Estimating Correlations and Reading the News by Candlesticks*

Tim Bollerslev† Jia Li‡ Qiyuan Li§ Yifan Li¶

May 7, 2026

Abstract

We develop a new class of nonparametric spot correlation estimators based on high-frequency candlestick data comprising open, close, high, and low prices over short intraday time intervals. The new estimators augment the conventional return-based estimator with a signed asymmetry feature. We establish consistency and mixed normality in a traditional large-estimation-window in-fill asymptotic setting. We also develop a fixed-k asymptotic framework and confidence intervals that achieve near-nominal finite-sample coverage with very few candlesticks. Simulations further demonstrate substantial mean-squared-error reductions over return-based estimators and robustness to market microstructure noise. Relying on one-minute candlesticks, we find that FOMC announcements systematically raise stock-bond correlations, consistent with a discount-rate channel interpretation of the news, while President Trump's tariff-related posts affect more dispersed changes in the correlations. The empirical results highlight the value of the new candlestick-based spot estimators for capturing rapid changes in asset comovements around economic news and illuminating the economic transmission mechanisms at work.

Keywords: spot correlation estimation, high-frequency data, candlesticks, fixed-k inference, news announcements, stock-bond comovements, economic shocks.
JEL Codes: C14, C58, G12, G14.

*We would like to thank participants at various seminars and conferences for useful comments. Jia Li's research is partially supported by a Singapore Ministry of Education tier-2 grant (T2EP40124-0021).

†Department of Economics, Duke University, Durham, NC 27708, and NBER; e-mail: boller@duke.edu.

‡School of Economics, Singapore Management University; e-mail: jiali@smu.edu.sg.

§Faculty of Business and Economics, University of Hong Kong; e-mail: qiyuanli@hku.hk.

¶Accounting & Finance Division, University of Manchester; e-mail: yifan.li@manchester.ac.uk

1 Introduction

We present a new class of nonparametric, consistent spot-correlation estimators based on high-frequency candlesticks, comprising the open, close, high, and low prices over short intraday time intervals, along with a framework for conducting reliable inference about sudden changes in the estimated correlations. Estimating the stock-bond spot correlations around a select set of recent FOMC announcements, we find that the correlations generally increase significantly immediately after the press release, consistent with a traditional discount-rate interpretation of the news. By comparison, while some of the recent tariff-related posts by President Trump clearly make volatility great again, they affect more dispersed changes in the correlations, reflective of their more ambiguous economic interpretations.

Placing the paper in the broader context of the literature, correlations of asset returns have long been known to play a critical role in portfolio construction, risk management, and derivatives pricing alike (see, e.g., the foundational works by Markowitz (1952), Sharpe (1964), and Ross (1976)). Correlations have also proven crucially important for interpreting macroeconomic news and understanding economic transmission mechanisms (see, e.g., the early work by Chen et al. (1986) and Campbell (1991)). A large body of empirical research has also documented that the conditional correlations of asset returns are strongly time-varying (see, e.g., the extensive evidence in Engle (2009)), and that correlations can change sharply in response to certain policy shocks and news announcements (see, e.g., Cutler et al. (1989) and Rigobon and Sack (2004)). While much of the earlier work on time-varying covariances and correlations relied on parametric models and daily or coarser frequency data (see, e.g., Bollerslev et al. (1988)), the increased availability of high-frequency intraday data has changed the approach to measuring volatility-type objects, with so-called realized measures now commonly used for consistently estimating the ex-post latent integrated covariances and correlations over longer non-trivial time intervals, typically a day (e.g., Barndorff-Nielsen and Shephard, 2004; Hayashi and Yoshida, 2005; Andersen et al., 2006; Aït-Sahalia et al., 2010; Aït-Sahalia and Xiu, 2016; Varneskov, 2016; Shephard and Xiu, 2017). However, when analyzing policy events and their economic impact, the most informative variation is often concentrated within a few minutes around the news releases (as first documented by Edrington and Lee (1993) and Andersen et al. (2007) among others). On the other hand, daily correlation estimates invariably average co-movements across pre- and post-event windows, thereby attenuating precisely the variation one would like to measure. To address this issue, a more recent growing body of work has sought to establish high-frequency-return-based procedures for consistently estimating spot covariances and spot correlations at a given point in time (e.g., Fan and Wang, 2008; Kristensen, 2010; Jacod and Protter, 2012; Vetter, 2012; Bibinger and Reiß, 2014; Bibinger et al., 2019; Jacod et al., 2021; Liu and Liu, 2022; Hansen and Luo, 2023). Our new class of spot correlation estimators improves on these existing procedures by incorporating the additional information inherent in high-frequency candlesticks.

Candlestick data have been available at daily and lower frequencies for centuries. They are

now also readily available from a range of different trading platforms and data vendors at high intraday frequencies for a host of different assets (see, e.g., the discussion in Bollerslev et al. (2026)). Intuitively, the information encoded in the high and low candlestick prices provides additional insights about the realized intraperiod price paths that should be useful for improved estimation of higher moments. This idea also underlies the classical range-based volatility estimators of Parkinson (1980), Garman and Klass (1980), and Rogers and Satchell (1991). Recent work by Li et al. (2024), Bollerslev et al. (2024a), and Li et al. (2025) has also further extended this same logic to high-frequency candlestick-based estimation of univariate spot and integrated variances. However, an analogous formal high-frequency inference framework for spot correlations has hitherto been lacking.

The existing candlestick-based correlation statistics of Rogers and Zhou (2008) and Popov (2016) are derived within a parametric Brownian model calibrated to daily data with constant volatility and constant correlation. As such, these existing procedures are not applicable and do not allow for reliable inference in a high-frequency dynamic setting. Most closely related to the present paper, concurrent work by Simsek (2026) has recently studied candlestick-based spot covariance estimation in a high-frequency setting. However, in contrast to the estimator developed here, which is explicitly designed for consistently estimating spot correlations, the analysis in Simsek (2026) is primarily focused on estimating spot betas. The analysis does not provide a consistent estimator for the scale-free spot correlation coefficient, nor does it address the question of how to conduct reliable high-frequency inference for spot correlations.

Our new class of candlestick-based spot correlation estimators is based on a simple convex combination of the usual open-to-close return cross-product and a cross-product that summarizes the shapes of the candlesticks in terms of their upper and lower wicks, along with a nonlinear bias correction that converts the raw statistics into a consistent spot correlation estimator under a standard infill asymptotic setting. In addition to proving consistency, we also establish a large-k, where k denotes the number of candlesticks used in the local estimation window, asymptotic Gaussian limit theory under a general Itô semimartingale model allowing for stochastic and possibly rough volatility, time-varying correlations, leverage effects, and jumps. We also identify a (nearly) minimax-variance estimator within our general class of consistent spot estimators.

To allow for more reliable inference in practical empirical applications with relatively few candlesticks and narrow estimation windows, we further extend the conventional large-k asymptotic theory to a fixed-k setting. Our fixed-k setting still entails the notion of an increasing sampling frequency, but the number of candlesticks used for the estimation is treated as fixed. The resulting fixed-k limit is nonstandard, with the limiting distribution depending nonlinearly on the unknown correlation. This contrasts with earlier fixed-k analyses in the high-frequency literature, where the limiting distribution typically only depends on the unknown nuisance parameter through a scale normalization and can therefore be pivotalized. To address this challenge, we instead construct Anderson–Rubin type confidence sets by inverting two-sided

tests of the spot correlation parameter.

Our finite-sample simulation results confirm that the resulting fixed-k confidence intervals attain near-nominal coverage, even when the local estimation window contains very few candlesticks. By comparison, the more conventional large-k Gaussian confidence intervals require substantially longer estimation windows to approach their nominal levels. The simulations also clearly demonstrate a sizable reduction in the mean-square-error for the new candlestick-based estimator compared to the traditional return-based estimator, along with much more accurate identification of the correct sign. Even though our formal theoretical analyses abstract from market frictions, the simulations also further demonstrate that our proposed new candlestick-based estimator remains remarkably robust to empirically realistic levels of market microstructure noise.

To illustrate the practical usefulness and additional insights afforded by the new estimator, we revisit the much-studied correlation between stocks and bonds. More specifically, we analyze the comovements of U.S. stocks and bonds around a select set of recent monetary and trade-policy-related announcements, using one-minute candlesticks for the S&P 500 SPY ETF and the long-duration Treasury Bond TLT ETF. For the FOMC announcements, we find that the estimated spot correlations tend to be close to zero before the release time and then typically become significantly positive shortly thereafter. This systematic pattern naturally points to a discount-rate channel at work, where revisions to the expected policy path move equities and long-duration bonds in the same direction.[1] By comparison, the unscheduled tariff-related posts by President Trump have a more heterogeneous impact on the stock-bond correlations depending on whether the changes in the tariffs are seen as escalating or de-escalating the trade wars, and in turn, which type(s) of economic shock(s) dominate.

The remainder of the paper is organized as follows. Section 2 introduces the new class of consistent estimators, develops corresponding large-k asymptotic distributions and fixed-k confidence intervals, and establishes our proposed minimax variance estimator. Section 3 reports the results from a series of simulation experiments designed to analyze the finite sample performance of the estimators and robustness to market microstructure noise. Section 4 presents our empirical estimates of the stock-bond spot correlations around select FOMC announcements and tariff changes. Section 5 concludes. All proofs are collected in the Appendix. Additional supportive material is provided in an Online Supplemental Appendix.

[1] This interpretation is also consistent with Bauer and Swanson (2023), who argue that much of the apparent information content of FOMC surprises reflects a "Fed response to news" channel. The findings do not support the information channel emphasized by Nakamura and Steinsson (2018) and others, under which favorable news about real fundamentals would tend to induce negative stock-bond comovements.

2 Candlestick-Based Inference for Spot Correlation

2.1 Econometric setting and assumptions

We assume that the bivariate log-price process $P = (P_t)_{t\geq 0}$ is defined on a filtered probability space $\big(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}\big)$ and takes the form of a bivariate Itô semimartingale:

$$P_t = P_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + J_t, \tag{1}$$

where $W = (W^{(1)}, W^{(2)})^\top$ is a bivariate standard Brownian motion, $b = (b_t)_{t\geq 0}$ is the bivariate spot drift process, $\sigma = (\sigma_t)_{t\geq 0}$ is the 2-by-2 spot covolatility matrix process, and J is a bivariate pure-jump process collecting all jumps in P. The processes b and σ are assumed to be càdlàg and adapted. We denote the spot covariance matrix as:

$$\Sigma_t = \sigma_t \sigma_t^\top = \begin{pmatrix} (\sigma_t^{(1)})^2 & \bullet \\ \sigma_t^{(1)}\sigma_t^{(2)}\rho_t & (\sigma_t^{(2)})^2 \end{pmatrix},$$

where $\sigma_t^{(j)}$ and ρ_t, respectively, denote the spot volatility for asset $j \in \{1, 2\}$ and the spot correlation of the two assets. The object of interest is the spot correlation at a given point in time τ, ρ_τ. Define $\text{atanh}(x) = \frac{1}{2}\log\big(\frac{1+x}{1-x}\big)$ as the inverse hyperbolic tangent function. We maintain the following regularity conditions on the underlying processes.

Assumption 1. *Suppose P is of the form in* (1) *and that there exists a sequence $(T_m)_{m\geq 1}$ of stopping times increasing to infinity such that the following conditions hold for each $m \geq 1$: (i) for all $t \in [0, T \wedge T_m]$, $\|b_t\| + \sum_{j=1}^2 \big(|\sigma_t^{(j)}| + |(\sigma_t^{(j)})^{-1}|\big) + |\text{atanh}(\rho_t)| + F_t(\mathbb{R}^2 \setminus \{0\}) < K_m$, where $F_t(z)$ is the spot Lévy measure of J_t; (ii) for some constant $\kappa > 0$, $\mathbb{E}[\|\sigma_{s\wedge T_m} - \sigma_{t\wedge T_m}\|^2] < K_m|t-s|^{2\kappa}$ for all $t, s \in [0, T]$, where $\|\cdot\|$ denotes the Euclidean norm.*

Assumption 1 naturally generalizes the univariate assumptions in Bollerslev et al. (2024a) and Li et al. (2024) to a multivariate setting. It is also similar to the conditions in Bollerslev et al. (2024b) used in the context of spot regressions. These mild regularity conditions imply that: (1) the spot volatility processes $\sigma^{(j)}$ and the spot correlation process ρ are locally κ-Hölder continuous under the L_2-norm; and (2) ρ is locally bounded away from ± 1, precluding perfectly correlated asset pairs. These conditions provide sufficient regularity conditions for the spot estimation of these quantities under infill asymptotics. They can be readily verified for Itô semimartingale and long-memory stochastic volatility models. The parameter κ controls the "roughness" of the spot volatility processes, with $\kappa < 1/2$ allowing for "rough paths," as advocated by Gatheral et al. (2018) among others.

We consider a standard infill asymptotic setting under which high-frequency observations of P are sampled at discrete times $0, \Delta_n, 2\Delta_n, \ldots$, over a fixed interval $[0, T]$, with $\Delta_n \to 0$ as $n \to \infty$. The ith candlestick observation consists of the component-wise open, close, high and low prices, formally represented as $P_{(i-1)\Delta_n}, P_{i\Delta_n}, \sup_{t\in\mathcal{I}_{n,i}} P_t$, and $\inf_{t\in\mathcal{I}_{n,i}} P_t$, respectively,

where $\mathcal{I}_{n,i} \equiv [(i-1)\Delta_n, i\Delta_n]$. We further denote the corresponding open-to-close, high, and low returns as:

$$C_{n,i} \equiv P_{i\Delta_n} - P_{(i-1)\Delta_n}, \quad H_{n,i} \equiv \sup_{t\in\mathcal{I}_{n,i}} P_t - P_{(i-1)\Delta_n}, \quad L_{n,i} \equiv \inf_{t\in\mathcal{I}_{n,i}} P_t - P_{(i-1)\Delta_n}.$$

To construct the spot correlation estimator, we consider a local estimation window $\mathcal{I}_n = \bigcup_{i=1}^{k} \mathcal{I}_{n,i_n(\tau)+i}$ comprising k observations immediately following time τ, where $i_n(\tau)$ refers to the interval index such that $\tau \in \mathcal{I}_{n,i_n(\tau)}$. For notational convenience, we normalize $i_n(\tau) \equiv 0$, so that the inference hinges on the candlestick observations $(C_{n,i}, H_{n,i}, L_{n,i})_{1\leq i\leq k}$, where the window size k plays the role of a "bandwidth" parameter. We analyze two distinct asymptotic frameworks: a large-k setting in which $k \to \infty$, and a fixed-k setting where k is treated as a constant. To explicitly distinguish these two scenarios, we denote k_n instead of k under the large-k setting, thereby emphasizing its dependence on the asymptotic stage n.

The large-k scheme has conventionally been used for nonparametric inference about spot covariances and correlations (see, e.g., Foster and Nelson (1996); Jacod and Protter (2012)). The assumption that $k_n \to \infty$ allows us to employ classical laws of large numbers and central limit theorems to establish the consistency and asymptotic mixed normality for our new general class of spot correlation estimators. This in turn allows us to construct asymptotically valid confidence intervals for ρ_τ as well. In practice, however, the window size k is often not "large enough" for the large-k asymptotic theory to "kick in," resulting in nontrivial finite-sample size distortions for corresponding test statistics and confidence intervals. We therefore augment the conventional asymptotic characterization of the estimator with one based on a fixed-k scheme, in which we directly exploit the finite-sample distributions of candlestick features, obviating the need to invoke asymptotic mixed normality through the aggregation of a "large" number of candlestick observations. The resulting fixed-k confidence intervals substantially improve finite-sample coverage properties. We begin our discussion by considering the more conventional large-k setting.

2.2 A class of consistent spot correlation estimators

The key idea behind our new class of estimators is to augment the conventional return-based correlation statistic with information encoded in candlestick shapes. To this end, define the signed asymmetry feature:

$$A_{n,i}^{(j)} \equiv H_{n,i}^{(j)} + L_{n,i}^{(j)} - C_{n,i}^{(j)}, \tag{2}$$

which measures the difference between the lengths of the upper and lower wicks of the i-th candlestick for asset j. Inspired by the daily candlestick-based covariance estimator proposed by Rogers and Zhou (2008), define:

$$\hat{\gamma}_n^{(j,l)}(\lambda) \equiv \frac{1}{k_n}\sum_{i=1}^{k_n}\big\{(1-\lambda)C_{n,i}^{(j)}C_{n,i}^{(l)} + \lambda A_{n,i}^{(j)}A_{n,i}^{(l)}\big\}, \qquad j,l \in \{1,2\}, \tag{3}$$

where $\lambda \in [0,1]$. The $\hat{\gamma}_n^{(1,2)}(\lambda)$ statistic, based on the convex combination of the return cross-product and the cross-product of the asymmetry features, is naturally interpreted as a covariance estimator. In fact, the daily covariance estimator proposed by Rogers and Zhou (2008) is obtained by fixing the weight parameter $\lambda_{RZ} \equiv (4 - 4\log 2)^{-1} \approx 0.815$ and $k_n \equiv 1$. Now consider the "correlation-like" statistic:

$$\hat{\varrho}_n(\lambda) \equiv \frac{\hat{\gamma}_n^{(1,2)}(\lambda)}{\sqrt{\hat{\gamma}_n^{(1,1)}(\lambda)\hat{\gamma}_n^{(2,2)}(\lambda)}}. \tag{4}$$

Since $\hat{\gamma}_n^{(j,j)}(\lambda) > 0$ almost surely for $\lambda \in [0,1]$, $\hat{\varrho}_n(\lambda) \in [-1,1]$ by construction. Accordingly, $\hat{\varrho}_n(\lambda)$ could seemingly be used as a correlation estimator. However, as discussed further below, unlike the traditional return-based correlation estimator obtained for $\lambda = 0$, for $\lambda \neq 0$ $\hat{\varrho}_n(\lambda)$ is generally asymptotically biased, necessitating the need for a bias adjustment.

Although the estimators defined above are structurally close to the daily statistics previously proposed in the literature, the econometric setting is fundamentally different. The estimator in Rogers and Zhou (2008), and the subsequent refinements by Popov (2016), are based on a fully parametric Brownian model with constant volatility and a constant correlation coefficient. Correspondingly, the estimation problem amounts to determining a single parameter within that specific parametric setting. By contrast, we target the spot correlation ρ_τ process at an arbitrary point in time τ within a general Itô semimartingale model that admits stochastic and possibly rough volatility, time-varying correlation, leverage, and jumps. We also rely on the asymptotic notion of increasingly finer sampled candlesticks to more accurately mimic a high-frequency setting.[2]

To appreciate the aforementioned bias of the "natural" correlation estimator defined in (4), consider a "toy" model in which $(B_t^{(1)}, B_t^{(2)})_{t\geq 0}$ is a bivariate Brownian motion with unit volatility and constant correlation ρ. For each $i = 1, \ldots, k_n$ and $j \in \{1,2\}$ define:

$$\zeta_{C,i}^{(j)} \equiv B_i^{(j)} - B_{i-1}^{(j)}, \qquad \zeta_{H,i}^{(j)} \equiv \sup_{t\in[i-1,i]} B_t^{(j)} - B_{i-1}^{(j)}, \qquad \zeta_{L,i}^{(j)} \equiv \inf_{t\in[i-1,i]} B_t^{(j)} - B_{i-1}^{(j)}, \tag{5}$$

and

$$\zeta_{A,i}^{(j)} \equiv \zeta_{H,i}^{(j)} + \zeta_{L,i}^{(j)} - \zeta_{C,i}^{(j)}. \tag{6}$$

It follows from Rogers and Shepp (2006) that

$$\mathbb{E}[\zeta_{H,1}^{(1)}\zeta_{H,1}^{(2)}] = c(\rho) \equiv \cos\big(\alpha(\rho)\big) \int_0^\infty \frac{\cosh\big(\nu\alpha(\rho)\big)}{\sinh(\nu\pi/2)} \tanh\big(\nu\gamma(\rho)\big) d\nu, \tag{7}$$

[2]Further along these lines, the Brownian-motion identities that we rely on in equations (7) and (8) below should be read as local infill approximations around τ, not as exact moment conditions of the data-generating process. As a consequence, the bias correction we extract from these identities operates within a high-frequency nonparametric estimation problem, and the consistency, mixed-normal limiting distribution, and fixed-k inference procedures developed below simply have no counterpart in the daily parametric setting.

