which represents the natural choice for* N*. Moreover, if* $N = n/2$*, one achieves the optimal asymptotic variance with regards to the term* $(4/3)\sigma^4(t)$*, which is smaller than its counterpart* $2\sigma^4(t)$ *in [Aït-Sahalia and Jacod, 2014, Th. 8.6]. This is due to the presence of the Fejér kernel in the Fourier-Fejér*

inversion formula, see also Remark 3.13 (iv) in Cuchiero and Teichmann [2015], where the variance reduction implied by the Fourier method is related to the James–Stein methodology of shrinkage to improve estimator variances. Similarly, letting $\beta = \sqrt{T}/a$*, even the second addend of the asymptotic variance in the efficient-rate case (6), that is,* $T/3a^2\gamma^2(t)$*, is smaller than the corresponding term* $T/a^2\gamma^2(t)$*.*

Based on the previous remarks, in the absence of microstructure noise, we conclude that, with the choice $N = n/2$, the Fourier estimator has the same rate of convergence and asymptotic variance of the Fejér-kernel based spot volatility estimator in (4), studied in Kristensen [2010] and Mancini et al. [2015], and the same rate and smaller variance with respect to the estimator by Aït-Sahalia and Jacod [2014]. Moreover, from Theorem 3.1 and Remark 3.3, the following result easily follows.

Corollary 3.4 *Let assumptions* $(A.I)$*,* $(A.II)$ *and* $(A.III)$ *hold. Moreover, assume that* $\lim_{n,N\to\infty} N/n = 1/2$ *and* $\lim_{n,M\to\infty} Mn^{-1/\tau} = a > 0$*, for* $1 < \tau < 2$*. Then, for any fixed* $t \in (0,T)$*, the following stable convergence in law holds, as* $n, N, M \to \infty$*:*

$$n^{1/2}M^{-1/2}\left(\widehat{\sigma}^2_{nNM}(t) - \sigma^2(t)\right) \to \mathcal{N}\left(0, \frac{4}{3}\,\sigma^4(t)\right).$$

A feasible version of the Theorem 3.1 requires a consistent estimator of the asymptotic variance, in particular of the spot quarticity and volatility of volatility. To this aim, we propose to exploit again the Fourier methodology. In fact, a consistent estimator of the integrated quarticity (namely, the 0-th Fourier coefficient of the fourth power of the volatility) is obtained in [Livieri et al., 2019, Th. 2], and reads as

$$c_0(\sigma^4_{nNS}) := \sum_{|k|\le S} c_k(\sigma^2_{nN})c_{-k}(\sigma^2_{nN}),$$

while a consistent estimator of the integrated volatility of volatility is obtained in [Toscano

et al., 2024, Th. 3.1], and reads as

$$c_0(\gamma^2_{nNS}) := \frac{T}{S+1}\sum_{|k|\leq S}\left(1-\frac{|k|}{S+1}\right)k^2 c_k(\sigma^2_{nN})c_{-k}(\sigma^2_{nN}) - \frac{1}{3}a^2\left(1+2K(2c)\right)c_0(\sigma^4_{nNS}). \tag{8}$$

These results, which consider the 0-th Fourier coefficient of the processes $\sigma^4(t)$ and $\gamma^2(t)$, can be extended to any k-th Fourier coefficient $c_k(\sigma^4_{nNS})$ and $c_k(\gamma^2_{nNS})$, respectively, as detailed in [Toscano et al., 2024, Remark 3.3]. Then, the uniform convergence in probability for $t \in (0,T)$ of the spot quantity estimators

$$\widehat{\sigma}^4_{nNSL}(t) := \sum_{|k|\leq L}\left(1-\frac{|k|}{L+1}\right)\mathrm{e}^{\mathrm{i}k\frac{2\pi}{T}t}c_k(\sigma^4_{nNS}) \tag{9}$$

and

$$\widehat{\gamma}^2_{nNSL}(t) := \sum_{|k|\leq L}\left(1-\frac{|k|}{L+1}\right)\mathrm{e}^{\mathrm{i}k\frac{2\pi}{T}t}c_k(\gamma^2_{nNS}) \tag{10}$$

is ensured by the Fourier-Fejér inversion formula.

The following feasible CLT with optimal rate follows.

Theorem 3.5 *Let assumptions (A.I), (A.II) and (A.III) hold. Moreover, assume that* $\lim_{n,N\to\infty} Nn^{-1} = c > 0$, $\lim_{n,M\to\infty} Mn^{-1/2} = a > 0$, $\lim_{n,S\to\infty} Sn^{-1/2} = \overline{a} > 0$ *and* $\lim_{L,S\to\infty} LS^{-1} = 0$. *Then, for any fixed* $t \in (0,T)$, *the following stable convergence in law holds, as* $n, N, M, S, L \to \infty$*:*

$$n^{1/2}M^{-1/2}\frac{\widehat{\sigma}^2_{nNM}(t) - \sigma^2(t)}{\sqrt{\frac{4}{3}(1+2K(2c))\,\widehat{\sigma}^4_{nNSL}(t) + \frac{T}{3a^2}\,\widehat{\gamma}^2_{nNSL}(t)}} \to \mathcal{N}(0,1). \tag{11}$$

3.2 Asymptotic Normality in the presence of microstructure noise

In this Section, we study the asymptotic normality with the rate of convergence $n^{1/8}$ for the estimator (1) when the price process is contaminated by microstructure noise. The following assumption on the noise process η is made.

(N.I) $(\eta_{t_i})_{i\geq 0}$ is a family of i.i.d. random variables, independent of the log-price process p, with zero mean and finite second and fourth moments. The noise variance is denoted as ξ.

Denote by $\widetilde{\sigma}^2_{nNM}(t)$ the Fourier spot volatility estimator (1) in the presence of noise, that is,

$$\widetilde{\sigma}^2_{nNM}(t) := \sum_{|k|\leq M} \left(1 - \frac{|k|}{M+1}\right) e^{\mathrm{i}k\frac{2\pi}{T}t} c_k(\widetilde{\sigma}^2_{nN}), \tag{12}$$

where $c_k(\widetilde{\sigma}^2_{nN})$ is the estimator of the k-th Fourier coefficient of the volatility process, obtained via eq. (2) using the discrete Fourier coefficients of the increments of the noisy log-price process $\widetilde{p}_{t_i} := p_{t_i} + \eta_{t_i}$, where we refer to p as the efficient log-price. Analogously to (3), the estimator (12) with noisy log-returns can be rewritten as

$$\frac{1}{T}\sum_{i=0}^{n-1}\sum_{j=0}^{n-1} F_M(t-t_j)D_N(t_j-t_i)\left(\delta_i(p)\delta_j(p) + \varepsilon_i\delta_j(p) + \delta_i(p)\varepsilon_j + \varepsilon_i\varepsilon_j\right), \tag{13}$$

with $\varepsilon_i := \eta_{t_{i+1}} - \eta_{t_i}$. The CLT in the presence of noise reads as follows.

