7-2024

Reading the candlesticks: An OK estimator for volatility

Jia LI
Singapore Management University, jiali@smu.edu.sg

Dishen WANG

Qiushi ZHANG
University of International Business and Economics

Citation

LI, Jia; WANG, Dishen; and ZHANG, Qiushi. Reading the candlesticks: An OK estimator for volatility. (2024). *Review of Economics and Statistics*. 106, (4), 1114-1128.
Available at: https://ink.library.smu.edu.sg/soe_research/2565

Reading the Candlesticks: An OK Estimator for Volatility

Jia Li* Dishen Wang† Qiushi Zhang‡

April 21, 2021

Abstract

Academic research on nonparametric "spot" volatility inference often relies on high-quality transaction data that are not available to an average investor. Most investors, however, have free access to intraday candlestick charts through their online trading applications. Based on such data, we propose an *Optimal candlesticK (OK)* estimator for the spot volatility at a given time point. Under a standard infill asymptotic setting for Itô semimartingale price process, we show that the OK estimator is asymptotically unbiased and has minimal asymptotic variance within a class of linear estimators. In addition, its estimation error can be coupled by a Brownian functional, whose distribution is pivotal and known in finite-sample. Optimal confidence intervals can be constructed using the highest density interval of the (nonstandard) coupling distribution. Our theoretical and numerical results suggest that the proposed candlestick-based estimator is much more accurate than the conventional spot volatility estimator based on high-frequency returns. An empirical illustration is provided, which documents the intraday spot volatility dynamics of various assets during the Fed Chairman's recent congressional testimony.

Keywords: high-frequency data, nonparametric inference, semimartingale, volatility.
JEL Codes: C14, C22, C32.

*Department of Economics, Duke University, Durham, NC 27708; e-mail: jl410@duke.edu.

†Haas School of Business, University of California at Berkeley, Berkeley, CA, 94720; email: dishen_wang@mfe.berkeley.edu

‡School of Banking and Finance, University of International Business and Economics, Beijing, China; email: qiushizhang@uibe.edu.cn.

1 Introduction

Realized volatility measures computed using high-frequency asset returns have been extensively studied and applied in econometrics and statistics. Following the pioneer work of Andersen and Bollerslev (1998), Barndorff-Nielsen and Shephard (2002), and Andersen, Bollerslev, Diebold, and Labys (2003), a large literature in econometrics has emerged to study integrated volatility over a fixed period of time, typically a trading day. As high-quality asset price data become increasingly more available at higher frequencies, researchers have also rekindled their interest on the nonparametric estimation and inference for the spot volatility at a given time point (see Foster and Nelson (1996) and Comte and Renault (1998) for early contributions), which is particularly useful for studying news-induced volatility shocks such as those triggered by macroeconomic news announcements (Bollerslev, Li, and Xue (2018), Nakamura and Steinsson (2018)).[1] The classical spot volatility estimator is constructed as a localized version of the realized variance, but conducting reliable spot inference is more challenging, fundamentally because of the scarcity of local information, which may lead to finite-sample distortions in conventional asymptotic approximations. Hence, the spot inference naturally demands "better" high-frequency data, typically in the form of asset price returns sampled at higher frequencies.[2]

Despite the aforementioned impressive methodological advance, it might have been difficult for an average investor to benefit directly from this academic development. We recognize the lack of data as a major obstacle. Indeed, published papers in this area often rely on intraday high-frequency returns data obtained from commercial databases (e.g., the TAQ database) through institutional subscriptions, which is simply beyond the reach of an average investor. We also note that the majority of academic research has focused on the U.S. market, benefiting from its highly developed data service infrastructure. It is conceivable that the data-related obstacle may be more severe in less developed financial markets outside the U.S.[3]

Set against this background, our goal in this paper is to develop an easy-to-implement inference method for spot volatility based on high-frequency data that are readily accessible to retail investors. We observe that most, if not all, online trading applications display "candlestick" charts for asset prices, with each candlestick containing the open, close, high, and low prices within a

[1] Also see Kristensen (2010), Chapter 13.3 of Jacod and Protter (2012), and Bollerslev, Li, and Liao (2021) for some recent innovations and additional references.

[2] Price data sampled at "ultra" high frequency (e.g., tick-by-tick data) are often deemed to be contaminated by microstructure noise. To mitigate the effect of noise, researchers often sparsely sample the data at a much lower frequency, typically on the level of a few minutes. For more comprehensive discussions, see, for example, Zhang, Mykland, and Aït-Sahalia (2005), Bandi and Russell (2008), and the recent paper by Li and Linton (2020).

[3] A useful publicly available source of information on high-frequency volatility measures is the Realized Library of Oxford-Man Institute of Quantitative Finance, which offers the general public daily updates of realized volatility measures for 31 market indices.

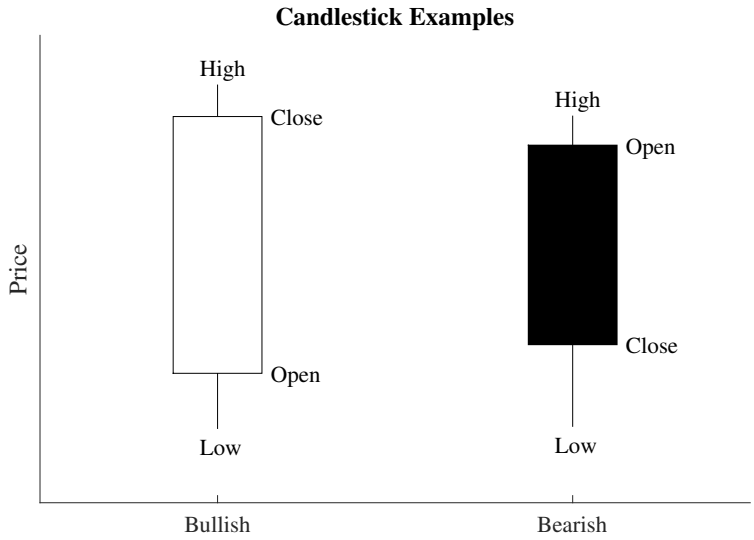

Figure 1: The figure illustrates the generic form of candlesticks. In the bullish (resp. bearish) case, the close price is higher (resp. lower) than the open price during the trading session.

short trading session.[4] For instance, the Robinhood mobile app, which is a popular trading application among retail investors, updates the candlesticks for stocks and ETFs listed in the U.S. market every 10 minutes. As a concrete illustration, we plot two candlesticks (bullish and bearish) in Figure 1. Each candlestick contains a vertical rectangle box (i.e., real-body) determined by the open and close prices during the trading period (say, 10 minutes), with its color signifying the direction of the price movement; the upper and lower ends of the candlestick indicate the highest and lowest prices within that period, respectively.

We propose an estimator for the spot volatility at a given time point based on the associated candlestick, and study its asymptotic property under a standard Itô semimartingale price model (which features stochastic drift and volatility, intraday seasonality, leverage effect, and price and volatility jumps) using infill asymptotics.[5] In its basic form, the proposed estimator is constructed as a linear combination of the high-low range and the absolute open-close return, which is more precisely given by

$$\frac{0.811 \times (\text{High} - \text{Low}) - 0.369 \times |\text{Close} - \text{Open}|}{(\text{Duration of the Trading Session})^{1/2}}.^{6} \tag{1.1}$$

Here, the specific numerical weights are chosen to achieve a type of optimality, that is, the resulting estimator is asymptotically unbiased and attains the minimal asymptotic variance within a class of linear estimators. We refer to this estimator as the *Optimal candlesticK (OK)* estimator.

[4]The candlestick chart has been commonly used in asset markets over the world for centuries; see Nison (2001).

[5]See Chapter 2 of Jacod and Protter (2012) for a comprehensive discussion on the Itô semimartingale model.

[6]The volatility estimate computed using the dollar-denominated prices is also in dollar unit, which may be transformed into relative terms by using the open price as a normalization. Almost equivalently, one may compute the volatility estimate in percentage terms using log-transformed prices.

Since the OK estimator is based on a single candlestick observation, one evidently cannot use the conventional law of large numbers or central limit theorem to derive its asymptotic property. Our asymptotic analysis instead relies on an "approximate finite-sample" approach in the same spirit as Bollerslev, Li, and Liao (2021). The underlying idea is simple: Within each "short" trading session, the Itô semimartingale price process can be approximated by a Brownian motion scaled by a constant volatility. Guided by this intuition, we show that the estimation error of the OK estimator may be strongly approximated, or "coupled," by a certain Brownian functional, whose (nonstandard) distribution is known in finite-sample. This result can then be used to construct confidence intervals (CI) for the spot volatility. In fact, we can optimize the CI (in terms of minimal length) by using the highest density interval (HDI) of the distribution of the coupling variable as the critical value. Our theory suggests that the resulting candlestick-based CI is much more efficient than the conventional CI based only on high-frequency returns. Our numerical findings from both simulated and real data also support this theoretical claim.

We stress that the OK estimator and the associated optimal CI are very easy to compute and, importantly, they only rely on readily accessible data. Our proposal thus offers the average investor an "affordable" option, both in terms of computational feasibility and data availability, for conducting reliable inference on spot volatility in real time. This exactly fulfills our main goal mentioned above. To guide practical applications, we provide an empirical illustration in Section 4 and show that the OK estimator can indeed provide economically sensible spot volatility estimates with adequate precision for a variety of different assets (bond, equity, gold, and cryptocurrency). We also note that, although we intentionally "tilt" our main discussion towards practicality, the underlying econometric idea can be extended in many directions for more complicated econometric settings. For example, when candlesticks are available at higher frequencies, we may aggregate them within a short window to further improve the estimation efficiency. Moreover, following the influential work of Barndorff-Nielsen and Shephard (2004), we may also consider a bipower version of the OK estimator by simultaneously using adjacent candlesticks, which may make the estimator more robust with respect to price jumps. Additional ideas and more detailed discussions are provided in Section 2.3.

Finally, we clarify the relation between the present paper and prior work in the literature. Our work is inspired by the classical papers of Parkinson (1980) and Garman and Klass (1980). Based on the probabilistic result of Feller (1951), Parkinson (1980) develops an estimator for asset return variance by aggregating a large number of observations of the high-low price range, and makes the important observation regarding the efficiency gain from using range data. Modeling the price process as a scaled Brownian motion, Garman and Klass (1980) construct a more efficient unbiased estimator for variance by exploiting the information of the entire candlestick. In contrast to this

prior work, we directly focus on the estimation of volatility (rather than variance), and conduct an infill asymptotic analysis under the far more general Itô semimartingale model. More importantly, we develop valid optimal CIs for the spot volatility whereas Parkinson (1980) and Garman and Klass (1980) only focus on estimation.

From a technical point of view, our paper is more closely related to the high-frequency econometrics literature on volatility inference; see the many papers cited above, and Andersen and Bollerslev (2018) and Jacod and Protter (2012) for a more complete list of references.[7] We highlight an important distinction between our method and the existing ones. In essentially all prior studies, the integrated variance and spot volatility are estimated by aggregating a "large" number of return or range observations. The consistency of the estimator can be claimed by invoking a law of large numbers, and the asymptotic Gaussian-based inference can be justified using a central limit theorem. In sharp contrast, we consider estimators formed using a *fixed* number of candlesticks, which better mimics the finite-sample environment for spot inference. Unlike conventional work, we do not pursue a consistency claim, but instead focus on the estimator's "approximate finite-sample" properties. Specifically, we show that the estimator is asymptotically unbiased with a well-defined sense of optimality, and its asymptotic distribution is captured by a nonstandard coupling variable, resulting in CIs that are quite distinct from conventional asymptotic Gaussian-based CIs. Our coupling-based approach is in the same spirit as Bollerslev, Li, and Liao (2021). However, that prior work only focuses on estimators formed using high-frequency returns. By exploiting the richer information from candlesticks in an optimal way, our estimator and CI substantially outperform those proposed by Bollerslev, Li, and Liao (2021). This efficiency improvement is crucial for our practical goal, that is, to provide an average investor with a reliable inference method in a data-scarce environment. In our empirical illustration, we show that in the benchmark scenario with 10-minute data, spot volatility estimates based only on returns are simply too noisy to be economically meaningful, whereas the OK volatility estimates are much more accurate and interpretable with respect to the underlying news flow.