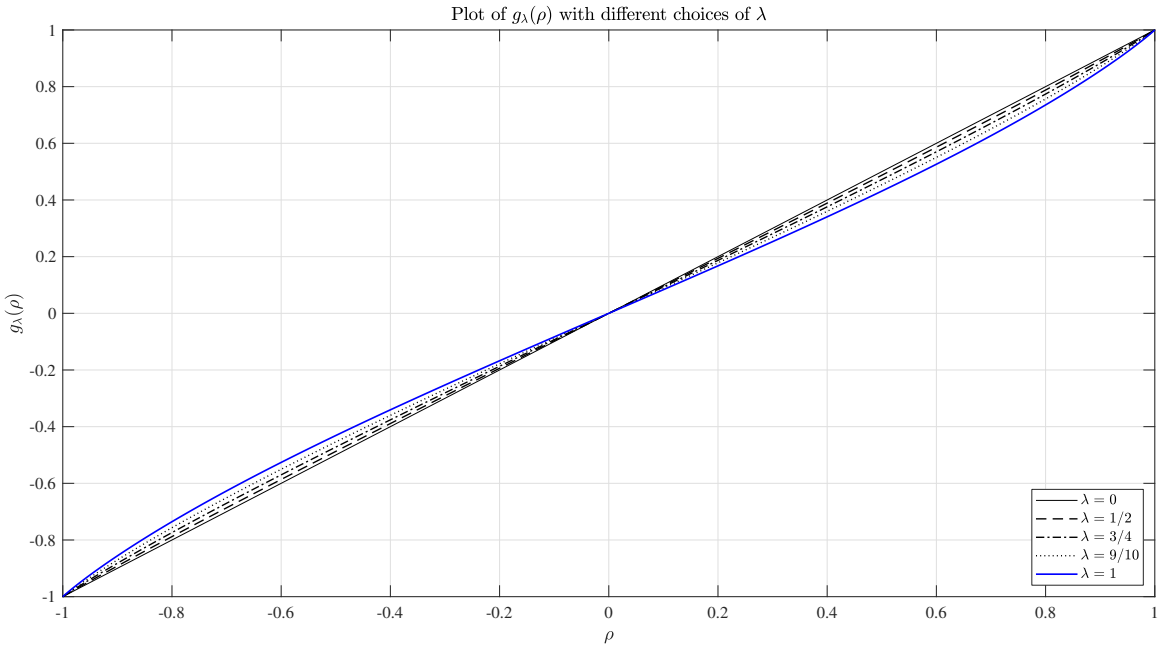

Figure 1: The figure plots $g_\lambda(\rho)$ for $\rho \in [-1, 1]$, and $\lambda \in \{0, 1/2, 3/4, 9/10, 1\}$.

where $\alpha(\rho) = \arcsin(\rho) \in (-\pi/2, \pi/2)$ and $\gamma(\rho) = \alpha(\rho)/2 + \pi/4$. Further using results from Rogers and Zhou (2008) and Popov (2016), it is possible to show:

$$g_\lambda(\rho) \equiv \frac{(1-\lambda)\mathbb{E}[\zeta_{C,1}^{(1)}\zeta_{C,1}^{(2)}] + \lambda\mathbb{E}[\zeta_{A,1}^{(1)}\zeta_{A,1}^{(2)}]}{(1-\lambda)\mathbb{E}[(\zeta_{C,1}^{(1)})^2] + \lambda\mathbb{E}[(\zeta_{A,1}^{(1)})^2]} = \frac{\rho + 2\lambda\big(c(\rho) - c(-\rho) - \rho\big)}{1 + 2\lambda(1 - 2\log 2)}. \tag{8}$$

Hence, under this simple "limit experiment," $\hat{\varrho}_n(\lambda)$ converges to $g_\lambda(\rho)$, rather than the true correlation ρ. To illustrate, Figure 1 plots the $g_\lambda(\cdot)$ functions for various values of λ. As the figure shows, $\hat{\varrho}_n(\lambda)$ is generally biased toward zero, with the bias depending on the underlying correlation in a nonlinear fashion. The maximum bias obtains for $\rho \approx \pm 0.6$.

To correct for this bias, in place of $\hat{\varrho}_n(\lambda)$ we instead propose the bias-corrected estimator:[3]

$$\hat{\rho}_n(\lambda) \equiv g_\lambda^{-1}\big(\hat{\varrho}_n(\lambda)\big). \tag{9}$$

Theorem 1, below, establishes the consistency and asymptotic mixed normality of this estimator. For ease of presentation, we define:

$$V_\lambda(\rho) \equiv \mathsf{Var}\Big(\zeta_\lambda^{(1,2)} - \frac{g_\lambda(\rho)}{2}(\zeta_\lambda^{(1,1)} + \zeta_\lambda^{(2,2)})\Big)$$

where

$$\zeta_\lambda^{(j,l)} \equiv \frac{(1-\lambda)\zeta_C^{(j)}\zeta_C^{(l)} + \lambda\zeta_A^{(j)}\zeta_A^{(l)}}{1 + 2\lambda(1 - 2\log 2)}.$$

[3]Note that for all $\lambda \in [0, 1]$, $g_\lambda(\cdot)$ is a strictly increasing odd bijection of $[-1, 1]$. It is also continuously differentiable on $(-1, 1)$ with derivatives bounded above and away from 0, see Lemma A.2. As a result, $g_\lambda^{-1}(\cdot)$ and hence $\hat{\rho}_n(\lambda)$ are well-defined.

Theorem 1. *Let $\lambda \in [0,1]$ and $k_n \asymp \Delta_n^{-\beta}$ for some $\beta \in (0,1)$. Under Assumption 1, it holds that*

$$\hat{\rho}_n(\lambda) \xrightarrow{\mathbb{P}} \rho_\tau. \tag{10}$$

If we further have $\beta < 2\kappa/(2\kappa+1)$, then

$$\sqrt{k_n}\big(\hat{\rho}_n(\lambda) - \rho_\tau\big) \xrightarrow{\mathcal{L}|\mathcal{F}_\tau} g'_\lambda(\rho_\tau)^{-1}\sqrt{V_\lambda(\rho_\tau)}Z, \tag{11}$$

where $\xrightarrow{\mathcal{L}|\mathcal{F}_\tau}$ denotes convergence in $\mathcal{F}_\tau$-conditional law (see Definition 1 in Appendix), $g'_\lambda(\rho_\tau)$ is the derivative of $g_\lambda(\rho)$ evaluated at ρ_τ, and Z is a standard Gaussian random variable defined on an extension of $\big(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}\big)$ independent of $\mathcal{F}$.

The central limit theorem in (11) is formulated in terms of convergence in $\mathcal{F}_\tau$-conditional law. This is a stronger convergence than ordinary convergence in distribution (see, e.g., Eagleson (1975), Barndorff-Nielsen et al. (2008, Definition 5), and Li and Xiu (2016, Definition A1) for formal discussions). This stronger mode of convergence is needed because the $\mathcal{F}_\tau$-conditional asymptotic variance is itself random, reflecting the randomness of the spot correlation ρ_τ. However, as an immediate corollary to the theorem, we obtain the following feasible central limit theorem.

Corollary 1. *Under the same conditions as in Theorem 1, we have:*

$$\frac{\sqrt{k_n}\big(\hat{\rho}_n(\lambda) - \rho_\tau\big)}{g'_\lambda\big(\hat{\rho}_n(\lambda)\big)^{-1}\sqrt{V_\lambda\big(\hat{\rho}_n(\lambda)\big)}} \xrightarrow{\mathcal{L}} \mathcal{N}(0,1).$$

Corollary 1 readily enables the construction of asymptotically valid tests and confidence intervals for ρ_τ. Specifically, for any $\alpha \in (0,1)$, a $(1-\alpha)$-level confidence interval for ρ_τ may be constructed as:

$$\left[\hat{\rho}_n(\lambda) - z_{1-\alpha/2}\frac{\sqrt{V_\lambda\big(\hat{\rho}_n(\lambda)\big)}}{\sqrt{k_n}g'_\lambda\big(\hat{\rho}_n(\lambda)\big)},\ \ \hat{\rho}_n(\lambda) + z_{1-\alpha/2}\frac{\sqrt{V_\lambda\big(\hat{\rho}_n(\lambda)\big)}}{\sqrt{k_n}g'_\lambda\big(\hat{\rho}_n(\lambda)\big)}\right], \tag{12}$$

where $z_{1-\alpha/2}$ denotes the $1-\alpha/2$ quantile of the standard normal distribution. The functions $g'_\lambda(\cdot)$ and $V_\lambda(\cdot)$ can both easily be evaluated numerically.

2.3 A (nearly) minimax estimator

The limiting results in Theorem 1 and Corollary 1 hold for any value of $\lambda \in [0,1]$, raising the question of how to choose λ in practice. A natural answer is to choose the value of λ that minimizes the conditional asymptotic variance of $\hat{\rho}_n(\lambda)$:

$$\widetilde{V}(\lambda;\rho_\tau) \equiv g'_\lambda(\rho_\tau)^{-2}\cdot V_\lambda(\rho_\tau).$$

However, since $\widetilde{V}(\lambda;\rho_\tau)$ depends not only on λ, but also on the unknown spot correlation ρ_τ, the resulting optimal choice of λ would depend on the unknown ρ_τ as well. A possible solution

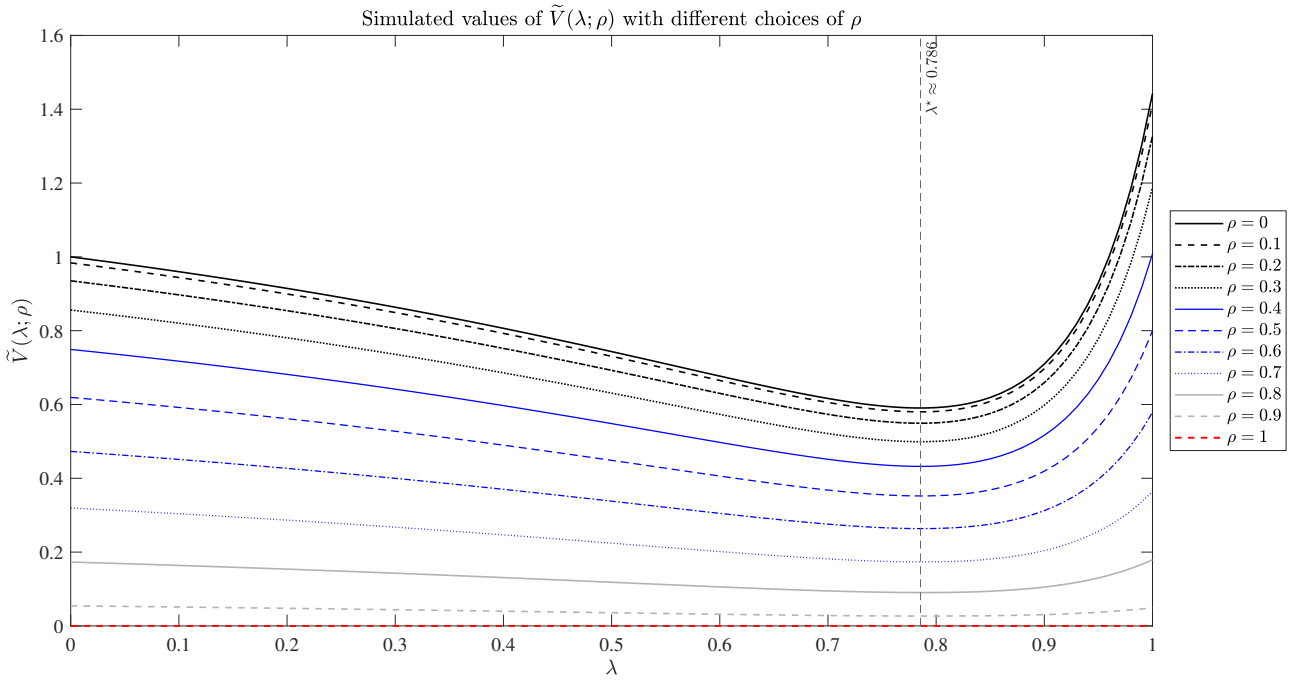

Figure 2: The figure shows the simulated values of $\widetilde{V}(\lambda;\rho)$ for $\lambda \in [0,1]$ and different values of ρ ranging from 0 to 1. The vertical dashed line indicates $\lambda^* = \arg\min_{\lambda\in[0,1]} \widetilde{V}(\lambda;0) \approx 0.786$.

would be to use a plug-in rule. That is, select λ to minimize the asymptotic variance $\widetilde{V}(\lambda;\rho_\tau)$ after replacing the unknown spot correlation by a preliminary estimate. Interestingly, such a preliminary step turns out to be largely unnecessary, because the optimal choice of λ is not very sensitive to the value of the spot correlation.

To illustrate this lack of sensitivity to ρ_τ, Figure 2 plots the numerically evaluated functions $\lambda \mapsto \widetilde{V}(\lambda;\rho)$ for $\rho \in [0,1]$.[4] As the figure shows, the curves corresponding to different values of ρ attain their minima in nearly the same region of λ. Consequently, as a simple rule of thumb, we propose choosing λ by minimizing $\widetilde{V}(\lambda;\rho)$ at $\rho = 0$. That is:

$$\widetilde{V}(\lambda;0) = \frac{(1-\lambda)^2 + \lambda^2(3-4\log 2)^2}{\big(1+2\lambda(2c'(0)-1)\big)^2},$$

where $c'(0) = 1/2 - 2/\pi^2 \approx 0.297$ denotes the derivative of $c(\rho)$ evaluated at $\rho = 0$ (see (A.5) in the Appendix). The resulting minimizer may be conveniently expressed as:

$$\lambda^* \equiv \underset{\lambda\in[0,1]}{\arg\min}\, \widetilde{V}(\lambda;0) = \frac{4c'(0)-1}{4\big((2\log 2-1)(2\log 2-2)+c'(0)\big)} \approx 0.786. \tag{13}$$

Plugging into the above formula, the corresponding minimized asymptotic variance for $\rho = 0$ equals $\widetilde{V}(\lambda^*;0) \approx 0.590$. By comparison, the return-based estimator has asymptotic variance equal to $\widetilde{V}(0;0) = 1$ for $\rho = 0$. Hence, incorporating candlestick information reduces the asymptotic variance by about 40% relative to the return-based spot correlation estimator. As shown in Figure 2, this same value of λ also approximately minimizes $\widetilde{V}(\lambda;\rho)$ for other values of ρ, supporting its use as a simple universal tuning parameter.

[4]Restricting attention to $\rho \in [0,1]$ is without loss of generality, since $\widetilde{V}(\lambda;\rho)$ is an even function of ρ.

The numerical results in Figure 2 further suggest that, for a given value of λ, the asymptotic variance $\widetilde{V}(\lambda;\rho)$ decreases monotonically in ρ. Assuming that this monotonicity is indeed true, the worst-case asymptotic variance for any given λ would be attained at $\rho = 0$.[5] This in turn suggests that λ^* is the minimax choice of λ, in the sense that:

$$\min_{\lambda\in[0,1]} \max_{\rho\in[-1,1]} \widetilde{V}(\lambda;\rho) = \min_{\lambda\in[0,1]} \widetilde{V}(\lambda;0) = \widetilde{V}(\lambda^*;0). \tag{14}$$

In other words, λ^* minimizes the worst-case asymptotic variance over all $\rho \in [-1,1]$, further supporting this simple uniform rule-of-thumb choice for λ and the resulting spot correlation estimator $\hat{\rho}_n^* \equiv \hat{\rho}_n(\lambda^*)$.[6]

Meanwhile, the practical implementation of $\hat{\rho}_n(\lambda^*)$ still requires the evaluation of $g_{\lambda^*}^{-1}(\rho)$. Doing so numerically may seem computationally burdensome. Instead, as a simple-to-implement estimator, we propose using the following third-order polynomial approximation:

$$g_{\lambda^*}^{-1}(\rho) \approx \tilde{g}_{\lambda^*}^{-1}(\rho) \equiv \rho\big(1 + 0.104(1-\rho^2) - 0.023(1-\rho^2)^2\big). \tag{15}$$

Higher-order polynomials could, of course, be used to achieve even more accurate approximations. However, numerically we find that $\sup_{\rho\in[-1,1]}\big|g_{\lambda^*}^{-1}(\rho) - \tilde{g}_{\lambda^*}^{-1}(\rho)\big| < 10^{-3}$, so the estimator in (15) is practically indistinguishable from the true minimax estimator.

2.4 Statistical inference in the fixed-k regime

Section 2.2 characterizes the behavior of the general class of spot correlation estimators $\hat{\rho}_n(\lambda)$ under a conventional large-k asymptotic regime. While this framework results in analytically convenient standard Gaussian-based inference procedures, it may provide poor approximations and be of limited practical use when the local estimation window contains only a moderate number of candlesticks. To address this issue and obtain more reliable approximations for practical empirical applications, we next study the behavior of $\hat{\rho}_n(\lambda)$ under a fixed-k regime, in which the number of candlesticks used for the estimation is held fixed as $\Delta_n \to 0$.

To characterize this alternative fixed-k limiting distribution, define:

$$\tilde{\rho}(\lambda;\rho) = g_\lambda^{-1}\left(\frac{\sum_{i=1}^k\big\{(1-\lambda)\zeta_{C,i}^{(1)}\zeta_{C,i}^{(2)} + \lambda\zeta_{A,i}^{(1)}\zeta_{A,i}^{(2)}\big\}}{\sqrt{\sum_{i=1}^k\big\{(1-\lambda)(\zeta_{C,i}^{(1)})^2 + \lambda(\zeta_{A,i}^{(1)})^2\big\}\sum_{i=1}^k\big\{(1-\lambda)(\zeta_{C,i}^{(2)})^2 + \lambda(\zeta_{A,i}^{(2)})^2\big\}}}\right), \tag{16}$$

where $(\zeta_{C,i}^{(j)}, \zeta_{A,i}^{(j)})_{j\in\{1,2\},\, i\in\{1,\ldots,k\}}$ are previously defined in (5) and (6), with ρ governing the correlation of the underlying bivariate Brownian motion $(B^{(1)}, B^{(2)})$. We then obtain the following theorem.

[5] While the numerical calculations in Figure 2 clearly support this conjecture, since the function $V_\lambda(\rho)$ does not admit a tractable closed-form expression, at present we have been unable to verify this conjecture analytically.

[6] Additional arguments in Section SA of the Online Supplemental Appendix further establish that not only is $\hat{\rho}_n(\lambda^*)$ (asymptotically) minimax in the linear class of estimators defined in (3), it is also minimax in the wider class of consistent and scale-invariant spot correlation estimators defined as an arbitrary continuously differentiable function of the quadratic candlestick features.