Theorem 3.6 *Let assumptions (A.I), (A.II), (A.III) and (N.I) hold. Moreover, assume that* $\lim_{N,n\to\infty} Nn^{-1/2} = c > 0$. *Then, for any fixed* $t \in (0,T)$, *the following stable convergence in law holds, as* $n, N, M \to \infty$:

- *if* $\lim_{N,M\to\infty} M\,N^{-1/\tau} = a > 0$, *for* $1 < \tau < 2$,

$$n^{1/4}M^{-1/2}\left(\widetilde{\sigma}^2_{nNM}(t) - \sigma^2(t)\right) \to \mathcal{N}\left(0, \frac{1}{c}\frac{2}{3}\sigma^4(t) + c\frac{4T}{9}\sigma^2(t)\xi + c^3\frac{T^2}{15}\xi^2\right); \tag{14}$$

- *if* $\lim_{N,M\to\infty} M\,N^{-1/2} = a > 0$,

$$n^{1/4}M^{-1/2}\left(\widetilde{\sigma}^2_{nNM}(t) - \sigma^2(t)\right) \to \mathcal{N}\left(0, \frac{1}{c}\frac{2}{3}\sigma^4(t) + \frac{1}{a^2c}\frac{T}{3}\gamma^2(t) + c\frac{4T}{9}\sigma^2(t)\xi + c^3\frac{T^2}{15}\xi^2\right). \tag{15}$$

Remark 3.7 *(Convergence rate)*

In the presence of microstructure noise, the convergence rate attained by the Fourier estimator in (15) is the optimal one, that is, $n^{1/8}$. *This rate is also found for the kernel-type volatility estimators studied in [Aït-Sahalia and Jacod, 2014, Section 8.7] and [Figueroa-López and Wu, 2022, Theorem 2.2]. However, while data pre-averaging and a bias correction are needed for these estimators, the Fourier estimator in (1) does not need any*

modification. Note that in (14), which assumes $MN^{-1/\tau} \sim a$, for $1 < \tau < 2$, the estimator reaches a slightly sub-optimal rate; however, the resulting asymptotic variance is smaller and does not depend on the volatility of volatility.

Remark 3.8 *(Asymptotic Variance)*

The asymptotic variance in (15) depends on the spot quarticity, the spot volatility of volatility, and the variance of the noise. The same quantities with different constants appear in the asymptotic variance of the pre-averaging and bias-corrected kernel-type estimators by Aït-Sahalia and Jacod [2014] and Figueroa-López and Wu [2022]. The latter authors consider general kernels to determine which one minimizes the asymptotic variance, and find that, in the optimal rate case $n^{1/8}$, the optimal kernel is the two-sided exponential kernel. Finally, notice that the constant $K(2c)$ in (7), arising from the discretization of the Dirichlet kernel, disappears, as the condition $N/n \to 0$ is in force (see Lemma 2.5 in the supplementary material).

The role of the frequency N appears clearly by comparing Theorems 3.6 and 3.1. In Theorem 3.6, N is chosen smaller than in the noise-less case, in line with the findings of Mancino and Sanfelici [2008]. In fact, the role of N is that of filtering out the noise: the high-frequency, or short-run, noise is ignored by cutting the highest frequencies in the construction of the Fourier coefficients of the volatility $c_k(\sigma^2_{nN})$.

Finally, we address the issue of proving a feasible version of Theorem 3.6, in particular for the optimal-rate case (15). Again, this result is obtained through consistent estimators of the spot quarticity and of the volatility of volatility in the presence of microstructure noise. They are defined as in (9) and (10) but considering as building blocks the k-th Fourier coefficients $c_k(\widetilde{\sigma}^2_{nN})$, obtained via Eq. (2) from the Fourier coefficients of the increments of the noisy log-price process $\widetilde{p}$.

For the feasible CLT in the presence of noise, we specify the assumption on the noise

process as follows:

(N.II) $(\eta_{t_i})_{i\geq 0}$ is a family of i.i.d. random variables, independent of the log-price process p, centered and symmetric,[2] and finite moments of order h, with $h = 2, 4, 6, 8$.

The following Proposition proves the consistency of the Fourier estimator - defined in Eq. (16) - of the k-th Fourier coefficient of the volatility of volatility $c_k(\gamma^2)$.

Proposition 3.9 *Let assumptions (A.I), (A.II), (A.III) and (N.II) hold. Define*

$$c_k(\widetilde{\gamma}^2_{nNS}) := \frac{T}{S+1}\sum_{|h|\leq S}\left(1-\frac{|h|}{S+1}\right)h(h-k)c_h(\widetilde{\sigma}^2_{nN})c_{k-h}(\widetilde{\sigma}^2_{nN}). \tag{16}$$

Then, under the conditions $\lim_{N,n\to\infty} Nn^{-1/2} = c > 0$ *and* $\lim_{N,S\to\infty} S\,N^{-1/2} = \overline{c} > 0$*, the estimator (16) converges in probability to* $c_k(\gamma^2)$*.*

Note that the estimator of the k-th Fourier coefficient of the volatility of volatility in the presence of noise in Eq. (16) differs from the one in Eq. (8) in the absence of noise, because (16) does not require a correction term. The same happens in [Toscano et al., 2024, Th. 3.6], and is due to the fact that the frequency S appearing in the convolution (16) is chosen to be $O(n^{1/4})$ instead of $O(n^{1/2})$, as in (8). Then, for the Fourier-Fejér inversion formula, a consistent estimator of the volatility of volatility in the presence of noise is given, for any fixed $t \in (0, T)$, by

$$\widetilde{\gamma}^2_{nNSL}(t) := \sum_{|k|\leq L}\left(1-\frac{|k|}{L+1}\right)\mathrm{e}^{\mathrm{i}k\frac{2\pi}{T}t}c_k(\widetilde{\gamma}^2_{nNS}), \tag{17}$$

provided that the parameter L is chosen such that $\lim_{S,L\to\infty} LS^{-1} = 0$.