The remainder of the paper is organized as follows. Section 2 presents our inference procedure and the corresponding asymptotic theory. Section 3 reports the finite-sample performance of the proposed method in a Monte Carlo experiment. An empirical illustration is provided in Section 4. Section 5 concludes. The Appendix collects all proofs.

[7] A paper worth highlighting is Christensen and Podolskij (2007), which establishes the infill asymptotic property of Parkinson's estimator as an estimator of integrated variance.

2 Candlestick-based inference for spot volatility

Section 2.1 introduces the theoretical setting and briefly reviews some existing methods for spot volatility inference. Section 2.2 describes a basic version of our candlestick-based inference procedure and establishes its theoretical validity. Further extensions are discussed in Section 2.3. In Section 2.4, we consider an efficiency comparison between the proposed method and commonly used benchmarks in a thought experiment. Below, all limits are for $n \to \infty$.

2.1 The setting and background

Suppose that the (log) price process P is an Itô semimartingale defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ written as

$$P_t = P_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + J_t, \tag{2.1}$$

where b is the drift process, σ is the stochastic volatility process, W is a standard Brownian motion, and J is a pure-jump process driven by a Poisson random measure. Our econometric interest is on the estimation and inference for the *spot volatility* σ_t at a given time point t. The method can be trivially extended to the joint inference for a finite collection of time points.

The classical nonparametric estimator for spot volatility, which is first proposed by Foster and Nelson (1996), can be constructed based on a localized version of the realized variance (Barndorff-Nielsen and Shephard (2002), Andersen, Bollerslev, Diebold, and Labys (2003)). Under the classical theory, we assume that the price process P is observed on a high-frequency time grid $0, \Delta_n, 2\Delta_n \ldots$ over a fixed time interval $[0, T]$, where the sampling interval $\Delta_n \to 0$ asymptotically. Let $r_i \equiv P_{i\Delta_n} - P_{(i-1)\Delta_n}$ denote the ith return. To conduct spot estimation, the user may choose a bandwidth sequence k_n and divide the returns into non-overlapping blocks, with the ith block containing $\left\{r_{(i-1)k_n+j} : 1 \leq j \leq k_n\right\}$. For any $t \in [(i-1)\, k_n\Delta_n, ik_n\Delta_n]$, the estimator for the spot variance σ_t^2 is given by

$$\hat{v}_t\left(k_n\right) \equiv \frac{1}{k_n\Delta_n} \sum_{j=1}^{k_n} r_{(i-1)k_n+j}^2. \tag{2.2}$$

Under mild regularity conditions, the conventional theory suggests that $\hat{v}_t\left(k_n\right)$ is a consistent estimator of σ_t^2, provided that the bandwidth sequence satisfies $k_n \to \infty$ and $k_n\Delta_n \to 0$.[8] The two

[8] It is useful to note that the spot variance estimator at any fixed time is robust to the presence of Poisson-type price jumps, even without using the multipower (Barndorff-Nielsen and Shephard (2004) and Barndorff-Nielsen, Shephard, and Winkel (2006)) or truncation (Mancini (2001)) techniques. These more advanced techniques are needed to achieve jump-robustness in the analysis of integrated variance over a non-degenerate time span. In contrast, the spot estimation concerns a shrinking estimation window, during which jumps occur with asymptotically negligible probability. This phenomenon is well understood in the literature on spot estimation; see, for example, Theorem 13.3.3 in Jacod and Protter (2012).

conditions on the tuning sequence play distinct roles in the asymptotic theory, which are intuitively easy to understand: The $k_n \to \infty$ condition permits the use of a law of large numbers to establish consistency, and the $k_n\Delta_n \to 0$ condition ensures that the $[(i-1)\,k_n\Delta_n, ik_n\Delta_n]$ estimation window "collapses" to the time point t, so as to give the local average estimator a nonparametric "spot" interpretation. Under additional restrictions, one may push the asymptotic analysis one step further to get the following feasible central limit theorem (see Theorem 2 in Foster and Nelson (1996) for an early contribution and Theorem 13.3.3 in Jacod and Protter (2012) for a more general result):

$$\frac{k_n^{1/2}\left(\hat{v}_t\left(k_n\right)-\sigma_t^2\right)}{\sqrt{2}\hat{v}_t\left(k_n\right)} \xrightarrow{d} \mathcal{N}\left(0,1\right). \tag{2.3}$$

By the delta method, an analogous result for the spot volatility σ_t also holds in the form of

$$\frac{k_n^{1/2}\left(\sqrt{\hat{v}_t\left(k_n\right)}-\sigma_t\right)}{\sqrt{\hat{v}_t\left(k_n\right)/2}} \xrightarrow{d} \mathcal{N}\left(0,1\right).$$

Consequently, for $\alpha \in (0,1)$, a two-sided $1-\alpha$ level confidence interval (CI) for σ_t may be constructed as

$$CI_{1-\alpha}^{G} \equiv \left[B_{\alpha-}^{G}\sqrt{\hat{v}_t\left(k_n\right)}, B_{\alpha+}^{G}\sqrt{\hat{v}_t\left(k_n\right)}\right], \tag{2.4}$$

where $B_{\alpha\pm}^{G} = 1 \pm z_{1-\alpha/2}/\sqrt{2k_n}$ and $z_{1-\alpha/2}$ is the $1-\alpha/2$ quantile of the $\mathcal{N}\left(0,1\right)$ distribution. Below, we refer to $CI_{1-\alpha}^{G}$ as the *Gaussian CI*.

Since the Gaussian CI is based on an asymptotic normal approximation (which requires $k_n \to \infty$), it may suffer from nontrivial size distortion when k_n is relatively small (e.g., $k_n \leq 10$). Meanwhile, it would be unwise to naively use a large k_n, because that will increase the bias stemming from the difference between the local average of the spot variance process and its spot value. Choosing the "proper" k_n to achieve good finite-sample coverage is a difficult practical matter. To the best of our knowledge, the existing literature does not offer a satisfactory answer to this question.

In a recent paper, Bollerslev, Li, and Liao (2021) propose an alternative "fixed-k" inference method for spot volatility. They treat the window size k_n as a fixed constant k in the asymptotic analysis, and suggest using a relatively small value of k in practice. When k is small, the bias becomes less of an issue. However, without the power of the law of large numbers, the spot estimator $\hat{v}_t\left(k\right)$ is no longer consistent and the standard asymptotic Gaussian-based inference cannot be justified using a central limit theorem. Nevertheless, Bollerslev, Li, and Liao (2021) show that feasible inference for σ_t can still be conducted based on a coupling result:

$$\frac{\sqrt{\hat{v}_t\left(k\right)}}{\sigma_t} = \sqrt{\frac{1}{k}\sum_{j=1}^{k}\left(\frac{W_{(i-1)k+j}-W_{(i-1)k+j-1}}{\Delta_n^{1/2}}\right)^2} + o_p(1). \tag{2.5}$$

Importantly, the leading term on the right-hand side of (2.5), which strongly approximates the multiplicative estimation error $\sqrt{\hat{v}_t(k)}/\sigma_t$, has the same finite-sample distribution as the square-root of a scaled chi-squared random variable with degree of freedom k. One can then exploit this distributional knowledge to construct CIs for σ_t as follows. Let $\bar{\chi}^2_k$ denote a generic copy of the scaled chi-squared variable with degree of freedom k. For $\alpha \in (0,1)$, let $B^F_{\alpha-}$ and $B^F_{\alpha+}$ be two constants such that $\mathbb{P}(B^F_{\alpha-} \leq (\bar{\chi}^2_k)^{-1/2} \leq B^F_{\alpha+}) = 1-\alpha$. The corresponding *fixed-k CI* for σ_t at confidence level $1-\alpha$ may be constructed as

$$CI^F_{1-\alpha} \equiv \left[B^F_{\alpha-}\sqrt{\hat{v}_t(k)}, B^F_{\alpha+}\sqrt{\hat{v}_t(k)}\right]. \tag{2.6}$$

This CI is not only asymptotically valid, but also finite-sample exact in the "limiting" model with the price process P being a scaled Brownian motion. In their simulation study, Bollerslev, Li, and Liao (2021) show that the fixed-k CI controls size more reliably than the Gaussian CI, and the former attains close-to-exact finite-sample coverage when k is small.

Although the Gaussian and fixed-k inference methods, especially the former, are well known in the high-frequency econometrics literature, we recognize that in practice they may not be easily applied by an average investor who is computationally constrained and/or does not have the access to the requisite high-frequency returns data.[9] Our first goal in this paper is to offer such investor an (extremely) easy-to-implement inference method based on readily accessible data. We observe that most, if not all, online trading applications update candlestick charts of asset prices at the 10-minute or higher frequencies in real time. To fix ideas, we denote the time interval associated with the ith candlestick by $I_{n,i} \equiv [(i-1)\Delta_n, i\Delta_n]$, with Δ_n being its duration (e.g., 10 minutes). The open, close, high, and low prices during this trading period can be expressed as, respectively,

$$O_i \equiv P_{(i-1)\Delta_n}, \quad C_i \equiv P_{i\Delta_n}, \quad H_i \equiv \sup_{t\in I_{n,i}} P_t, \quad L_i \equiv \inf_{t\in I_{n,i}} P_t.$$

The open-close return and the high-low range are then denoted by

$$r_i \equiv C_i - O_i, \quad w_i \equiv H_i - L_i.$$

We start by considering a class of linear estimators for σ_t that take the following form:

$$\hat{\sigma}_t(\lambda) \equiv \frac{\lambda_1 |r_i| + \lambda_2 w_i}{\Delta_n^{1/2}}, \quad t \in I_{n,i}, \tag{2.7}$$

[9]Published papers often rely on the TAQ database, which can only be obtained by researchers in relatively resourceful institutions. Even the TAQ database has a very limited scope, as it mainly covers the U.S. equity market. To obtain the other types of high-frequency data (e.g., bond, commodity futures, exchange rates, options, and international financial assets), one often has to resort to much more costly commercial databases such as Tick Data. Computationally more sophisticated investors may use application programming interface to automate data acquisition, but these investors are excluded from our definition of the "average" investor.

where the constant $\lambda = (\lambda_1, \lambda_2)$ will be chosen according to certain optimality requirement as detailed in Section 2.2, below. We then construct asymptotically valid CIs for σ_t using the optimal estimator. At first glance, it might be surprising that valid feasible inference can be carried out in the present setting, because the $\hat{\sigma}_t(\lambda)$ estimator uses only "one observation," precluding the traditional law of large numbers and central limit theorems from playing any role. Our inference is indeed justified using a different approach, which is in the same spirit as the coupling approach first proposed by Bollerslev, Li, and Liao (2021). In fact, with $\lambda_1 = 1$ and $\lambda_2 = 0$, $\hat{\sigma}_t(\lambda) = \Delta_n^{-1/2}|r_i|$ coincides with $\sqrt{\hat{v}_t(1)}$, and Bollerslev et al.'s fixed-k method can be directly used to construct CIs. In its more general form (2.7), the $\hat{\sigma}_t(\lambda)$ estimator also exploits the information from the high-low range w_i, and the estimator's coupling variable will take a more complicated form as a Brownian functional with a nonstandard distribution as we will show below.

Our estimator is clearly inspired by the classical estimators of Parkinson (1980) and Garman and Klass (1980). When the log price is modeled as a Brownian motion scaled by constant volatility, Parkinson's (1980) range-based estimator is unbiased for the variance. Under the same setting, Garman and Klass (1980) propose an improved unbiased estimator for the variance by optimally exploiting information from the candlestick. Garman and Klass also propose a "more practical" version of their optimal estimator which is popular among traders and is given by

$$\hat{v}_t^{GK} \equiv \frac{0.5w_i^2 - (2\log 2 - 1)r_i^2}{\Delta_n}. \tag{2.8}$$

Unlike this prior work, we focus on the (asymptotically) unbiased estimation of volatility rather than variance, which may be more directly useful in certain applications.[10] We also note two important theoretical differences between our analysis and those two papers. Firstly, we analyze our estimator under the Itô semimartingale model (2.1) under an infill asymptotic setting, which is more general than the baseline Brownian motion model. Secondly and more importantly, we construct asymptotically valid (and optimal) CIs for the spot volatility, whereas the aforementioned work only focuses on estimation. Our inference is formally justified by the aforementioned approximate finite-sample approach that involves a nonstandard limiting distribution. We now turn to the details.

