Estimation and Inference for Spot Variance of Microstructure Noise*

Z. Merrick Li† Haotian Zhang‡

November 13, 2026

Abstract

We develop estimation and inference for the spot variance of market microstructure noise, the discrepancy between observed transaction prices and underlying efficient prices arising from trading frictions. The noise is a stationary, serially dependent sequence scaled by the square root of a positive semimartingale that may jump at prescheduled times. Using transaction prices alone, we construct one-sided estimators of the left and right limits of the noise variance from a multi-horizon differencing scheme; averaging over a growing set of horizons renders them asymptotically equivalent to the oracle that observes the noise. We establish a joint stable central limit theorem and provide consistent asymptotic-variance estimators and a test for scheduled jumps. In SPY transactions over 2016–2025, the noise variance jumps at macroeconomic releases and closing-auction deadlines with event-specific recovery dynamics.

1 Introduction

High-frequency transaction prices measure an underlying efficient price with error. Bid–ask bounce, inventory control, and the discreteness of trading all induce a gap between the recorded price and the frictionless value that would prevail in an idealized continuous market (Hasbrouck, 2007). The standard additive decomposition writes the observed price as

$$Y = X + \varepsilon, \tag{1}$$

*We gratefully acknowledge the insightful discussions with Yacine Aït-Sahalia, Torben Andersen, Federico Bandi, Giuseppe Cavaliere, Carsten Chong, Wenhao Cui, Jia Li, Yingying Li, Oliver Linton, Mark Podolskij, Roberto Renò, Shuping Shi, Viktor Todorov, Jun Yu, Xinghua Zheng, and Haibin Zhu. We appreciate the valuable feedback from participants at the 19th International Symposium on Econometric Theory and Applications (SETA 2025), the 17th Annual SoFiE Conference in Paris, and the 2025 World Congress of the Econometric Society (ESWC 2025) in Seoul, Korea.

†Department of Economics, The Chinese University of Hong Kong, Shatin, Hong Kong SAR, China. Email: merrickli@cuhk.edu.hk.

‡Department of Economics, The Chinese University of Hong Kong, Shatin, Hong Kong SAR, China. Email: htzhang@link.cuhk.edu.hk.

where X is an Itô semimartingale, identified with an arbitrage-free efficient price (Delbaen and Schachermayer, 1994), and ε is the measurement error. Because ε arises from market microstructure frictions, the literature calls it *microstructure noise*. The variance of ε is of interest in itself, for economic and statistical reasons. Economically, it is a proxy for transaction costs and market quality (Lehmann, 1990; Hasbrouck, 1993). Statistically, it is the main nuisance in inference on the parameters and processes of X, and denoising methods must estimate it to remove its effect (Aït-Sahalia and Jacod, 2014).

Early work mostly assumes that the noise is i.i.d. Under this assumption, realized volatility computed at the highest frequencies identifies the noise variance (Zhang, Mykland, and Aït-Sahalia, 2005; Bandi and Russell, 2006, 2008), and quasi-maximum likelihood estimates the noise variance jointly with volatility (Xiu, 2010; Aït-Sahalia, Fan, and Xiu, 2010). Later work allows the noise to have a stochastic scale and serial dependence. It estimates integrated moments of the noise by local averaging (Jacod, Li, and Zheng, 2017) or by differencing (Li and Linton, 2022), or it obtains them as a by-product of quasi-maximum likelihood (Da and Xiu, 2021, 2026). These estimators remain valid when the noise scale varies within the day. However, they estimate the average or integrated moments of the scale over a fixed interval, not the scale at a given point in time.

The economics of financial markets, however, calls for local measurement. Some forces move microstructure noise gradually and randomly over the day: large traders split their orders (Chordia and Subrahmanyam, 2004), and market makers manage their inventories (Hendershott and Menkveld, 2014). Other forces move frictions abruptly, often at times fixed in advance. Macroeconomic announcements arrive at published times and change trading activity and liquidity (Fleming and Remolona, 1999; Bollerslev, Li, and Xue, 2018; Crego, 2020). We also find that exchange trading rules impose a second calendar. Before the closing auction, exchanges stop accepting or cancelling on-close orders and begin publishing order imbalances, all at fixed times, and the noise variance jumps at these deadlines. Moreover, these scheduled events are frequent.[1] Treating the noise variance as (almost surely) continuous at these times therefore distorts the intraday pattern of microstructure noise.

We model the noise variance as a positive semimartingale, as in Jacod, Li, and Zheng (2017), Da and Xiu (2021), and Li and Linton (2022), and extend their models by allowing the noise variance to jump at prescheduled times. We then propose a new spot inference method for the noise variance. The method is asymptotically equivalent to an infeasible oracle that observes the noise directly and therefore can compute its sample variance. Specifically, we construct one-sided estimators of the left and right limits of the noise variance, using observations in shrinking windows on the respective sides of a given time. These estimators recover the pre- and post-event noise variances separately, allowing the two limits to differ at scheduled times. We establish joint stable convergence in law of the appropriately normalized estimation errors at any fixed finite collection

[1] The closing-auction deadlines occur in every trading session, and about 64% of the sessions in our sample contain at least one scheduled macroeconomic release at 10:00, 14:00, or 15:00.

of times, with a conditionally Gaussian limit. The limiting variance is random and combines the sampling error of the serially dependent noise with the local variation of the noise variance, in proportions set by the bandwidth. We construct two-scale estimators of this variance in the spirit of Mykland and Zhang (2017), using only observations on the respective side of the given time, and show that they are consistent in every admissible bandwidth regime. The central limit theorem is therefore feasible: it yields confidence intervals for the left and right limits of the noise variance without requiring knowledge of the regime. The difference between the right- and left-limit estimators yields a jump test with correct asymptotic size and consistency against fixed alternatives. Together with the estimates of jump size, this allows us to detect and quantify abrupt changes in trading frictions at prespecified event times.

Our application to SPY transactions over 2016–2025 shows that measuring noise locally matters economically. At scheduled macroeconomic releases, the average noise variance after the release is roughly nine to sixteen times the average before it. Most of the increase after the 10:00 releases disappears within a minute. FOMC announcements, by contrast, show a buildup before the release and a longer-lasting rise afterward. Apart from news, late-session bursts recur at the closing-auction deadlines and imbalance disclosures, and their timing shifts when exchange rules change. Intraday trading frictions therefore reflect both the arrival of information and the organization of trading. Spot measurement shows how large these frictions are and how long they last, which integrated noise estimates cannot resolve.

Our paper contributes to several strands of the literature. The first is the literature on the general modeling and estimation of microstructure noise (Jacod, Li, and Zheng, 2017; Da and Xiu, 2021; Li and Linton, 2022). This literature targets integrated moments of the noise, and its models keep the noise variance almost surely continuous at any fixed time. Because an integrated moment pools the noise variance over the whole interval, it cannot tell the variance before a scheduled event apart from the variance after it. Our generalized model, combined with spot inference, identifies both. In terms of estimation method, the closest work is Li and Linton (2022). While Li and Linton (2022) use a single-horizon differencing scheme, we average the differencing statistic over a growing set of multiple horizons.[2] The multi-horizon approach has two advantages. First, it is more stable than an estimator based on a single horizon. Second, it improves efficiency; indeed, it is essential for asymptotic equivalence to the oracle estimator. Both advantages matter most for spot inference, where samples are typically much smaller than in integrated inference. Two recent papers study structured noise. Chong, Delerue, and Li (2025) model the noise as a rough fractional process with continuous stochastic volatility and estimate its roughness and integrated volatility. Cui (2026) splits the noise into a component explained by observable trading information and a residual, and measures how much noise variance the former explains. Both keep the noise variance continuous in time. Our target is instead the joint estimation and inference on the spot variance of the total noise, estimated nonparametrically from prices alone, which can jump at

[2]Our approach is distinct from that of Li and Yang (2026), who hold fixed both the horizon length and the number of horizons. Their main theory also serves a different purpose: they apply it to the semimartingale efficient price process to test whether noise is present.

prescheduled times.