Analogously, we obtain a consistent estimator of the spot quarticity in the presence of noise as follows

$$\widetilde{\sigma}^4_{nNSL}(t) := \sum_{|k|\leq L}\left(1-\frac{|k|}{L+1}\right)\mathrm{e}^{\mathrm{i}k\frac{2\pi}{T}t}c_k(\widetilde{\sigma}^4_{nNS}), \tag{18}$$

where $c_k(\widetilde{\sigma}^4_{nNS}) := \sum_{|h|\leq S} c_{k-h}(\widetilde{\sigma}^2_{nN})c_h(\widetilde{\sigma}^2_{nN})$ is a consistent estimator of the k-th Fourier coefficient of the quarticity in the presence of noise.

[2]The symmetric assumption can be relaxed at the cost of burdensome computations.

Finally, according to Bandi and Russell [2006], a consistent estimator of the noise variance ξ is given by

$$\hat{\xi}_n := \frac{1}{2n} \sum_{i=0}^{n-1} \delta_i^2(\widetilde{p}). \tag{19}$$

We can now state the following feasible CLT.

Theorem 3.10 *Let assumptions (A.I), (A.II), (A.III) and (N.II) hold. Moreover, assume that* $\lim_{N,n\to\infty} Nn^{-1/2} = c > 0$, $\lim_{N,M\to\infty} M\,N^{-1/2} = a > 0$, $\lim_{N,S\to\infty} S\,N^{-1/2} = \overline{c} > 0$, *and* $\lim_{L,S\to\infty} LS^{-1} = 0$. *Then, for any fixed* $t \in (0,T)$, *the following stable convergence in law holds, as* $n, N, M, S, L \to \infty$*:*

$$n^{1/4}M^{-1/2}\left(\frac{\widetilde{\sigma}^2_{nNM}(t) - \sigma^2(t)}{\sqrt{\frac{2}{3c}\widetilde{\sigma}^4_{nNSL}(t) + \frac{T}{3a^2c}\widetilde{\gamma}^2_{nNSL}(t) + c\frac{4T}{9}\widetilde{\sigma}^2_{nNM}(t)\hat{\xi}_n + c^3\frac{T^2}{15}\hat{\xi}^2_n}}\right) \to \mathcal{N}(0,1). \tag{20}$$

A numerical study that provides support for the Normality of the standardized estimation error in (20) with finite samples is detailed in the supplementary material.

4 Finite-sample properties

In this Section, we investigate the finite-sample properties of the Fourier estimator of the spot volatility defined in (12) with simulated prices. Firstly, we study the (unfeasible) optimal selection of the cutting frequencies N and M. Secondly, we introduce a feasible adaptive method for selecting N and M, which optimizes the integrated asymptotic variance in the rate-efficient case, see (15) in Theorem 3.6, and measure its relative efficiency with respect to the unfeasible case. Finally, we compare the finite-sample performance of the Fourier estimator with those of alternative noise-robust spot volatility estimators.

4.1 Simulation design

We consider four simulated scenarios, characterized by progressively more realistic features of the efficient log-price and the microstructure noise. For each scenario, we simulate 10^4

daily trajectories of noisy log-price observations. As we use the trading day as the unit of time, the horizon T of the simulated trajectories is equal to 1. For simulating scenarios I, II and III, we use the equally-spaced sampling grid with mesh $\rho(n) = T/n$ of 1 second. As we assume the trading day to be 6.5-hour long, n equals 23400. For the simulation of scenarios IV and V, we employ the random sampling scheme detailed in Section 4.1.4.

4.1.1 Scenario I

In the first scenario, we simulate the efficient log-price using the model by Heston [1993], which satisfies (A.I), (A.II) and (A.III). The model reads as

$$dp(t) = \sigma(t)dW_t + \left(\mu - \frac{1}{2}\sigma^2(t)\right)dt, \qquad d\sigma^2(t) = \gamma\sigma(t)dZ_t + \theta\left(\alpha - \sigma^2(t)\right)dt, \tag{21}$$

where W and Z are Brownian motions with correlation ϕ and $\mu, p(0) \in \mathbb{R}, \gamma, \theta, \alpha, \sigma^2(0) > 0$. For the simulation, we set $(\mu, \theta, \alpha, \gamma, \phi) = (0.001, 1, 0.3, 0.2, -0.8)$. Furthermore, we set $p(0) = \ln(100)$, while $\sigma^2(0)$ is drawn, trajectory-wise, from the stationary distribution of σ^2, that is, from $\Gamma\left(\frac{2\theta\alpha}{\gamma^2}, \frac{\gamma^2}{2\theta}\right)$. For simulating the additive noise, for any i, we assume that η_{t_i} is independent of p and such that $\eta_{t_i} \sim_{i.i.d.} \mathcal{N}(0, \xi)$ and $\sqrt{\xi} = \zeta\sqrt{\alpha\rho(n)}$, where ζ denotes the noise-to-signal ratio. In the simulation, we adopt three increasing values of ζ, namely, $\zeta = 1, 2, 3$.

4.1.2 Scenario II

In the second scenario, we simulate the efficient log-price using a model with stochastic leverage and volatility of volatility. The model reads as

$$\begin{gathered} dp(t) = \sigma(t)dW_t + \left(\mu - \frac{1}{2}\sigma^2(t)\right)dt, \qquad d\sigma^2(t) = \gamma(t)dZ_t + \theta(\alpha - \sigma^2(t))dt, \\ d\gamma^2(t) = \beta\gamma(t)dY_t + \theta_\gamma(\alpha_\gamma - \gamma^2(t))dt, \end{gathered} \tag{22}$$

where Y is another Brownian motion and $\theta_\gamma, \alpha_\gamma, \beta, \gamma^2(0) > 0$. The correlations of the pairs (W, Z) and (Z, Y) satisfy the model by Veraart and Veraart [2012], that is,

$$dW_t = \phi_1(t)dZ_t + \sqrt{1-\phi_1^2(t)}dZ_t^{\perp}, \qquad d\phi_1(t) = \kappa_1\sqrt{1-\phi_1^2(t)}dB_t^{(1)} + \lambda_1(\nu_1 - \phi_1(t))dt, \quad (23)$$

$$dZ_t = \phi_2(t)dY_t + \sqrt{1-\phi_2^2(t)}dY_t^{\perp}, \qquad d\phi_2(t) = \kappa_2\sqrt{1-\phi_2^2(t)}dB_t^{(2)} + \lambda_2(\nu_2 - \phi_2(t))dt, \quad (24)$$

where: $Z^{\perp}$ and $Y^{\perp}$ are Brownian motions independent of, respectively, Z and Y; $B^{(1)}$ and $B^{(2)}$ are Brownian motions independent of each other and independent of W, Z and Y; for $i = 1, 2$, $k_i \in \mathbb{R}, \lambda_i > 0, \nu_i \in [-1, 1]$. For the simulation, we set $(\mu, \theta_1, \alpha_1, \beta, \theta_2, \alpha_2) = (0.001, 8, 0.3, 0.5, 10, 0.1)$ and $(\kappa_1, \lambda_1, \nu_1, \kappa_2, \lambda_2, \nu_2) = (0.1, 0.2, -0.6, 0.6, 1, 0.5)$. The initial conditions are set as $p(0) = \ln(100), \sigma^2(0) = \alpha_1, \gamma^2(0) = \alpha_2, \phi_1(0) = \nu_1$ and $\phi_2(0) = \nu_2$. The additive noise is simulated as in scenario I.