Theorem 2. *Suppose that Assumption 1 holds. Then, as $\Delta_n \to 0$, for every $\lambda \in [0,1]$,*

$$\hat{\rho}_n(\lambda) \xrightarrow{\mathcal{L}|\mathcal{F}_\tau} \tilde{\rho}(\lambda;\rho_\tau). \tag{17}$$

In contrast to earlier work, which has relied on a similar fixed-k approach in a high-frequency setting (see, e.g., Bollerslev et al. (2021), Li et al. (2024), and Bollerslev et al. (2024a)), the limiting variable $\tilde{\rho}(\lambda;\rho_\tau)$ in Theorem 2 is not pivotal, but depends on ρ_τ in a nonstandard and highly nonlinear fashion. As a result, the usual confidence interval construction strategies developed in the existing literature are not directly applicable here.

Instead, we propose the use of Anderson and Rubin (1949) type confidence intervals via test inversion. Specifically, for a given significance level α and any candidate value $\rho_0 \in (-1,1)$, we can test the null hypothesis $\mathbb{H}_0\colon \rho_\tau = \rho_0$ using a two-sided test with equal-tailed acceptance region by:

$$I_{1-\alpha}(\lambda;\rho_0) \equiv [F_\lambda^{-1}(\alpha/2;\rho_0), F_\lambda^{-1}(1-\alpha/2;\rho_0)],$$

where $F_\lambda^{-1}(u;\rho)$ denotes the u-quantile of $\tilde{\rho}(\lambda;\rho)$. A $(1-\alpha)$-level confidence set for ρ_τ is thus readily obtained by inverting this test:

$$\mathrm{CS}_{n,1-\alpha}(\lambda) \equiv \{\rho \in [-1,1] : \hat{\rho}_n(\lambda) \in I_{1-\alpha}(\lambda;\rho)\}. \tag{18}$$

Utilizing the result in Theorem 2, the following corollary establishes the correct asymptotic coverage of $\mathrm{CS}_{n,1-\alpha}(\lambda)$.

Corollary 2. *Under Assumption 1 and the additional assumption that $\mathcal{F}_\tau$-conditional law of $\tilde{\rho}(\lambda;\rho_\tau)$ has no atoms almost surely,*[7] *it holds that:* $\mathbb{P}\big(\rho_\tau \in \mathrm{CS}_{n,1-\alpha}(\lambda) \mid \mathcal{F}_\tau\big) \xrightarrow{\mathbb{P}} 1-\alpha$.

The fixed-k confidence intervals in (18) complement the large-k confidence intervals previously established in (12). Importantly, as shown in the simulation study discussed in the next section, the fixed-k intervals also provide much more accurate coverage for our recommended $\hat{\rho}_n^* = \hat{\rho}_n(\lambda^*)$ spot correlation estimator when fewer than 20-30 candlesticks are used in the estimation.

3 Simulation Study

We follow the simulation design of Bollerslev et al. (2020) and consider the following bivariate one-factor stochastic volatility model for the asset log-prices P_t for $t \in [0,1]$, where we normalize a typical trading day to the unit interval:

$$\begin{aligned} dP_t^{(j)} &= \sigma_t^{(j)}\big(\sqrt{\rho_t}\, dB_t + \sqrt{1-\rho_t}\, dW_t^{(j)}\big), \\ \sigma_t^{(j)} &= \exp\{\beta_0 + \beta_1 v_t^{(j)}\}, \quad dv_t^{(j)} = \alpha v_t^{(j)} dt + dW_t^{(j)}, \end{aligned}$$

[7] The additional non-atomicity assumption formally rules out degenerate fixed-k cases, such as $k=1$ with $\lambda \in \{0,1\}$, where the coupling variable reduces to a sign statistic and hence has atoms. In such a case, the confidence set is still well-defined, but the equal-tailed construction may have incorrect limiting coverage.

and B, $W^{(1)}$, and $W^{(2)}$ denote independent standard Brownian motions. We consider parameter values $(\beta_0, \beta_1, \alpha) = (-5/16, 1/8, -1/40)$, with the initial values $v_0^{(j)}$ drawn from the stationary distribution $\mathcal{N}\big(0, (-2\alpha)^{-1}\big)$. These parameter choices ensure that $\mathbb{E}\big[\int_0^1 (\sigma_t^{(j)})^2 dt\big] = 1$ for each $j \in \{1, 2\}$. We consider two specifications for the spot correlation process ρ_t. The first treats the correlation as constant $\rho_t \equiv \rho$, with $\rho \in \{0, 0.3, 0.6, 0.9\}$. The second relies on the stochastic correlation model previously studied by Barndorff-Nielsen and Shephard (2004), in which $\rho_t = \tanh(Y_t)$ and Y_t follows the GARCH diffusion:

$$dY_t = -0.03(Y_t - 0.64)dt + 0.118 Y_t dB'_t,$$

where B' is independent of B, $W^{(1)}$, and $W^{(2)}$. We initialize the process from its stationary distribution $Y_0 \sim \text{InvGamma}(5.31, 2.76)$. Under this second specification, $\rho_t > 0$ with long-run mean approximately equal to 0.53.

We follow Bollerslev et al. (2024a) and simulate all the continuous-time processes using a dense Euler scheme with 10,000 steps per minute. The bivariate candlesticks are constructed from the resulting Euler-discretized price paths over non-overlapping one-minute intervals of the form $[(i-1)\Delta_n, i\Delta_n]$ with $1 \leq i \leq k$ and $\Delta_n = 1/390$. We rely on a total of 10,000 replications. Based on the k candlesticks from each of these simulated paths, we then construct our new proposed estimator $\hat{\rho}_n^* \equiv \hat{\rho}_n(\lambda^*)$. For comparison purposes, we also consider two additional consistent spot correlation estimators: the conventional return-based estimator, $\hat{\rho}_{C,n} \equiv \hat{\rho}_n(0)$, and the candlestick-based estimator proposed by Popov (2016) adapted to our high-frequency setting, $\hat{\rho}_{P,n} \equiv \big(\hat{\rho}_n(0) + \hat{\rho}_n(1)\big)/2$.[8] We begin by considering a "frictionless" setting in which all of the observed return and candlestick pairs are constructed from the simulated continuous price paths as described above.

3.1 Simulation results in the frictionless setting

Table 1 reports the finite-sample bias, variance, and relative MSE for each of the three estimators, under each of the five different specifications for ρ_t. Several patterns emerge. First, all three estimators are essentially unbiased when $\rho_t \equiv 0$, reflecting their symmetry under sign reversals. When ρ_t is positive, however, all the estimators exhibit a downward bias, with the magnitude of the bias generally increasing with the level of ρ_t and decreasing with k. Meanwhile, $\hat{\rho}_n^*$ always exhibits the lowest bias.

Second, the results also clearly evidence the large gains available from incorporating candlestick information into the estimation. Across all the different designs for ρ_t and values of k, the variance of both $\hat{\rho}_n^*$ and $\hat{\rho}_{P,n}$ is substantially lower than the variance of $\hat{\rho}_{C,n}$. Correspondingly, $\hat{\rho}_n^*$ uniformly dominates $\hat{\rho}_{C,n}$ in terms of MSE, with efficiency gains ranging from roughly 37%

[8] The consistency of $\hat{\rho}_{P,n}$ in the high-frequency setting follows from Theorem 1. Popov (2016) also claims consistency and asymptotic normality of the estimator in the daily constant correlation setting, but does so without verifying the requisite conditions on $g_1(\rho)$ for the continuous mapping and delta method arguments to hold.

	Bias			Variance $\times 100$			Relative MSE	
	$\hat{\rho}_n^*$	$\hat{\rho}_{C,n}$	$\hat{\rho}_{P,n}$	$\hat{\rho}_n^*$	$\hat{\rho}_{C,n}$	$\hat{\rho}_{P,n}$	$\hat{\rho}_n^*$	$\hat{\rho}_{P,n}$
Panel 1: $\rho_t \equiv 0$								
$k = 5$	0.001	0.001	-0.002	11.651	20.068	11.436	0.581	0.570
$k = 10$	0.000	0.000	-0.001	5.949	10.022	5.969	0.594	0.596
$k = 30$	-0.001	-0.001	-0.002	2.017	3.309	2.057	0.610	0.622
$k = 60$	-0.001	-0.001	0.000	1.001	1.656	1.034	0.605	0.624
Panel 2: $\rho_t \equiv 0.3$								
$k = 5$	-0.018	-0.030	-0.036	9.918	17.606	10.174	0.562	0.582
$k = 10$	-0.010	-0.019	-0.020	4.955	8.703	5.097	0.568	0.588
$k = 30$	-0.004	-0.009	-0.007	1.649	2.813	1.691	0.585	0.601
$k = 60$	-0.002	-0.004	-0.004	0.798	1.366	0.821	0.584	0.601
Panel 3: $\rho_t \equiv 0.6$								
$k = 5$	-0.024	-0.039	-0.053	5.843	11.267	6.781	0.517	0.618
$k = 10$	-0.011	-0.017	-0.026	2.690	4.928	3.021	0.545	0.623
$k = 30$	-0.003	-0.005	-0.009	0.897	1.477	0.945	0.607	0.644
$k = 60$	-0.002	-0.003	-0.004	0.428	0.701	0.445	0.610	0.636
Panel 4: $\rho_t \equiv 0.9$								
$k = 5$	-0.011	-0.023	-0.028	0.748	1.947	1.274	0.380	0.677
$k = 10$	-0.005	-0.010	-0.012	0.285	0.567	0.376	0.499	0.676
$k = 30$	-0.002	-0.003	-0.004	0.083	0.137	0.090	0.604	0.662
$k = 60$	-0.001	-0.002	-0.002	0.041	0.065	0.042	0.630	0.647
Panel 5: $\rho_t = \tanh Y_t$								
$k = 5$	-0.021	-0.035	-0.046	9.635	15.318	10.241	0.541	0.604
$k = 10$	-0.011	-0.017	-0.024	6.196	8.627	6.448	0.565	0.617
$k = 30$	-0.003	-0.006	-0.007	4.019	4.750	4.061	0.593	0.618
$k = 60$	-0.002	-0.004	-0.005	3.514	3.888	3.526	0.599	0.617

Table 1: Finite-sample bias, variance, and relative mean squared error of $\hat{\rho}_n^*$, $\hat{\rho}_{C,n}$, and $\hat{\rho}_{P,n}$ under different specifications of ρ_t in the frictionless setting. For each estimator, the relative MSE is calculated relative to that of the benchmark return-based estimator $\hat{\rho}_{C,n}$. All moments are simulated based on 10,000 Monte Carlo replications.

to 62%. The new $\hat{\rho}_{P,n}$ estimator also systematically improves on $\hat{\rho}_{C,n}$, but the relative gains are naturally smaller.

To further underscore the superiority of the new estimator, Figure 3 reports the probability that each of the estimators correctly identifies the sign of the true spot correlation under the stochastic correlation model $\rho_t = \tanh(Y_t) > 0$. As the figure shows, the simulated probabilities of correctly identifying the sign increase monotonically with k for all three estimators. Meanwhile, the proposed estimator $\hat{\rho}_n^*$ uniformly delivers the highest sign accuracy, with its advantage most pronounced for smaller values of k. In other words, $\hat{\rho}_n^*$ may be especially useful in empirical applications where the main inferential objective centers on the direction, rather than the exact magnitude, of changes in the spot correlation.

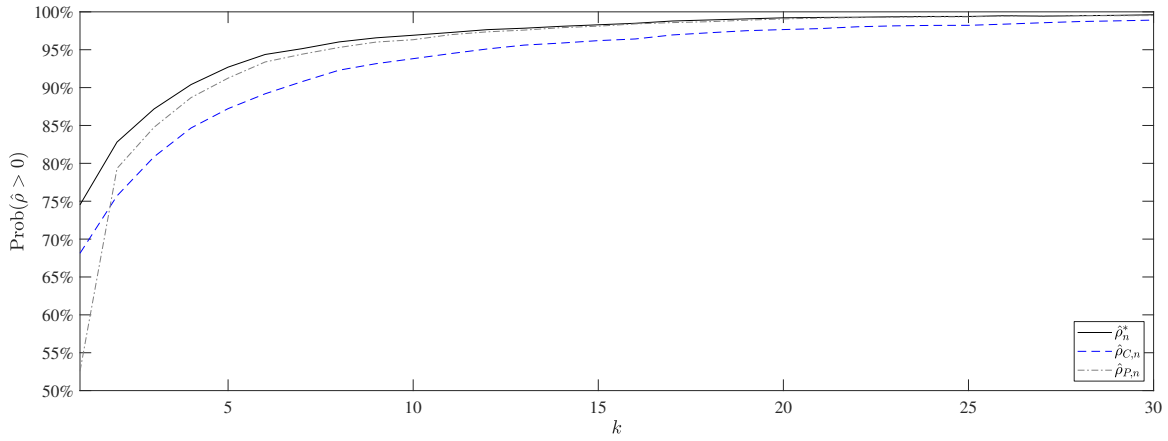

Figure 3: Finite-sample sign accuracies of $\hat{\rho}_n^*$, $\hat{\rho}_{C,n}$, and $\hat{\rho}_{P,n}$ under the stochastic correlation model in the frictionless setting. The vertical axis reports the simulated probability that each estimator correctly identifies the sign of $\rho_0 > 0$.

Finally, to examine the finite-sample coverage of the confidence intervals associated with $\hat{\rho}_n^*$, Figure 4 reports the simulated coverage rates for $\alpha \in \{10\%, 5\%, 1\%\}$ under the stochastic correlation model for both the large-k Gaussian confidence intervals based on (12) and the fixed-k Anderson–Rubin type confidence intervals constructed according to (18). As seen from the figure, the fixed-k confidence intervals achieve coverage rates close to their nominal levels for all values of k, with nearly exact coverage for $k \leq 30$. In contrast, the large-k confidence intervals provide coverage far below the nominal level when k is small, underscoring the advantage of using the fixed-k asymptotic approximations in empirical applications with moderately sized estimation windows.

3.2 Simulation results under market frictions

To further examine the performance of the estimators in the presence of empirically realistic market frictions, we construct candlesticks from asynchronously observed "frictionless" prices contaminated by additive measurement errors, or "market microstructure noise." In so doing, we rely on the stochastic correlation model $\rho_t = \tanh(Y_t)$ for all the simulations. Following

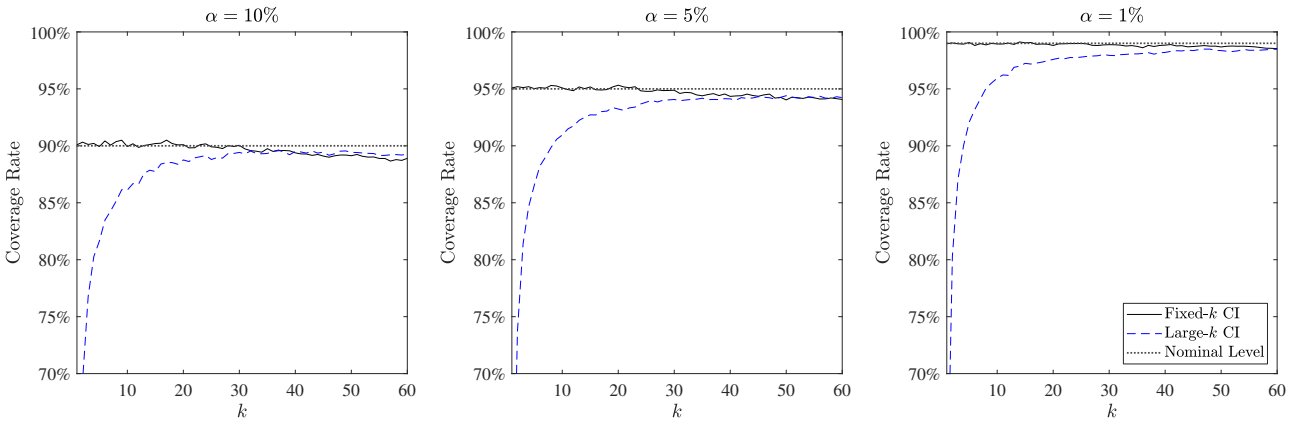

Figure 4: Finite-sample coverage rates of $(1-\alpha)$-level large-k and fixed-k confidence intervals under the stochastic correlation model in the frictionless setting.

Barndorff-Nielsen et al. (2011), we model the asynchronous observation arrivals of $(P_t^{(1)}, P_t^{(2)})$ using two independent Poisson processes for generating the observation times $\{t_m^{(1)}\}$ and $\{t_m^{(2)}\}$. We set the per-minute intensities to $(\lambda^{(1)}, \lambda^{(2)}) \in \{(1000, 1000), (1000, 500), (1000, 100)\}$, so the observations of $P^{(1)}$ arrive, on average, 1000 times per minute, while the average number of observations for $P^{(2)}$ varies among 1000, 500, and 100, over each one-minute interval. We denote the total number of observations for the j-th asset over the one-hour window $[0, 60\Delta_n]$ by $N^{(j)} = \sum_{m\geq 1} \mathbb{1}_{\{t_m^{(j)} \leq 60\Delta_n\}}$.

Each observation time $t_m^{(j)}$ is then associated with an observed price $X^{(j)}_{t_m^{(j)}}$, contaminated by an additive measurement error:

$$X^{(j)}_{t_m^{(j)}} = P^{(j)}_{t_m^{(j)}} + \omega_m^{(j)} \chi_m^{(j)},$$

where $(\chi_m^{(1)}, \chi_m^{(2)})$ denotes a sequence of i.i.d. bivariate Gaussian random variables with zero mean and identity variance-covariance matrix. Following Christensen et al. (2022), the scale of the noise is specified as $\omega_m^{(j)} = \gamma \cdot \sigma^{(j)}_{t_m^{(j)}} \sqrt{60\Delta_n / N^{(j)}}$, where γ effectively represents a noise-to-signal ratio, approximating the ratio of the standard deviation of the measurement error to that of the corresponding efficient-price return. We set $\gamma \in \{0, 0.5, 1, 2\}$, corresponding to no contamination, moderate contamination, high contamination, and extreme contamination, respectively (see Christensen et al. (2014, Table 3)). Note that even though γ is held fixed across both assets, the absolute magnitude of the measurement error depends on $N^{(j)}$, and thus effectively increases as the number of observations decreases. Finally, we construct the candlesticks for asset $j \in \{1, 2\}$ on the i-th one-minute interval $\mathcal{I}_{n,i} = [(i-1)\Delta_n, i\Delta_n]$ for $1 \leq i \leq k$, based on the asynchronously observed prices $\{X^{(j)}_{t_m^{(j)}} : t_m^{(j)} \in \mathcal{I}_{n,i}\}$, with the convention that the opening price of the i-th minute equals the closing price of the $(i-1)$-th minute.

We consider the same three estimators $\hat{\rho}_n \in \{\hat{\rho}_n^*, \hat{\rho}_{C,n}, \hat{\rho}_{P,n}\}$ and values of $k \in \{5, 10, 30, 60\}$ analyzed in Section 3.1 above. To conserve space, we focus our discussion on the results for the moderate-asynchronicity design with $(\lambda^{(1)}, \lambda^{(2)}) = (1000, 500)$ and $\gamma \in \{0, 0.5, 1, 2\}$ reported in Table 2. The results for the balanced and extremely asynchronous designs, $(\lambda^{(1)}, \lambda^{(2)}) \in$

$\{(1000, 1000), (1000, 100)\}$, are reported in Tables S1 and S2 in the Online Supplemental Appendix.