Our paper also contributes to the literature on nonparametric spot inference from high-frequency data. Much of this literature concerns the volatility, covariance, and regression coefficients of the efficient price.[3] Building our new method on local noise estimation, we develop feasible joint inference for the left and right limits of microstructure-noise variance. To our knowledge, this is the first general inference theory that accommodates serially dependent noise, irregular sampling, and jumps at prescheduled times.[4] We also contribute to the literature on the feasibility of spot inference over a spectrum of bandwidth regimes. Specifically, at the balanced bandwidth, sampling uncertainty and local parameter variation enter the limiting distribution at the same order, and feasible inference is difficult.[5] We develop a local, one-sided two-scale construction that consistently estimates the two sources of uncertainty. The resulting variance estimators are consistent throughout the admissible bandwidth regimes, including the balanced regime yielding the $n^{1/4}$ rate. Combined with oracle equivalence and the joint stable central limit theorem, this delivers feasible confidence intervals for the noise variance immediately before and after scheduled events, together with tests for its discontinuities, without requiring undersmoothing to eliminate local variance dynamics.

Our findings also relate to the growing literature on intraday variation analysis. Andersen, Thyrsgaard, and Todorov (2021) document systematic variation in the cross-sectional dispersion of market betas and its responses to news, while Andersen, Riva, Thyrsgaard, and Todorov (2023) extend this analysis to distributions of exposures to multiple risk factors. Andersen, Su, Todorov, and Zhang (2024) develop functional inference for average intraday volatility curves and document shifts in these curves over time. We complement these studies with a formal framework for inference on the intraday noise variance. Using this framework, we document abrupt changes in the noise variance at scheduled announcements and exchange deadlines. These results highlight three considerations when estimating intraday profiles of other high-frequency quantities, such as volatility or factor exposures. First, profiles should be conditioned on the announcement calendar. Scheduled releases are frequent, yet their effects are absent on non-announcement days, so pooling all days mixes event responses with regular intraday variation. Second, profiles should respect trading-rule regimes, because changes in exchange rules can shift the times at which the underlying process jumps. Third, measurement around scheduled events requires a fine time grid to separate anticipation, immediate adjustment, and recovery.

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 develops the oracle and feasible estimators and the joint limit theory. Section 4 reports simulations;

[3] A partial list of studies is Foster and Nelson (1996); Kristensen (2010); Zu and Boswijk (2014); Bollerslev, Li, and Liao (2021); Bollerslev, Li, and Ren (2024); Li and Linton (2026).

[4] Several studies estimate the noise variance locally, including Kalnina and Linton (2008); Jacod, Li, and Zheng (2019); Zhang, Li, and Bollerslev (2022); Li and Linton (2026). Their methods differ, and their settings are less general than ours. In these studies, the local noise-variance estimates serve only as intermediate inputs to volatility estimators. None of them develops an inference theory for the estimates.

[5] In their discussion of spot-volatility inference, Aït-Sahalia and Jacod (2014, Section 8.1.3, p. 271) describe the corresponding spot volatility-of-volatility quantity as "nearly impossible to estimate."

Section 5 contains the empirical application. Asymptotic conditions and variance estimators are collected in the Appendix. All proofs are given in Section SA of the supplementary materials (Li and Zhang, 2026), and Section SB documents the data and the empirical implementation and reports additional empirical results; Section SC reports additional simulation results under different settings.

2 Model Setting

Let Z be a generic Itô semimartingale defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$. It can be represented as (Jacod and Protter, 2012):

$$Z_t = Z_0 + \int_0^t b_s^Z \mathrm{d}s + \int_0^t \sigma_s^Z \mathrm{d}W_s^Z + \left(\delta^Z \mathbf{1}_{\{|\delta^Z|\leq 1\}}\right) \star \left(\mathfrak{p} - \mathfrak{q}\right)_t + \left(\delta^Z \mathbf{1}_{\{|\delta^Z|>1\}}\right) \star \mathfrak{p}_t, \tag{2}$$

where W^Z is a Brownian motion, $\mathfrak{p}$ is a Poisson random measure on $\mathbb{R}_+ \times \mathbb{R}$ with deterministic compensator $\mathfrak{q}(\mathrm{d}s, \mathrm{d}z) = \mathrm{d}s \otimes \lambda(\mathrm{d}z)$, where λ is a given σ-finite measure. The processes b^Z and σ^Z are optional and locally bounded, while the function δ^Z on $\Omega \times \mathbb{R}_+ \times \mathbb{R}$ is predictable. Moreover, we assume the following standard regularity conditions hold for all Itô semimartingales considered in this paper:

Assumption (H). *Let Z be a semimartingale represented by (2). The volatility process σ^Z is adapted and càdlàg. There exists a sequence of stopping times $\{\tau_m\}_m$ increasing to infinity almost surely, a sequence of real numbers $\{w_m\}_m$, and for each m a deterministic nonnegative function J_m^Z on $\mathbb{R}$ so that the following hold for all (ω, t, z) with $t \leq \tau_m(\omega)$:*

$$\left|b_t^Z(\omega)\right| \leq w_m, \quad \left|\sigma_t^Z(\omega)\right| \leq w_m, \quad \left|\delta^Z(\omega, t, z)\right| \wedge 1 \leq J_m^Z(z), \quad \int_{\mathbb{R}} \left(J_m^Z(z)\right)^2 \lambda(\mathrm{d}z) < \infty.$$

Specifically, we assume that the efficient price process X is an Itô semimartingale satisfying Assumption (H). The semimartingale property is necessary for a market to be free of arbitrage (Delbaen and Schachermayer, 1994).

We work under infill asymptotics over a fixed observation interval, normalized to $[0, 1]$, with increasingly frequent observations as $n \to \infty$. Our analysis of microstructure noise is based on high-frequency transaction prices, which are naturally recorded at random time points. To accommodate this feature, we work under a general framework of random observation times introduced in Jacod, Li, and Zheng (2017, 2019). Specifically, for each $n \in \mathbb{N}_+$, we consider a sequence of strictly increasing observation times $\{t_i^n\}_{i\geq 0}$ with $t_0^n = 0$ and $t_i^n \to \infty$ as $i \to \infty$. We denote the random spacing between two consecutive observation times by $\Delta(n, i) := t_i^n - t_{i-1}^n$ for $i = 1, 2, \ldots$. Throughout, for any process V we write $V_i^n := V_{t_i^n}$. For example, X_i^n and Y_i^n represent the latent efficient price and the observed price, respectively, at time t_i^n. Moreover, for any process V with one-sided limits, we write $V_{t-} := \lim_{s\uparrow t} V_s$ for $t > 0$ and $V_{t+} := \lim_{s\downarrow t} V_s$ for $t \geq 0$. If V is càdlàg, then $V_{t+} = V_t$. Let $\Delta_n := 1/n \to 0$ denote the nominal spacing between observations. The local

sampling frequency is governed by a positive intensity process α: by (3) below, the expected spacing after t^n_{i-1} is approximately Δ_n/α^n_{i-1}, so about α_t/Δ_n observations per unit of time arrive near t. We assume that the observation times satisfy the following condition:

Assumption (O). *α is an Itô semimartingale defined on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ satisfying Assumption (H). We further assume there is a localizing sequence $\{\tau_m\}_m$ of stopping times and positive constants $\eta_{m,1}$, $\kappa_{m,p}$ and η, such that:*

(i) For $t < \tau_m$, we have $\frac{1}{\eta_{m,1}} \leq \alpha_{t-} \leq \eta_{m,1}$.