4.1.3 Scenario III

In the third scenario, we employ the model of scenario II for simulating the efficient log-price, while the additive noise is simulated using the model proposed by Jacod et al. [2017], which allows for auto-correlation and dependence with the efficient log-price. The latter assumes that $\eta_{t_i} = \psi_{t_i}\chi_{t_i}$ for any i, where ψ_{t_i} is sampled from $d\psi(t) = \theta_\psi(\alpha_\psi - \psi(t))dt + \sigma_\psi dW_t$, while χ_{t_i} is independent of W and such that $\chi_{t_i} = a\chi_{t_{i-1}} + e_{t_i}$, where $|a| < 1$ and e_{t_i} is an i.i.d. sequence of t-distributed random variables with $\nu > 4$ degrees of freedom. For the simulation, we set $(a, \nu) = (0.6, 5)$ and consider three alternatives for $(\alpha_\psi, \theta_\psi, \sigma_\psi, \psi(0))$, namely $\mathrm{v}_1 := (0.0001, 10, 0.01, 0.0001)$, $\mathrm{v}_2 := (0.0002, 10, 0.02, 0.0002)$ and $\mathrm{v}_3 := (0.0003, 10, 0.03, 0.0003)$. By imposing that $\sqrt{\bar{\xi}} = \zeta\sqrt{\alpha\rho(n)}$, where $\bar{\xi} = \frac{\nu}{\nu-2}\frac{1}{1-a^2} \times \left(\alpha_\psi^2 + \frac{\sigma_\psi^2}{2\theta_\psi}\right)$ denotes the steady-state variance of η, we note that $\mathrm{v}_1, \mathrm{v}_2$ and v_3 correspond to a noise-to-signal ratio ζ (approximately) equal to $1, 2$ and 3, respectively.

4.1.4 Scenario IV

The fourth scenario differs from the third only for the sampling scheme employed, which assumes that log-price observation times follow a non-homogeneous Poisson process. Specifically, following Li and Linton [2023], we set the rate ι as $\iota(t) = 23400\left(1 + \frac{1}{2}\cos\left(\frac{2\pi}{T}t\right)\right)$. Such specification mimics the clustering of trades at the beginning and the end of a trading day, which generates the typical intraday volatility U-shape.

4.2 Unfeasible optimal selection of N and M

In this Section, we study the unfeasible optimal selection of the cutting frequencies N and M based on the minimization of the mean integrated squared error (MISE). The latter is defined as $E\left[\frac{T}{m}\sum_{j=0}^{m-1}\left(\widetilde{\sigma}^2_{nNM}(t_j) - \sigma^2(t_j)\right)^2\right]$, where m is the number of equally-spaced estimates on the grid $t_j = j\frac{T}{m}$, $j \in \{0, ..., m-1\}$, and we approximate it with the average across all simulated trajectories.

For implementing the Fourier spot volatility estimator, we use the entire sample of available high-frequency prices on $[0, T]$, corresponding to $n = 23400$ in the case of regular sampling[3]. Based on the conditions of the rate-efficient CLT, see (15) in Theorem 3.6, we set the cutting frequencies as $N = \lfloor c\sqrt{n}\rfloor$ and $M = \lfloor a\sqrt{N}\rfloor$ and optimize the constants c and a. Following [Mancino and Recchioni, 2015, Section 4], we reconstruct the spot volatility on the equally-spaced grid with mesh equal to 1 minute, so that $m = 390$. Table 1 illustrates the values of the pair (c, a) that minimize the MISE and the resulting optimized MISE values.

Table 1 suggests that the optimal choice of c is smaller in correspondence with a larger noise-to-signal ratio. Additionally, it is worth noting that the optimal values of c are approximately halved when moving from an i.i.d. noise specification (scenarios I and II) to

[3]Numerical results included in the supplementary material suggest that the use of the sampling window $[a, b]$, $a > 0, b < T$, for the estimation of $\sigma^2(t)$, $t \in (a, b)$, leads to sub-optimal finite-sample efficiency.

scenario	ζ	optimal (c, a)	MISE	scenario	ζ	optimal (c, a)	MISE
I	1	$(8, 0.3)$	$1.1 \cdot 10^{-3}$	III	1	$(4, 0.4)$	$2.2 \cdot 10^{-3}$
	2	$(5, 0.3)$	$1.3 \cdot 10^{-3}$		2	$(2, 0.4)$	$3.3 \cdot 10^{-3}$
	3	$(3, 0.3)$	$1.9 \cdot 10^{-3}$		3	$(2, 0.4)$	$3.8 \cdot 10^{-3}$
II	1	$(8, 0.4)$	$1.5 \cdot 10^{-3}$	IV	1	$(4, 0.3)$	$2.5 \cdot 10^{-3}$
	2	$(5, 0.4)$	$2.2 \cdot 10^{-3}$		2	$(2, 0.3)$	$3.3 \cdot 10^{-3}$
	3	$(3, 0.4)$	$2.6 \cdot 10^{-3}$		3	$(1, 0.3)$	$4.0 \cdot 10^{-3}$

Table 1: MISE-optimal selection of (c, a) and resulting MISE values in the different scenarios and for the different values of the noise-to-signal ratio ζ. The constants c and a were optimized, respectively, on the equally-spaced grids $\{1, 2, ..., 10\}$ and $\{0.1, 0., 2, ..., 0.5\}$.

an auto-correlated, price-dependent noise specification (scenarios III and IV). Instead, the optimal choice of a appears to be noise- and scenario-insensitive, oscillating between 0.3 and 0.4. These findings align with the numerical results in [Mancino and Recchioni, 2015, Section 4].