2.2 Inference on spot volatility via candlesticks

In this subsection, we establish the asymptotic property of the $\hat{\sigma}_t(\lambda)$ estimator defined in (2.7) and then propose CIs for the spot volatility. Readers who are mainly interested in applications may skip the theoretical discussion and jump directly to the last paragraph of this subsection for practical guidance. We start by introducing some regularity conditions.

[10] By Jensen's inequality, the square-root of an unbiased estimator of the variance is guaranteed to underestimate the volatility on average.

Assumption 1. *Suppose that the price process P has the form (2.1) and there exists a sequence $(T_m)_{m\geq 1}$ of stopping times increasing to infinity and a sequence $(K_m)_{m\geq 1}$ of constants such that the following conditions hold for each $m \geq 1$: (i) for all $t \in [0, T_m]$, $|b_t|+|\sigma_t|+|\sigma_t|^{-1}+F_t(\mathbb{R}\backslash\{0\}) \leq K_m$, where F_t denotes the spot Lévy measure of J; (ii) for some constant $\kappa > 0$, $\mathbb{E}[|\sigma_{t\wedge T_m} - \sigma_{s\wedge T_m}|^2] \leq K_m|t-s|^{2\kappa}$ for all $t, s \in [0, T]$.*

Assumption 1 entails relatively mild regularity conditions that allow for leverage effect, intraday periodicity, and price and volatility jumps. Condition (i), in particular, imposes local boundedness on various processes, while condition (ii) states that the volatility process is locally κ-Hölder continuous under the L_2 norm. Note that the Hölder index κ is allowed to be arbitrarily small, and we do not need to know its value for the conduct of inference. When $\kappa = 1/2$, condition (ii) can be readily verified if σ is an Itô semimartingale or a long-memory process driven by a fractional Brownian motion (see, e.g., Comte and Renault (1998)). In particular, this condition accommodates essentially unrestricted volatility jumps driven by a (possibly compensated) Poisson random measure. Condition (ii) even allows the volatility to have "rough" paths, corresponding to $\kappa \in (0, 1/2)$.

Theorem 1, below, establishes a coupling result for the spot volatility estimator $\hat{\sigma}_t(\lambda)$.

Theorem 1. *Suppose that Assumption 1 holds. Then, for each $i \geq 1$ and $t \in I_{n,i}$,*

$$\frac{\hat{\sigma}_t(\lambda)}{\sigma_t} = \lambda_1\zeta_{1,i} + \lambda_2\zeta_{2,i} + o_p(1), \tag{2.9}$$

where $\zeta_{1,i} \equiv \Delta_n^{-1/2}\left|W_{i\Delta_n} - W_{(i-1)\Delta_n}\right|$ and $\zeta_{2,i} \equiv \Delta_n^{-1/2}\sup_{s,t\in I_{n,i}}|W_s - W_t|$.

Theorem 1 shows that the multiplicative estimation error $\hat{\sigma}_t(\lambda)/\sigma_t$ can be strongly approximated, or "coupled," by $\lambda_1\zeta_{1,i}+\lambda_2\zeta_{2,i}$. Since the coupling variable is non-degenerate for any $\lambda \neq 0$, $\hat{\sigma}_t(\lambda)$ is not a consistent estimator for σ_t, which is hardly surprising in the present "small-sample" scenario. Nevertheless, the distribution of the coupling variable is known in finite-sample, which permits the construction of feasible inference. To see this more clearly, let $\widetilde{W}$ be a generic copy of the standard Brownian motion on $[0, 1]$, and set $\xi_1 \equiv |\widetilde{W}_1 - \widetilde{W}_0|$ and $\xi_2 \equiv \sup_{s,t\in[0,1]}|\widetilde{W}_s - \widetilde{W}_t|$. By the scaling property of the Brownian motion, we see that $\lambda_1\zeta_{1,i} + \lambda_2\zeta_{2,i}$ has the same distribution as $\xi(\lambda) \equiv \lambda_1\xi_1 + \lambda_2\xi_2$, which can be easily computed via Monte Carlo simulation.

We may apply Theorem 1 to construct CIs for σ_t as follows. By the continuous mapping theorem, (2.9) implies that

$$\frac{\sigma_t}{\hat{\sigma}_t(\lambda)} \xrightarrow{d} \frac{1}{\xi(\lambda)}. \tag{2.10}$$

For $\alpha \in (0, 1)$, we can pick constants $B^C_{\alpha-}(\lambda)$ and $B^C_{\alpha+}(\lambda)$ such that

$$\mathbb{P}\left(B^C_{\alpha-}(\lambda) \leq \xi(\lambda)^{-1} \leq B^C_{\alpha+}(\lambda)\right) = 1 - \alpha. \tag{2.11}$$

The convergence in (2.10) then implies that

$$CI^C_{1-\alpha}(\lambda) \equiv \left[B^C_{\alpha-}(\lambda)\,\hat{\sigma}_t(\lambda), B^C_{\alpha+}(\lambda)\,\hat{\sigma}_t(\lambda)\right] \tag{2.12}$$

is a CI for σ_t with asymptotic level $1-\alpha$, that is, $\mathbb{P}\left(\sigma_t \in CI^C_{1-\alpha}(\lambda)\right) \to 1-\alpha$. We refer to $CI^C_{1-\alpha}(\lambda)$ as a *candlestick CI*.

The candlestick CI is asymptotically valid for any choice of $\lambda \neq 0$. But the arbitrariness in λ is clearly undesirable in practice. Following classical statistical principles, we pin down $\lambda = (\lambda_1, \lambda_2)$ so that (i) the candlestick estimator is asymptotically unbiased and (ii) its asymptotic variance is minimized within the class of linear estimators defined by (2.7). In view of Theorem 1, this amounts to setting λ as

$$\lambda^* = \underset{\lambda}{\operatorname{argmin}}\, Var\left(\xi(\lambda)\right), \quad \text{s.t.} \quad \mathbb{E}\left[\xi(\lambda)\right] = 1. \tag{2.13}$$

Denoting $\mu_1 \equiv \mathbb{E}[\xi_1]$ and $\mu_2 \equiv \mathbb{E}[\xi_2]$, the solution to this minimization problem can be written in explicit form as

$$\lambda_1^* = -\frac{1}{\mu_1}\frac{Cov\left(\frac{\xi_1}{\mu_1}-\frac{\xi_2}{\mu_2}, \frac{\xi_2}{\mu_2}\right)}{Var\left(\frac{\xi_1}{\mu_1}-\frac{\xi_2}{\mu_2}\right)} \approx -0.369, \quad \lambda_2^* = \frac{1}{\mu_2}\frac{Cov\left(\frac{\xi_1}{\mu_1}-\frac{\xi_2}{\mu_2}, \frac{\xi_1}{\mu_1}\right)}{Var\left(\frac{\xi_1}{\mu_1}-\frac{\xi_2}{\mu_2}\right)} \approx 0.811.^{11} \tag{2.14}$$

We refer to the resulting estimator, $\hat{\sigma}_t(\lambda^*)$, as the *Optimal candlesticK (OK)* estimator, which is what we have recommended (recall (1.1)) in the Introduction. It is interesting to note that the OK estimator assigns a nontrivial *negative* weight on $|r_i|$, although the absolute return is quite commonly used as a proxy of volatility in empirical work. A useful practical implication is: if two candlesticks have the same length (measured by the high-low range), the one with a shorter real-body (i.e., smaller absolute open-close return) actually indicates a higher level of volatility, and vice versa.

Equipped with the optimal weights, we need to further select the $B^C_{\alpha\pm}(\lambda^*)$ constants in order to compute the CIs described in (2.12). Note that the width of the CI is determined by $B^C_{\alpha+}(\lambda^*) - B^C_{\alpha-}(\lambda^*)$. Therefore, to obtain the shortest CI that satisfies the coverage restriction (2.11), we can simply set $\left[B^C_{\alpha-}(\lambda^*), B^C_{\alpha+}(\lambda^*)\right]$ as the highest density interval (HDI) of the random variable $\xi(\lambda^*)^{-1}$. To visualize this construction, we plot in Figure 2 the probability density function of $\xi(\lambda^*)^{-1}$ along with its 90% HDI. For comparison, we also plot the distribution and the 90% HDI of $\xi(\lambda)^{-1}$ with (λ_1, λ_2) being $\left(\mu_1^{-1}, 0\right)$ (resp. $\left(0, \mu_2^{-1}\right)$), which corresponds to the asymptotically unbiased estimator formed using only the absolute open-close return (resp. the high-low range). Looking at the left panel of Figure 2, we see that the distribution of $\xi(\lambda^*)^{-1}$ is much tighter

[11] The numerical values of λ_1^* and λ_2^* are obtained by computing the moments using 100 million simulated paths of $\widetilde{W}$ on the unit interval $[0,1]$, which is discretized with mesh size 10^{-7} in each simulation.

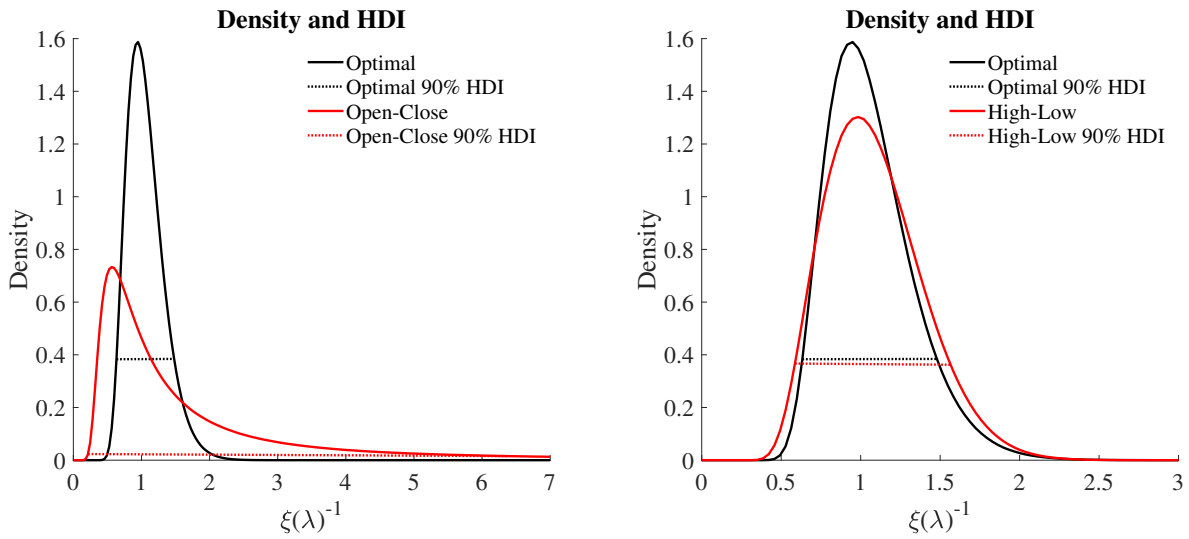

Figure 2: The figure plots the probability density function of the $\xi(\lambda)^{-1}$ variable evaluated at the optimal weight λ^* and the associated 90% highest density interval (HDI). For comparison, the density and associated 90% HDI are also plotted for $\lambda = (\mu_1^{-1}, 0)$ (resp. $\lambda = (0, \mu_2^{-1})$) on the left (resp. right) panel, corresponding to the "degenerate" candlestick estimator that only depends on the open-close return (resp. high-low range).

than that of $\xi\left(\mu_1^{-1}, 0\right)^{-1}$ and, in particular, the former's 90% HDI is considerably shorter than the latter's. This suggests that CIs constructed using the OK estimator are much more accurate than those based only on the absolute return. Similarly, the right panel of the figure also reveals the superiority of the OK estimator relative to the estimator solely based on the high-low range, although the contrast is not as striking as what we see on the left panel.