(ii) Let $(\mathcal{F}^n_t)_{t\geq 0}$ be the smallest filtration containing $(\mathcal{F}_t)_{t\geq 0}$ for which all t^n_i, $i = 0, 1, 2, \ldots$, are stopping times, and write $\mathcal{I}^n_i := \mathcal{F}^n_{t^n_i}$. For each $i = 1, 2, \ldots$, $\Delta(n, i)$ is, conditionally on $\mathcal{I}^n_{i-1}$, independent of $\mathcal{F}_\infty := \bigvee_{t\geq 0} \mathcal{F}_t$.

(iii) With the restriction $\{t^n_{i-1} < \tau_m\}$, and for all $p > 0$,

$$\left|\mathbb{E}\left(\Delta(n,i)|\,\mathcal{I}^n_{i-1}\right) - \frac{\Delta_n}{\alpha^n_{i-1}}\right| \leq \kappa_{m,1}\Delta_n^{1+\eta}, \quad \mathbb{E}\left((\Delta(n,i))^p|\,\mathcal{I}^n_{i-1}\right) \leq \kappa_{m,p}\Delta_n^p. \tag{3}$$

Next, we present the model for the market microstructure noise. We first specify the noise variance process Γ as a positive semimartingale consisting of an Itô semimartingale component and possible additional jumps at finitely many deterministic times. Scheduled announcements, such as FOMC decisions and corporate earnings releases, can alter liquidity provision, as reflected in changes in bid–ask spreads and market depth around their release times (Lee, Mucklow, and Ready, 1993; Fleming and Remolona, 1999; Crego, 2020). The empirical literature often takes the dispersion of microstructure noise as a measure of illiquidity or market quality (Hasbrouck, 1993; Aït-Sahalia and Yu, 2009). This motivates allowing the spot noise variance to jump at such times. Moreover, the process Γ may depend on the efficient price process, for example, as a function of its volatility. Our assumption extends the noise specifications in Jacod, Li, and Zheng (2017); Da and Xiu (2021); Li and Linton (2022), under which, for every fixed deterministic time, the noise scale and hence its variance are continuous at that time almost surely.

Assumption (V). *Let $\widetilde{\mathcal{T}} := \{\widetilde{t}_\ell \in (0,1) : \ell = 1, 2, \ldots, |\widetilde{\mathcal{T}}|\}$ be a finite deterministic set. For each ℓ, let φ_ℓ be a finite real-valued, $\mathcal{F}_{\widetilde{t}_\ell}$-measurable random variable. The spot variance process Γ satisfies*

$$\Gamma_t = \Gamma'_t + \sum\nolimits_{\ell=1}^{|\widetilde{\mathcal{T}}|} \varphi_\ell \mathbf{1}_{\{\widetilde{t}_\ell \leq t\}}, \qquad \inf_{0\leq s\leq 1} \Gamma_s > 0 \quad \textit{almost surely}, \tag{4}$$

where Γ' satisfies Assumption (H).

Following Jacod, Li, and Zheng (2017) and Li and Linton (2022), we model the latent microstructure noise by a multiplicative form as:

$$\varepsilon^n_i := \gamma^n_i \cdot \chi_i, \tag{5}$$

where $\gamma_t := \sqrt{\Gamma_t}$ and the spot variance process Γ satisfies Assumption (V), while χ is a zero-mean unit-variance but potentially autocorrelated sequence reflecting the dynamics of order flow. We assume the χ process is independent of the efficient price process. The serial dependence of χ is restricted through a ρ-mixing condition with a polynomial decay rate.

Assumption (N). *Let $\{\chi_i\}_{i\in\mathbb{Z}}$ be a stationary ρ-mixing random sequence with the mixing coefficient ρ_k decaying at a polynomial rate, i.e., $\rho_k \le Kk^{-v}, \forall k \in \mathbb{N}_+$, for some constants $K > 0$ and $v > 0$. We further assume that $\{\chi_i\}_{i\in\mathbb{Z}}$ is centered at 0 with variance 1 and finite moments of all orders, and is independent of $\mathcal{F}^{(0)} := \mathcal{F}_\infty \vee \sigma(t_i^n : i \ge 0,\ n \ge 1)$.*

3 Spot Estimation and Inference

Let $\mathcal{T} := \widetilde{\mathcal{T}} \cup \widetilde{\mathcal{T}}'$ be a finite collection of fixed time points taking values in $(0, 1)$, where $\widetilde{\mathcal{T}}$ is a set of prescheduled time points described in Assumption (V), and $\widetilde{\mathcal{T}}'$ is a finite set of fixed times at which no prescheduled jump occurs. For any $t \in (0, 1]$, write $n_t := \max\{i \ge 0 : t_i^n \le t\}$. The targets are the right and left limits Γ_{t+} $(= \Gamma_t)$ and Γ_{t-} of the spot noise variance, for $t \in \mathcal{T}$; we write $\Gamma_{t\pm}$ for either of them.

3.1 An infeasible oracle estimation strategy

Since ε is latent, practical estimation from observed transaction prices presents a challenge. To develop our estimation strategy and establish theoretical benchmarks, we first consider an infeasible "oracle" scenario in which ε is treated as observable. In this hypothetical setting, the natural estimator for $\Gamma_{t\pm}$ at time $t \in \mathcal{T}$ would be the local one-sided sample variance. Given a local bandwidth l_n, let $\widehat{\mathbb{I}}_{t\pm}^n := \{n_t \pm j : j = 1, \dots, l_n\}$ and define

$$\widetilde{\Gamma}_{t\pm}^n := \frac{1}{l_n} \sum\nolimits_{i\in\widehat{\mathbb{I}}_{t\pm}^n} (\varepsilon_i^n)^2. \tag{6}$$

Under very mild regularity conditions, we can establish the consistency (Theorem 3.1) and asymptotic normality (Theorem 3.2) of this oracle estimator.

Theorem 3.1. *Assume that Condition A.1 holds. For any $t \in \mathcal{T}$, we have $\widetilde{\Gamma}_{t\pm}^n \xrightarrow{\mathbb{P}} \Gamma_{t\pm}$.*

Theorem 3.2. *Let $\widehat{Z}_{t\pm}^n := \widetilde{\Gamma}_{t\pm}^n - \Gamma_{t\pm}$, $c_n := l_n^2\Delta_n \to c \in [0, \infty]$, and $u_n = \sqrt{\frac{l_n}{1+c_n}}$. Under Condition A.2, we have the joint $\mathcal{F}_\infty$-stable convergence in law*

$$u_n \left(\widehat{Z}_{t+}^n, \widehat{Z}_{t-}^n\right)_{t\in\mathcal{T}} \xrightarrow{\mathcal{L}_s} (Z_{t+}, Z_{t-})_{t\in\mathcal{T}}\,.$$

Conditionally on $\mathcal{F}$, all $2|\mathcal{T}|$ coordinates of the limit vector are independent centered Gaussian random variables, with conditional variance $\mathfrak{S}_{t\pm} := \frac{\Gamma_{t\pm}^2\mathfrak{S}}{1+c} + \frac{c\widetilde{\sigma}_{t\pm}^2}{3\alpha_{t\pm}(1+c)}$, where $\mathfrak{S} := \sum_{k=-\infty}^{\infty} \left(\mathbb{E}(\chi_0^2\chi_k^2) - 1\right)$ and $\widetilde{\sigma}$ is the volatility process of Γ'.