4.3 Feasible method for selecting N and M

In this Section, we introduce a feasible method for selecting the pair (N, M) which optimizes the asymptotic variance in Theorem 3.6 in the rate-efficient case, see equation (15). More specifically, as in [Figueroa-López and Wu, 2022, Section 3], the method allows performing the selection of (N, M) globally on the interval $[0, T]$ by optimizing the integrated

asymptotic variance[4], that is,

$$\Psi(N,M) := \frac{M}{N}\frac{2}{3}\int_0^T \sigma^4(t)\,dt + \frac{1}{M}\frac{T}{3}\int_0^T \gamma^2(t)\,dt + \frac{MN}{n}\frac{4T}{9}\xi\int_0^T \sigma^2(t)\,dt + \frac{N^3M}{n^2}\frac{T^3}{15}\xi^2. \tag{25}$$

Based on the consistency results derived in Section 3.2 for proving the feasible Theorem 3.5, a consistent estimator of $\Psi(N,M)$ is provided by

$$\widehat{\Psi}(N,M) := \frac{M}{N}\frac{2}{3}Tc_0(\widetilde{\sigma}^4_{nNS}) + \frac{1}{M}\frac{T^2}{3}c_0(\widetilde{\gamma}^2_{nNS}) + \frac{MN}{n}\frac{4T^2}{9}\widehat{\xi}_n\,c_0(\widetilde{\sigma}^2_{nN}) + \frac{N^3M}{n^2}\frac{T^3}{15}\widehat{\xi}^2_n, \tag{26}$$

where $S = \lfloor \bar{c}\sqrt{N} \rfloor$[5].

For a given sample size n, the feasible selection of (N,M) amounts then to solving the constrained optimization problem

$$\min_{(N,M)\in R}\ \widehat{\Psi}(N,M),$$

where $R = \left[\underline{N},\overline{N}\right] \times \left[\underline{M},\overline{M}\right] \subset \mathbb{R}^+ \times \mathbb{R}^+$ is such that $M < N < n$. To do so, we propose to use the following adaptive method, based on the gradient-descent algorithm.

[4]This method could also be applied to optimize the asymptotic variance at any specific time $t \in (0,T)$. However, if one is interested in estimating the entire volatility path, simulations suggest that the computational cost is substantial, while the gain in accuracy is not. Similar circumstances are experienced for the selection of the bandwidth parameter in Figueroa-López and Wu [2022], see Remark 4.1 therein.

[5]Simulations suggest selecting a relatively small S, that is, $\bar{c}$, for accurately estimating the integrated quarticity and volatility of volatility. Accordingly, we choose $\bar{c} = 0.1$ in our study.

Adaptive method for the selection of the pair (N, M)

Given the initial condition $(N_0 = \underline{N}, M_0 = \underline{M})$, for $k = 1, ..., K$, follow the update rule[a]

$$N_k = \left(\underline{N} \vee N_{k-1} - \lambda \frac{\partial \widehat{\Psi}}{\partial N}(N_{k-1}, M_{k-1}) \right) \wedge \overline{N},$$

$$M_k = \left(\underline{M} \vee M_{k-1} - \lambda \frac{\partial \widehat{\Psi}}{\partial M}(N_{k-1}, M_{k-1}) \right) \wedge \overline{M},$$

where λ is a positive constant, typically referred to as the *learning rate* parameter. The optimal value (N^*, M^*) is achieved as soon as the marginal absolute change of the objective function falls below a given threshold ϑ, i.e., in correspondence of the smallest k such that

$$\delta_k := \frac{\left| \widehat{\Psi}(N_k, M_k) - \widehat{\Psi}(N_{k-1}, M_{k-1}) \right|}{\widehat{\Psi}(N_{k-1}, M_{k-1})} < \vartheta.$$

Instead, if $\delta_k > \vartheta \ \forall k$, make the final selection $(N^*, M^*) = (N_K, M_K)$.

[a]The notation $\frac{\partial \widehat{\Psi}}{\partial x}(u, v)$ is shorthand for $\frac{\partial \widehat{\Psi}(N,M)}{\partial x}\Big|_{N=u, M=v}$.

The implementation of the method requires the selection of tuning parameters, namely the constraint region R and the vector of optimization parameters (λ, ϑ, K). In this regard, for the largest n available, we set $R = \left[\lfloor 2^{-1} n^{1/4} \rfloor, \lfloor 10 n^{1/4} \rfloor\right] \times \left[\lfloor n^{1/2} \rfloor, \lfloor 20 n^{1/2} \rfloor\right]$ and $(\lambda, \vartheta, K) = (\propto \hat{\xi}^{-1}, 10^{-4}, 10^4)$. The choice of R is rather conservative, as it allows for a fairly large range of variation for N and M. The selection of ϑ and K is also conservative since it ensures that the stopping rule $\delta_k < \vartheta$ is always satisfied for some $k \leq K$ in the simulation study. In general, the choice of the learning rate λ is known to play a crucial role in determining the accuracy of the gradient-descent algorithm, see Ruder [2017] and the references therein. In this regard, simulations suggest that selecting a value for λ which is inversely proportional to the (estimated) variance of the noise may yield satisfactory estimation accuracy. Specifically, letting $\lambda = c_\lambda \, \hat{\xi}^{-1}$, simulations suggest that the MISE is optimized if c_λ is chosen in a neighborhood of $c_\lambda = 5$ in the case of scenarios I and II, or $c_\lambda = 1$ in the case of scenarios III and IV. As for robustness, it is worth noting that MISE

values appear to be rather stable for $c_\lambda \in [3, 7]$ (respectively, $c_\lambda \in [0.5, 1.5]$) in scenarios I and II (respectively, in scenarios III and IV).

The efficiency of the feasible adaptive method for selecting N and M is summarized in Table 2. Compared to the unfeasible results reported in Table 1, the feasible results illustrated in Table 2 appear satisfactory. Indeed, based on the values of Δ_{MISE}, which denotes the relative change in the MISE compared to the corresponding unfeasible MISE reported in Table 1, the loss of efficiency ranges between 5% and 9%. We note that the feasible adaptive method tends to produce values of c and a that are relatively close, on average, to the unfeasible optimal values in Table 1. Finally, note that the feasible method for selecting N and M provides a satisfactory performance also when the assumptions of Theorem 3.6 are partially violated, that is, in scenarios II, III and IV.

scenario	ζ	$(\bar{c}, \bar{a})$	MISE	Δ_{MISE}	scenario	ζ	$(\bar{c}, \bar{a})$	MISE	Δ_{MISE}
I	1	$(7.5, 0.3)$	$1.2 \cdot 10^{-3}$	+9%	III	1	$(4.1, 0.4)$	$2.4 \cdot 10^{-3}$	+9%
	2	$(4.4, 0.3)$	$1.4 \cdot 10^{-3}$	+8%		2	$(2.4, 0.3)$	$3.5 \cdot 10^{-3}$	+6%
	3	$(2.9, 0.4)$	$2.0 \cdot 10^{-3}$	+5%		3	$(2.1, 0.4)$	$4.1 \cdot 10^{-3}$	+8%
II	1	$(7.4, 0.4)$	$1.6 \cdot 10^{-3}$	+7%	IV	1	$(4.0, 0.3)$	$2.7 \cdot 10^{-3}$	+8%
	2	$(4.2, 0.4)$	$2.3 \cdot 10^{-3}$	+5%		2	$(2.3, 0.3)$	$3.5 \cdot 10^{-3}$	+6%
	3	$(2.8, 0.4)$	$2.8 \cdot 10^{-3}$	+8%		3	$(1.3, 0.4)$	$4.3 \cdot 10^{-3}$	+8%