	Bias			Variance $\times 100$			Relative MSE	
	$\hat{\rho}_n^*$	$\hat{\rho}_{C,n}$	$\hat{\rho}_{P,n}$	$\hat{\rho}_n^*$	$\hat{\rho}_{C,n}$	$\hat{\rho}_{P,n}$	$\hat{\rho}_n^*$	$\hat{\rho}_{P,n}$
Panel 1: $\gamma = 0, (\lambda^{(1)}, \lambda^{(2)}) = (1000, 500)$								
$k = 5$	-0.024	-0.036	-0.049	9.997	15.748	10.622	0.543	0.608
$k = 10$	-0.013	-0.018	-0.027	6.430	8.990	6.637	0.568	0.610
$k = 30$	-0.006	-0.008	-0.011	4.054	4.843	4.097	0.587	0.613
$k = 60$	-0.004	-0.005	-0.007	3.478	3.862	3.496	0.600	0.622
Panel 2: $\gamma = 0.5, (\lambda^{(1)}, \lambda^{(2)}) = (1000, 500)$								
$k = 5$	-0.020	-0.035	-0.043	9.598	15.485	10.126	0.534	0.595
$k = 10$	-0.011	-0.015	-0.023	6.222	8.745	6.418	0.573	0.617
$k = 30$	-0.007	-0.007	-0.012	4.064	4.794	4.125	0.600	0.636
$k = 60$	-0.005	-0.004	-0.007	3.518	3.884	3.545	0.602	0.623
Panel 3: $\gamma = 1, (\lambda^{(1)}, \lambda^{(2)}) = (1000, 500)$								
$k = 5$	-0.027	-0.038	-0.055	9.778	15.807	10.422	0.526	0.596
$k = 10$	-0.017	-0.019	-0.030	6.280	8.877	6.522	0.563	0.615
$k = 30$	-0.010	-0.009	-0.015	4.063	4.800	4.118	0.608	0.637
$k = 60$	-0.009	-0.006	-0.011	3.487	3.876	3.512	0.614	0.638
Panel 4: $\gamma = 2, (\lambda^{(1)}, \lambda^{(2)}) = (1000, 500)$								
$k = 5$	-0.036	-0.042	-0.060	9.428	15.571	10.065	0.537	0.613
$k = 10$	-0.028	-0.022	-0.041	6.049	8.557	6.264	0.586	0.644
$k = 30$	-0.019	-0.011	-0.025	3.808	4.661	3.870	0.602	0.641
$k = 60$	-0.018	-0.008	-0.020	3.310	3.764	3.342	0.633	0.661

Table 2: Finite-sample bias, variance, and relative mean squared error of $\hat{\rho}_n^*$, $\hat{\rho}_{C,n}$, and $\hat{\rho}_{P,n}$ under the stochastic correlation model with moderate observation asynchronicity and measurement error. For each estimator, the relative MSE is calculated relative to that of the benchmark return-based estimator $\hat{\rho}_{C,n}$. All moments are based on 10,000 Monte Carlo replications. The parameter γ controls the noise-to-signal ratio of the measurement error, while $(\lambda^{(1)}, \lambda^{(2)})$ denotes the per-minute average number of Poisson observation arrivals.

Comparing the results in Table 2 (and Tables S1 and S2) to the frictionless results in Panel 5 of Table 1 shows that asynchronicity and market microstructure noise both tend to generate slightly larger downward biases in all the spot correlation estimators. This, of course, is in line with the so-called Epps effect (Epps, 1979), and the idea that common price movements are only partially accounted for when the two assets are not observed at the same time. Similarly, the addition of independent measurement errors effectively adds idiosyncratic variation to the

observed prices and candlestick features, in turn resulting in a similar attenuation effect. As a result, increasing the degree of asynchronicity and/or the noise level tends to push the estimated correlation downward relative to the true correlation.

Meanwhile, the results in Table 2 again show that $\hat{\rho}_n^*$ uniformly dominates the conventional $\hat{\rho}_{C,n}$ return-based estimator and the $\hat{\rho}_{P,n}$ candlestick-based estimator. Only under extreme market frictions (i.e., $(\lambda^{(1)}, \lambda^{(2)}) = (1000, 100)$, $\gamma = 2$) and for the largest window size (i.e., k=60), as reported in Table S2, do the relative improvements in MSE appear noticeable lower.

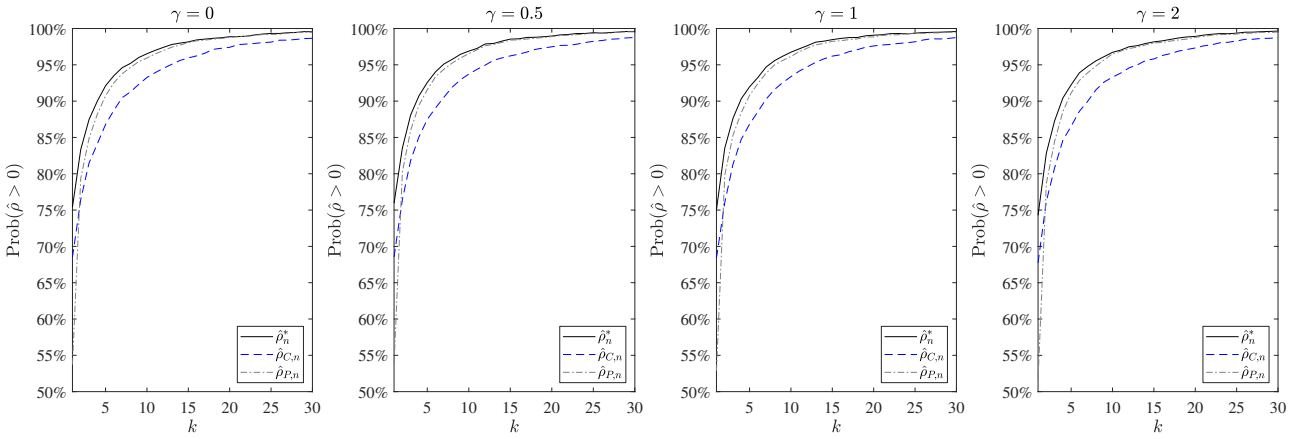

Figure 5: Finite-sample sign accuracies of $\hat{\rho}_n^*$, $\hat{\rho}_{C,n}$, and $\hat{\rho}_{P,n}$ under the stochastic correlation model with moderate observation asynchronicity, $(\lambda^{(1)}, \lambda^{(2)}) = (1000, 500)$, and measurement error levels $\gamma \in \{0, 0.5, 1, 2\}$. The vertical axis reports the simulated probability that each estimator correctly identifies the sign of $\rho_0 > 0$.

The results for the sign accuracy in Figure 5 similarly corroborate the earlier results for the frictionless setting in Figure 3. The new $\hat{\rho}_n^*$ estimator continues to deliver the most accurate identification of the sign of the true spot correlation, especially when using relatively narrow estimation windows, or smaller values of k. The fact that the accuracy of $\hat{\rho}_{P,n}$ fairly closely matches that of $\hat{\rho}_n^*$ merely underscores the value of the additional information provided by the candlesticks for accurately signing the spot correlation. At a broader level, this also echoes the robustness to market microstructure noise of the subsampled quadrant-based realized correlation estimator studied by Hansen and Luo (2023).

Lastly, paralleling the results for the frictionless setting in Figure 4, Figure 6 plots the coverage rates for nominal 95% confidence intervals for the fixed-k and large-k asymptotic approximations based on simulated "contaminated" prices.[9] The figure again shows that for small to moderate noise levels (i.e., $\gamma = 0.5, 1$), the coverage rates for the fixed-k intervals remain uniformly close to the nominal level. Meanwhile, in contrast to the findings for the frictionless setting, for the highest noise level (i.e., $\gamma = 2$) and larger estimation windows (i.e., $k \geq 30$), the large-k confidence intervals tend to work slightly better than their fixed-k counterparts.

In sum, the results discussed above confirm that the new candlestick-based spot correlation estimator $\hat{\rho}_n^*$ continues to perform well in the presence of empirically realistic levels of market

[9]The results for other confidence levels are qualitatively similar and omitted for brevity.

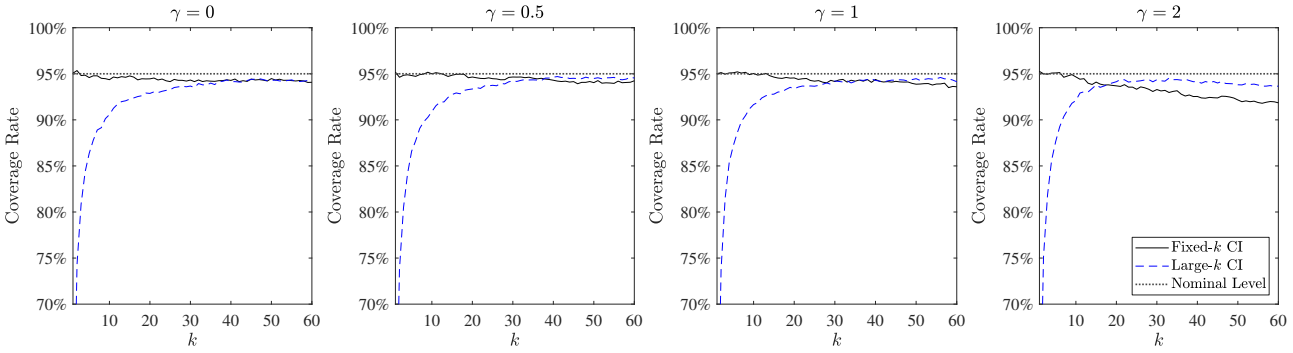

Figure 6: Finite-sample coverage rates of 95%-level large-k and fixed-k confidence intervals under the stochastic correlation model with moderate observation asynchronicity, $(\lambda^{(1)}, \lambda^{(2)}) = (1000, 500)$. The parameter γ controls the noise-to-signal ratio of the measurement error.

frictions. We turn next to an actual empirical application of the estimator, highlighting its practical usefulness for identifying significant changes in the correlation between stocks and bonds at the exact times of important economic news.

4 Stock-Bond Spot Correlations and Economic News

The correlation between equity and long-term Treasury returns is a central input into portfolio construction, risk management, and the interpretation of macroeconomic news. Following a prolonged period in which long-term bonds served as a reliable hedge against equity risk, this correlation has become less stable in recent years, with inflation and discount-rate shocks playing an increasingly important role in determining the comovement of the two asset classes (see, e.g., Campbell et al., 2020; Cieslak and Pang, 2021; Duffee, 2023). Recent evidence further suggests that Treasuries may switch between safe-haven and risky-asset behavior even within a single trading day (Hu et al., 2025). These observations in turn underscore the importance of measuring stock-bond comovements over intraday horizons for properly understanding the economic mechanisms at work. The new candlestick-based estimator $\hat{\rho}_n^*$ provides a practical tool for doing so.[10]

In general, if an important news event is priced primarily through a discount-rate or inflation-expectations channel, equities and long-duration bonds tend to move in the same direction, leading to a higher post-event correlation. If, instead, the event is priced primarily through a growth or flight-to-safety channel, equity cash-flow expectations and Treasury prices tend to move in opposite directions, leading to a lower, and potentially negative, post-event correlation. When several channels are active at the same time, the spot correlation effectively summarizes their net effect over the event window.

[10] We do not argue that candlestick data are preferable to tick-level data when the latter are available. However, the use of candlesticks makes high-frequency correlation analysis feasible for a much broader set of markets and assets, for which tick-level data are costly, incomplete, difficult to access, or simply unavailable.

We proxy the returns on the two asset classes with their most actively traded ETFs, namely the SPDR S&P 500 ETF (SPY) for equities and the iShares 20+ Year Treasury Bond ETF (TLT) for long-duration Treasuries.[11] We form our spot correlation estimates based on one-minute candlesticks over a 30-minute local estimation window, corresponding to $k = 30$. To better trace the intraday evolution of the correlation around each event, we plot rolling estimates obtained by advancing the 30-minute estimation window in 15-minute increments, so that consecutive $\hat{\rho}_n^*$ estimates share half of their underlying candlesticks. We do so for two different types of 2025 economic news events. The first set consists of highly structured monetary-policy news in the form of precisely scheduled FOMC announcements. The second set consists of unscheduled tariff-related posts by President Trump, which are less structured, more discretionary, and harder to interpret ex ante. We begin our discussion by considering the correlations observed around the FOMC announcements.

4.1 FOMC Announcements

FOMC announcements are directly related to monetary policy. However, the policy signals in the statements are multidimensional, with each statement conveying the current decision on the target rate, the expected policy path, and the committee's assessment of the economic outlook. These different components have different implications for stock-bond comovement. The intraday correlation observed around the announcement time helps distinguish between these different channels.

To illustrate, Figure 7 plots the estimated spot correlation, together with 95% confidence intervals, around the eight FOMC announcements in 2025. A clear pattern emerges. Before the 14:00 EST release, the spot correlation is typically close to zero and often mildly negative. After the announcement, it shifts upward and, on several dates, becomes significantly positive. This pattern is much sharper than what is visible from the daily average, suggesting that the most informative variation in stock-bond comovement is concentrated around the announcement window and is largely averaged out at the daily frequency. Correspondingly, for investors assessing hedging demand and portfolio risk around such events, daily correlations can be very misleading.

Meanwhile, the systematic upward shift in the estimated spot correlations observed in Figure 7 clearly favors a traditional discount-rate interpretation of the news. In other words, markets appear to price the announcements mainly through revisions to the expected policy path, rather than as news about real fundamentals. This reading is in line with Bauer and Swanson (2023), who argue that much of the apparent information content of FOMC surprises is best understood as a "Fed response to news" channel. By contrast, under the "information channel" interpretation emphasized by Nakamura and Steinsson (2018) among others, an an-

[11] Over the recent five-year window 2021–2025, the daily-return correlation of these ETFs averages only 0.06. As we show below, however, the spot correlation can move by several tenths of a percentage point within a narrow window around major policy events.

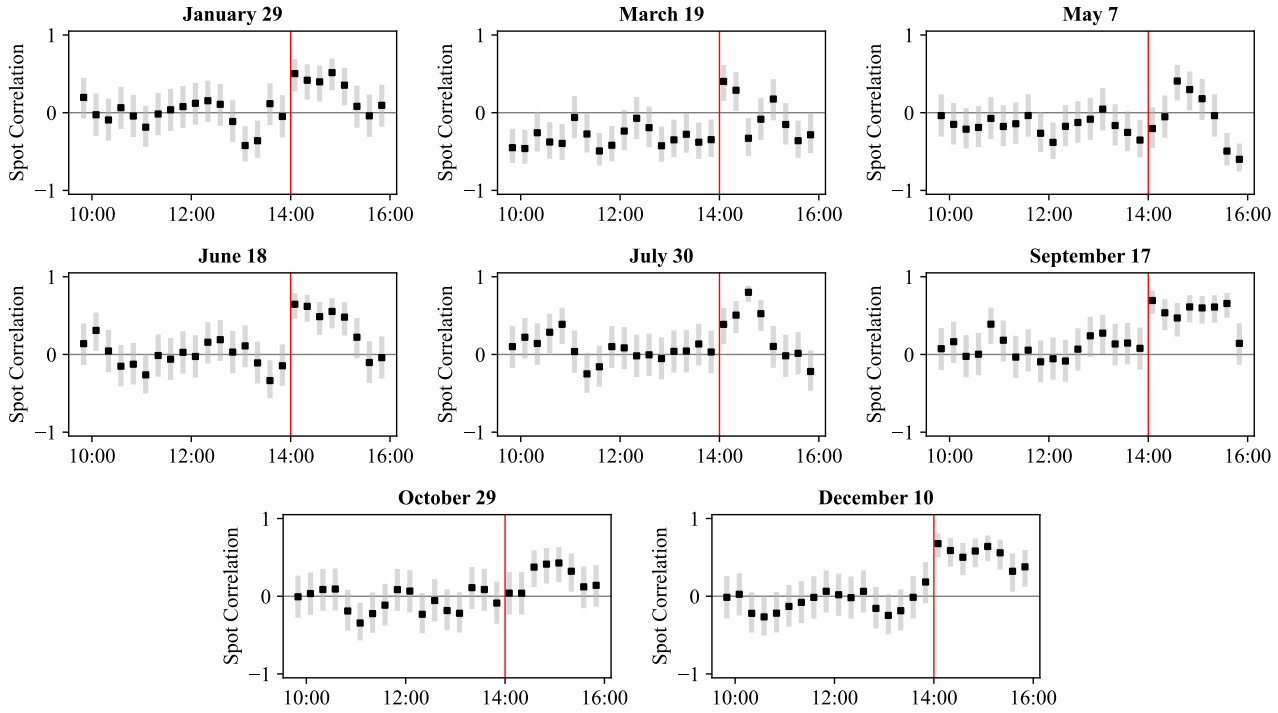

Figure 7: The figure shows the intraday patterns of spot correlation on eight FOMC announcement days in 2025. Each panel corresponds to one event day. Gray bands represent 95% confidence intervals. The vertical lines mark the exact announcement times.

nouncement interpreted as good news about real economic fundamentals should raise equity prices while lowering bond prices, thereby inducing a negative stock-bond correlation.

The four announcements in March, June, September, and December, which also include the Fed's summaries of economic projections (SEP), provide a useful check on this interpretation. By explicitly communicating its projections for growth, inflation, unemployment, and the intended future policy path, these FOMC announcements provide precisely the type of forward-looking information through which the Fed information channel would naturally operate. Yet the estimated correlations for these four announcements again tend to shift upward at the announcement times, reinforcing the view that markets primarily interpreted FOMC news in terms of its implications for future discount rates.

4.2 Trump Trade Policy Shocks

In contrast to the highly structured FOMC announcements analyzed in the previous section, we next study the impact of a set of discretionary and unscheduled tariff-related posts by President Trump. A single post may naturally affect expected growth, inflation, supply-chain risk, and beliefs about future monetary-policy responses. This ambiguity makes spot correlations especially useful for identifying the market's dominant interpretation of the news. We consider three such posts from the 2025 tariff cycle: the April 2 "Liberation Day" announcement, the April 9 restructuring of tariff pressure across countries, and the July 2 transition toward

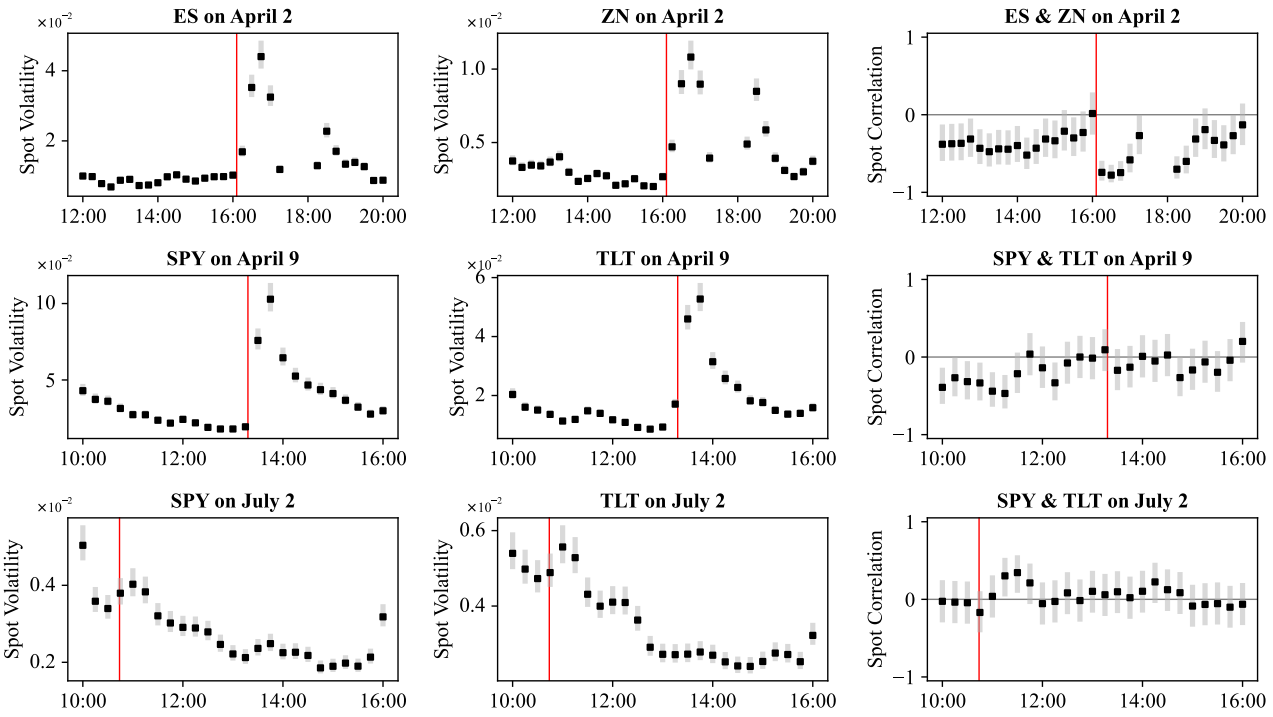

Figure 8: The figure shows the intraday patterns of spot volatility and spot correlation on three distinct dates of the 2025 tariff cycle. Each row corresponds to one event day. The first two columns report the estimated spot volatilities. The third column reports the estimated spot correlation. Gray bands represent 95% confidence intervals. The vertical lines mark the exact times of the posts.

negotiated deals.