To facilitate applications, we tabulate in Table 1 the aforementioned HDI-based critical values $B^C_{\alpha\pm}(\lambda^*)$ associated with the OK estimator for various confidence levels. Since we consider a "data-scarce" environment, we include some lower-than-conventional confidence levels which may be relevant for investors who seek sharper (but more "aggressive") inference. For simplicity, we refer to the resulting CI as the *optimal candlestick confidence interval (OKCI)*. In addition, to gauge more precisely the statistical efficiency of the CI, we report the width $B^C_{\alpha+}(\lambda^*) - B^C_{\alpha-}(\lambda^*)$. We also report the analogous quantities for the open-close and the high-low estimators (i.e., those corresponding to $\lambda = (\mu_1^{-1}, 0)$ and $(0, \mu_2^{-1})$). From the table, we see that the OKCI is always tighter than the other two alternatives across all confidence levels. Specifically, at the 90% level, the OKCI is $6.150/0.849 \approx 7.24$ (resp. $0.979/0.849 \approx 1.15$) times as efficient as the CI based on the open-close return (resp. high-low range).

We summarize the above theoretical discussion more concisely as a "user's guide." We recom-

Table 1: Critical Values for Candlestick Confidence Intervals

	Optimal			Open-Close			High-Low		
Level	Lower	Upper	Width	Lower	Upper	Width	Lower	Upper	Width
50%	0.793	1.135	0.341	0.337	1.234	0.897	0.792	1.202	0.411
60%	0.762	1.189	0.427	0.307	1.561	1.255	0.749	1.260	0.511
70%	0.727	1.255	0.528	0.279	2.102	1.824	0.704	1.331	0.627
80%	0.688	1.343	0.656	0.249	3.173	2.924	0.654	1.424	0.770
90%	0.636	1.485	0.849	0.216	6.366	6.150	0.587	1.565	0.979

Note: The table reports the critical values of alternative candlestick confidence intervals (CIs) for the spot volatility. The confidence level $1-\alpha$ ranges from 50% to 90%. The optimal, open-close, and high-low CIs correspond to λ being λ^*, $(\mu_1^{-1}, 0)$, and $(0, \mu_2^{-1})$, respectively. For each case, we report the lower bound $B^C_{\alpha-}(\lambda)$ and the upper bound $B^C_{\alpha+}(\lambda)$ computed as the highest density interval (HDI) of the distribution of $\xi(\lambda)^{-1}$, along with the width of the HDI. The numerical values are computed by simulation using 100 million Monte Carlo draws, for which the Brownian motion $\widetilde{W}$ on the unit interval is simulated under a Euler scheme with mesh size 10^{-7}.

mend estimating the spot volatility using the OK estimator[12]

$$\hat{\sigma}_t^* \equiv \hat{\sigma}_t(\lambda^*) = \frac{0.811 w_i - 0.369\,|r_i|}{\sqrt{\Delta_n}}, \tag{2.15}$$

and gauging its sampling variability using the 90%-level OKCI

$$CI^*_{90\%} \equiv [0.636\,\hat{\sigma}_t^*, 1.485\,\hat{\sigma}_t^*], \tag{2.16}$$

where the critical values are obtained from Table 1. At this confidence level, the OKCI provides a reasonable balance between statistical confidence and precision.[13] The OK estimator and the OKCI are obviously very easy to compute. In fact, the calculation can be carried out manually on

[12] In the U.S. stock market, each regular trading day contains 390 minutes. If each candlestick spans a 10-minute trading session, the user should set $\Delta_n = 10/390 = 1/39$ to obtain volatility estimates quoted in daily terms. The length of a trading day and the duration of the candlestick may vary across different markets and trading applications, so the value of Δ_n should be adjusted accordingly. If the user takes the duration of the candlestick (e.g., 10 minutes) as the unit of time, they can simply set $\Delta_n = 1$.

[13] If the practitioner is willing to trade some statistical confidence for a sharper CI, they may use, for example, the 50% OKCI given by $[0.793\,\hat{\sigma}_t^*, 1.135\,\hat{\sigma}_t^*]$.

Table 2: Critical Values for Optimal Confidence Intervals based on k Candlesticks

	$k=3$			$k=5$			$k=10$		
	Lower	Upper	Width	Lower	Upper	Width	Lower	Upper	Width
50%	0.892	1.087	0.195	0.917	1.069	0.151	0.944	1.051	0.107
60%	0.870	1.114	0.244	0.900	1.089	0.189	0.931	1.064	0.133
70%	0.846	1.147	0.301	0.882	1.114	0.233	0.917	1.081	0.164
80%	0.818	1.191	0.373	0.858	1.146	0.288	0.899	1.103	0.203
90%	0.779	1.259	0.480	0.826	1.197	0.370	0.875	1.136	0.261

Note: The table reports the critical values for optimal confidence intervals based on the k-candlestick estimator defined in (2.17). At confidence level $1-\alpha$, the critical values are given by the highest density interval of the distribution of $k/\sum_{j=1}^{k}(\lambda_1^*\zeta_{1,i+j}+\lambda_2^*\zeta_{2,i+j})$. The numerical values are computed by simulation under the same scheme as in Table 1.

a basic calculator within a couple of minutes. Our method may thus enable an ordinary investor to make inference on spot volatility in real time as the candlestick chart updates in their trading application; see Section 4 for a concrete empirical demonstration. This baseline estimator may be further improved in various ways, to which we now turn.

2.3 Extensions

The OK estimator proposed in Section 2.2 concerns the inference for spot volatility using a single candlestick. This basic method is most relevant in a data-scarce environment in which the user only has coarsely sampled candlesticks, say, for example, at the 10-minute frequency. We have intentionally focused on this scenario so far due to its simplicity and practicality. Needless to say, the underlying econometric idea may be extended to many different and more complicated settings. In this subsection, we discuss a few possibilities, which may be mixed and matched, in order to guide future research in this direction.

Firstly, we note that the single-candlestick estimator (2.7) can be easily improved when higher-frequency candlesticks are available. For example, if a researcher has access to 1-minute candlesticks, they may construct a local average version of the candlestick volatility estimator for a

10-minute estimation window by using, with $\Delta_n = 1$ minute and $k = 10$,

$$\hat{\sigma}_t(k,\lambda) \equiv \frac{1}{k}\sum_{j=1}^{k} \frac{\lambda_1 |r_{i+j}| + \lambda_2 w_{i+j}}{\Delta_n^{1/2}}, \quad \text{for} \quad t \in [(i-1)k\Delta_n, ik\Delta_n]. \tag{2.17}$$

This construction exactly parallels the conventional return-based spot variance estimator (2.2). By harnessing higher-frequency candlestick data (if they are available), the new k-candlestick estimator displayed above will evidently be more efficient than that based on a single 10-minute candlestick.[14] Treating the bandwidth k as a fixed constant (as in Bollerslev, Li, and Liao (2021)), we can easily extend Theorem 1 to establish an analogous coupling result for $\hat{\sigma}_t(k,\lambda)$, namely,

$$\frac{\hat{\sigma}_t(k,\lambda)}{\sigma_t} = \frac{1}{k}\sum_{j=1}^{k} (\lambda_1 \zeta_{1,i+j} + \lambda_2 \zeta_{2,i+j}) + o_p(1).$$

Since the $(\zeta_{1,i}, \zeta_{i,2})$ variables are independent across i, it is also easy to see that λ^* determined by (2.14) is still the optimal choice for minimizing the asymptotic variance of $\hat{\sigma}_t(k,\lambda)$ while maintaining asymptotic unbiasedness, and the aggregation leads to a k-fold reduction of the asymptotic variance. Moreover, the distribution of the coupling variable is known in finite-sample, and we can use the $1-\alpha$ level HDI of the distribution of $k/\sum_{j=1}^{k}(\lambda_1^* \zeta_{1,i+j} + \lambda_2^* \zeta_{2,i+j})$, denoted $\left[B_{\alpha-}^C(k,\lambda^*), B_{\alpha+}^C(k,\lambda^*)\right]$, to construct the optimal CI given by

$$CI_{1-\alpha}^*(k) \equiv \left[B_{\alpha-}^C(k,\lambda^*)\,\hat{\sigma}_t(k,\lambda^*), B_{\alpha+}^C(k,\lambda^*)\,\hat{\sigma}_t(k,\lambda^*)\right].$$

We tabulate the $B_{\alpha\pm}^C(k,\lambda^*)$ critical values in Table 2 for various confidence levels and k values. As expected, when k increases, the OKCIs become tighter approximately along the $k^{-1/2}$ scale.[15] For example, at the 90% confidence level, the length of the 10-candlestick CI is $0.261/0.849 \approx 31\%$ of that of the single-candlestick CI.

Secondly, we note that the result in Theorem 1 can be used to construct a formal test for detecting "large" volatility jumps. Consider two time points s and t in two distinct trading sessions $I_{n,i}$ and $I_{n,j}$, respectively. We may test the null hypothesis $H_0 : \sigma_t = \sigma_s$ against a one-sided alternative $H_a : \sigma_t > \sigma_s$ using the test statistic $\log(\hat{\sigma}_t^*/\hat{\sigma}_s^*) = \log(\hat{\sigma}_t^*) - \log(\hat{\sigma}_s^*)$, where the log transformation is employed to "symmetrize" the roles of the two volatility estimators in the test statistic. Under the null hypothesis, Theorem 1 and the continuous mapping theorem[16] imply that

$$\log(\hat{\sigma}_t^*) - \log(\hat{\sigma}_s^*) = \log\left(\frac{\lambda_1^* \zeta_{1,j} + \lambda_2^* \zeta_{2,j}}{\lambda_1^* \zeta_{1,i} + \lambda_2^* \zeta_{2,i}}\right) + o_p(1).$$

[14] Since the estimators span the same estimation window, they are subject to essentially the same level of nonparametric bias.

[15] The reduction in the width of CI is consistent with the k-fold reduction in the asymptotic variance. But these two effects are not exactly the same because the OKCIs are not based on the asymptotic Gaussian approximation.

[16] Since the OK estimators are not consistent for the spot volatility, the delta method cannot be applied to derive the asymptotic distributions of their nonlinear transformations.

For a significance level $\alpha \in (0,1)$, the $1-\alpha$ quantile of the coupling variable $\log\left(\frac{\lambda_1^*\zeta_{1,j}+\lambda_2^*\zeta_{2,j}}{\lambda_1^*\zeta_{1,i}+\lambda_2^*\zeta_{2,i}}\right)$, which is known in finite-sample, may be used as the critical value, and we reject the null hypothesis when $\log(\hat{\sigma}_t^*) - \log(\hat{\sigma}_s^*)$ is greater than the critical value. A two-sided test may be constructed similarly. Since the coupling variable is nondegenerate, this test is not consistent against fixed, but "small," derivation from the null. This is not surprising given the asymptotically "fixed" amount of information being exploited. That said, the test has valid asymptotic size control and is asymptotically unbiased (i.e., it has nontrivial asymptotic power under the alternative). It can be used to detect "large" moves in volatility across different time points, for example, around important macroeconomic news announcements as studied in Lucca and Moench (2015), Bollerslev, Li, and Xue (2018), and Nakamura and Steinsson (2018), among others. The power of the test may be further improved by aggregating multiple higher-frequency candlesticks.

Thirdly, our coupling-based inference can be readily adapted to construct valid inference for the other types of estimators. For example, if one aims to make inference directly on the spot variance σ_t^2, a good candidate estimator is the Garman–Klass estimator defined in (2.8). Using a similar argument as Theorem 1, we can show that the Garman–Klass estimator is an asymptotically unbiased estimator for σ_t^2 and it admits the following coupling:

$$\frac{\hat{v}_t^{GK}}{\sigma_t^2} = 0.5\zeta_{2,i}^2 - (2\log 2 - 1)\zeta_{1,i}^2 + o_p(1).$$

CIs for σ_t^2 may be constructed as $\left[B_{\alpha-}^{GK}\hat{v}_t^{GK}, B_{\alpha+}^{GK}\hat{v}_t^{GK}\right]$, whose length can be minimized by taking $\left[B_{\alpha-}^{GK}, B_{\alpha+}^{GK}\right]$ as the HDI of the distribution of $(0.5\zeta_{2,i}^2 - (2\log 2 - 1)\zeta_{1,i}^2)^{-1}$.