Table 2: MISE values produced by the feasible adaptive method for selecting N and M, in the different scenarios and for the different values of the noise-to-signal ratio ζ considered. Letting $N = c\sqrt{n}$ and $M = a\sqrt{N}$, $\bar{c}$ and $\bar{a}$ denote the sample averages of, respectively, the values of c and a produced by the feasible method. Moreover, Δ_{MISE} denotes the relative change in the MISE compared to unfeasible values reported in Table 1.

4.4 Performance comparison

We compare the finite-sample performance of the Fourier spot volatility estimator with that of four alternative noise-robust spot estimators, namely the Smoothed Two-Scale estimator (STS) by Mykland et al. [2019], the Pre-Averaging Kernel estimator by Figueroa-López and Wu [2022] (PAK), the Pre-Averaging ReMeDI (PAR) estimator by Li and Linton [2023] and the Local Method of Moments (LMM) estimator by Bibinger et al. [2019]. We refer the reader to the cited papers for the estimators' definitions. The selection of the tuning parameters and auxiliary functions involved in the implementation thereof was performed using as guidance, respectively, [Chen et al., 2020, Section 6.3] (STS), [Figueroa-López and Wu, 2022, Section 3] (PAK), [Li and Linton, 2023, Section 5] (PAR) and [Bibinger et al., 2019, Section 3.4] (LMM).

The MISE values resulting from the application of the STS, PAK, PAR and LMM estimators on the 1-minute estimation grid are shown in Table 3. Based on the comparison of Table 3 with Table 2, the feasible application of the Fourier estimator yields the lowest MISE in every scenario and for every noise-to-signal ratio considered. These findings provide further support for the finite-sample robustness of the Fourier spot volatility estimator, suggesting that it may represent a competitive alternative to other spot volatility estimators not only in the first scenario with i.i.d. noise but also in scenarios where the assumptions of Theorem 3.6 are partially violated.

5 Empirical study

The accuracy of spot volatility estimates can be investigated empirically by testing the distribution of high-frequency standardized log-returns, see [Mancino and Recchioni, 2015, Section 4]. We briefly recall the testing procedure. Define the standardized log-return on $[t, t+h]$, $h > 0$, as $z_h(t) := r_h(t) \left(\sigma^2(t)\, h\right)^{-1/2}$, where $r_h(t) := p(t+h) - p(t)$. Moreover,

sc.	ζ	STS	PAK	PAR	LMM	sc.	ζ	STS	PAK	PAR	LMM
I	1	$5.1 \cdot 10^{-3}$	$1.8 \cdot 10^{-3}$	$3.0 \cdot 10^{-3}$	$1.8 \cdot 10^{-3}$	III	1	$8.7 \cdot 10^{-3}$	$4.7 \cdot 10^{-3}$	$5.3 \cdot 10^{-3}$	$4.2 \cdot 10^{-3}$
	2	$5.0 \cdot 10^{-3}$	$2.2 \cdot 10^{-3}$	$3.2 \cdot 10^{-3}$	$2.3 \cdot 10^{-3}$		2	$8.7 \cdot 10^{-3}$	$7.9 \cdot 10^{-3}$	$8.2 \cdot 10^{-3}$	$4.9 \cdot 10^{-3}$
	3	$4.9 \cdot 10^{-3}$	$2.4 \cdot 10^{-3}$	$4.4 \cdot 10^{-3}$	$3.1 \cdot 10^{-3}$		3	$9.5 \cdot 10^{-3}$	$8.1 \cdot 10^{-3}$	$8.8 \cdot 10^{-3}$	$5.5 \cdot 10^{-3}$
II	1	$7.5 \cdot 10^{-2}$	$3.5 \cdot 10^{-3}$	$4.4 \cdot 10^{-3}$	$2.2 \cdot 10^{-3}$	IV	1	$8.9 \cdot 10^{-3}$	$5.2 \cdot 10^{-3}$	$6.5 \cdot 10^{-3}$	$4.4 \cdot 10^{-3}$
	2	$8.2 \cdot 10^{-2}$	$4.1 \cdot 10^{-3}$	$4.7 \cdot 10^{-3}$	$2.8 \cdot 10^{-3}$		2	$9.0 \cdot 10^{-3}$	$8.7 \cdot 10^{-3}$	$8.5 \cdot 10^{-3}$	$6.2 \cdot 10^{-3}$
	3	$9.2 \cdot 10^{-2}$	$4.4 \cdot 10^{-3}$	$5.2 \cdot 10^{-3}$	$5.0 \cdot 10^{-3}$		3	$9.3 \cdot 10^{-2}$	$9.0 \cdot 10^{-3}$	$8.6 \cdot 10^{-3}$	$7.4 \cdot 10^{-3}$

Table 3: MISE values produced by the STS, PAK, PAR and LMM estimators in the different scenarios and for the different values of the noise-to-signal ratio ζ.

denote by $\hat{z}_h(t)$ the estimated standardized log-return, obtained by replacing $\sigma^2(t)$ with an estimate thereof. Under Assumption (A.I), if h is sufficiently small, the sequence $\{z_h(t)\}_{t\in\tau}$, $\tau = \{0, h, 2h, ..., (\lfloor T/h \rfloor - 1)h\}$, approximates a sequence of i.i.d. standard Normal random variables. Thus, by testing the empirical distribution of the estimated standardized log-returns $\{\hat{z}_h(t)\}_{t\in\tau}$ for Gaussianity, one can obtain insight into the accuracy of spot volatility estimates.

We empirically implement the Fourier spot volatility estimator in (12) using high-frequency prices of three US stocks, Intel (INTC), Comcast (CMCSA) and Starbucks (SBUX), sampled over the period December 1, 2022 - December 30, 2022, which corresponds to 21 trading days. During this period, the average trading volume of INTC, CMCSA and SBUX is approximately equal to, respectively, 14, 7 and 2 million shares, reflecting three different liquidity levels.