The April 2 "Liberation Day" announcement essentially transformed bilateral tariff threats into a broader and more systemic tariff regime.[12] Because the announcement occurred at 16:06 EST after the cash-market close at 16:00 EST, we rely on data for the E-mini S&P 500 futures (ES) and the 10-Year U.S. Treasury Note futures (ZN) for capturing the returns on stocks and long-duration bonds.[13] The first row of Figure 8 shows an immediate sharp increase in the resulting stock and bond spot volatility estimates at the 16:06 EST announcement time. At the same time, the estimated spot correlations also turn distinctly negative. This negative response in the correlation suggests that investors placed greater weight on the adverse growth implications of a systemic tariff regime than on the inflation or discount-rate channel. In other words, equities repriced downward on cash-flow and growth concerns, while long-duration Treasuries benefited from flight-to-safety demand. The negative estimates after the CME daily maintenance halt from 17:00 EST to 18:00 EST (visible as a gap in the figure) further indicate that the market continued to respond to and price the news primarily through the growth and

[12]See `https://www.presidency.ucsb.edu/documents/remarks-announcing-additional-united-states-tariff-actions-foreign-imports`.

[13]Since there is no standard 20-year Treasury futures contract, ZN is the closest available benchmark. Using 30-year Treasury bond futures (ZB) leads to qualitatively similar results.

flight-to-safety channels.

The subsequent April 9 post represents the first major restructuring of the new tariff regime.[14] The reciprocal tariff on Chinese imports was raised from 84% to 125%, while the country-specific reciprocal tariffs for most other countries were suspended for 90 days and replaced by a flat 10% baseline tariff. The second row of Figure 8 shows that volatility again jumped sharply in both markets, underscoring the economic significance of the news. The spot correlation, however, shows no persistent directional shift and remains close to zero throughout the event window. This contrast is informative. The April 9 announcement contained two major policy signals pointing in opposite directions: a sharp escalation toward China and a partial de-escalation toward the rest of the world, rendering the net implication of the news ambiguous, regardless of the precise economic channels at work.

The final of our three posts on July 2 in turn marks a transition in the trade war from escalation to negotiated resolution.[15] Although the post was explicitly about a new trade deal reached with Vietnam, the market likely interpreted the post as a template for a negotiated resolution with other countries. Accordingly, the third row of Figure 8 shows a different pattern from the two previous posts. In particular, the spot volatility appears to decline in both markets, consistent with a resolution-of-uncertainty interpretation of the news. The spot correlation, by contrast, moves upward and turns mildly positive. A natural interpretation is that the post and its likely wider implications for reduced tariff pass-through into U.S. consumer prices, lowered inflation expectations and Treasury yields, while increasing equity valuations, resulting in positive stock-bond comovement.

Taken together, the three events show that tariff-related shocks, and by implication, other important economic news announcements, can generate rich patterns of stock-bond comovement, and that it is difficult to infer the patterns from the news headlines alone. A systemic trade war escalation can drive the correlation negative through growth concerns, as on April 2; a mixed escalation-plus-de-escalation can generate a large volatility response with little change in comovement, as on April 9; and a move toward negotiated resolution can shift the correlation upward by reducing expected tariff pass-through, as on July 2. Such patterns would be difficult, if not impossible, to discern from lower-frequency daily correlations.

5 Conclusion

We develop a new econometric framework for estimating and conducting inference on spot correlations using high-frequency candlestick data. The new estimator blends the standard return cross-product with a cross-product of signed candlestick asymmetries together with a nonlinear bias correction. We establish consistency and asymptotic normality within a conventional

[14]The post was published on Truth Social at 13:18 EST, see `https://truthsocial.com/@realDonaldTrump/posts/114309144289505174`.

[15]The post was published on Truth Social at 10:44 EST, see `https://truthsocial.com/@realDonaldTrump/posts/114784170652465525`.

large-k infill asymptotic framework, along with more reliable confidence intervals when only a few candlesticks are employed in the estimation based on alternative fixed-k approximations. Within the general class of consistent estimators, we further determine the "optimal" estimator $\hat{\rho}_n(\lambda^*)$ that (approximately) minimizes the asymptotic variance across the full range of possible correlation values. Monte Carlo simulations demonstrate that this new estimator performs well in empirically realistic situations.

An empirical application to U.S. stock-bond comovements around recent policy events illustrates both the practical value of the estimator and the economic content it reveals. For FOMC announcements, the estimated spot correlation between SPY and TLT is typically near zero or mildly negative in the pre-announcement window, then shifts upward and often becomes significantly positive following the news release. This systematic pattern points to a dominant discount-rate channel in which revisions to the expected policy path move equities and long-duration bonds in the same direction. The evidence sits uneasily with an information-channel interpretation, under which favorable news about real fundamentals should generate negative stock-bond comovement. Tariff-related announcements tell a more heterogeneous story, with some generating an immediate and persistent negative correlation response, consistent with markets pricing the shock primarily through adverse growth expectations and flight-to-safety, while others producing a large volatility response in both markets, but no discernible directional shift in correlation, reflecting the offsetting economic signals embedded in the news. These different patterns would be largely obscured at the daily frequency, underscoring the value of intraday correlation measurement for interpreting the economic content of policy news in real time.

Several directions for future research naturally suggest themselves. On the theoretical side, a formal proof of the conjectured minimax optimality would provide a sharper analytical foundation for the proposed estimator. It would also be interesting to further expand the class of candlestick-based estimators by incorporating additional non-quadratic features into the formulation of the estimator. Extensions to candlestick-based estimation of spot betas, factor loadings, and other regression coefficients in a high-frequency panel context also represent promising avenues for future research.

On the empirical side, while the application in the paper focuses on U.S. stock-bond comovements around a small set of monetary and trade policy announcements, a more systematic study spanning a longer time horizon and a wider set of news could help further illuminate how the economic channels at work vary with the inflation regime, the level of interest rates, and the credibility of the central bank. The framework is equally well-suited to studying a broader range of asset pairs and news events. Applying the estimator to currency and commodity markets could shed light on how exchange rate and inflation expectations respond jointly to central bank communications, geopolitical shocks, or supply disruptions. Within equity markets, the estimator could also be used to trace the intraday dynamics of sector correlations around earnings announcements, or episodes of financial stress, providing a sharper lens through which to

distinguish between cash-flow, discount-rate, and risk-appetite channels of transmission.

More broadly, the results of the paper suggest that the recent widespread availability of high-frequency intraday candlestick data — across asset classes, geographies, and sampling frequencies — represents a still underutilized source of statistical information, and that principled econometric frameworks for exploiting such data could yield meaningful new insights for estimation and economic inference concerning questions that daily data have proven too coarse to answer cleanly. We look forward to further exploring these ideas in future work.

Appendix: Proofs

Following Bollerslev et al. (2024a) and Li et al. (2024), we shall assume that jumps are absent in intervals of order $O(k_n\Delta_n)$ near τ such that P is continuous on $\mathcal{I}_n \equiv \bigcup_{i=1}^{k_n} \mathcal{I}_{n,i}$, as the probability of observing a jump in this interval converges to zero as $k_n\Delta_n \to 0$. We also use the strengthened version of Assumption 1 by assuming that the conditions hold with $T_1 = \infty$ using a standard localization procedure (see the discussion in Li et al. (2024)). The positive constants K, K' are generic and may vary from line to line, independent of n. For a generic random variable X, we define $\|X\|_p = (\mathbb{E}[|X|^p])^{1/p}$ as its L_p-norm for $1 \leq p < \infty$.

As a preliminary, we shall give the definition of convergence in $\mathcal{F}$-conditional law in the vein of Li and Xiu (2016, Definition A1) and a useful lemma.

Definition 1. *Let (X_n) be a sequence of random variables defined on $\big(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}\big)$, and let X denote a random variable defined on an extension of $\big(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}\big)$. We say that X_n converges in $\mathcal{F}$-conditional law to X, written as $X_n \xrightarrow{\mathcal{L}|\mathcal{F}} X$, if for every bounded continuous function f, it holds that $\mathbb{E}[f(X_n)\,|\,\mathcal{F}] \xrightarrow{\mathbb{P}} \mathbb{E}[f(X)\,|\,\mathcal{F}]$, where the expectation is taken with respect to the extended probability space if necessary.*

The convergence in $\mathcal{F}$-conditional law clearly nests the mere convergence in law since $\mathbb{E}[f(X_n)] \to \mathbb{E}[f(X)]$ follows from $\mathbb{E}[f(X_n) \mid \mathcal{F}] \xrightarrow{\mathbb{P}} \mathbb{E}[f(X) \mid \mathcal{F}]$ using the boundedness of f. The following lemma is repeatedly used in the proof.

Lemma A.1. *For two sequences of random variables (X_n) and (Y_n) and a random variable X with law $\mathbb{P}_X$, if $X_n \xrightarrow{\mathcal{L}|\mathcal{F}} X$ and $|Y_n - X_n| \xrightarrow{\mathbb{P}} 0$, then*

(i) (Continuous mapping.) For any $\mathbb{P}_X$-almost everywhere continuous function f, it holds that $|f(Y_n) - f(X_n)| \xrightarrow{\mathbb{P}} 0$.

(ii) (Asymptotic equivalence.) $Y_n \xrightarrow{\mathcal{L}|\mathcal{F}} X$.

Proof. For (i), by $X_n \xrightarrow{\mathcal{L}} X$ and $|X_n - Y_n| \xrightarrow{\mathbb{P}} 0$, we have the joint convergence $(X_n, Y_n) \Rightarrow (X, X)$ by the Slutsky's Theorem. Then $\big(f(X_n), f(Y_n)\big) \Rightarrow \big(f(X), f(X)\big)$ and hence $|f(X_n) - f(Y_n)| \xrightarrow{\mathcal{L}} 0$ follows by the continuous mapping theorem. But since the limit is degenerate, we conclude that $|f(X_n) - f(Y_n)| \xrightarrow{\mathbb{P}} 0$.

For (ii), take any bounded continuous function φ, it holds from (i) that $|\varphi(Y_n) - \varphi(X_n)| \xrightarrow{\mathbb{P}} 0$. With the boundedness of φ, we further deduce $\|\varphi(Y_n) - \varphi(X_n)\|_1 \to 0$. The required claim follows from

$$\begin{aligned}\|\mathbb{E}[\varphi(Y_n) \mid \mathcal{F}] - \mathbb{E}[\varphi(X) \mid \mathcal{F}]\|_1 &\leq \|\mathbb{E}[\varphi(Y_n) - \varphi(X_n) \mid \mathcal{F}]\|_1 + \|\mathbb{E}[\varphi(X_n) \mid \mathcal{F}] - \mathbb{E}[\varphi(X) \mid \mathcal{F}]\|_1 \\ &\leq \|\varphi(Y_n) - \varphi(X_n)\|_1 + \|\mathbb{E}[\varphi(X_n) \mid \mathcal{F}] - \mathbb{E}[\varphi(X) \mid \mathcal{F}]\|_1 \to 0,\end{aligned}$$

where the last estimate is by the assumption that $X_n \xrightarrow{\mathcal{L}|\mathcal{F}} X$. The above L_1 convergence implies that $\mathbb{E}[\varphi(Y_n) \mid \mathcal{F}] \xrightarrow{\mathbb{P}} \mathbb{E}[\varphi(X) \mid \mathcal{F}]$, which completes the proof. □

We shall prove some analytical properties for the function $g_\lambda(\cdot)$ defined in (8).

Lemma A.2. *For every $\lambda \in [0,1]$, the function $g_\lambda : [-1,1] \to [-1,1]$ is an odd, continuous, and strictly increasing bijection. Moreover, $g_\lambda(\cdot)$ is differentiable on $(-1,1)$ with $g'_\lambda(\cdot)$ continuous on $(-1,1)$ and satisfies*

$$0 < \frac{1 - 8/\pi^2}{3 - 4\log 2} \le g'_\lambda(\rho) \le \frac{1}{3 - 4\log 2}, \qquad \rho \in (-1,1).$$

Consequently, $g_\lambda^{-1}(\cdot)$ exists, is Lipschitz and continuously differentiable on $(-1,1)$, with

$$0 < 3 - 4\log 2 \le (g_\lambda^{-1})'(u) = \frac{1}{g'_\lambda\big(g_\lambda^{-1}(u)\big)} \le \frac{3 - 4\log 2}{1 - 8/\pi^2}, \qquad u \in (-1,1).$$

Proof. Let $B = (B_t)_{0 \le t \le 1}$ be a standard Brownian motion and define the functionals $C(B) \equiv B_1$, $H(B) \equiv \sup_{0 \le t \le 1} B_t$, $L(B) \equiv \inf_{0 \le t \le 1} B_t$, and $A(B) \equiv H(B) + L(B) - C(B)$. One should verify that:

$$C(-B) = -C(B), \quad H(-B) = -L(B), \quad L(-B) = -H(B),$$

so we have $A(-B) = -A(B)$, hence $A(B)$ is odd in B with $\mathbb{E}[A(B)] = 0$.

For $\rho \in [-1,1]$, let $B^\perp$ be an independent copy of B, and define

$$B^\rho \equiv \rho B + \sqrt{1-\rho^2} B^\perp, \tag{A.1}$$

so that (B, B^ρ) is a bivariate Brownian motion with correlation ρ. Define $g_A(\rho) \equiv \mathbb{E}[A(B)A(B^\rho)]$, and note that $g_A(\rho) = 2(c(\rho) - c(-\rho)) - \rho$. With this notation, we can rewrite $g_\lambda(\rho)$ as

$$g_\lambda(\rho) = \frac{(1-\lambda)\rho + \lambda g_A(\rho)}{(1-\lambda) + \lambda(3 - 4\log 2)}. \tag{A.2}$$

We shall first prove oddity and continuity of $g_\lambda(\rho)$ on $[-1,1]$, which immediately follows from the oddity and continuity of $g_A(\rho)$. First, notice that $(B, B^{-\rho}) \overset{d}{=} (-B, B^\rho)$, and therefore, for all $\rho \in [-1,1]$:

$$-g_A(\rho) = \mathbb{E}[-A(B)A(B^\rho)] = \mathbb{E}[A(-B)A(B^\rho)] = \mathbb{E}[A(B)A(B^{-\rho})] = g_A(-\rho),$$

so g_A is odd. To prove the continuity of g_A, take any sequence $\rho_n \to \rho$ on $[-1,1]$. The construction of B^ρ in (A.1) implies that $\sup_{0 \le t \le 1} |B_t^{\rho_n} - B_t^\rho| \to 0$ almost surely, so B^ρ is continuous in ρ under the sup-norm. The candlestick maps $B \mapsto C(B)$, $B \mapsto H(B)$, $B \mapsto L(B)$, and hence $B \mapsto A(B)$ are continuous under the sup-norm, thus $A(B)A(B^{\rho_n}) \to A(B)A(B^\rho)$, almost surely. Now, write $M \equiv \sup_{0 \le t \le 1} |B_t| + \sup_{0 \le t \le 1} |B_t^\perp|$, and it is well-known that Brownian suprema have finite moments of all orders, i.e., $\mathbb{E}[M^k] < \infty, \forall k \in \mathbb{N}$. Uniformly over $\rho \in [-1,1]$, it holds that

$$|C(B^\rho)| \le M, \quad |H(B^\rho)| \le M, \quad |L(B^\rho)| \le M,$$

and hence $|A(B^\rho)| \le 3M$. This further implies that $|A(B)A(B^\rho)| \le 9M^2$. Since $\mathbb{E}[M^2] < \infty$, by the dominated convergence theorem, we conclude that

$$g_A(\rho_n) = \mathbb{E}[A(B)A(B^{\rho_n})] \to \mathbb{E}[A(B)A(B^\rho)] = g_A(\rho),$$

proving the continuity of g_A on $[-1,1]$.

To provide estimates for the derivative of g_A, we use Malliavin calculus techniques. On the Wiener space over $[0,1]$, let $\mathcal{H} = L^2([0,1])$. To highlight the dimension, denote $\|f\|_{\mathcal{H}(q)} \equiv \big(\int_{[0,1]^q} f^2(x)dx\big)^{1/2}$ for $f : [0,1]^q \to \mathbb{R}$ defined on q-dimensional cube. The Malliavin derivative D is first defined on smooth cylindrical random variables and then closed as an operator from $L^2(\Omega)$ into $L^2(\Omega;\mathcal{H})$. The first-order Malliavin–Sobolev space $\mathbb{D}^{1,2}$ is the domain of this closed operator under the norm

$$\|F\|_{1,2}^2 = \mathbb{E}[F^2] + \mathbb{E}[\|DF\|_{\mathcal{H}(1)}^2],$$

where D denotes the Malliavin derivative as an $\mathcal{H}$-valued operator, while D_sF denotes its coordinate at time s, so that $DF = (D_sF)_{0\le s\le 1}$, see Nualart (2006, Section 1.2).

Let τ_H and τ_L denote the almost surely unique times at which B attains its maximum and minimum on $[0,1]$. The Malliavin derivative identity for the Brownian maximum gives $H \in \mathbb{D}^{1,2}$ and $D_sH = \mathbb{1}_{\{s\le\tau_H\}}$, see Nualart (2006, Exercise 1.2.11 and Proposition 2.1.10). Applying the same result to $-B$ yields $L \in \mathbb{D}^{1,2}$ and $D_sL = \mathbb{1}_{\{s\le\tau_L\}}$. Also, $C(B) = B_1 \in \mathbb{D}^{1,2}$ and $D_sC \equiv 1$. Therefore,

$$A = H + L - C \in \mathbb{D}^{1,2}, \quad D_sA = \mathbb{1}_{\{s\le\tau_H\}} + \mathbb{1}_{\{s\le\tau_L\}} - 1.$$

In particular, $|D_sA| \le 1$ for all $s \in [0,1]$, so $\mathbb{E}[\|DA\|_{\mathcal{H}(1)}^2] \le 1$.