Remark 3.1. *The convergence rate depends on the asymptotic order of the local bandwidth l_n relative to Δ_n, or equivalently, n. When $l_n \asymp n^a$ with $a < 1/2$, the rate is $\sqrt{l_n}$, and the statistical errors are $O_p(1/\sqrt{l_n})$. In this regime, Γ is locally constant to first order on the indicated side of t, and the asymptotic variance is dominated by the long-run variance $\mathfrak{S}$ of χ_i^2, scaled by $\Gamma_{t\pm}^2$. For $a > 1/2$, the rate becomes $\sqrt{1/(l_n\Delta_n)}$. Within this broader local window, the variation in Γ_s ($s \neq t$) dominates, and the asymptotic variance reflects the volatility of Γ at t. The optimal convergence rate of $n^{1/4}$ is attained when $a = \frac{1}{2}$, where the sampling error of the noise and the local variation of Γ contribute to the asymptotic variance at the same order.*

Remark 3.2. *Assumption (N) requires χ to be globally stationary. A natural extension is to allow its serial dependence to vary over time. Under the unit-variance normalization, the spot noise variance $\Gamma_{t\pm}$ is determined by the scale γ, while the dependence structure of χ affects estimation uncertainty (via $\mathfrak{S}$). Since our estimators use shrinking one-sided windows, we expect analogous results under suitable local-stationarity conditions that allow for time-varying serial dependence.*

3.2 A feasible estimation strategy

We now introduce a feasible estimator constructed from transaction prices alone that attains asymptotic equivalence to the oracle estimator. By recovering the spot variance of the unobservable transitory component ε from observed transaction prices Y, our approach effectively treats the latent process as if it were observable.

Inspired by Li and Linton (2022), we employ a differencing method originally designed for *integrated* moment estimation of the noise. Extending this approach to one-sided spot estimation requires additional analysis because estimation is performed over shrinking time windows. The local bandwidth governs the relative contributions of sampling variability and stochastic variation in the noise variance, while the differencing horizons must be chosen to make the feasible–oracle approximation error negligible at the spot convergence rate. The resulting joint limits at distinct fixed times are conditionally independent Gaussian variables, giving the finite-dimensional white-noise structure characteristic of local spot estimation. See Chapter 8 of Aït-Sahalia and Jacod (2014) and Chapter 13 of Jacod and Protter (2012) for detailed discussions.

Beyond this adaptation, we develop a novel multi-horizon differencing scheme, improving substantially upon the single-horizon approach of Li and Linton (2022). This innovation enhances efficiency and mitigates instability from selecting a single differencing horizon, making it particularly advantageous for spot estimation with small local sample sizes.

Now we formally present our multi-horizon differencing estimation strategy. For any integer $k \geq 1$, we let $f(k)_i^n := (Y_i^n - Y_{i-k}^n)(Y_i^n - Y_{i+k}^n)$. Our multi-horizon differencing method is governed by two sequences of positive integers $\{k_n\}_n$ and $\{p_n\}_n$. Let $k_{n,p} := k_n + p$, $\overline{k}_n = k_n + p_n$, and define the post-$t$ and pre-t index sets jointly as $\mathbb{I}_{t\pm}^n := \{n_t \pm j : j = \overline{k}_n + 1, \ldots, l_n - \overline{k}_n\}$.

We define:

$$\widehat{\Gamma}^n_{t\pm} := \frac{1}{l_n p_n} \sum\nolimits_{p=0}^{p_n-1} \sum\nolimits_{i \in \mathbb{I}^n_{t\pm}} f(k_{n,p})^n_i. \tag{7}$$

The intuition of our multi-horizon differencing method is as follows: increments of the efficient price are martingale differences up to the drift, and $f(k)^n_i$ is constructed to leverage this property. Specifically, $\mathbb{E}(f(k_{n,p})^n_i) \approx \mathbb{E}((\varepsilon^n_i - \varepsilon^n_{i-k_{n,p}})(\varepsilon^n_i - \varepsilon^n_{i+k_{n,p}}))$.[6] When multiple horizons are considered, we have $\frac{1}{p_n}\sum_{p=0}^{p_n-1} f(k_{n,p})^n_i \approx (\varepsilon^n_i)^2 - \varepsilon^n_i \overline{\varepsilon} - \varepsilon^n_i \overline{\varepsilon}' + \overline{\varepsilon}''$, where

$$\overline{\varepsilon} := \frac{1}{p_n} \sum\nolimits_{p=0}^{p_n-1} \varepsilon^n_{i+k_{n,p}}, \quad \overline{\varepsilon}' := \frac{1}{p_n} \sum\nolimits_{p=0}^{p_n-1} \varepsilon^n_{i-k_{n,p}}, \quad \overline{\varepsilon}'' := \frac{1}{p_n} \sum\nolimits_{p=0}^{p_n-1} \varepsilon^n_{i+k_{n,p}} \varepsilon^n_{i-k_{n,p}}.$$

These averages of ε are asymptotically negligible when p_n and $k_{n,p}$ are large. Consequently, $\frac{1}{p_n}\sum_p f(k_{n,p})^n_i \approx (\varepsilon^n_i)^2$ on either side, which implies $\widehat{\Gamma}^n_{t\pm} \approx \widetilde{\Gamma}^n_{t\pm}$.

The following theorem formalizes the preceding intuition by establishing the asymptotic equivalence between the proposed feasible estimators and the infeasible oracle estimators introduced in Section 3.1. Moreover, it provides the asymptotic distribution of our feasible estimators, enabling valid statistical inference.

Theorem 3.3. *Suppose Condition A.3 holds. For any $t \in \mathcal{T}$, we have*

$$\widehat{\Gamma}^n_{t\pm} \xrightarrow{\mathbb{P}} \Gamma_{t\pm}. \tag{8}$$

Furthermore, under Condition A.4, $u_n(\widehat{\Gamma}^n_{t\pm} - \widetilde{\Gamma}^n_{t\pm}) \xrightarrow{\mathbb{P}} 0$ for every $t \in \mathcal{T}$. The joint stable convergence in Theorem 3.2 holds with $\widetilde{\Gamma}^n_{t\pm}$ replaced by $\widehat{\Gamma}^n_{t\pm}$. Consequently, if in addition Condition A.5 holds, for any $t \in \mathcal{T}$ and side $\pm$ for which $\mathfrak{S}_{t\pm} > 0$ almost surely,

$$\frac{u_n\left(\widehat{\Gamma}^n_{t\pm} - \Gamma_{t\pm}\right)}{\sqrt{\widehat{\mathfrak{S}}^n_{t\pm}}} \xrightarrow{\mathcal{L}_s} \mathcal{N}(0,1), \tag{9}$$

where $\widehat{\mathfrak{S}}^n_{t\pm}$ is the consistent estimator of the asymptotic variance $\mathfrak{S}_{t\pm}$ defined in Appendix B.