Spot volatility paths are estimated on the daily horizon, that is, on the interval from 9:30 a.m. to 4:00 p.m. We measure time in trading days, hence we set $T = 1$. The Fourier estimator is implemented using tick-by-tick log-returns. Hence, the sample size n employed changes from day to day, reflecting the number of ticks available on each day. For the given

n available, on each day we select N and M using the adaptive method described in Section 4.3. The rectangle S and the optimization parameters λ, ϑ, K are chosen as in Section 4.3. Specifically, we set $c_\lambda = 1$.

Price jumps are identified using the test by Lee and Mykland [2012] and removed from the sample employed for the estimation[6]. Overnight returns, that is, the difference between the opening log-price and the closing log-price of the previous day, are not employed in the estimation, as they often contain jumps due to the fact that market-related news is typically released when the market is closed.

To compute the sequence of the standardized log-returns, one needs to select a value for the interval h which is small enough to satisfactorily approximate their true distribution and, at the same time, large enough to make the impact of microstructure noise negligible. Based on a common practice adopted in the literature, we select h equal to 5 minutes. Accordingly, Fourier estimates of the spot volatility are obtained on the equally-spaced grid with a mesh size equal to 5 minutes. The estimated daily volatility paths of INTC, CMCSA and SBUX are displayed in Figure 1.

To evaluate the accuracy of estimates we perform the Jarque-Bera and Anderson-Darling tests at the 95% confidence level on the 21 daily sub-samples of estimated standardized log-returns. Table 4, which illustrates the null-hypothesis rejection rates and the average p-values, suggests that the daily empirical distributions of $\hat{z}_h(t)$ approximate a Gaussian distribution quite well for each of the three stocks considered, thereby suggesting that the Fourier spot volatility estimates obtained are accurate under the different liquidity

[6] As in [Lee and Mykland, 2012, Section 6], we apply the jump test at the 99% confidence level on seven intraday sub-intervals (namely, 9.30-10, 10-11, 11-12, 12-13, 13-14, 14-15 and 15-16) and, if the null hypothesis of absence of jumps is rejected on a given sub-interval, we replace the highest log-return in absolute value on that interval with a zero log-return. For each sub-interval the procedure is iterated until the test does not reject the null. Overall, we find evidence of the presence of jumps for 9 days (CMCSA), 9 days (INTC) and 12 days (SBUX).

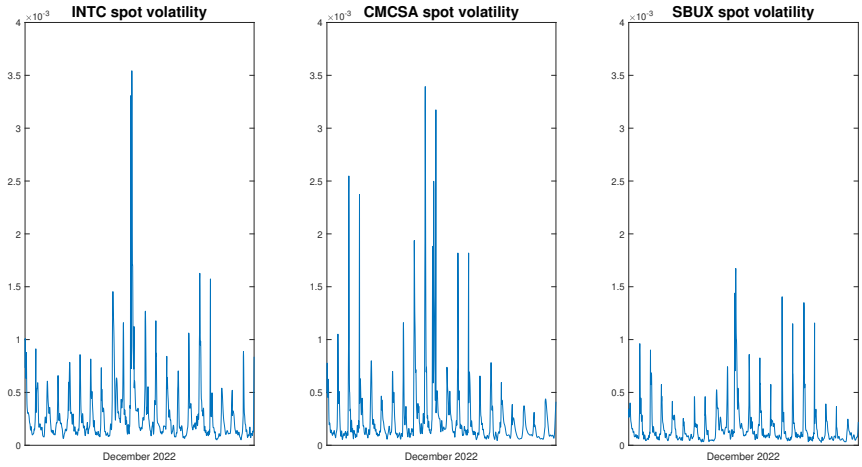

Figure 1: Sequences of the estimated daily volatility paths of CMCSA, INTC and SBUX in December 2022, on the 5-minute grid.

levels considered.

	JB test		AD test	
stock	rej. rate	av. p-value	rej. rate	av. p-value
INTC	2/21	0.525	1/21	0.589
CMCSA	1/21	0.515	1/21	0.496
SBUX	1/21	0.482	2/21	0.446

Table 4: Rejection rates and average p-values of the Jarque-Bera (JB) and Anderson-Darling (AD) tests, performed at the 95% confidence level on the 21 daily samples of 5-minute standardized returns covering the period December 1, 2022 - December 30, 2022.

6 Conclusions

The main contribution of this paper is providing a (pointwise) rate-optimal CLT for the Fourier spot volatility estimator in the presence of additive noise. Besides, we obtain

a rate-optimal CLT in the absence of noise, thereby completing the asymptotic theory in Mancino and Recchioni [2015]. Moreover, we prove the associated feasible CLTs by showing the consistency of the Fourier estimators of spot volatility of volatility and quarticity in the presence of noise. Finally, we propose a numerical feasible method for selecting the cutting frequencies defining the Fourier estimator in the presence of noise, based on the optimization of the asymptotic variance of the estimation error. The accuracy and robustness of this method are supported by an extended simulation study and an empirical exercise, which employs tick-by-tick prices of three US stocks with different liquidity levels.

As noticed in the Introduction, the Fourier estimation method is a global one, because it is designed to easily provide an estimation of the entire volatility function $t \mapsto \sigma^2(t)$ on the interval of interest $(0, T)$, using the Fourier-Fejér inversion theorem. For a continuous volatility function, the same theorem implies the uniform convergence in probability, that is, $\sup_{t\in(0,T)} \|\widehat{\sigma}^2_{nNM}(t) - \sigma^2(t)\| \xrightarrow{p} 0$, given the convergence in probability of the Fourier coefficients of the volatility function (see also Malliavin and Mancino [2002], Mancino and Recchioni [2015], Park et al. [2016]). For global estimators, besides the pointwise convergence rate studied in this paper, it is of interest to assert the quality of the estimation by using global criteria, as done in Fan and Wang [2008], Hoffmann [1999], Jacod et al. [2021]. The speed of convergence for the functional global inference of the Fourier estimator is an open problem, left for future research.

Acknowledgements

G. Toscano gratefully acknowledges financial support from the Institute Louis Bachelier 2022 grant 'Risk management in times of unprecedented geo-political volatility: a machine learning approach'. M.E. Mancino and G. Toscano were partially supported by INdAM-GNAMPA. The authors are also thankful to the anonymous Associate Editor and Referees

for their comments and suggestions, which contributed to improving the quality of the manuscript.

Conflicts of interests

The authors report there are no competing interests to declare.

SUPPLEMENTARY MATERIAL

Proofs and Auxiliary Results: contains the proofs of theorems, auxiliary lemmas on the Fejér and Dirichlet kernels and additional numerical results.

References

Y. Aït-Sahalia and J. Jacod. *High-frequency financial econometrics*. Princeton University Press, 2014.

T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys. Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625, 2003.