Because $A \in \mathbb{D}^{1,2} \subset L^2(\Omega)$ and A is measurable with respect to the Brownian path, the Wiener–Itô chaos decomposition gives a unique expansion $A = \mathbb{E}[A] + \sum_{q=1}^{\infty} I_q(f_q)$, where $f_q \in L^2([0,1]^q)$ and $I_q(f_q)$ denotes the q-th multiple Wiener integral, see Nualart (2006, Section 1.1). Recall that A is odd with $\mathbb{E}[A] = 0$, and the reflection $B \mapsto -B$ changes the sign of $I_q(f_q)$ by the factor $(-1)^q$. By uniqueness of the Wiener–Itô expansion, all even-order chaos kernels must vanish. Hence

$$A = \sum_{q\in\mathbb{N}_{\text{odd}}} I_q(f_q).$$

By the standard derivative formula for Wiener chaos expansions, we further have $\mathbb{E}[\|DA\|_{\mathcal{H}(1)}^2] = \sum_{q\in\mathbb{N}_{\text{odd}}} q q!\|f_q\|_{\mathcal{H}(q)}^2 \le 1$, see Nualart (2006, Proposition 1.2.2).

We compute the first chaos coefficient explicitly. Since $D_sI_q(f_q) = qI_{q-1}\big(f_q(\cdot,s)\big)$, taking expectations in the derivative expansion gives $f_1(s) = \mathbb{E}[D_sA]$. By the arcsine law for the Brownian maximizer, $\mathbb{P}(\tau_H \le s) = \frac{2}{\pi}\arcsin\sqrt{s}$, for $s \in [0,1]$, see Mörters and Peres (2010, Theorem 5.26). The same formula holds for τ_L, since τ_L is the maximizer of $-B$. Hence $f_1(s) = \mathbb{P}(\tau_H \ge s) + \mathbb{P}(\tau_L \ge s) - 1 = 1 - \frac{4}{\pi}\arcsin\sqrt{s}$. Consequently,

$$\|f_1\|_{\mathcal{H}(1)}^2 = \int_0^1 \Big(1 - \frac{4}{\pi}\arcsin\sqrt{s}\Big)^2 ds = 1 - \frac{8}{\pi^2}.$$

We now use the chaos expansion to differentiate g_A. By Mehler's formula (Nualart, 2006, Section 1.4, Equation 1.67) for the Ornstein–Uhlenbeck semigroup $\{T_t, t \ge 0\}$, for $\rho = e^{-t} \in (0,1]$ and a square-integrable Brownian functional F,

$$T_tF(B) = \mathbb{E}[F(\rho B + \sqrt{1-\rho^2}B^{\perp}) \mid B] = \mathbb{E}[F(B^{\rho}) \mid B].$$

Taking $F(B) = I_q^B(f_q)$, and using the fact that the q-th Wiener chaos is an eigenspace of T_t with eigenvalue e^{-qt}, we obtain for $\rho \in (0,1]$,

$$\mathbb{E}[I_q^{B^\rho}(f_q) \mid B] = T_t I_q^B(f_q) = e^{-qt} I_q^B(f_q) = \rho^q I_q^B(f_q). \tag{A.3}$$

For $\rho \in [-1,0)$, we may take $r = -\rho = e^{-t}$ and applying the previous argument with $-B$ in place of B, which gives $I_q^{-B}(f_q) = (-1)^q I_q^B(f_q)$ and hence $\mathbb{E}[I_q^{B^\rho}(f_q) \mid B] = r^q(-1)^q I_q^B(f_q) = \rho^q I_q^B(f_q)$. For $\rho = 0$, the identity in (A.3) follows from independence, since $\mathbb{E}[I_q^{B^\perp}(f_q) \mid B] = 0$, for $q \geq 1$. So (A.3) holds for all $\rho \in [-1,1]$. Therefore, by (A.3), the law of iterated expectation, and the orthogonality of Wiener chaoses, we have

$$\mathbb{E}[I_p^B(f_p) I_q^{B^\rho}(f_q)] = \rho^q \mathbb{E}[I_p^B(f_p) I_q^B(f_q)] = \mathbb{1}_{\{p=q\}} \rho^q q! \|f_q\|^2_{\mathcal{H}(q)},$$

for all $\rho \in [-1,1]$. Plugging the chaos expansion of A into the expression of g_A and simplifying using the above relation yields

$$g_A(\rho) = \mathbb{E}[A(B)A(B^\rho)] = \sum_{q \in \mathbb{N}_{\text{odd}}} \rho^q q! \|f_q\|^2_{\mathcal{H}(q)}.$$

The series is absolutely and uniformly convergent on $[-1,1]$, since

$$\sum_{q \in \mathbb{N}_{\text{odd}}} |\rho^q| q! \|f_q\|^2_{\mathcal{H}(q)} \leq \sum_{q \in \mathbb{N}_{\text{odd}}} q! \|f_q\|^2_{\mathcal{H}(q)} = \mathbb{E}[A(B)^2] = 3 - 4\log 2,$$

where the last estimate follows from Popov (2016). Moreover, since

$$\sum_{q \in \mathbb{N}_{\text{odd}}} q|\rho^{q-1}| q! \|f_q\|^2_{\mathcal{H}(q)} \leq \sum_{q \in \mathbb{N}_{\text{odd}}} q q! \|f_q\|^2_{\mathcal{H}(q)} = \mathbb{E}[\|DA\|^2_{\mathcal{H}(1)}] \leq 1,$$

the ρ-derivative series is also absolutely and uniformly convergent on $[-1,1]$. Hence g_A is differentiable on $(-1,1)$, with

$$g_A'(\rho) = \sum_{q \in \mathbb{N}_{\text{odd}}} q \rho^{q-1} q! \|f_q\|^2_{\mathcal{H}(q)} \leq 1. \tag{A.4}$$

Since its partial sums are continuous functions of ρ, the uniform limit theorem implies that g_A' is continuous. Consequently, g_A' is continuous on $(-1,1)$. For odd q, $q-1$ is even, and therefore $\rho^{q-1} \geq 0$ for every $\rho \in [-1,1]$. Thus we also have

$$g_A'(\rho) \geq g_A'(0) = 1! \|f_1\|^2_{\mathcal{H}(1)} = 1 - \frac{8}{\pi^2}, \tag{A.5}$$

and it is worth noting that $c'(0) = \big(g_A'(0) + 1\big)/4 = 1/2 - 2/\pi^2$. By (A.2) we have

$$g_\lambda'(\rho) = \frac{(1-\lambda) + \lambda g_A'(\rho)}{(1-\lambda) + \lambda(3 - 4\log 2)}.$$

The continuity of g_A' on $(-1,1)$ thus implies the continuity of g_λ' on $(-1,1)$. Moreover, we conclude from (A.4) that

$$g_\lambda'(\rho) \leq \frac{(1-\lambda) + \lambda}{(1-\lambda) + \lambda(3 - 4\log 2)} \leq \frac{1}{3 - 4\log 2}.$$

Also, it follows from (A.5) that

$$g'_\lambda(\rho) \geq \frac{(1-\lambda)+\lambda(1-8/\pi^2)}{(1-\lambda)+\lambda(3-4\log 2)}.$$

Since $1-8/\pi^2 < 3-4\log 2$, the right-hand side is decreasing in $\lambda \in [0,1]$, and hence is bounded below by its value at $\lambda = 1$. Thus

$$g'_\lambda(\rho) \geq \frac{1-8/\pi^2}{3-4\log 2} \equiv K_g > 0.$$

The above result shows that g'_λ is continuous on $(-1,1)$, and hence $g_\lambda \in C^1([-1,1])$. Moreover, since $g'_\lambda(\rho) \geq K_g > 0$, g_λ is strictly increasing on $[-1,1]$. At $\rho = 1$, we have $g_A(1) = 3 - 4\log 2$, and hence $g_\lambda(1) = 1$. By oddness, $g_\lambda(-1) = -1$. Therefore g_λ is a continuous strictly increasing map from $[-1,1]$ onto $[-1,1]$, and is consequently a bijection.

As g_λ is bijective, its inverse g_λ^{-1} exists. Using the bounds of g'_λ, the inverse function theorem implies that $g_\lambda^{-1} \in C^1([-1,1])$, with $(g_\lambda^{-1})'(u) = g'_\lambda\big(g_\lambda^{-1}(u)\big)^{-1}$ for all $u \in (-1,1)$, see Rudin (1976, Theorem 9.24). Using the bounds on g'_λ, we obtain

$$3 - 4\log 2 \leq (g_\lambda^{-1})'(u) \leq K_g^{-1}.$$

By the mean value theorem, we conclude that $g_\lambda^{-1}(\cdot)$ is Lipschitz on $[-1,1]$ with the constant K_g^{-1}, which completes the proof. □

Proof of Theorem 1. We begin with a coupling result in the spirit of Bollerslev et al. (2024a) and Li et al. (2024) which holds for every $1 \leq i \leq k_n$. To this end, let us first define the linearly combined Brownian motion for $t \in \bigcup_{i=1}^{k_n} \mathcal{I}_{n,i}$,

$$\widetilde{W}_t^{(m)} = \sigma_\tau^{(m,1)} W_t^{(1)} + \sigma_\tau^{(m,2)} W_t^{(2)}, \tag{A.6}$$

such that $\widetilde{W}_t = (\widetilde{W}_t^{(1)}, \widetilde{W}_t^{(2)})^\top$ satisfies $\widetilde{W}_t = \sigma_\tau W_t$ for $t \in \bigcup_{i=1}^{k_n} \mathcal{I}_{n,i}$. Define the following coupled high and low returns over $\mathcal{I}_{n,i}$,

$$\widetilde{H}_{n,i}^{(m)} = \Delta_n^{-1/2}\Big(\sup_{s\in\mathcal{I}_{n,i}} \widetilde{W}_s^{(m)} - \widetilde{W}_{\tau_n+(i-1)\Delta_n}^{(m)}\Big), \quad \widetilde{L}_{n,i}^{(m)} = \Delta_n^{-1/2}\Big(\inf_{s\in\mathcal{I}_{n,i}} \widetilde{W}_s^{(m)} - \widetilde{W}_{\tau_n+(i-1)\Delta_n}^{(m)}\Big),$$

and the coupled return $\widetilde{C}_{n,i}^{(m)} = \Delta_n^{-1/2}(\widetilde{W}_{\tau_n+i\Delta_n}^{(m)} - \widetilde{W}_{\tau_n+(i-1)\Delta_n}^{(m)})$, where $\tau_n \equiv i^n(\tau)\Delta_n$ is the left endpoint of the local window $(\mathcal{I}_{n,i})_{1\leq i\leq k_n}$ with $\tau_n - \tau \in [0,\Delta_n), \forall n$.

We prove the following L_2-bound of the coupling error.

Lemma A.3. *Under Assumption 1, it holds for all $1 \leq i \leq k_n$, $Z \in \{C,H,L\}$, and $m \in \{1,2\}$ that*

$$\|\Delta_n^{-1/2} Z_{n,i}^{(m)} - \widetilde{Z}_{n,i}^{(m)}\|_2 \leq K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa}. \tag{A.7}$$

Proof. Note that the strengthened version of Assumption 1 implies the following κ-Hölder continuity in L_2-norm for all elements of σ

$$\|\sigma_t^{(m,m')} - \sigma_s^{(m,m')}\|_2 \leq K|t-s|^\kappa,$$

for all $m, m' \in \{1,2\}$ and $t, s \in [0,T]$. We now verify Equation (A.7). Starting with the returns $Z = C$, by the triangle inequality

$$\begin{aligned}
\|\Delta_n^{-1/2} C_{n,i}^{(m)} - \widetilde{C}_{n,i}^{(m)}\|_2 &\leq K\Delta_n^{-1/2}\left(\left\|\int_{\mathcal{I}_{n,i}} |b_s^{(m)}|ds\right\|_2 + \sum_{\ell=1}^{2}\left\|\int_{\mathcal{I}_{n,i}} (\sigma_t^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_s^{(\ell)}\right\|_2\right) \\
&\leq K\Delta_n^{1/2} + \Delta_n^{-1/2}\sum_{\ell=1}^{2}\left\|\int_{\mathcal{I}_{n,i}} (\sigma_t^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_s^{(\ell)}\right\|_2,
\end{aligned}$$

where the last line follows from the boundedness of b. For the last term above, by the Itô isometry, the κ-Hölder continuity of $\sigma_t^{(m,m')}$, and the fact that $|t-\tau| \leq (k_n+1)\Delta_n \leq Kk_n\Delta_n$ for all $t \in \bigcup_{i=1}^{k_n} \mathcal{I}_{n,i}$, we deduce that, for $m, \ell \in \{1,2\}$,

$$\left\|\int_{\mathcal{I}_{n,i}} (\sigma_t^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_t^{(\ell)}\right\|_2^2 = \int_{\mathcal{I}_{n,i}} \|\sigma_t^{(m,\ell)} - \sigma_\tau^{(m,\ell)}\|_2^2 dt \leq K\Delta_n(k_n\Delta_n)^{2\kappa},$$

and hence $\|\Delta_n^{-1/2} C_{n,i}^{(m)} - \widetilde{C}_{n,i}^{(m)}\|_2 \leq K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa}$ with $k_n \asymp \Delta_n^{-\beta}$, as desired. Now for the high return H, by repeated application of the triangle inequality,

$$\begin{aligned}
|\Delta_n^{-1/2} H_{n,i}^{(m)} - \widetilde{H}_{n,i}^{(m)}| &\leq \Delta_n^{-1/2}\sup_{t\in\mathcal{I}_{n,i}}\left|\int_{\tau_n+(i-1)\Delta_n}^{t} |b_s^{(m)}|ds + \sum_{\ell=1}^{2}\int_{\tau_n+(i-1)\Delta_n}^{t} (\sigma_s^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_s^{(\ell)}\right| \\
&\leq \Delta_n^{-1/2}\int_{t\in\mathcal{I}_{n,i}} |b_s^{(m)}|ds + \Delta_n^{-1/2}\sum_{\ell=1}^{2}\sup_{t\in\mathcal{I}_{n,i}}\left|\int_{\tau_n+(i-1)\Delta_n}^{t} (\sigma_s^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_s^{(\ell)}\right| \\
&\leq K\Delta_n^{1/2} + \Delta_n^{-1/2}\sum_{\ell=1}^{2}\sup_{t\in\mathcal{I}_{n,i}}\left|\int_{\tau_n+(i-1)\Delta_n}^{t} (\sigma_s^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_s^{(\ell)}\right|,
\end{aligned}$$

which implies that

$$\|\Delta_n^{-1/2} H_{n,i}^{(m)} - \widetilde{H}_{n,i}^{(m)}\|_2 \leq K\Delta_n^{1/2} + \Delta_n^{-1/2}\sum_{\ell=1}^{2}\left\|\sup_{t\in\mathcal{I}_{n,i}}\left|\int_{\tau_n+(i-1)\Delta_n}^{t} (\sigma_s^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_s^{(\ell)}\right|\right\|_2.$$

For the last term above, by the Burkholder–Davis–Gundy inequality, we have for $m, \ell \in \{1,2\}$

$$\left\|\sup_{t\in\mathcal{I}_{n,i}}\left|\int_{\tau_n+(i-1)\Delta_n}^{t} (\sigma_s^{(m,\ell)} - \sigma_\tau^{(m,\ell)})dW_s^{(\ell)}\right|\right\|_2^2 \leq K\int_{\mathcal{I}_{n,i}} \|\sigma_t^{(m,\ell)} - \sigma_\tau^{(m,\ell)}\|_2^2 dt \leq K\Delta_n(k_n\Delta_n)^{2\kappa}.$$

Therefore, we also have $\|\Delta_n^{-1/2} H_{n,i}^{(m)} - \widetilde{H}_{n,i}^{(m)}\|_2 \leq K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa}$, similar to the derivation for the returns. Finally, for the low return $\|\Delta_n^{-1/2} L_{n,i}^{(m)} - \widetilde{L}_{n,i}^{(m)}\|_2$, it suffices to notice that for a generic process X, $\inf_{t\in\mathcal{I}_{n,i}} X_t = \sup_{t\in\mathcal{I}_{n,i}}\{-X_t\}$, and the required expectation bound follows by replacing inf with sup using essentially the same steps as in the proof for the high returns. This completes the proof. □

Write $\widetilde{A}^{(m)}_{n,i} = \widetilde{H}^{(m)}_{n,i} + \widetilde{L}^{(m)}_{n,i} - \widetilde{C}^{(m)}_{n,i}$. By the triangle inequality, one can also verify that $\|\Delta_n^{-1/2} A^{(m)}_{n,i} - \widetilde{A}^{(m)}_{n,i}\|_2 \leq K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa}$ by Lemma A.3. We now identify the finite sample joint distribution of the tuple $(\widetilde{C}^{(1)}_{n,i}, \widetilde{C}^{(2)}_{n,i}, \widetilde{A}^{(1)}_{n,i}, \widetilde{A}^{(2)}_{n,i})_{1\leq i\leq k_n}$ for all n.

Lemma A.4. *Conditioning on $\mathcal{F}_\tau$, it holds that*

$$(\widetilde{C}^{(1)}_{n,i}, \widetilde{C}^{(2)}_{n,i}, \widetilde{A}^{(1)}_{n,i}, \widetilde{A}^{(2)}_{n,i})_{1\leq i\leq k_n} \overset{d}{=} (\sigma^{(1)}_\tau \zeta^{(1)}_{C,i}, \sigma^{(2)}_\tau \zeta^{(2)}_{C,i}, \sigma^{(1)}_\tau \zeta^{(1)}_{A,i}, \sigma^{(2)}_\tau \zeta^{(2)}_{A,i})_{1\leq i\leq k_n}, \tag{A.8}$$

where $(\zeta^{(1)}_{C,i}, \zeta^{(2)}_{C,i}, \zeta^{(1)}_{A,i}, \zeta^{(2)}_{A,i})_{1\leq i\leq k_n}$ are defined according to Equation (5) and the underlying Brownian motion B satisfies $\mathbb{E}[B^{(1)}_t B^{(2)}_t \mid \mathcal{F}_\tau] = \rho_\tau t, \forall t > 0$.

Proof. Recall from the definition of Σ_τ that

$$\sigma^{(m)}_\tau = \sqrt{(\sigma^{(m,1)}_\tau)^2 + (\sigma^{(m,2)}_\tau)^2}, \quad \rho_\tau = \frac{\sigma^{(1,1)}_\tau \sigma^{(2,1)}_\tau + \sigma^{(1,2)}_\tau \sigma^{(2,2)}_\tau}{\sigma^{(1)}_\tau \sigma^{(2)}_\tau}.$$

By the construction of $\widetilde{W}^{(m)}$ in Equation (A.6), conditioning on $\mathcal{F}_\tau$ we have the following equality in law by the Markov property and the scaling property of the Brownian motion for all n,

$$\Delta_n^{-1/2}(\widetilde{W}^{(1)}_{\tau_n+s\Delta_n} - \widetilde{W}^{(1)}_{\tau_n}, \widetilde{W}^{(2)}_{\tau_n+s\Delta_n} - \widetilde{W}^{(2)}_{\tau_n})_{s\in[0,k_n]} \overset{d}{=} (\sigma^{(1)}_\tau B^{(1)}_s, \sigma^{(2)}_\tau B^{(2)}_s)_{s\in[0,k_n]},$$

which can be verified by computing the $\mathcal{F}_\tau$-conditional variance-covariance matrix for each s on both sides. Now, notice that the following two mappings are identical

$$\begin{aligned}
\Delta_n^{-1/2}(\widetilde{W}^{(1)}_{\tau_n+s\Delta_n} - \widetilde{W}^{(1)}_{\tau_n}, \widetilde{W}^{(2)}_{\tau_n+s\Delta_n} - \widetilde{W}^{(2)}_{\tau_n})_{s\in[i-1,i]} &\mapsto (\widetilde{C}^{(1)}_{n,i}, \widetilde{C}^{(2)}_{n,i}, \widetilde{A}^{(1)}_{n,i}, \widetilde{A}^{(2)}_{n,i}) \\
(\sigma^{(1)}_\tau B^{(1)}_s, \sigma^{(2)}_\tau B^{(2)}_s)_{s\in[i-1,i]} &\mapsto (\sigma^{(1)}_\tau \zeta^{(1)}_{C,i}, \sigma^{(2)}_\tau \zeta^{(2)}_{C,i}, \sigma^{(1)}_\tau \zeta^{(1)}_{A,i}, \sigma^{(2)}_\tau \zeta^{(2)}_{A,i}).
\end{aligned}$$

Therefore, under the pushforward measure and by combining the results for all $1 \leq i \leq k_n$ exploiting independence across different i, we deduce Equation (A.8) conditional on $\mathcal{F}_\tau$, and the proof is complete. □

We proceed to study the sample mean over k_n blocks for the original and the coupled candlestick features

$$\begin{aligned}
S^{(m,m')}_n &= \frac{1}{k_n\Delta_n}\sum_{i=1}^{k_n}\big\{(1-\lambda)C^{(m)}_{n,i}C^{(m')}_{n,i} + \lambda A^{(m)}_{n,i}A^{(m')}_{n,i}\big\}, \\
\widetilde{S}^{(m,m')}_n &= \frac{1}{k_n}\sum_{i=1}^{k_n}\big\{(1-\lambda)\widetilde{C}^{(m)}_{n,i}\widetilde{C}^{(m')}_{n,i} + \lambda \widetilde{A}^{(m)}_{n,i}\widetilde{A}^{(m')}_{n,i}\big\}.
\end{aligned}$$

We shall also need the following existence of moments result for Brownian candlesticks, and recall the definition of $\zeta^{(m)}_{C,i}$, $\zeta^{(m)}_{H,i}$, $\zeta^{(m)}_{L,i}$ and $\zeta^{(m)}_{A,i}$ in Equations (5) and (6).