Remark 3.3. *A local adaptation of the single-horizon ReMeDI estimator of Li and Linton (2022) is obtained by averaging $\left(Y^n_i - Y^n_{i-2k_n}\right)\left(Y^n_i - Y^n_{i+k_n}\right)$ over each one-sided window. Under the model assumptions and the restrictions on υ, κ and λ in Condition A.4, this estimator has a centered mixed-Gaussian limit under the same normalization u_n as in Theorem 3.2, with conditional asymptotic variance $\frac{\Gamma^2_{t\pm}\mathfrak{S}'}{1+c} + \frac{c\widetilde{\sigma}^2_{t\pm}}{3\alpha_{t\pm}(1+c)}$, where $\mathfrak{S}' := \mathfrak{S} + 3\sum_{k=-\infty}^{\infty} r_k^2$ and $r_k := \mathbb{E}(\chi_0\chi_k)$. For $c < \infty$, this variance exceeds that of both the oracle estimator in Theorem 3.2 and the multi-horizon estimator in Theorem 3.3 by $3\Gamma^2_{t\pm}\sum_{k=-\infty}^{\infty} r_k^2/(1+c) > 0$. Thus, the ReMeDI approach is less efficient. Besides this asymptotic efficiency gain, our multi-horizon differencing method is more robust to the choice of differencing horizons in finite samples, as the simulation*

[6] Over a differencing horizon $k_{n,p}$, the drift increment is $O_p(k_{n,p}\Delta_n)$; its contribution to the estimator is asymptotically negligible.

study in Section 4 shows.

3.3 Testing for scheduled jumps in spot noise variance

Scheduled announcements may change the dispersion of the transitory price component abruptly. Comparing the noise variance immediately before and after an announcement isolates this discontinuity from continuous intraday variation. The calendar time is specified in advance; the occurrence, sign, and size of the jump remain random.

For a fixed $t \in \mathcal{T}$, write $\Delta\Gamma_t := \Gamma_t - \Gamma_{t-}$ and consider

$$\mathbb{H}_{0,t} : \Delta\Gamma_t = 0, \qquad \mathbb{H}_{1,t} : \Delta\Gamma_t \neq 0. \tag{10}$$

At a scheduled time $t = \widetilde{t}_\ell$, $\Delta\Gamma_t = \varphi_\ell$, so the null hypothesis is directly $\varphi_\ell = 0$. Let $V_t^n := \widehat{\mathfrak{S}}_{t+}^n + \widehat{\mathfrak{S}}_{t-}^n$ and $V_t := \mathfrak{S}_{t+} + \mathfrak{S}_{t-}$. Define[7]

$$T_t^n := \frac{u_n(\widehat{\Gamma}_{t+}^n - \widehat{\Gamma}_{t-}^n)}{\sqrt{V_t^n}}, \qquad V_t^n > 0. \tag{11}$$

For a significance level $a \in (0,1)$, the test rejects $\mathbb{H}_{0,t}$ when $|T_t^n| > z_{1-a/2}$, where z_p is the p-quantile of $\mathcal{N}(0,1)$.

Corollary 3.1 (Spot jump test)**.** *Under Conditions A.4 and A.5, fix $t \in \mathcal{T}$ and $a \in (0,1)$. Assume that $V_t > 0$ almost surely. Then*

$$\mathbb{P}(\Delta\Gamma_t = 0,\ |T_t^n| > z_{1-a/2}) \longrightarrow a\,\mathbb{P}(\Delta\Gamma_t = 0), \tag{12}$$

$$\mathbb{P}(\Delta\Gamma_t \neq 0,\ |T_t^n| \leq z_{1-a/2}) \longrightarrow 0. \tag{13}$$

4 Simulation Study

In this section, we conduct Monte Carlo simulations to evaluate the performance of the proposed spot noise variance estimator. We show that the multi-horizon differencing method is more efficient than the single-horizon version of (7). We also show that our method is robust to the choice of tuning parameters, and that price jumps and volatility jumps have little effect on the performance of our test for jumps in the spot noise variance.

4.1 Simulation setup

We suppose the efficient price X follows a jump-diffusion process with stochastic volatility:

$$\mathrm{d}X_t = \kappa_1(\mu_1 - X_t)\mathrm{d}t + \sigma_t \mathrm{d}W_t + \xi_t \mathrm{d}N_t, \quad \mathrm{d}\sigma_t^2 = \kappa_2(\mu_2 - \sigma_t^2)\mathrm{d}t + \eta\sigma_t \mathrm{d}B_t,$$

[7] Set $T_t^n := 0$ when $V_t^n \leq 0$. This convention is asymptotically non-binding: by Theorem B.1, $V_t^n \xrightarrow{\mathbb{P}} V_t$, so $\mathbb{P}(V_t^n \leq 0) \to 0$ whenever $V_t > 0$ almost surely, as assumed in Corollary 3.1.

where W and B are standard Brownian motions with $\mathrm{Corr}(W, B) = \rho_{\mathrm{WB}}$, and N is a Poisson process with intensity λ_J. The jump sizes ξ follow an asymmetric double-exponential distribution and are negative with probability p_ξ, that is, $\xi \sim -\mathrm{Exp}(\mu_-^\xi)\mathbf{1}_{\{U \le p_\xi\}} + \mathrm{Exp}(\mu_+^\xi)\mathbf{1}_{\{U > p_\xi\}}$, where U is uniformly distributed on $[0, 1]$. The parameters for the above dynamics are given by:

$$\kappa_1 = 0.5,\ \mu_1 = 3.6,\ \kappa_2 = 5/252,\ \mu_2 = 0.04/252,\ \eta = 0.05/252;\ \rho_{\mathrm{WB}} = -0.45,$$
$$\lambda_J = 1,\ \mu_+^\xi = 0.0075,\ \mu_-^\xi = 1.5 \cdot \mu_+^\xi,\ p_\xi = 0.65.$$

We consider a random observation scheme in which the observation times t_i^n follow an inhomogeneous Poisson process with rate $n\alpha_t$, where $\alpha_t = 1 + 0.5\cos(2\pi t)$ and $n = 234,000$. This setting implies 10 observations per second on average over a regular trading day and closely aligns with the empirical data used in our applications. We model the transitory component with a discrete part and a signed multiplier, $\varepsilon_i^n = \gamma_i^n \chi_i$. The multiplier takes the form $\gamma_t = K_\gamma \widetilde{\gamma}_t$ with $\mathrm{d}\widetilde{\gamma}_t = -\kappa_3(\widetilde{\gamma}_t - \mu_t^{(\gamma)})\mathrm{d}t + \sigma_\gamma \mathrm{d}W_t$. The component $\widetilde{\gamma}_t$ captures a stochastic intraday pattern, where $\kappa_3 = 10$, $\mu_t^{(\gamma)} := 1 + 0.1\cos(2\pi t)$ and $\sigma_\gamma = 0.25$. The parameter K_γ sets the amplitude and hence the noise-to-signal ratio, which critically affects estimator performance. We set $K_\gamma = 5 \times 10^{-4}$. The discrete part $\{\chi_i\}_i$ is set as a stationary AR(1) process: $\chi_i = \rho_\chi \chi_{i-1} + \sqrt{1 - \rho_\chi^2} e_i$ with $\{e_i\}_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$.

4.2 Efficiency gains

Figure 1 reports the relative root mean squared error (RMSE) of $\widehat{\Gamma}_t^n$ at 9:45 for $k_n \in \{2, 4, 8\}$, $p_n \in \{1, \ldots, 11\}$, and $\rho_\chi \in \{\pm 0.1, \pm 0.4\}$. The figure shows that the multi-horizon estimator outperforms the single-horizon version of (7). When $p_n = 1$, the estimator uses only the horizon k_n. As p_n increases, the estimator combines more horizons and the RMSE falls in every scenario. This reflects gains in both efficiency and robustness. The serial dependence of the noise also matters. When the dependence is strong, estimators based on small horizons have relatively large RMSE, and most of this error comes from the bias that small horizons leave. When the dependence is weak, the choice of k_n matters less, but using multiple horizons still lowers the RMSE. Section SC of the supplementary materials (Li and Zhang, 2026) repeats this comparison with smaller noise amplitudes K_γ. Combining horizons still lowers the RMSE, although the gains level off when the noise is weak.