Finally, we note that all aforementioned procedures may be generalized to a bipower (or multipower) version in the spirit of Barndorff-Nielsen and Shephard (2004) and Barndorff-Nielsen, Shephard, and Winkel (2006). In an influential paper, Barndorff-Nielsen and Shephard (2004) propose the bipower method to elegantly achieve jump-robust estimation of the integrated variance. Although the candlestick-based spot estimators are robust to Poisson-type jumps (because with probability approaching 1, $[t-\Delta_n, t+\Delta_n]$ does not contain any jump for each fixed t), it is conceivable that using the bipower construction can further improve the estimator's robustness with respect to jumps. Specifically, we may consider a bipower extension of our single-candlestick estimator (2.7) as

$$\hat{\sigma}_t^{Bipower}(\lambda) = \frac{\lambda_1 |r_i|^{1/2} |r_{i+1}|^{1/2} + \lambda_2 w_i^{1/2} w_{i+1}^{1/2}}{\Delta_n^{1/2}}.$$

Similar to Theorem 1, it can be shown that

$$\frac{\hat{\sigma}_t^{Bipower}(\lambda)}{\sigma_t} = \frac{\lambda_1 |\zeta_{1,i}|^{1/2} |\zeta_{1,i+1}|^{1/2} + \lambda_2 |\zeta_{2,i}|^{1/2} |\zeta_{2,i+1}|^{1/2}}{\Delta_n^{1/2}} + o_p(1).$$

The constants λ_1 and λ_2 can then be determined to achieve certain optimality requirement (e.g., unbiasedness and minimal variance). As mentioned above, the advantage of the bipower version is

that it will be more robust with respect to price jumps. But there is no free lunch. The bipower version also implicitly requires the volatility to be "nearly" identical across the two consecutive candlesticks on $I_{n,i}$ and $I_{n,i+1}$ time intervals. From a finite-sample point of view, the latter requirement is more plausible for higher-frequency data (say, 1 minute), but might be somewhat restrictive on a longer time scale (say, 10 minutes) especially when news flows quickly in the market. A comprehensive investigation on the bipower/multipower extension may be an interesting topic for a separate research, but is beyond the scope of the present paper.

The above list of possible extensions is by no means exhaustive. We do not investigate them in full details in order to remain focused on our main goal, namely, developing a simple practical method for spot volatility inference that can be adopted by an average, perhaps computationally constrained, investor. Further extensions are left for future research.

2.4 Efficiency comparison in a thought experiment

So far, we have shown how to optimally exploit information from candlesticks for volatility inference, which may allow a retail investor to make reasonably sharp inference on spot volatility using easy-to-access data. But the retail investor is apparently at a disadvantage compared to an institutional researcher who has access to "professional grade" data, such as those commonly seen in the recent high-frequency econometrics literature. One may naturally wonder: How severe is this disadvantage?

We shed some light on this question by considering a thought experiment. Imagine two market observers, a "retail trader" and an "econometrician," who are both interested in making inference for an asset's spot volatility over a 10-minute trading session.[17] The retail trader observes a 10-minute candlestick from their mobile trading app and conducts inference using the OKCI. The econometrician, on the other hand, has access to returns data at the 1-minute frequency. Like Bollerslev, Li, and Liao (2021), the econometrician may compute CIs for the spot volatility using either the conventional asymptotic Gaussian-based method or the fixed-k method, setting the bandwidth $k_n = 10$; recall the discussion in Section 2.1. Since the econometrician employs higher-frequency data and uses state-of-the-art econometric methods from the existing literature, it is natural to expect that they will produce a tighter CI than the retail trader. However, following our proposal, the retail trader exploits the additional information in the candlestick, which may compensate their disadvantage from using lower-frequency data. We analyze whether, and to which extent, the econometrician can draw a sharper inference than the retail trader.

We perform the comparison as follows. From (2.4), (2.6), and (2.12), we observe that the

[17]The trader may be interested in the spot volatility for several reasons, such as gauging the prevailing liquidity (say using a volatility-volume ratio), accessing the risk of getting a margin call, or making decisions on vega-related option trading.

Table 3: Critical Values in Thought Experiment

	Candlestick (10 min)			Gaussian (1 min $\times 10$)			Fixed-k (1 min $\times 10$)		
	Lower	Upper	Width	Lower	Upper	Width	Lower	Upper	Width
50%	0.793	1.135	0.341	0.849	1.151	0.302	0.822	1.126	0.304
60%	0.762	1.189	0.427	0.812	1.188	0.376	0.794	1.176	0.382
70%	0.727	1.255	0.528	0.768	1.232	0.464	0.762	1.239	0.476
80%	0.688	1.343	0.656	0.713	1.287	0.573	0.727	1.326	0.600
90%	0.636	1.485	0.849	0.632	1.368	0.736	0.679	1.476	0.797

Note: The table reports the critical values of alternative confidence intervals (CI) for the spot volatility. The left panel reports $B^C_{\alpha\pm}(\lambda^*)$ for the optimal candlestick CI based on a single 10-minute candlestick. The middle (resp. right) panel reports the $B^G_{\alpha\pm}$ (resp. $B^F_{\alpha\pm}$) critical values for the Gaussian (resp. fixed-k) CI based on $k_n = 10$ returns sampled at the 1-minute frequency.

Gaussian, fixed-k, and candlestick CIs are all proportional to their corresponding spot volatility estimates. We thus compare the CIs' critical values $B^G_{\alpha\pm}$, $B^F_{\alpha\pm}$, and $B^C_{\alpha\pm}(\lambda^*)$, and use the width between the upper and lower bounds to gauge their relative efficiency. Table 3 reports these numbers for various confidence levels. As expected, we find that the OKCI based on a single 10-minute candlestick is indeed wider than the Gaussian and fixed-k CIs based on ten 1-minute returns. The conventional Gaussian CI is also tighter than the fixed-k CI. However, the seemingly better performance of the Gaussian CI should be taken with a grain of salt: As shown in the simulation study of Bollerslev, Li, and Liao (2021), the Gaussian CI may suffer from nontrivial size distortion when k_n is small. This is quite intuitive, because the asymptotic Gaussian approximation does not "kick in" sufficiently well when the averaging is done over only ten observations. On the other hand, Bollerslev, Li, and Liao (2021) show that the fixed-k CIs have almost exact coverage for small k_n's, and hence, may serve as a more relevant benchmark for our comparison here.

Compared to the fixed-k CI, the proposed OKCI is only moderately less efficient. For example, at the 90% confidence level, the OKCI is $0.797/0.849 \approx 94\%$ as tight as the fixed-k CI. This finding is remarkable, as it suggests that the retail trader, whose candlestick observation is updated at the (slow) 10-minute frequency, may be able to make formal inference about the spot volatility at nearly the same statistical accuracy as in many academic research papers (e.g., Bollerslev,

Li, and Xue (2018), Bollerslev, Li, and Liao (2021)) based on returns data sampled at the 1-minute frequency. Of course, if one implements the multiple-candlestick estimator $\hat{\sigma}_t(k, \lambda^*)$ by aggregating ten 1-minute candlesticks, the width of the resulting 90%-level OKCI is only one third of that of the fixed-k CI. Overall, these findings further highlight the usefulness of the proposed candlestick-based method in applied work.

3 Monte Carlo simulations

We examine the finite-sample properties of the proposed inference method in a Monte Carlo experiment. Following Bollerslev and Todorov (2011), we simulate the (log) price process from a two-factor stochastic volatility model. Specifically, with the unit time interval normalized to "one day," we generate the process P according to

$$\begin{aligned} dP_t &= \sigma_t dW_t, \quad \sigma_t^2 = V_{1,t} + V_{2,t}, \\ dV_{1,t} &= 0.0128(0.4068 - V_{1,t})dt + 0.0954\sqrt{V_{1,t}}\left(\rho dW_t + \sqrt{1-\rho^2}dB_{1,t}\right), \\ dV_{2,t} &= 0.6930(0.4068 - V_{2,t})dt + 0.7023\sqrt{V_{2,t}}\left(\rho dW_t + \sqrt{1-\rho^2}dB_{2,t}\right), \end{aligned}$$

where W, B_1, and B_2 denote independent standard Brownian motions. The $\rho = -0.7$ parameter captures the well-documented negative correlation between price and volatility shocks (i.e., the "leverage" effect). The V_1 volatility factor is highly persistent with a half-life of 2.5 months, while the V_2 volatility factor is quickly mean-reverting with a half-life of only one day. For ease of discussion, we fix $V_{1,0} = V_{2,0} = 0.5$, so that $\sigma_0 = 1$. We simulate the "continuous-time processes" using a Euler scheme with mesh size being 10^{-7} minute. The candlesticks used in the calculations are then constructed on 10-minute and 1-minute intervals. The estimand σ_t is sampled at the mid-point of each 10-minute estimation window.[18] All numerical results reported below are based on 10,000 Monte Carlo replications.

We consider six inference methods in total. The first three are based on a single 10-minute candlestick. Specifically, we compute three spot volatility estimators:

$$\hat{\sigma}_t^* = \frac{0.811 w_i - 0.369\,|r_i|}{\sqrt{\Delta_n}}, \quad \hat{\sigma}_t^r = \frac{|r_i|}{\mu_1\sqrt{\Delta_n}}, \quad \hat{\sigma}_t^w = \frac{w_i}{\mu_2\sqrt{\Delta_n}},$$

where $\hat{\sigma}_t^*$ is the recommended OK estimator, and $\hat{\sigma}_t^r$ and $\hat{\sigma}_t^w$ are asymptotically unbiased estimators based only on the open-close return and the high-low range, respectively. Using the critical values reported in Table 1, we further compute their associated 90%-level CIs respectively as

$$CI_{90\%}^* = [0.636\,\hat{\sigma}_t^*, 1.485\,\hat{\sigma}_t^*], \quad CI_{90\%}^r = [0.216\hat{\sigma}_t^r, 6.366\hat{\sigma}_t^r], \quad CI_{90\%}^w = [0.587\hat{\sigma}_t^w, 1.565\hat{\sigma}_t^w].$$

[18] Recall that Theorem 1 holds for any $t \in I_{n,i}$. We have also considered σ_t sampled at the start, end, and a random position of the 10-minute estimation window. The results are very similar, and hence, are omitted for brevity.

Table 4: Simulation Results

Method	Bias	RMSE	Coverage Rate	CI Width
Panel A. Methods based on one 10-minute observation				
OK 10-min	0.002	0.253	0.896	0.849
Open-Close	-0.004	0.755	0.896	6.120
High-Low	0.001	0.300	0.896	0.978
Panel B. Methods based on ten 1-minute observations				
OK 1-min	0.001	0.081	0.892	0.261
Gaussian	-0.023	0.224	0.861	0.718
Fixed-k	-0.023	0.224	0.898	0.777

Note: The table reports the relative biases and root-mean-squared-errors (RMSE) of $\hat{\sigma}_t^*$ (OK 10-min), $\hat{\sigma}_t^r$ (Open-Close), $\hat{\sigma}_t^w$ (High-Low), $\hat{\sigma}_t(10, \lambda^*)$ (OK 1-min) and $\sqrt{\hat{v}_t(k_n)}$ (Gaussian and Fixed-k), and the coverage rates and average widths of their associated 90%-level CIs. These numbers are calculated based on 10,000 Monte Carlo replications.

The remaining three methods are implemented using ten observations sampled at the 1-minute frequency. We use 1-minute candlesticks to compute the local average OK estimator $\hat{\sigma}_t(k, \lambda^*)$ with $k = 10$; recall (2.17). The corresponding 90%-level optimal CI is given by

$$CI^*_{90\%}(10) = \left[0.875\,\hat{\sigma}_t(10, \lambda^*), 1.136\,\hat{\sigma}_t(10, \lambda^*)\right].$$

Finally, we compute the conventional spot volatility estimator $\sqrt{\hat{v}_t(k_n)}$ using 1-minute returns and $k_n = 10$. Using the critical values reported in Table 3, we construct the associated Gaussian and fixed-k CIs as

$$CI^G_{90\%} = \left[0.632\sqrt{\hat{v}_t(10)}, 1.368\sqrt{\hat{v}_t(10)}\right], \quad CI^F_{90\%} = \left[0.679\sqrt{\hat{v}_t(10)}, 1.476\sqrt{\hat{v}_t(10)}\right].$$

Table 4 summarizes the finite-sample performance of these spot volatility estimators and CIs. Specifically, we report the bias and root-mean-squared-error (RMSE) of the spot volatility estimates, both in relative term (i.e., estimation errors are normalized by the true value of σ_t). We also report the coverage rates of the 90%-level CIs, along with their widths averaged across all simulations.