Lemma A.5. *For any $k \in \mathbb{N}$, $\lambda \in [0,1]$, $m, m' \in \{1,2\}$, and $\rho \in [-1,1]$, it holds that:*

$$\mathbb{E}\big[\big|(1-\lambda)\zeta^{(m)}_{C,1}\zeta^{(m')}_{C,1} + \lambda\zeta^{(m)}_{A,1}\zeta^{(m')}_{A,1}\big|^k\big] < \infty.$$

Proof. We begin with the elementary fact that, for univariate Brownian motions $m \in \{1,2\}$, it holds that $\mathbb{E}[|\zeta_{C,1}^{(m)}|^k] = \mathbb{E}[|\zeta_{H,1}^{(m)}|^k] = \mathbb{E}[|\zeta_{L,1}^{(m)}|^k] < \infty$ for all $k \in \mathbb{N}$, where we used $|\zeta_{C,1}^{(m)}| \overset{d}{=} \zeta_{H,1}^{(m)} \overset{d}{=} -\zeta_{L,1}^{(m)}$ by the reflection principle and the reflection symmetry of Brownian motion. Therefore, by the elementary inequality $|a+b+c|^k \le K_k(|a|^k + |b|^k + |c|^k)$ for some $K_k \ge 3^{k-1}$, we have:

$$\mathbb{E}[|\zeta_{A,1}^{(m)}|^k] \le K_k(\mathbb{E}[|\zeta_{C,1}^{(m)}|^k] + \mathbb{E}[|\zeta_{H,1}^{(m)}|^k] + \mathbb{E}[|\zeta_{L,1}^{(m)}|^k]) < \infty.$$

Then, by the triangle inequality and the Cauchy–Schwarz inequality:

$$\begin{aligned}
&\mathbb{E}\big[\big|(1-\lambda)\zeta_{C,1}^{(m)}\zeta_{C,1}^{(m')} + \lambda\zeta_{A,1}^{(m)}\zeta_{A,1}^{(m')}\big|^k\big] \\
&\quad \le K_k\big(\mathbb{E}[\|\zeta_{C,1}^{(m)}\zeta_{C,1}^{(m')}\|^k] + \mathbb{E}[\|\zeta_{A,1}^{(m)}\zeta_{A,1}^{(m')}\|^k]\big) \\
&\quad \le K_k\big(\|(\zeta_{C,1}^{(m)})^k\|_2\|(\zeta_{C,1}^{(m')})^k\|_2 + \|(\zeta_{A,1}^{(m)})^k\|_2\|(\zeta_{A,1}^{(m')})^k\|_2\big) < \infty,
\end{aligned}$$

as desired. This completes the proof. □

Lemmas A.3-A.5 lead to the following coupling relation.

Lemma A.6. *For $m, m' \in \{1,2\}$ and $\lambda \in [0,1]$, it holds that*

$$\big|S_n^{(m,m')} - \widetilde{S}_n^{(m,m')}\big| = O_p(\Delta_n^{1/2\wedge(1-\beta)\kappa}). \tag{A.9}$$

Proof. By the triangle inequality and that $\lambda \vee (1-\lambda) \le 1$,

$$\begin{aligned}
&\big|S_n^{(m,m')} - \widetilde{S}_n^{(m,m')}\big| \\
&= \left|\frac{1}{k_n\Delta_n}\sum_{i=1}^{k_n}\big\{(1-\lambda)C_{n,i}^{(m)}C_{n,i}^{(m')} + \lambda A_{n,i}^{(m)}A_{n,i}^{(m')}\big\} - \frac{1}{k_n}\sum_{i=1}^{k_n}\big\{(1-\lambda)\widetilde{C}_{n,i}^{(m)}\widetilde{C}_{n,i}^{(m')} + \lambda\widetilde{A}_{n,i}^{(m)}\widetilde{A}_{n,i}^{(m')}\big\}\right| \\
&\le \frac{1}{k_n}\sum_{i=1}^{k_n}\big|\Delta_n^{-1}C_{n,i}^{(m)}C_{n,i}^{(m')} - \widetilde{C}_{n,i}^{(m)}\widetilde{C}_{n,i}^{(m')}\big| + \frac{1}{k_n}\sum_{i=1}^{k_n}\big|\Delta_n^{-1}A_{n,i}^{(m)}A_{n,i}^{(m')} - \widetilde{A}_{n,i}^{(m)}\widetilde{A}_{n,i}^{(m')}\big|.
\end{aligned}$$

For the first sum above, using the triangle inequality, the identity $ab - cd = c(b-d) + d(a-c) + (a-c)(b-d)$, the Cauchy–Schwarz inequality, and Lemma A.3, we have

$$\begin{aligned}
&\left\|\frac{1}{k_n}\sum_{i=1}^{k_n}\big|\Delta_n^{-1}C_{n,i}^{(m)}C_{n,i}^{(m')} - \widetilde{C}_{n,i}^{(m)}\widetilde{C}_{n,i}^{(m')}\big|\right\|_1 \\
&\le \frac{1}{k_n}\sum_{i=1}^{k_n}\Big(\big\|\widetilde{C}_{n,i}^{(m)}(\Delta_n^{-1/2}C_{n,i}^{(m')} - \widetilde{C}_{n,i}^{(m')})\big\|_1 + \big\|\widetilde{C}_{n,i}^{(m')}(\Delta_n^{-1/2}C_{n,i}^{(m)} - \widetilde{C}_{n,i}^{(m)})\big\|_1 \\
&\qquad + \big\|(\Delta_n^{-1/2}C_{n,i}^{(m')} - \widetilde{C}_{n,i}^{(m')})(\Delta_n^{-1/2}C_{n,i}^{(m)} - \widetilde{C}_{n,i}^{(m)})\big\|_1\Big) \\
&\le (K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa})(1 + K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa}) \\
&\le K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa}.
\end{aligned} \tag{A.10}$$

where the penultimate line follows from the fact that $\|\widetilde{C}_{n,i}^{(m)}\|_2^2 = \|\sigma_\tau^{(m)}\|_2^2 \le K$ by the boundedness of $\sigma^{(m)}$. Using exactly the same decomposition and that $\|\widetilde{A}_{n,i}^{(m)}\|_2^2 \le K$ again by the boundedness of $\sigma^{(m)}$ and Lemmas A.4 and A.5, we conclude that

$$\left\|\frac{1}{k_n}\sum_{i=1}^{k_n}\big|\Delta_n^{-1}A_{n,i}^{(m)}A_{n,i}^{(m')} - \widetilde{A}_{n,i}^{(m)}\widetilde{A}_{n,i}^{(m')}\big|\right\|_1 \le K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge(1-\beta)\kappa}. \tag{A.11}$$

Combining Equations (A.10) and (A.11), we deduce (A.9) by the Markov's inequality, and the proof is complete. □

Lemma A.6 suggests that $\big|S_n^{(m,m')} - \widetilde{S}_n^{(m,m')}\big| = o_p(1)$ for all $\beta < 1$, and if $\beta < 2\kappa/(2\kappa+1)$, then we further have $\sqrt{k_n}\big|S_n^{(m,m')} - \widetilde{S}_n^{(m,m')}\big| = o_p(1)$. As a result, we only need to prove the convergence in probability and the convergence in distribution for $\widetilde{S}_n^{(m,m')}$, as the same convergence also holds for $S_n^{(m,m')}$ by Lemma A.1.

We proceed to study the asymptotic behavior of $\widetilde{S}_n^{(m,m')}$ as $k_n \to \infty$. We begin by identifying the limit of $\widetilde{S}_n^{(m,m')}$.

Lemma A.7. *Let $\theta_\lambda = 1 + 2\lambda(1 - 2\log 2)$. It holds as $k_n \to \infty$ that*

$$\widetilde{S}_n^{(m,m)} \xrightarrow{\mathbb{P}} \theta_\lambda (\sigma_\tau^{(m)})^2, \quad \widetilde{S}_n^{(1,2)} \xrightarrow{\mathbb{P}} \theta_\lambda \sigma_\tau^{(1)} \sigma_\tau^{(2)} g_\lambda(\rho_\tau).$$

Proof. We shall only prove the convergence for $\widetilde{S}_n^{(1,2)}$, as the same proof applies to $\widetilde{S}_n^{(m,m)}$ in the case of perfect correlation, i.e., $g_\lambda(1) = 1$. We begin with the following fact, which holds for all i and n,

$$\mathbb{E}[(1-\lambda)\widetilde{C}_{n,i}^{(1)}\widetilde{C}_{n,i}^{(2)} + \lambda \widetilde{A}_{n,i}^{(1)}\widetilde{A}_{n,i}^{(2)} \mid \mathcal{F}_{\tau_n + (i-1)\Delta_n}] = \theta_\lambda \sigma_\tau^{(1)}\sigma_\tau^{(2)} g_\lambda(\rho_\tau).$$

This follows directly from Lemma A.4 and the Brownian moment identity in (8):

$$\mathbb{E}[(1-\lambda)\zeta_{C,1}^{(1)}\zeta_{C,1}^{(2)} + \lambda \zeta_{A,1}^{(1)}\zeta_{A,1}^{(2)}] = \theta_\lambda g_\lambda(\rho). \tag{A.12}$$

Also, by Lemmas A.4 and A.5, we deduce

$$\mathbb{E}\big[\big((1-\lambda)\widetilde{C}_{n,i}^{(1)}\widetilde{C}_{n,i}^{(2)} + \lambda \widetilde{A}_{n,i}^{(1)}\widetilde{A}_{n,i}^{(2)}\big)^2 \mid \mathcal{F}_{\tau_n + (i-1)\Delta_n}\big] < \infty,$$

and hence

$$\frac{1}{k_n^2}\mathbb{E}\big[\big((1-\lambda)\widetilde{C}_{n,i}^{(1)}\widetilde{C}_{n,i}^{(2)} + \lambda \widetilde{A}_{n,i}^{(1)}\widetilde{A}_{n,i}^{(2)}\big)^2 \mid \mathcal{F}_{\tau_n + (i-1)\Delta_n}\big] \to 0.$$

Then, by Jacod and Protter (2012, Lemma 2.2.11), the triangular array $(z_i^n)_{1 \le i \le k_n}$ defined by

$$z_i^n = \frac{1}{k_n}\big((1-\lambda)\widetilde{C}_{n,i}^{(1)}\widetilde{C}_{n,i}^{(2)} + \lambda \widetilde{A}_{n,i}^{(1)}\widetilde{A}_{n,i}^{(2)} - \theta_\lambda \sigma_\tau^{(1)}\sigma_\tau^{(2)} g_\lambda(\rho_\tau)\big),$$

is asymptotically negligible, which implies that $\sum_{i=1}^{k_n} z_i^n \xrightarrow{\mathbb{P}} 0$. This completes the proof. □

By Lemma A.2, $g_\lambda^{-1}(\cdot)$ is continuous on $[-1, 1]$. Therefore, Lemma A.7 and the continuous mapping theorem lead to

$$\tilde{\rho}_n(\lambda) \equiv g_\lambda^{-1}\left(\frac{\widetilde{S}_n^{(1,2)}}{\sqrt{\widetilde{S}_n^{(1,1)}\widetilde{S}_n^{(2,2)}}}\right) \xrightarrow{\mathbb{P}} \rho_\tau. \tag{A.13}$$

Also, using the Lipschitz property of $g_\lambda^{-1}(\cdot)$ on $[-1,1]$ by Lemma A.2, Lemmas A.6 and A.7 give the following estimate

$$\begin{aligned}
|\hat{\rho}_n(\lambda) - \tilde{\rho}_n(\lambda)| &= \left| g_\lambda^{-1}\left(\frac{S_n^{(1,2)}}{\sqrt{S_n^{(1,1)} S_n^{(2,2)}}} \right) - \tilde{\rho}_n(\lambda) \right| \\
&\leq K \left| \frac{\widetilde{S}_n^{(1,2)} + o_p(1)}{\sqrt{\widetilde{S}_n^{(1,1)} \widetilde{S}_n^{(2,2)}} + o_p(1)} - \frac{\widetilde{S}_n^{(1,2)}}{\sqrt{\widetilde{S}_n^{(1,1)} \widetilde{S}_n^{(2,2)}}} \right| \\
&\leq K \left| \frac{(\sqrt{\widetilde{S}_n^{(1,1)} \widetilde{S}_n^{(2,2)}} - \widetilde{S}_n^{(1,2)}) \cdot o_p(1)}{\widetilde{S}_n^{(1,1)} \widetilde{S}_n^{(2,2)} + \sqrt{\widetilde{S}_n^{(1,1)} \widetilde{S}_n^{(2,2)}} \cdot o_p(1)} \right| = o_p(1).
\end{aligned} \tag{A.14}$$

Combining Equations (A.13) with (A.14) yields (10).

Moving on to the proof of (11). From Lemma A.7 and under the additional condition $\beta < 2\kappa/(2\kappa+1)$, we can replace all the $o_p(1)$ terms by $o_p(k_n^{-1/2})$ which improves the estimate in (A.14) to $\sqrt{k_n}|\hat{\rho}_n(\lambda) - \tilde{\rho}_n(\lambda)| = o_p(1)$. Therefore, it suffices to prove the same $\mathcal{F}_\tau$-conditional convergence for $\tilde{\rho}_n(\lambda)$ to conclude (11) using the asymptotic equivalence lemma. To prove this, consider the following vector-valued triangular array $\eta_{n,i} = (\eta_{n,i}^{(1,2)}, \eta_{n,i}^{(1,1)}, \eta_{n,i}^{(2,2)})^\top$, where

$$\begin{aligned}
\eta_{n,i}^{(1,2)} &= \frac{1}{\sqrt{k_n}} \left(\frac{(1-\lambda)\widetilde{C}_{n,i}^{(1)} \widetilde{C}_{n,i}^{(2)} + \lambda \widetilde{A}_{n,i}^{(1)} \widetilde{A}_{n,i}^{(2)}}{\theta_\lambda \sigma_\tau^{(1)} \sigma_\tau^{(2)}} - g_\lambda(\rho_\tau) \right), \\
\eta_{n,i}^{(m,m)} &= \frac{1}{\sqrt{k_n}} \left(\frac{(1-\lambda)\widetilde{C}_{n,i}^{(m)} \widetilde{C}_{n,i}^{(m)} + \lambda \widetilde{A}_{n,i}^{(m)} \widetilde{A}_{n,i}^{(m)}}{\theta_\lambda (\sigma_\tau^{(m)})^2} - 1 \right), \quad m \in \{1,2\},
\end{aligned}$$

which is, by Equation (A.12), a vector-valued martingale difference sequence for each n with all moments finite. By Lemma A.4, we see that

$$\mathbb{E}[\eta_{n,i}\eta_{n,i}^\top \mid \mathcal{F}_{\tau_n + (i-1)\Delta_n}] = \frac{1}{k_n}\Sigma_\lambda(\rho_\tau), \quad \forall 1 \leq i \leq k_n,$$

where $\Sigma_\lambda(\rho) = \left(\Sigma_\lambda^{(m,m')}(\rho)\right)_{m,m' \in \{1,2,3\}}$ is a 3-by-3 symmetric matrix with the following entries:

$$\begin{aligned}
\Sigma_\lambda^{(1,1)}(\rho) &= \mathsf{Var}(\zeta_\lambda^{(1,2)}), \quad \Sigma_\lambda^{(1,2)}(\rho) = \Sigma_\lambda^{(1,3)}(\rho) = \mathsf{Cov}(\zeta_\lambda^{(1,2)}, \zeta_\lambda^{(1,1)}), \\
\Sigma_\lambda^{(2,2)}(\rho) &= \Sigma_\lambda^{(3,3)}(\rho) = \mathsf{Var}(\zeta_\lambda^{(1,1)}), \quad \Sigma_\lambda^{(2,3)}(\rho) = \mathsf{Cov}(\zeta_\lambda^{(1,1)}, \zeta_\lambda^{(2,2)}),
\end{aligned}$$

in which $\zeta_\lambda^{(m,m')} = \left((1-\lambda)\zeta_{C,1}^{(m)}\zeta_{C,1}^{(m')} + \lambda\zeta_{A,1}^{(m)}\zeta_{A,1}^{(m')}\right)/\theta_\lambda$. Also, for any fixed $\varepsilon > 0$ and $a \in \mathbb{R}^3$,

$$\begin{aligned}
&\sum_{i=1}^{k_n} \mathbb{E}[(a^\top \eta_{n,i})^2 \mathbb{1}_{\{|a^\top \eta_{n,i}| > \varepsilon\}} \mid \mathcal{F}_{\tau_n + (i-1)\Delta_n}] \leq \frac{1}{\varepsilon^2} \sum_{i=1}^{k_n} \mathbb{E}[(a^\top \eta_{n,i})^4 \mid \mathcal{F}_{\tau_n + (i-1)\Delta_n}] \\
\leq &\frac{K\|a\|_\infty^4}{\varepsilon^2} \sum_{i=1}^{k_n} k_n^{-2} \mathbb{E}[(\zeta_\lambda^{(1,1)})^4 + (\zeta_\lambda^{(2,2)})^4 + (\zeta_\lambda^{(1,2)})^4] \xrightarrow{\mathbb{P}} 0,
\end{aligned} \tag{A.15}$$

where $\|a\|_\infty$ is the sup norm of a, and the last convergence follows by Lemma A.5. Therefore, by Jacod and Protter (2012, Lemma 2.2.11) we deduce:

$$\sum_{i=1}^{k_n} \mathbb{E}[\eta_{n,i}\eta_{n,i}^\top \mid \mathcal{F}_{\tau_n + (i-1)\Delta_n}] \xrightarrow{\mathbb{P}} \Sigma_\lambda(\rho_\tau) \in \mathcal{F}_\tau. \tag{A.16}$$

By the Cramér–Wold device and using the corollary of Theorem 3 of Eagleson (1975) with the two conditions in Equations (A.15) and (A.16) satisfied, we conclude that:

$$\sqrt{k_n}\left[\begin{pmatrix}(\theta_\lambda\sigma_\tau^{(1)}\sigma_\tau^{(2)})^{-1}\widetilde{S}_n^{(1,2)}\\ \left(\theta_\lambda(\sigma_\tau^{(1)})^2\right)^{-1}\widetilde{S}_n^{(1,1)}\\ \left(\theta_\lambda(\sigma_\tau^{(2)})^2\right)^{-1}\widetilde{S}_n^{(2,2)}\end{pmatrix}-\begin{pmatrix}g_\lambda(\rho_\tau)\\ 1\\ 1\end{pmatrix}\right]\equiv\sum_{i=1}^{k_n}\eta_{n,i}\xrightarrow{\mathcal{L}|\mathcal{F}_\tau}\Sigma_\lambda(\rho_\tau)^{1/2}Z_3, \tag{A.17}$$

where Z_3 is a 3-dimensional spherical standard Gaussian vector defined on an extension of $\big(\Omega,\mathcal{F},(\mathcal{F}_t)_{t\geq 0},\mathbb{P}\big)$ independent of $\mathcal{F}$, and $\Sigma_\lambda(\rho_\tau)^{1/2}$ is the (unique) square-root matrix of $\Sigma_\lambda(\rho_\tau)$.