4.3 Robustness to tuning parameters and to price and volatility jumps

We also evaluate the feasible CLT in (9) at $t = 9{:}45$, $12{:}00$, and $13{:}30$ in Figure 2. The noise uses $\rho_\chi = 0.1$. The spot estimator and the asymptotic variance estimator both use the same differencing horizons: $k_{n,p} \in \{3, 4, 5\}$. The spot estimator uses local window bandwidth $l_n \in \{300, 500, 800\}$, and the asymptotic variance estimator uses $(L_{n,1}, q_{n,1}, d_{n,1}) = (\lfloor 5\sqrt{n} \rfloor, 20, \lfloor 0.25\sqrt{n} \rfloor)$ and $(L_{n,2}, q_{n,2}, d_{n,2}) = (\lfloor 20\sqrt{n} \rfloor, 10, \lfloor 2\sqrt{n} \rfloor)$. The QQ plots of the standardized spot variance estimator are close to the 45-degree line for $l_n = 500$ and 800 at all three times. With $l_n = 300$, the upper

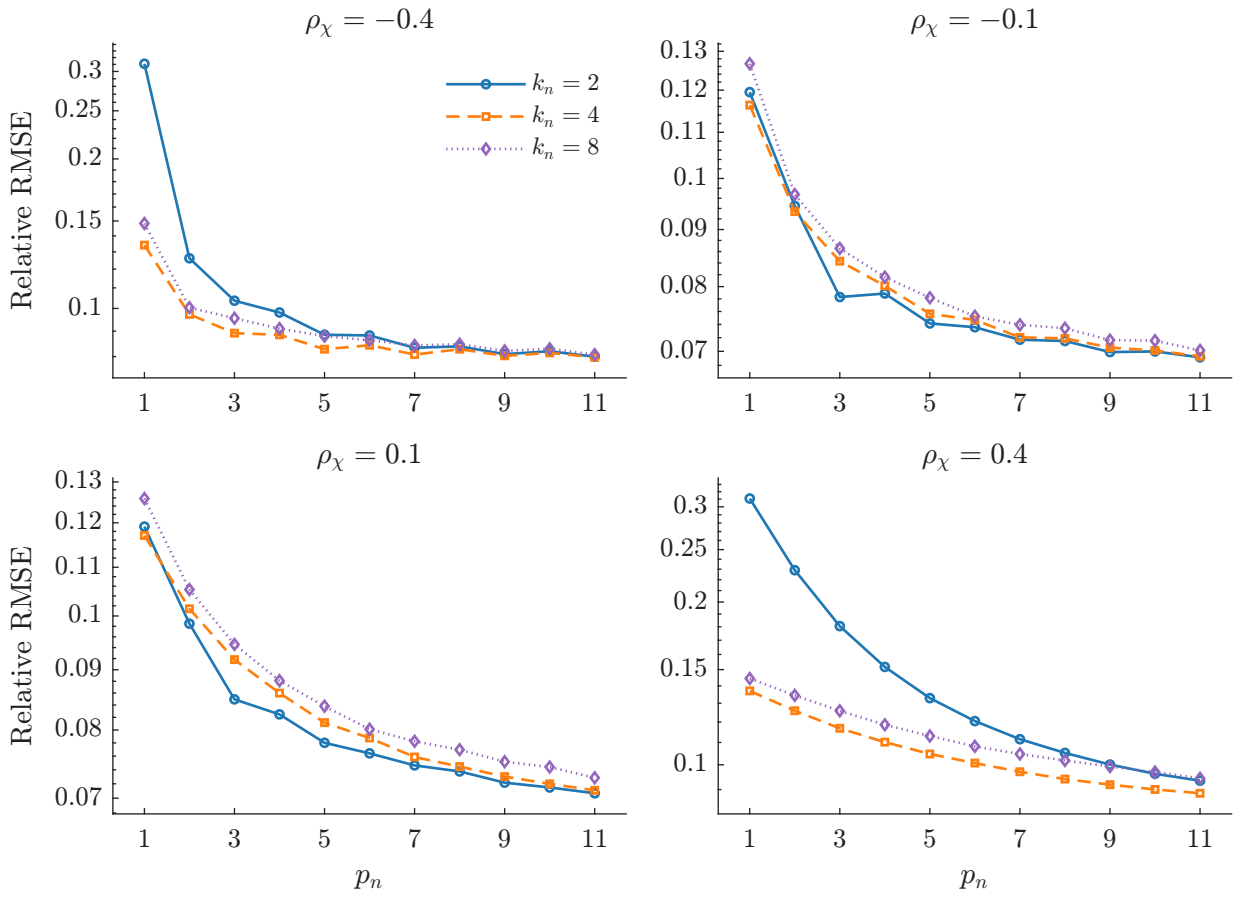


Figure 1: Relative RMSE of $\widehat{\Gamma}_t^n$ at $t = 9{:}45$ over 1,000 replications with $l_n = 500$. The panels vary the AR(1) coefficient ρ_χ of χ; $p_n = 1$ gives the single-horizon version of (7).

tail bends above the line at 9:45 and 12:00. This moderate distortion is expected given the small number of observations in each window, and it fades as l_n increases.[8]

We next evaluate the prescheduled jump test in Corollary 3.1 at $\tau = 14{:}00$, the time of the FOMC policy announcement that the empirical analysis below examines. In both designs, we add a permanent increment $\varphi \in [0, 4 \times 10^{-7}]$ to the baseline noise variance, so that $\Gamma_t^{(\varphi)} = \Gamma_t^{\text{base}} + \varphi \mathbf{1}_{\{t \geq \tau\}}$. The jump is therefore $\Delta\Gamma_\tau^{(\varphi)} = \varphi$, measured in squared log-price units. In Panel (a) of Figure 3, the scheduled jump in noise variance is the only change. Panel (b) adds simultaneous jumps in the efficient price and its volatility. This design tests whether announcement-time changes in these quantities could be mistaken for a jump in noise variance. Specifically, let Z follow the original efficient-price equation, but replace its diffusion coefficient with $\sigma_t^J = \sigma_t(1 + \mathbf{1}_{\{t \geq \tau\}})$, and set $X_t^J = Z_t + 0.01 \cdot \mathbf{1}_{\{t \geq \tau\}}$. The scheduled step in the log price is thus 0.01, a price increase of about 1%. Diffusion volatility doubles after the event, so diffusion variance quadruples. The original volatility process is the baseline to which this persistent multiplier applies.

Each panel reports 1,000 replications for each of 75 equally spaced jump sizes, with common

[8] Section SC of the supplementary materials repeats this check with smaller noise amplitudes, $K_\gamma \in \{2.5 \times 10^{-4}, 10^{-4}\}$, and stronger serial dependence.