Looking at Panel A of the table, which reports results based on the "coarse" 10-minute level data, we see that all three estimators are essentially unbiased, with their relative biases bounded by 0.4% in magnitude. Their CIs also have virtually exact size control. As expected, the OK estimator has the smallest RMSE and the OKCI has the shortest width, whereas the estimator based only on the open-close return is evidently the least accurate.

Panel B reports the performance of the local average OK estimator $\hat{\sigma}_t(10, \lambda^*)$ and the conventional spot volatility estimator using the "fine" 1-minute level data. Since these estimators are computed using more observations, they are, not surprisingly, more accurate than those in Panel A. We note that the "fine" OK estimator formed by aggregating ten 1-minute candlesticks is far more accurate than the conventional estimator, as evidenced by the former's smaller bias and RMSE. The OKCI is also tighter than conventional Gaussian and fixed-k CIs, and remains to exhibit excellent size control. It is also interesting to note that even the "coarse" OK estimator based on a single 10-minute candlestick is $0.224/0.253 \approx 89\%$ as efficient (in terms of RMSE) as the conventional estimator computed using more observations, and the former's CI is $0.777/0.849 \approx 92\%$ as tight as the fixed-k CI.[19]

In summary, these simulation results are consistent with our theoretical predictions. We see that the OK estimator based on the 10-minute candlestick has adequate accuracy, and is more efficient than candlestick-based estimators with suboptimal weights. In addition, the candlestick-based CIs all have almost exact finite-sample coverage. We also see that the OK method applied to 10-minute level data can achieve roughly 90% of the accuracy of the conventional method applied to 1-minute level data and, if the OK method is also applied to 1-minute candlestick observations, its accuracy can be further improved by a factor of three. Overall, these results suggest that the proposed method can be reliably used for making inference on spot volatility.

4 Empirical illustration

We illustrate the candlestick-based inference method in a case study for a recent event on February 23, 2021, when the Federal Reserve Chairman, Jerome Powell, delivered his semiannual monetary policy report to the U.S. Congress. The media coverage of this event started at 10:00 EST and the event lasted for approximately 2.5 hours. Till that day, the 10-year Treasury yield had experienced an 8-month liftoff from 0.54% in July, 2020, to 1.37%, whereas the short rate was kept at the zero lower bound. The steepening of the yield curve reflected the market's expectation for a strong economic recovery from the COVID-19 pandemic. Meanwhile, the rising 10-year yield also led to much anxiety in the equity market, in that the higher discount rate might lead to a major correction

[19] Consistent with the simulation findings of Bollerslev, Li, and Liao (2021), we see that the Gaussian-based CI is somewhat undersized, whereas the fixed-k CI has almost exact coverage.

on asset valuation. The Federal Reserve's economic outlook and its stance on monetary policy was the focal point of the financial market at the time. This event is thus ideal for our illustrative purpose for two reasons. Firstly, it corresponds to a high-stake economic and policy environment. Secondly, the intense flow of information together with the market's considerable attention paid to it also means that the price and volatility tend to fluctuate substantially, providing us with a challenging real-data scenario to gauge the practical performance of the proposed method.

We consider five assets from different asset classes including the 10-year Treasury note, a passive equity ETF for the S&P 500 index (SPY), an active equity ETF of technology stocks (ARKK), a gold ETF (GLD), and Bitcoin, which we use to represent the bond, equity, commodity, and cryptocurrency markets. Candlestick data for these assets are obtained from the Bloomberg Terminal at the 10-minute frequency.[20] We focus on the 10-minute data for two reasons. First, this sampling scheme is sufficiently sparse to guard against complications stemming from microstructure noise for the different asset classes studied here. Second, the 10-minute frequency also appears to be the lowest at which the candlestick chart is updated in popular online trading applications, and hence, allows us to examine how the proposed method will behave in a "worst case" scenario (in terms of data availability). In the discussion below, we sometimes associate the price and volatility dynamics with Powell's testimony, the video recording of which can be found on the website of the U.S. Senate Committee on Banking, Housing, and Urban Affairs.[21] By doing so, we aim to trace the market and the news flow simultaneously. We do not attempt to formally assert any causal link between them, as it is impossible to rule out all potential confounding factors.

Since Powell's testimony was about monetary policy, our primary focus, which was shared broadly by market participants at that time, is naturally on the 10-year Treasury yield. To set the stage, we plot the 10-minute candlestick chart for the 10-year Treasury yield on the top panel of Figure 3. The sample period is from 7:30 to 17:00 EST, and we highlight the time intervals for the regular trading hours of the equity market (i.e., 9:30 – 16:00) and Powell's testimony for ease of discussion. In the bottom panel of the figure, we plot the corresponding OK volatility estimates along with the 90% OKCIs, as described in (2.15) and (2.16), respectively.

Figure 3 reveals some interesting price and volatility dynamics in the 10-year yield. From the candlestick chart, we see that the 10-year yield gradually rose from 1.36% to 1.38% during the two hours before the stock market open, and then fluctuated violently between 9:30 and 10:00, as evidenced by the three long candlesticks during that half-hour trading session. The volatility

[20]Similar data are also available in real time to retail investors from various trading platforms (e.g., Charles Schwab, Fidelity, Interactive Broker, Robinhood, etc.) or public websites such as Yahoo Finance.

[21]The video is available at https://www.banking.senate.gov/hearings/02/12/2021/the-semiannual-monetary-policy-report-to-the-congress. The media coverage started at 10:00 AM EST, which corresponds the 14:47 timestamp in the video.

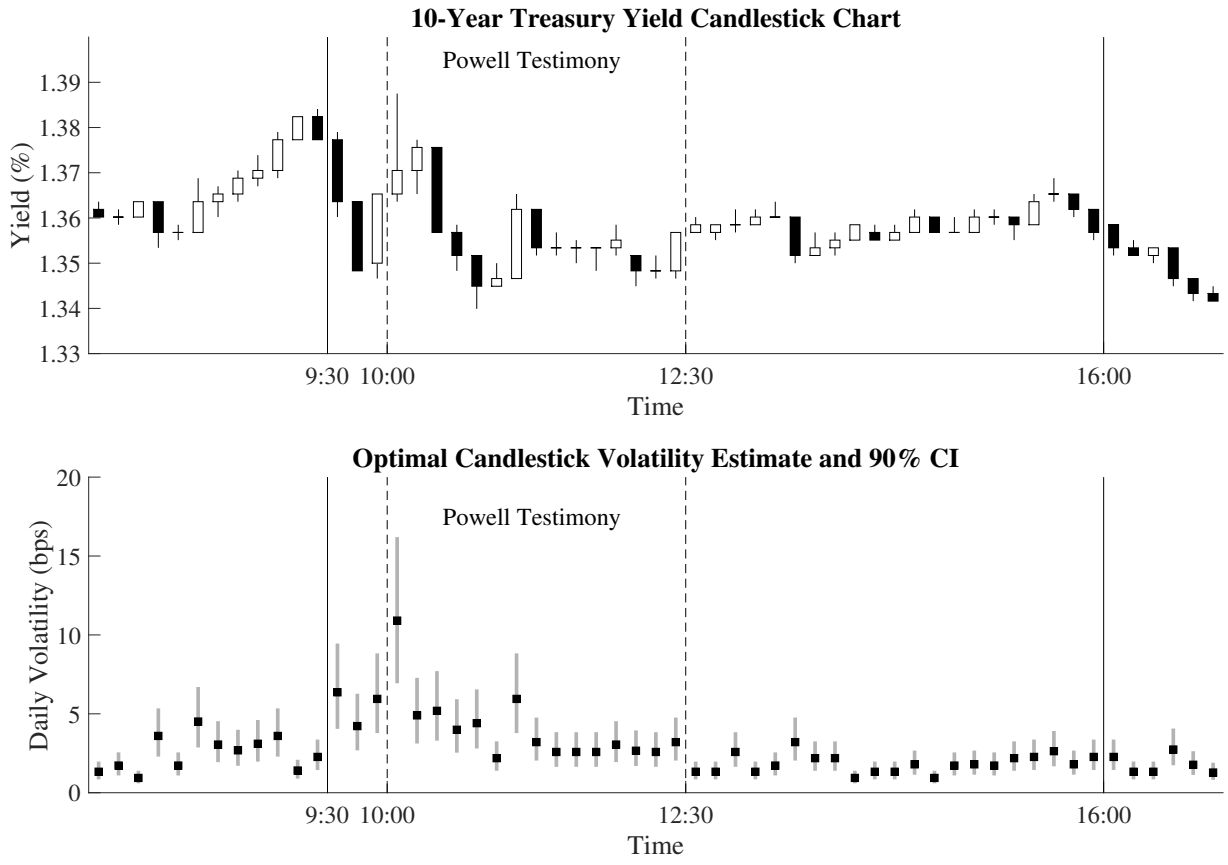

Figure 3: The top panel plots the 10-minute candlestick chart for the yield of the 10-year U.S. Treasury note from 7:30 to 17:00 EST on February 23, 2021. The solid and dashed lines highlight the regular trading hour and the Fed Chairman Powell's congressional testimony, respectively. The bottom panel shows the optimal candlestick volatility estimate (on daily horizon in basis points) of the yield for each 10-minute trading session, along with the associated 90% confidence interval.

jump at the market open can be seen more clearly from the spot volatility estimates plotted on the bottom panel, particularly in view of the fact that the volatility actually dipped in the 20-minute period before 9:30.

Powell's testimony started at 10:00. The opening statements given by the Committee Chairman Sherrod Brown (D-OH) and the ranking member Senator Pat Toomey (R-PA) during the first 15 minutes of the testimony reflected the partisan disagreement on a range of policy issues, including particularly the 1.9-trillion-dollar stimulus plan soon to be voted in the Senate. Interestingly, the estimated volatility of the 10-year yield also reached its daily maximum of 11 basis points as the two Senators spoke.

It is instructive to explain how this large volatility estimate is related to the corresponding candlestick. Looking at the candlestick immediately after the 10:00 timestamp, we see that it not only has a wide high-low range but also a relatively short real-body (i.e., open and close prices are

similar); indeed, the most salient feature of this candlestick is its very long upper shadow. Recall that the OK estimator assigns a negative weight on the absolute open-close return. This explains why the volatility estimate for the 10:00–10:10 trading session is much higher than those between 9:30 and 10:00, despite the fact that all these candlesticks have similar "sizes."

The Fed Chairman began his statement at around 10:15. The aforementioned volatility spike soon reverted to a lower level around 4–5 basis points and appeared to trend down during the rest of the first hour of the testimony. Meanwhile, the yield also dropped from 1.38% to 1.34%. These price and volatility movements occurred as Chairman Powell confirmed that the central bank would continue to be accommodative in various policy dimensions such as the federal funds rate, inflation target, and asset purchase. Immediately after the first hour, we see, quite interestingly, a 6-basis-point spike in volatility on the 11:00-11:10 interval. This might be attributed to the conversation between Senator Mike Rounds (R-SD) and Chairman Powell on whether the Federal Reserve would extend the temporary change to its supplementary leverage ratio (SLR) for bank holding companies. During the COVID-19 crisis, the SLR exclusion temporarily allowed banks to exclude U.S. Treasuries and deposits at the Federal Reserve from the SLR denominator. Without an extension, banks would have less capacity to own Treasuries, and the resulting selling pressure might push the long-term Treasury yield even higher. This was an important concern in the bond market. Chairman Powell responded to Senator Rounds by saying that the Fed was thinking about this issue and would make a decision "pretty soon." The unresolved policy uncertainty might have disappointed the market (in view of Powell's evidently accommodative tone to the other policy issues), which may explain the peculiar volatility spike during that 10-minute trading session.[22] But this volatility spike is short-lived: The volatility level soon dropped to a lower 2–3 basis points range for the remaining part of the testimony, and dropped even further afterwards.