Now, by Lemma A.2, $u\mapsto g_\lambda^{-1}(u)$ is continuously differentiable on $(-1,1)$ for all $\lambda\in[0,1]$ with derivative bounded and bounded away from zero. Since ρ_τ is bounded away from ±1 by the strengthened version of Assumption 1, $g_\lambda(\rho_\tau)$ is bounded away from ±1, and hence g_λ^{-1} is differentiable at $g_\lambda(\rho_\tau)$ with $0<(g_\lambda^{-1})'(g_\lambda(\rho_\tau))=g_\lambda'(\rho_\tau)^{-1}<\infty$. An application of the multivariate delta method to (A.17) under the mapping $(x,y,z)^\top\mapsto g_\lambda^{-1}(x/\sqrt{yz})$ yields:

$$\sqrt{k_n}(\tilde{\rho}_n(\lambda)-\rho_\tau)\equiv\sqrt{k_n}\left[g_\lambda^{-1}\left(\frac{\widetilde{S}_n^{(1,2)}}{\sqrt{\widetilde{S}_n^{(1,1)}\widetilde{S}_n^{(2,2)}}}\right)-\rho_\tau\right]\xrightarrow{\mathcal{L}|\mathcal{F}_\tau}g_\lambda'(\rho_\tau)^{-1}\sqrt{V_\lambda(\rho_\tau)}Z, \tag{A.18}$$

where, with $\Xi_\lambda(\rho_\tau)\equiv\zeta_\lambda^{(1,2)}-g_\lambda(\rho_\tau)(\zeta_\lambda^{(1,1)}+\zeta_\lambda^{(2,2)})/2$ and noting that $\mathbb{E}[\Xi_\lambda(\rho_\tau)]=0$,

$$\begin{aligned}V_\lambda(\rho_\tau)&=\mathsf{Var}(\Xi_\lambda(\rho_\tau))=\mathbb{E}[\Xi_\lambda(\rho_\tau)^2]\\&=\Sigma_\lambda^{(1,1)}(\rho_\tau)+\frac{g_\lambda(\rho_\tau)^2}{2}\big(\Sigma_\lambda^{(2,2)}(\rho_\tau)+\Sigma_\lambda^{(2,3)}(\rho_\tau)\big)-2g_\lambda(\rho_\tau)\Sigma_\lambda^{(1,2)}(\rho_\tau).\end{aligned}$$

We note that $V_\lambda(\rho_\tau)>0$ for all $\rho_\tau\in(-1,1)$, so the limit in Eq. (A.18) is non-degenerate. To see this, it suffices to prove that $\Xi_\lambda(\rho_\tau)$ is not almost surely 0 for $\rho_\tau\in(-1,1)$. Indeed, Lemma A.2 gives $g_\lambda(\rho)\in(-1,1)$. Now pick any deterministic continuous paths h that generate $\zeta_{C,1}^{(m)}=\zeta_{A,1}^{(m)}=1$ for $m\in\{1,2\}$, e.g., a linear interpolation of $h(0)=0$, $h(t)=2$ and $h(1)=1$ for $t\in(0,1)$. On the bivariate path (h,h), we have $\zeta_\lambda^{(m,m')}\equiv\theta_\lambda^{-1}$ and hence $\Xi_\lambda(\rho_\tau)=(1-g_\lambda(\rho_\tau))/\theta_\lambda>0$. By continuity of the candlestick functionals, there exists an open neighbourhood of (h,h) such that $\Xi_\lambda(\rho_\tau)>0$. Since $(B^{(1)},B^{(2)})$ has full support on $C_0([0,1])^2$, such an open neighbourhood has positive probability measure, and hence $\Xi_\lambda(\rho_\tau)$ is not almost surely zero. It is also easy to verify that $V_\lambda(\pm1)=0$, which is prohibited by the bound of ρ_τ in Assumption 1.

Combining Equations (A.18) and the estimate $\sqrt{k_n}|\hat{\rho}_n(\lambda)-\tilde{\rho}_n(\lambda)|=o_p(1)$ yields (11) by Lemma A.1(ii), and the proof is complete. □

Proof of Corollary 1. By the $\mathcal{F}_\tau$-conditional CLT in Theorem 1, the feasible plug-in CLT follows from the continuous mapping theorem and Slutsky's theorem provided that, on $(-1,1)$: (1) g_λ' and V_λ are positive, and (2) g_λ' and V_λ are continuous in ρ. Lemma A.2 proves that $g_\lambda'>0$ and the continuity of g_λ', and $V_\lambda>0$ is proved in the proof of Theorem 1.

It remains to prove that V_λ is continuous in ρ. Following the same technique for the continuity of g_λ in the proof of Lemma A.2, we take any sequence $\rho_n\to\rho\in(-1,1)$. Since $\Xi_\lambda(\rho)^2$ is

a continuous function of the candlestick functionals, which are in turn continuous in the Brownian paths that are continuous in ρ almost surely, we have $\Xi_\lambda(\rho_n)^2 \to \Xi_\lambda(\rho)^2$ almost surely as $\rho_n \to \rho$. Furthermore, recall $M \equiv \sup_{0\le t\le 1}|B_t| + \sup_{0\le t\le 1}|B_t^\perp|$. For $m, m' \in \{1,2\}$ and M constructed from the corresponding driving Brownian motions, we have the following bound independent of ρ and $\lambda \in [0,1]$:

$$|\zeta_\lambda^{(m,m')}| \le \frac{(1-\lambda)|\zeta_{C,1}^{(m)}\zeta_{C,1}^{(m')}| + \lambda|\zeta_{A,1}^{(m)}\zeta_{A,1}^{(m')}|}{\theta_\lambda} \le \frac{1+8\lambda}{\theta_\lambda}M^2 \le KM^2,$$

where the last estimate uses $(1+8\lambda)/\theta_\lambda \le 9/(3-4\log 2)$ for $\lambda \in [0,1]$. Therefore,

$$\Xi_\lambda(\rho)^2 \le \left(|\zeta_\lambda^{(1,2)}| + \frac{|g_\lambda(\rho_\tau)|}{2}(|\zeta_\lambda^{(1,1)}| + |\zeta_\lambda^{(2,2)}|)\right)^2 \le KM^4.$$

Recall from the that proof of Lemma A.2 that $\mathbb{E}[M^4] < \infty$. Applying the dominated convergence theorem and using the continuity of g_λ, we conclude that

$$V_\lambda(\rho_n) = \mathbb{E}[\Xi_\lambda(\rho_n)^2] \to \mathbb{E}[\Xi_\lambda(\rho)^2] = V_\lambda(\rho),$$

proving the continuity of V_λ on $(-1,1)$. This completes the proof. □

Proof of Theorem 2. Recall that under the fixed-k setting, we have $\beta = 0$. Therefore, Lemma A.3 gives

$$\begin{aligned}\|\Delta_n^{-1/2}C_{n,i}^{(m)} - \widetilde{C}_{n,i}^{(m)}\|_2 &\le K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge\kappa},\\ \|\Delta_n^{-1/2}A_{n,i}^{(m)} - \widetilde{A}_{n,i}^{(m)}\|_2 &\le K\Delta_n^{1/2} + K'\Delta_n^{1/2\wedge\kappa},\end{aligned}$$

which implies that, jointly for all $1 \le i \le k$,

$$\Delta_n^{-1/2}(C_{n,i}^{(1)}, C_{n,i}^{(2)}, A_{n,i}^{(1)}, A_{n,i}^{(2)}) = (\widetilde{C}_{n,i}^{(1)}, \widetilde{C}_{n,i}^{(2)}, \widetilde{A}_{n,i}^{(1)}, \widetilde{A}_{n,i}^{(2)}) + o_p(1).$$

By Lemma A.1(i), for all $\lambda \in [0,1]$, we further have

$$\hat{\rho}_n(\lambda) = \tilde{\rho}_n(\lambda) + o_p(1), \tag{A.19}$$

and recall $\tilde{\rho}_n(\lambda)$ defined in Equation (A.13).

It remains to notice that, by Lemma A.4 and conditioning on $\mathcal{F}_\tau$, we have the following $\mathcal{F}_\tau$-conditional equal in distribution independent of n

$$\tilde{\rho}_n(\lambda) \overset{\mathcal{L}|\mathcal{F}_\tau}{=} \tilde{\rho}(\lambda;\rho_\tau).$$

The above result directly leads to the desired $\mathcal{F}_\tau$-conditional convergence in law in Equation (17) by (A.19) and the asymptotic equivalence lemma. This completes the proof. □

Proof of Corollary 2. By the assumed $\mathcal{F}_\tau$-conditional non-atomicity of $\tilde{\rho}_\lambda(\rho_\tau)$, its boundary has zero conditional probability

$$\mathbb{P}\big(\tilde{\rho}_\lambda(\rho_\tau) \in \partial I_{1-\alpha}(\lambda;\rho_\tau) \mid \mathcal{F}_\tau\big) = 0, \quad \text{a.s.}$$

Hence $I_{1-\alpha}(\lambda;\rho_\tau)$ is a conditional continuity set for the limiting law. Therefore, Theorem 2 and the conditional portmanteau theorem give:

$$\mathbb{P}\big(\hat{\rho}_n(\lambda) \in I_{1-\alpha}(\lambda;\rho_\tau) \mid \mathcal{F}_\tau\big) \xrightarrow{\mathbb{P}} \mathbb{P}\big(\tilde{\rho}(\lambda;\rho_\tau) \in I_{1-\alpha}(\lambda;\rho_\tau) \mid \mathcal{F}_\tau\big) = 1-\alpha,$$

where the equality follows from the construction of $I_\lambda(\alpha;\rho_\tau)$ and the non-atomic condition. Now by the definition of the inverted confidence set

$$\{\rho_\tau \in \mathrm{CS}_{n,1-\alpha}(\lambda)\} = \{\hat{\rho}_n(\lambda) \in I_{1-\alpha}(\lambda;\rho_\tau)\}.$$

Therefore,

$$\mathbb{P}\big(\rho_\tau \in \mathrm{CS}_{n,1-\alpha}(\lambda) \mid \mathcal{F}_\tau\big) = \mathbb{P}\big(\hat{\rho}_n(\lambda) \in I_{1-\alpha}(\lambda;\rho_n) \mid \mathcal{F}_\tau\big) \xrightarrow{\mathbb{P}} 1-\alpha,$$

which completes the proof.

□

References

Aït-Sahalia, Y., J. Fan, and D. Xiu (2010): "High-Frequency Covariance Estimates With Noisy and Asynchronous Financial Data," *Journal of the American Statistical Association*, 105, 1504–1517.

Aït-Sahalia, Y. and D. Xiu (2016): "Increased correlation among asset classes: Are volatility or jumps to blame, or both?" *Journal of Econometrics*, 194, 205–219.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2007): "Real-Time Price Discovery in Global Stock, Bond and Foreign Exchange Markets," *Journal of International Economics*, 73, 251–277.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and G. Wu (2006): "Realized Beta: Persistence and Predictability," in *Econometric Analysis of Financial and Economic Time Series - Part B*, Elsevier, vol. 20 of *Advances in Econometrics*, 1–40.

Anderson, T. W. and H. Rubin (1949): "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *The Annals of Mathematical Statistics*, 20, 46–63.

Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2008): "Designing Realised Kernels to Measure the Ex-Post Variation of Equity Prices in the Presence of Noise," *Econometrica*, 76, 1481–1536.

——— (2011): "Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading," *Journal of Econometrics*, 162, 149–169.

Barndorff-Nielsen, O. E. and N. Shephard (2004): "Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics," *Econometrica*, 72, 885–925.

Bauer, M. D. and E. T. Swanson (2023): "An alternative explanation for the "fed information effect"," *American Economic Review*, 113, 664–700.

Bibinger, M., N. Hautsch, P. Malec, and M. Reiss (2019): "Estimating the Spot Covariation of Asset Prices–Statistical Theory and Empirical Evidence," *Journal of Business & Economic Statistics*, 37, 419–435.

Bibinger, M. and M. Reiss (2014): "Spectral Estimation of Covolatility from Noisy Observations Using Local Weights," *Scandinavian Journal of Statistics*, 41, 23–50.

Bollerslev, T., R. F. Engle, and J. M. Wooldridge (1988): "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy*, 96, 116–131.

Bollerslev, T., J. Li, and Q. Li (2024a): "Optimal nonparametric range-based volatility estimation," *Journal of Econometrics*, 238, 105548.

Bollerslev, T., J. Li, Y. Li, and Q. Zhang (2026): "Illuminating Important Economic News by Candlesticks: Optimal Testing Meets Technical Analysis," Working paper, Duke University.

Bollerslev, T., J. Li, and Z. Liao (2021): "Fixed-k inference for volatility," *Quantitative Economics*, 12, 1053–1084.

Bollerslev, T., J. Li, A. J. Patton, and R. Quaedvlieg (2020): "Realized Semicovariances," *Econometrica*, 88, 1515–1551.

Bollerslev, T., J. Li, and Y. Ren (2024b): "Optimal Inference for Spot Regressions," *American Economic Review*, 114, 678–708.

Campbell, J. Y. (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

Campbell, J. Y., C. Pflueger, and L. M. Viceira (2020): "Macroeconomic Drivers of Bond and Equity Risks," *Journal of Political Economy*, 128, 3148–3185.

Chen, N.-F., R. Roll, and S. A. Ross (1986): "Economic Forces and the Stock Market," *Journal of Business*, 59, 383–403.

Christensen, K., R. Oomen, and R. Renò (2022): "The drift burst hypothesis," *Journal of Econometrics*, 227, 461–497.

Christensen, K., R. C. A. Oomen, and M. Podolskij (2014): "Fact or friction: Jumps at ultra high frequency," *Journal of Financial Economics*, 114, 576–599.

Cieslak, A. and H. Pang (2021): "Common Shocks in Stocks and Bonds," *Journal of Financial Economics*, 142, 880–904.

Cutler, D. M., J. M. Poterba, and L. Summers (1989): "What Moves Stock Prices?" *Journal of Portfolio Management*, 15, 4–12.

Duffee, G. R. (2023): "Macroeconomic News and Stock-Bond Comovement," *Review of Finance*, 27, 1859–1882.

Eagleson, G. K. (1975): "Martingale Convergence to Mixtures of Infinitely Divisible Laws," *Annals of Probability*, 3, 557–562.

Edrington, L. H. and J. H. Lee (1993): "How Markets Process Information: News Releases and Volatility," *Journal of Finance*, 48, 1161–1191.

Engle, R. F. (2009): *Anticipating Correlations: A New Paradigm for Risk Management*, Princeton, NJ: Princeton University Press.

Epps, T. W. (1979): "Comovements in Stock Prices in the Very Short Run," *Journal of the American Statistical Association*, 74, 291–298.

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

Foster, D. P. and D. B. Nelson (1996): "Continuous Record Asymptotics for Rolling Sample Variance Estimators," *Econometrica*, 64, 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, 67–78.

Gatheral, J., T. Jaisson, and M. Rosenbaum (2018): "Volatility is rough," *Quantitative Finance*, 18, 933–949.

Hansen, P. R. and Y. Luo (2023): "Robust estimation of realized correlation: New insight about intraday fluctuations in market betas," *arXiv preprint arXiv:2310.19992*.

Hayashi, T. and N. Yoshida (2005): "On covariance estimation of non-synchronously observed diffusion processes," *Bernoulli*, 11, 359–379.

Hu, G. X., Z. Jin, and J. Pan (2025): "The Stock-Bond Correlation: A Tale of Two Days in the U.S. Treasury Market," Working paper.

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

Jacod, J. and P. E. Protter (2012): *Discretization of Processes*, New York: Springer-Verlag, 1st ed.

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

Li, J., D. Wang, and Q. Zhang (2024): "Reading the Candlesticks: An OK Estimator for Volatility," *The Review of Economics and Statistics*, 106, 1114–1128.

Li, J. and D. Xiu (2016): "Generalized Method of Integrated Moments for High-Frequency Data," *Econometrica*, 84, 1613–1633.

Li, Y., I. Nolte, S. Nolte, and S. Yu (2025): "Realized candlestick wicks," *Journal of Econometrics*, 250, 106014.

Liu, Q. and Z. Liu (2022): "Statistical Inference of Spot Correlation and Spot Market Beta under Infinite Variation Jumps," *Journal of Financial Econometrics*, 20, 612–654.

Markowitz, H. (1952): "Portfolio Selection," *The Journal of Finance*, 7, 77–91.

Mörters, P. and Y. Peres (2010): *Brownian Motion*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge: Cambridge University Press.

Nakamura, E. and J. Steinsson (2018): "High-frequency identification of monetary non-neutrality: the information effect," *The Quarterly Journal of Economics*, 133, 1283–1330.

Nualart, D. (2006): *The Malliavin Calculus and Related Topics*, Probability and Its Applications, Berlin: Springer, 2nd ed.

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

Popov, V. (2016): "Correlation estimation using components of Japanese candlesticks," *Quantitative Finance*, 16, 1615–1630.

Rigobon, R. and B. Sack (2004): "The Impact of Monetary Policy on Asset Prices," *Journal of Monetary Economics*, 51, 1553–1575.

Rogers, L. C. G. and S. E. Satchell (1991): "Estimating Variance From High, Low and Closing Prices," *Annals of Applied Probability*, 1, 504–512.

Rogers, L. C. G. and L. Shepp (2006): "The correlation of the maxima of correlated Brownian motions," *Journal of Applied Probability*, 43, 880–883.

Rogers, L. C. G. and F. Zhou (2008): "Estimating Correlation From High, Low, Opening and Closing Prices," *Annals of Applied Probability*, 18, 813–823.

Ross, S. A. (1976): "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, 13, 341–360.

Rudin, W. (1976): *Principles of Mathematical Analysis*, New York: McGraw-Hill, 3rd ed.

Sharpe, W. F. (1964): "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance*, 19, 425–442.

Shephard, N. and D. Xiu (2017): "Econometric analysis of multivariate realised QML: Estimation of the covariation of equity prices under asynchronous trading," *Journal of Econometrics*, 201, 19–42.

Simsek, Y. (2026): "Spot Regressions with Candlesticks," *Duke University Working Paper*.

Varneskov, R. T. (2016): "Flat-Top Realized Kernel Estimation of Quadratic Covariation With Nonsynchronous and Noisy Asset Prices," *Journal of Business and Economic Statistics*, 34, 1–22.

Vetter, M. (2012): "Estimation of Correlation for Continuous Semimartingales," *Scandinavian Journal of Statistics*, 39, 757–771.