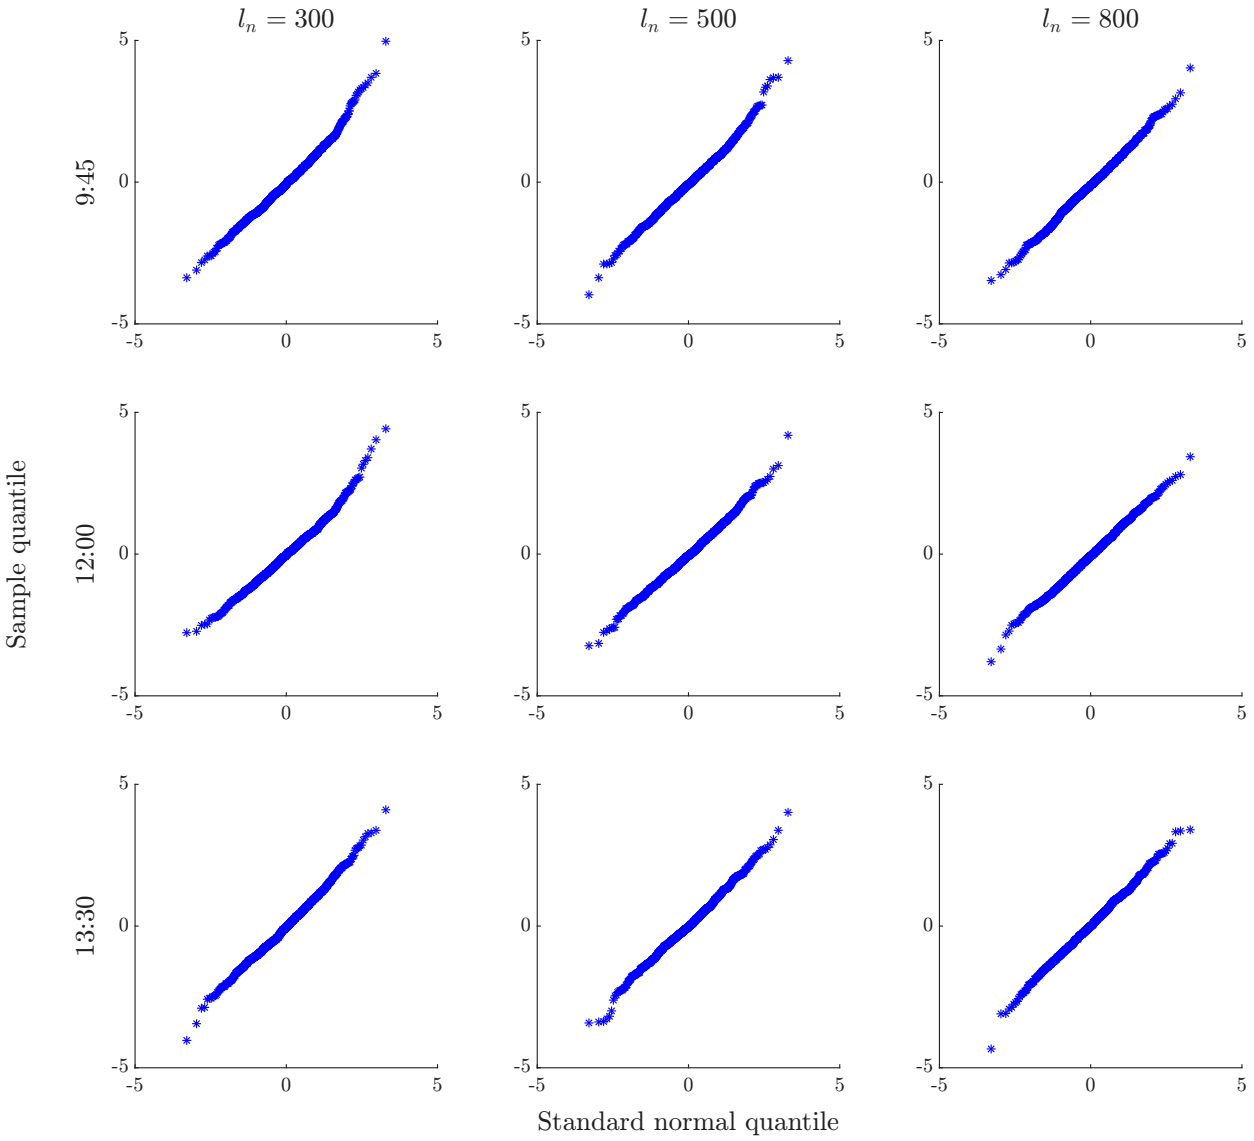


Figure 2: QQ plots of the standardized statistic in (9) over 1,000 replications, by target time (rows) and bandwidth l_n (columns).

draws across sizes within a panel and independent draws across panels. Both designs use $l_n = 300$, $k_{n,p} \in \{3, 4, 5\}$ for the spot variance estimator and asymptotic variance estimator, $\rho_\chi = 0.1$, and $K_\gamma = 5 \times 10^{-4}$. In panel (a), the rejection rates under the null are 10.2%, 4.6%, and 1.1% at the nominal 10%, 5%, and 1% levels. With simultaneous jumps in price and volatility, the rates are 9.5%, 4.7%, and 0.9%. Both sets of rates are close to the nominal levels, so this design shows no evidence of overrejection. Power is also similar across the two panels. At $\varphi = 2 \times 10^{-7}$, the rejection rates are 98.3%, 95.9%, and 85.5% in the baseline design and 97.7%, 95.2%, and 85.9% under simultaneous jumps. At $\varphi = 4 \times 10^{-7}$, all three rates exceed 99% in both panels.

This robustness follows from the construction of $f(k)_i^n$. Every observation used on a given side of τ lies strictly on that side, so the scheduled price step cancels from each price difference.

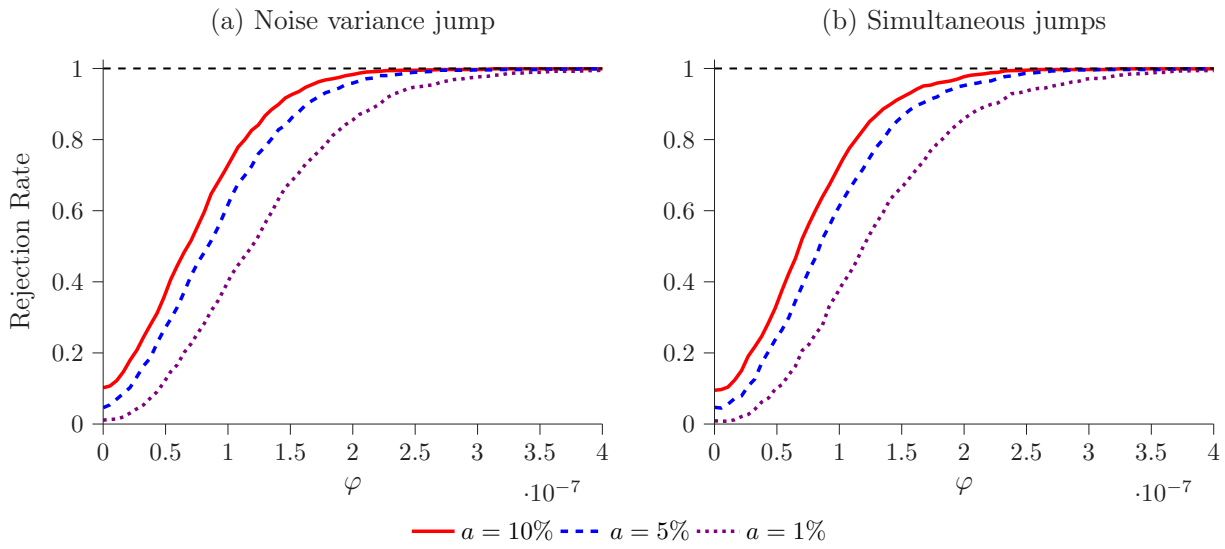


Figure 3: Rejection rates of the two-sided scheduled-jump test at $\tau = 14{:}00$ against the noise-variance jump φ, at nominal levels $a = 10\%, 5\%, 1\%$. Panel (b) adds a log-price step of 0.01 and doubles diffusion volatility at τ. Dashed lines mark unit power.