Our main goal of discussing the above real-data example is to demonstrate the empirical applicability of our proposed econometric method for studying asset price volatility. We find the estimated spot volatility path to be economically sensible, without any obviously erratic behavior that may be hard to interpret. By carefully going through the contemporaneous economic events, we also see that the occasional "spikes" in the estimated volatility tend to reflect "features" of what is happening, rather than "bugs" of the estimator. In addition, we note that the 90% OKCIs are reasonably tight to make meaningful inferential statements about the ups and downs of the volatility process. This is remarkable in view of the fact that each CI is only based on a single 10-minute candlestick.

To examine whether the proposed method works well on the other types of assets, we further

[22]The market's response turns out to be "rationalizable" with the benefit of hindsight, as the central bank announced the termination of the SLR exclusion later on March 19, 2021. The official announcement can be found at https://www.federalreserve.gov/newsevents/pressreleases/bcreg20210319a.htm.

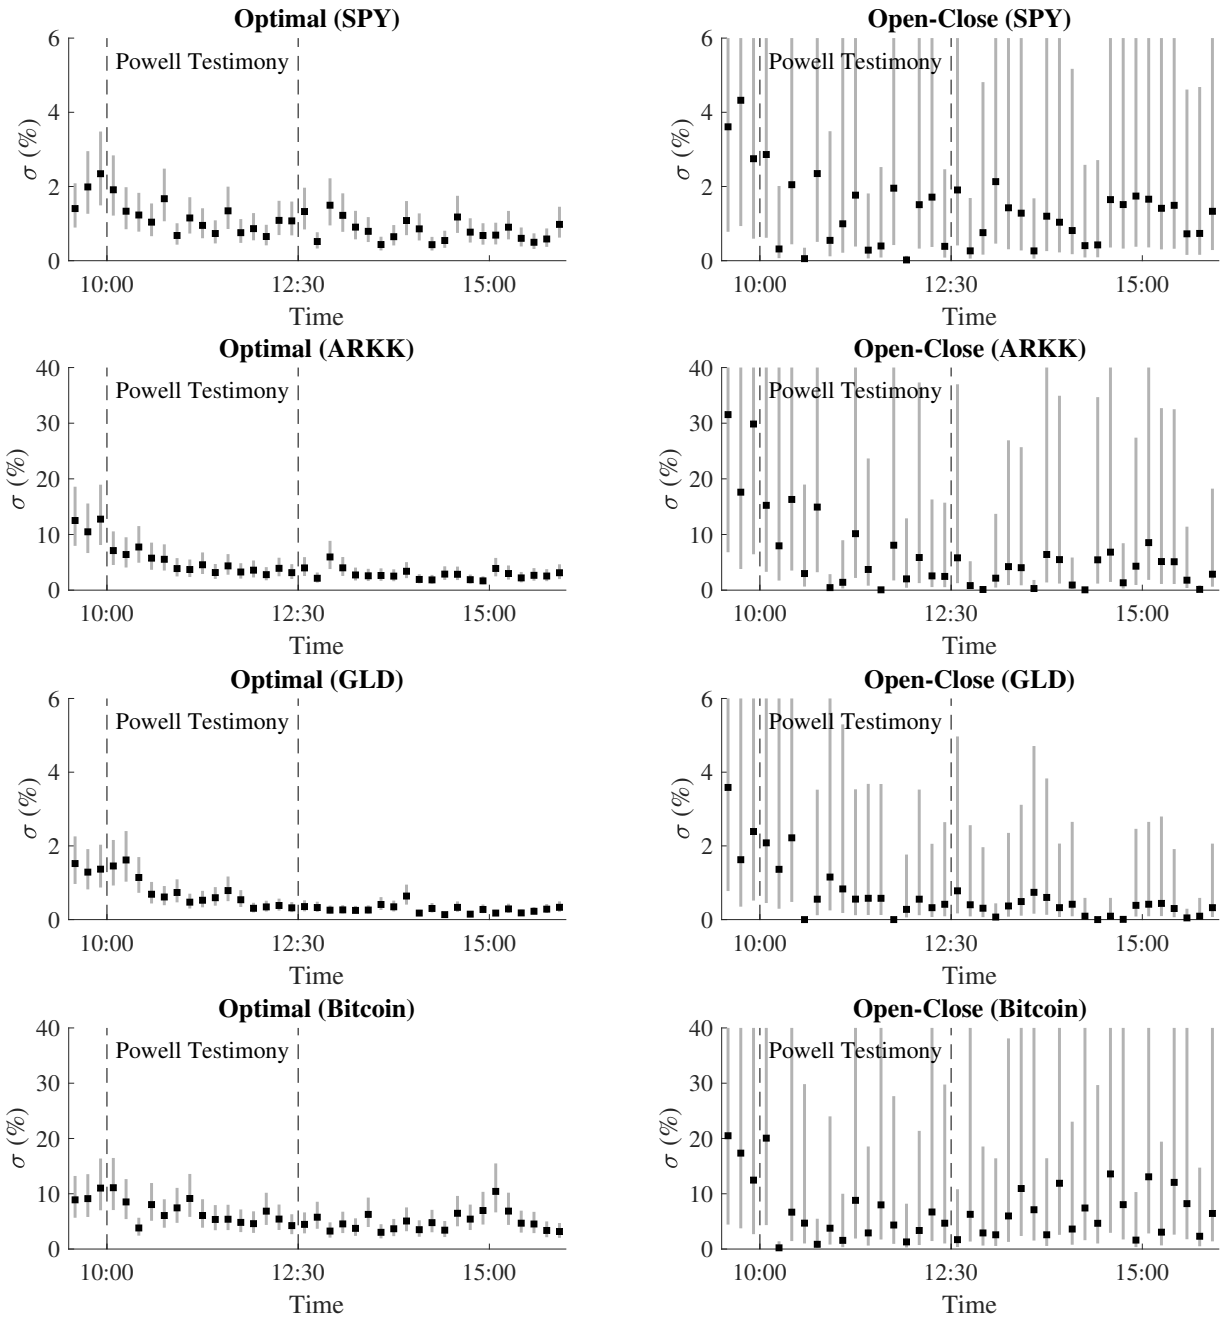

Figure 4: The figure plots volatility estimates (in daily percentage terms) of various assets on February 23, 2021, based on the optimal candlestick estimator (left) and the open-close estimator $\mu_1^{-1}|r_i|/\Delta_n^{1/2}$ (right) using data sampled at the 10-minute frequency. Confidence intervals are computed using the critical values in Table 1 at the 90% level.

compute the OK estimates and the 90% OKCIs for SPY, ARKK, GLD, and Bitcoin, and plot the results on the left column of Figure 4. For comparison, we also implement the estimator based only on the open-close return, that is, $\mu_1^{-1}|r_i|/\Delta_n^{1/2}$, which we refer to as the open-close estimator for simplicity. We remind the reader that the open-close estimator is also asymptotically unbiased, and its CI is asymptotically valid; however, it is far less efficient than the OK estimator as shown in Table 1. Numerical results for the open-close estimator are plotted on the right column of Figure 4. For ease of comparison, we plot the two methods' estimates under the same scale.

Looking at the left column of Figure 4, we see that the OK estimator generates sensible volatility estimates and relatively tight CIs for this broader collection of assets, like what we have seen for the 10-year Treasury yield in Figure 3. To demonstrate more clearly the merit of this optimal estimator, we compare it with the open-close estimator shown on the right column. The latter serves as a relevant benchmark because, in empirical work, it is not rare to see the absolute return being used as a proxy for volatility. The contrast between the two methods is striking. The estimated volatility paths based on the open-close estimator are clearly quite erratic. Indeed, volatility estimates for adjacent time intervals often bounce up and down in an apparently "random" fashion, and the CIs are generally very wide. This is very different from the much smoother estimated volatility path generated by the optimal estimator, even though smoothness across different time intervals is not forced upon either estimator as each candlestick is treated on its own in our estimation.

The erratic behavior of the open-close estimator may be largely explained by its statistical inefficiency, resulting in very noisy estimates. We also notice another undesirable feature of the open-close estimator, that is, it often generates implausibly low volatility estimate. For example, the open-close volatility estimates for SPY on the three 10-minute intervals between 10:20 to 10:50 are 2.049%, 0.055%, and 2.349% per day. In particular, the 10:30–10:40 estimate, 0.055% (or 0.869% in annualized terms), is very close to zero and is much lower than the adjacent ones. It is obvious that this estimate cannot be reflecting SPY's actual volatility level at that time. In contrast, the OK volatility estimates for these three intervals are 1.236%, 1.042%, and 1.677% (or 19.54%, 16.47%, and 26.51% in annualized terms), which are economically much more sensible and in line with SPY's typical volatility level.

To better understand the mechanics underlying this particular discrepancy between the open-close and the OK estimators, we examine the candlestick chart of SPY shown in Figure 5, where we highlight the 10:30–10:40 trading session for ease of discussion. The "†-shaped" candlestick for that trading session is indeed visibly distinct from the others: The open and close prices are virtually identical, whereas the high and low prices are far apart. This is why the open-close volatility estimate is close to zero, but the OK estimate remains to be large. This example highlights a "bug," in a practical sense, of the open-close volatility estimator, that is, it ignores essentially

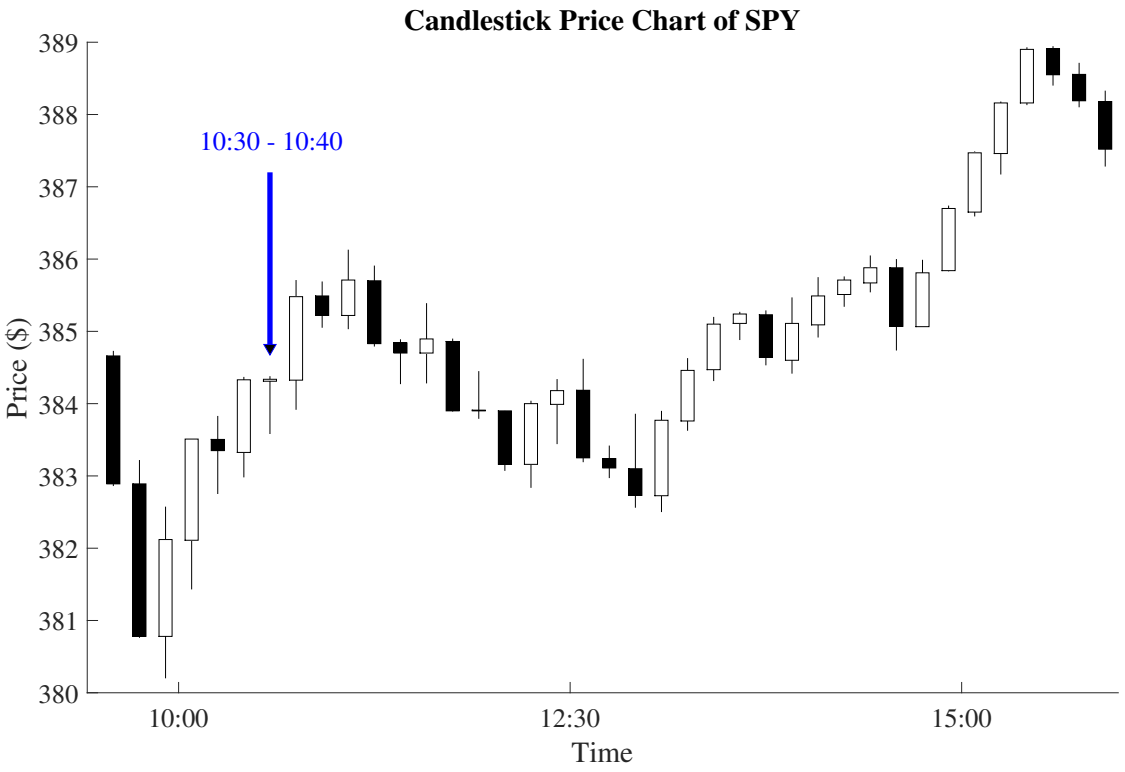

Figure 5: The figure plots the candlestick chart of the price of the SPY ETF on February 23, 2021, during regular trading hours. The arrow highlights a "dragonfly doji" pattern, which generates quite distinct volatility estimates for the optimal candlestick estimator and the open-close estimator.

all the price variation within the 10-minute trading session. Coincidentally, the cross-shaped candlestick is called a *doji* in technical analysis, which literally means "mistake" in Japanese (and the one highlighted in the figure is dubbed "dragonfly doji"). Although technical analysis is largely orthogonal to our econometric discussion, we do need to emphasize the practical relevance of the doji pattern for our candlestick-based volatility inference, in that it tends to generate insensible volatility estimate based only on the open-close return. The OK estimator, on the other hand, can better exploit the information from the whole candlestick and deliver more robust and efficient empirical estimates. The doji pattern is of course not unique to the SPY ETF, but is commonly seen in real data; for example, the 10-year Treasury yield charted in Figure 3 contains even more dojis.