F. M. Bandi and J. R. Russell. Separating microstructure noise from volatility. *Journal of Financial Economics*, 79(3):655–692, 2006.

O. E. Barndorff-Nielsen, S. E. Graversen, J. Jacod, M. Podolskij, and N. Shephard. A central limit theorem for realised power and bipower variations of continuous semimartingales. In *From stochastic calculus to mathematical finance*, pages 33–68. 2001.

O. E. Barndorff-Nielsen, P. R. Hansen, A. Lunde, and N. Shephard. Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica*, 76(6):1481–1536, 2008.

E. Barucci and R. Renò. On measuring volatility and the garch forecasting performance. *Journal of International Financial Markets, Institutions and Money*, 12(3):183–200, 2002.

M. Bibinger, N. Hautsch, P. Malec, and M. Reiss. Estimating the spot covariation of asset prices—statistical theory and empirical evidence. *Journal of Business & Economic Statistics*, 37(3):419–435, 2019.

D. Chen, P. A. Mykland, and L. Zhang. The five trolls under the bridge: principal component analysis with asynchronous and noisy high-frequency data. *Journal of the American Statistical Association*, 115(532):1960–1977, 2020.

E. Clement and A. Gloter. Limit theorems in the Fourier transform method for the estimation of multivariate volatility. *Stochastic Processes and their Applications*, 121(5): 1097–1124, 2011.

C. Cuchiero and J. Teichmann. Fourier transform methods for pathwise covariance estimation in the presence of jumps. *Stochastic Processes and their Applications*, 125(1): 116–160, 2015.

J. Fan and Y. Wang. Spot volatility estimation for high-frequency data. *Statistics and its interface*, 1(2):279–288, 2008.

J. E. Figueroa-López and C. Li. Optimal kernel estimation of spot volatility of stochastic differential equations. *Stochastic Processes and their Applications*, 130(8):4693–4720, 2020.

J. E. Figueroa-López and B. Wu. Kernel estimation of spot volatility with microstructure noise using pre-averaging. *Econometric Theory*, 40(3):558–607, 2022.

D. Foster and D. Nelson. Continuous record asymptotics for rolling sample variance estimators. *Econometrica*, 64(1):139–174, 1996.

P. R. Hansen and A. Lunde. Realized variance and market microstructure noise. *Journal of Business & Economic Statistics*, 24(2):127–161, 2006.

S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343, 1993.

M. Hoffmann. L_p estimation of the diffusion coefficient. *Bernoulli*, 5(3):447–481, 1999.

W. Huang, C.-A. Lehalle, and M. Rosenbaum. Simulating and analyzing order book data: the queue-reactive model. *Journal of the American Statistical Association*, 110(509): 107–122, 2015.

J. Jacod. Limit of random measures associated with the increments of a Brownian semimartingale. *Tech.rep., Université de Paris VI*, 1994.

J. Jacod and M. Rosenbaum. Quarticity and other functionals of volatility: efficient estimation. The Annals of Statistics, 41(3):1462–1484, 2013.

J. Jacod, Y. Li, and X. Zheng. Statistical properties of microstructure noise. *Econometrica*, 85(4):1133–1174, 2017.

J. Jacod, J. Li, and Z. Liao. Volatility coupling. *The Annals of Statistics*, 49(4):1982 – 1998, 2021.

D. Kristensen. Nonparametric filtering of the realized spot volatility: a kernel-based approach. *Econometric Theory*, 26(1):60–93, 2010.

S. S. Lee and P. A. Mykland. Jumps in equilibrium prices and market microstructure noise. *Journal of Econometrics*, 168(2):396–406, 2012.

Z. M. Li and O. B. Linton. Robust estimation of integrated and spot volatility. *Journal of Econometrics*, 2023.

G. Livieri, M. E. Mancino, and S. Marmi. Asymptotic results for the Fourier estimator of the integrated quarticity. *Decisions in Economics and Finance*, 42:471–502, 2019.

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

P. Malliavin and M. E. Mancino. A Fourier transform method for nonparametric estimation of multivariate volatility. *The Annals of Statistics*, 37(4):1983–2010, 2009.

C. Mancini, V. Mattiussi, and R. Renò. Spot volatility estimation using delta sequences. *Finance and Stochastics*, 19:261–293, 2015.

M. E. Mancino and M. C. Recchioni. Fourier spot volatility estimator: asymptotic normality and efficiency with liquid and illiquid high-frequency data. *PLOS ONE*, 10(9), 2015.

M. E. Mancino and S. Sanfelici. Robustness of Fourier estimator of integrated volatility in the presence of microstructure noise. *Computational Statistics & Data Analysis*, 52(6): 2966–2989, 2008.

T. Mariotti, F. Lillo, and G. Toscano. From zero-intelligence to queue-reactive: limit-order-book modeling for high-frequency volatility estimation and optimal execution. *Quantitative Finance*, 23(3):367–388, 2023.

P. A. Mykland, L. Zhang, and D. Chen. The algebra of two scales estimation, and the S-TSRV: high-frequency estimation that is robust to sampling times. *Journal of Econometrics*, 208(1):101–119, 2019.

M. Ø. Nielsen and P. Frederiksen. Finite-sample accuracy and choice of sampling frequency in integrated volatility estimation. *Journal of Empirical Finance*, 15(2):265–286, 2008.

S. Park, S. Y. Hong, and O. Linton. Estimating the quadratic covariation matrix for asynchronously observed high-frequency stock returns corrupted by additive measurement error. *Journal of Econometrics*, 191(2):325–347, 2016.

S. Ruder. An overview of gradient descent optimization algorithms. *arXiv:1609.04747*, 2017.

V. Todorov. Non parametric spot volatility from options. *Annals of Applied Probability*, 29(6):3590–3636, 2019.

V. Todorov and G. Tauchen. The realized Laplace transform of volatility. *Econometrica*, 80(3):1105–1127, 2012.

V. Todorov and Y. Zhang. Bias reduction in spot volatility estimation from options. *Journal of Econometrics*, 234(1):53–81, 2023.

G. Toscano, G. Livieri, M. E. Mancino, and S. Marmi. Volatility of volatility estimation: central limit theorems for the Fourier transform estimator and empirical study of the daily time series stylized facts. *Journal of Financial Econometrics*, 22(1):252–296, 2024.

A. E. Veraart and L. A. Veraart. Stochastic volatility and stochastic leverage. *Annals of Finance*, 8:205–233, 2012.

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

Y. Zu and H. P. Boswijk. Estimating spot volatility with high-frequency financial data. *Journal of Econometrics*, 181(2):117–135, 2014.