In summary, the above illustration highlights the empirical usefulness of the proposed volatility inference procedure based on the OK estimator. We see that, for a broad variety of assets, our method can deliver economically sensible volatility estimates with adequate statistical accuracy. Importantly, the method can be easily implemented based on data available to retail investors in real time, and hence, may help an ordinary investor better manage risk and make more informed

investment decisions. We also demonstrate the clear benefit of adopting the OK estimator relative to a commonly used benchmark based on the absolute return. Here, the take-home message is clear: replacing the absolute return with the OK estimator for an equally simple but substantially more accurate volatility proxy.

5 Conclusion

With the goal to help ordinary investors conduct reliable inference on spot volatility, we propose an easy-to-implement econometric procedure based on readily accessible candlestick data. The proposed optimal candlestick (OK) volatility estimator is asymptotically unbiased and minimizes the asymptotic variance within a class of linear estimators. Under an approximate finite-sample approach, we construct asymptotically valid CIs for the spot volatility based on a nonstandard limiting distribution. We show that the candlestick-based estimator and CI are much more accurate than those based on high-frequency returns alone, and demonstrate their ability to generate economically sensible volatility estimates in practically relevant empirical settings. The proposed inference method is user-friendly, as it can be carried out manually on a basic calculator in real time. Our proposal may thus offer the average investor an "affordable" way to conduct volatility inference.

Appendix: Proofs

Proof of Theorem 1. Throughout the proof, we fix some $i \geq 1$ and $t \in I_{n,i}$, and use K to denote a generic positive constant. By a standard localization procedure, we can strengthen Assumption 1 by assuming that the conditions hold with $T_1 = \infty$ without loss of generality; see Section 4.4.1 in Jacod and Protter (2012) for details on the localization procedure. Finally, we note that under Assumption 1(i), the probability that the interval $I_{n,i}$ contains at least one price jump is $O(\Delta_n)$. Therefore, with probability approaching 1, $I_{n,i}$ does not contain any price jump. Since our calculation concentrates on this one interval, we can and will assume in the subsequent analysis that there are no price jumps without loss of generality.

Turning to the proof, we start with introducing some notation and preliminary estimates. We

rewrite $w_i = u_i - l_i$, where

$$\begin{aligned} u_i &\equiv \sup_{t\in I_{n,i}} \left(\int_{(i-1)\Delta_n}^{t} b_s ds + \int_{(i-1)\Delta_n}^{t} \sigma_s dW_s \right), \\ l_i &\equiv \inf_{t\in I_{n,i}} \left(\int_{(i-1)\Delta_n}^{t} b_s ds + \int_{(i-1)\Delta_n}^{t} \sigma_s dW_s \right), \end{aligned}$$

are respectively the upper and lower shadow of the ith candlestick. For ease of notation, we further denote

$$\begin{aligned} u_i' &\equiv \sigma_{(i-1)\Delta_n} \sup_{t\in I_{n,i}} \left(W_t - W_{(i-1)\Delta_n} \right), \quad l_i' \equiv \sigma_{(i-1)\Delta_n} \inf_{t\in I_{n,i}} \left(W_t - W_{(i-1)\Delta_n} \right), \\ r_i' &\equiv \sigma_{(i-1)\Delta_n} (W_{i\Delta_n} - W_{(i-1)\Delta_n}), \quad w_i' \equiv u_i' - l_i', \quad \hat{\sigma}_t'(\lambda) \equiv \frac{\lambda_1 |r_i'| + \lambda_2 w_i'}{\Delta_n^{1/2}}. \end{aligned}$$

It is easy to see that

$$\left| \int_{(i-1)\Delta_n}^{i\Delta_n} b_s ds \right| \leq \int_{(i-1)\Delta_n}^{i\Delta_n} |b_s| \, ds = O_p(\Delta_n) = o_p(\Delta_n^{1/2}). \tag{A.1}$$

Moreover, by the Burkholder–Davis–Gundy inequality (see, e.g., (2.1.34) in Jacod and Protter (2012)), we have

$$\begin{aligned} \mathbb{E}\left[\sup_{t\in I_{n,i}} \left| \int_{(i-1)\Delta_n}^{t} \left(\sigma_s - \sigma_{(i-1)\Delta_n} \right) dW_s \right|^2 \right] &\leq K\mathbb{E}\left[\int_{(i-1)\Delta_n}^{i\Delta_n} \left(\sigma_s - \sigma_{(i-1)\Delta_n} \right)^2 ds \right] \\ &\leq K\Delta_n^{1+2\kappa}, \end{aligned}$$

where the second inequality follows from $\mathbb{E}[(\sigma_s - \sigma_{(i-1)\Delta_n})^2] \leq K\Delta_n^{2\kappa}$. This estimate further implies

$$\sup_{t\in I_{n,i}} \left| \int_{(i-1)\Delta_n}^{t} \left(\sigma_s - \sigma_{(i-1)\Delta_n} \right) dW_s \right| = O_p\left(\Delta_n^{1/2+\kappa} \right) = o_p(\Delta_n^{1/2}). \tag{A.2}$$

By the triangle inequality, (A.1), and (A.2),

$$\left| |r_i| - |r_i'| \right| \leq \left| \int_{(i-1)\Delta_n}^{i\Delta_n} b_s ds \right| + \left| \int_{(i-1)\Delta_n}^{i\Delta_n} \left(\sigma_s - \sigma_{(i-1)\Delta_n} \right) dW_s \right| = o_p(\Delta_n^{1/2}). \tag{A.3}$$

In addition, we note that

$$\begin{aligned} \left| u_i - u_i' \right| &= \left| \sup_{t\in I_{n,i}} \left(\int_{(i-1)\Delta_n}^{t} b_s ds + \int_{(i-1)\Delta_n}^{t} \sigma_s dW_s \right) - \sigma_{(i-1)\Delta_n} \sup_{t\in I_{n,i}} \left(W_t - W_{(i-1)\Delta_n} \right) \right| \\ &\leq \sup_{t\in I_{n,i}} \left| \int_{(i-1)\Delta_n}^{t} b_s ds + \int_{(i-1)\Delta_n}^{t} \left(\sigma_s - \sigma_{(i-1)\Delta_n} \right) dW_s \right| \\ &\leq \int_{(i-1)\Delta_n}^{i\Delta_n} |b_s| \, ds + \sup_{t\in I_{n,i}} \left| \int_{(i-1)\Delta_n}^{t} \left(\sigma_s - \sigma_{(i-1)\Delta_n} \right) dW_s \right| \\ &= o_p(\Delta_n^{1/2}), \end{aligned}$$

where the first three lines are obvious, and the last line follows from (A.1) and (A.2). Similarly, we can derive $|l_i - l_i'| = o_p(\Delta_n^{1/2})$. Since $w_i = u_i - l_i$ and $w_i' = u_i' - l_i'$, we further deduce

$$w_i - w_i' = o_p(\Delta_n^{1/2}). \tag{A.4}$$

Combining (A.3) and (A.4), we deduce

$$\hat{\sigma}_t(\lambda) = \hat{\sigma}_t'(\lambda) + o_p(1). \tag{A.5}$$

From the definition of $\hat{\sigma}_t'(\lambda)$, we see that

$$\frac{\hat{\sigma}_t'(\lambda)}{\sigma_{(i-1)\Delta_n}} = \lambda_1 \zeta_{1,i} + \lambda_2 \zeta_{2,i},$$

where $\zeta_{1,i}$ and $\zeta_{2,i}$ are defined in Theorem 1. Since the σ process is càdlàg and bounded away from zero, we have for any $t \in I_{n,i}$,

$$\frac{\sigma_{(i-1)\Delta_n}}{\sigma_t} \xrightarrow{\mathbb{P}} 1.$$

Hence,

$$\frac{\hat{\sigma}_t'(\lambda)}{\sigma_t} = \lambda_1 \zeta_{1,i} + \lambda_2 \zeta_{2,i} + o_p(1). \tag{A.6}$$

The assertion of Theorem 1 then readily follows from (A.5) and (A.6). □

References

Andersen, T., and T. Bollerslev (1998): "Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts," *International Economic Review*, 39(4), 885–905.

Andersen, T. G., and T. Bollerslev (2018): *Volatility*. Edward Elgar Publishing.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003): "Modeling and Forecasting Realized Volatility," *Econometrica*, 71(2), pp. 579–625.

Bandi, F. M., and J. R. Russell (2008): "Microstructure Noise, Realized Variance, and Optimal Sampling," *The Review of Economic Studies*, 75(2), 339–369.

Barndorff-Nielsen, O., and N. Shephard (2002): "Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society*, Series B, 64, 253–280.

Barndorff-Nielsen, O., and N. Shephard (2004): "Power and Bipower Variation with Stochastic Volatility and Jumps," *Journal of Financial Econometrics*, 2, 1–37.

Barndorff-Nielsen, O., N. Shephard, and M. Winkel (2006): "Limit Theorems for Multipower Variation in the Presence of Jumps in Financial Econometrics," *Stochastic Processes and Their Applications*, 116, 796–806.

Bollerslev, T., J. Li, and Z. Liao (2021): "Fixed-k Inference for Volatility," *Quantitative Economics, Forthcoming*.

Bollerslev, T., J. Li, and Y. Xue (2018): "Volume, Volatility, and Public News Announcements," *Review of Economic Studies*, 85(4), 2005–2041.

Bollerslev, T., and V. Todorov (2011): "Estimation of Jump Tails," *Econometrica*, 79(6), 1727–1783.

Christensen, K., and M. Podolskij (2007): "Realized Range-based Estimation of Integrated Variance," *Journal of Econometrics*, 141(2), 323–349.

Comte, F., and E. Renault (1998): "Long Memory in Continuous Time Stochastic Volatility Models," *Mathematical Finance*, 8(4), 291–323.

Feller, W. (1951): "The Asymptotic Distribution of the Range of Sums of Independent Random Variables," *The Annals of Mathematical Statistics*, 22(3), 427 – 432.

Foster, D. P., and D. B. Nelson (1996): "Continuous Record Asymptotics for Rolling Sample Variance Estimators," *Econometrica*, 64(1), 139–174.

Garman, M. B., and M. J. Klass (1980): "On the Estimation of Security Price Volatilities from Historical Data," *The Journal of Business*, 53(1), 67–78.

Jacod, J., and P. Protter (2012): *Discretization of Processes*. Springer.

Kristensen, D. (2010): "Nonparametric Filtering of the Realized Spot Volatility: A Kernel-Based Approach," *Econometric Theory*, 26(1), 60–93.

Li, Z. M., and O. B. Linton (2020): "A ReMeDI for Microstructure Noise," Discussion paper, University of Cambridge.

Lucca, D. O., and E. Moench (2015): "The Pre-FOMC Announcement Drift," *Journal of Finance*, 70(1), 329–371.

Mancini, C. (2001): "Disentangling the Jumps of the Diffusion in a Geometric Jumping Brownian Motion," *Giornale dell'Istituto Italiano degli Attuari*, LXIV, 19–47.

Nakamura, E., and J. Steinsson (2018): "Identification in Macroeconomics," *Journal of Economic Perspectives*, 32(3), 59–86.

Nison, S. (2001): *Japanese Candlestick Charting Techniques*. Prentice Hall Press, 2nd edn.

Parkinson, M. (1980): "The Extreme Value Method for Estimating the Variance of the Rate of Return," *The Journal of Business*, 53(1), 61–65.

Zhang, L., P. A. Mykland, and Y. Aït-Sahalia (2005): "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data," *Journal of the American Statistical Association*, 100(472), 1394–1411.