Moreover, the two factors of $f(k)_i^n$ are returns over adjacent, nonoverlapping intervals. Diffusive increments and price jumps therefore enter only through cross-products with mean zero. Doubling the volatility makes these terms more volatile but does not shift the center of the estimator.

5 Empirical Application

We apply the spot noise variance estimator to intraday transaction prices of the SPDR S&P 500 ETF (SPY) from the TAQ database over 2016–2025. The sample contains 2,493 regular trading sessions; the 21 early-close sessions, which end at 13:00, are excluded. We keep regular-session trades that have no correction flag and none of the excluded sale-condition codes, and we remove trades priced far from their neighbors by the screen of Brownlees and Gallo (2006).[9] The overall goal is to describe the intraday dynamics of noise variance and to study their relation to scheduled news releases and trading deadlines. Throughout, times are stated in U.S. Eastern Time.

5.1 Announcement Groups and Estimation

We compare four groups of trading days categorized by the types of news announcements released on those days.[10] The first group is defined by the 10:00 core-news releases of the ISM Manufacturing Index, ISM Services (formerly Non-Manufacturing) Index, and Conference Board

[9] This mild data filtration removes around 0.7% of the total observations; Section SB.1 of the supplementary materials (Li and Zhang, 2026) gives the details.

[10] Table SB.2 of the supplementary materials (Li and Zhang, 2026) gives the assignment rule and the number of days in each group.

Consumer Confidence Index (344 days). The second group comprises scheduled FOMC policy announcements at 14:00 (68 days). Days containing both a core-news release and a scheduled FOMC announcement are excluded from these two groups. The third group comprises preliminary and final University of Michigan consumer sentiment releases at 10:00 on days without either type of core announcement (233 days). The last group, the comparison group, contains days without recorded releases in our selected calendar at 10:00 or 14:00 (900 days). The four groups omit 11 joint core-news/FOMC days and 937 days with other recorded releases.

We compare the left and right noise variance estimates using the two-sided test of Corollary 3.1, with two-scale variance estimates and a nominal 5% level. Rejection frequencies report the fraction of eligible days on which the null is rejected. On each side of a grid time, the spot estimator (7) uses the fixed horizons $\{3, 4, 5\}$ and up to 300 trades within the adjacent grid interval; a side with fewer than 150 trades is omitted. The two-scale variance estimates use windows of at most five and 15 minutes on the same side, which do not cross recorded releases, the NYSE/Nasdaq closing deadlines, or session boundaries.[11]

5.2 Intraday Profiles across Announcement Groups

Figure 4 presents estimates on a five-minute grid from 9:30 to 15:55. The four columns correspond to the core-news, FOMC, Michigan sentiment, and comparison groups, from left to right. The top row displays the cross-day averages of the left and right noise variance estimates on a logarithmic scale, together with the cross-day average of the full-session ReMeDI estimate of Li and Linton (2022) as a benchmark. Because the benchmark averages over the whole session, including the open and the spikes, it lies above most of the spot averages. The bottom row displays the rejection frequencies at the nominal 5% level.

The announcement groups differ most clearly at their scheduled release times. At 10:00, the average right-side estimate is about 15.1 times the left-side estimate for core news and 9.3 times for Michigan sentiment. On FOMC days, the corresponding ratio at 14:00 is 16.4. At these times, the rejection frequencies rise to 46.8%, 31.8%, and 45.6%, respectively. The comparison group shows no comparable spike at either announcement time, with rejection frequencies of 7.7% at 10:00 and 7.0% at 14:00. Because large changes cluster at known release times, the data support including a prescheduled jump component in Γ in (4). Noise models that are continuous at fixed times (Jacod, Li, and Zheng, 2017; Da and Xiu, 2021; Li and Linton, 2022) rule out such a component.

A separate feature appears at 15:50 in every group. Rejection frequencies at this time range from 38.2% to 52.8% and equal 42.7% on comparison days. Those days have no recorded release at 15:50, so the feature may reflect closing-auction activity on the exchanges that list the stocks in SPY.[12] Table 1 summarizes the three rule changes. For each of the four periods they define, it also reports rejection frequencies across all sessions. As the timing of imbalance information and order

[11] Section SB.3 of the supplementary materials describes the implementation. Section SB.4 reports rejection frequencies under alternative fixed horizons and bandwidths.

[12] SPY itself is listed on NYSE Arca, whose closing auction restricts market-on-close and limit-on-close orders only from 15:59 (Rule 7.35-E(d)(2); see 90 FR 43676). Thus neither 15:50 nor 15:55 is a deadline in SPY's own auction.

restrictions changes, rejections shift between 15:50 and 15:55. In particular, the rejection frequency at 15:55 falls after Nasdaq begins releasing early imbalance information at 15:50, even though its cutoff for entering market-on-close orders remains at 15:55.

Table 1: Closing-auction rules and noise-variance jump rejections.

Sample period	Sessions	NYSE cutoff and imbalance start	Nasdaq first imbalance / MOC cutoff	Rejection (%) 15:50	Rejection (%) 15:55
Jan. 2016–26 Oct. 2018	707	15:45	15:50 / 15:50	23.9	12.4
29 Oct. 2018–29 Mar. 2019	102	15:45	15:55 / 15:55	9.8	46.1
1 Apr. 2019–1 Nov. 2019	150	15:50	15:55 / 15:55	38.7	55.3
4 Nov. 2019–Dec. 2025	1,534	15:50	15:50* / 15:55	56.4	18.0

Rejection frequencies use the baseline test on the five-minute grid. The NYSE column reports the market-on-close (MOC) and limit-on-close cutoff and the start of imbalance publication.
*Notes**: From November 4, 2019, Nasdaq publishes an early imbalance indicator and restricts cancellations at 15:50; full imbalance dissemination and the MOC entry cutoff remain at 15:55.[13]

5.3 One-Minute Event-Time Dynamics

The five-minute grid locates the sharp changes in Γ, but it cannot show how long an increase lasts or how Γ moves around it; in particular, the return to baseline by the next five-minute grid point could reflect either rapid decay in Γ' or a second, downward jump shortly after the release. Figure 5 examines the dynamics around the same announcements on a one-minute grid. The first three columns retain the announcement groups, while the fourth shows the minutes around 15:50 on the 1,534 sessions from November 4, 2019 onwards, which form the last rule period in Section 5.2.

Figure 5 shows that, at one-minute resolution, each event has a sharp move at the event time—consistent with the scheduled jump φ_ℓ in (4)—along with more gradual changes before and after it. For the 10:00 releases, the increase is short-lived. The right-side estimate surges in that minute and then reverts to its pre-release level in the next. Michigan sentiment days repeat this pattern at a smaller magnitude. FOMC announcements look different. Both sides rise in the minutes before 14:00, spike at 14:00, and then decline gradually, staying above their pre-event level at 14:10. The upward drift may reflect anticipation of the announcement and changes in liquidity provision. Both are plausible before what is arguably the most consequential scheduled economic event. Event-minute rejection frequencies are 46.8% (core news), 45.6% (FOMC), and 31.8% (Michigan sentiment), the same as the five-minute values in Figure 4. The last column sharpens the evidence

[13] *Sources:* Nasdaq Equity Trader Alert 2018-82 and 2019-66; NYSE Trader Update, March 26, 2019.

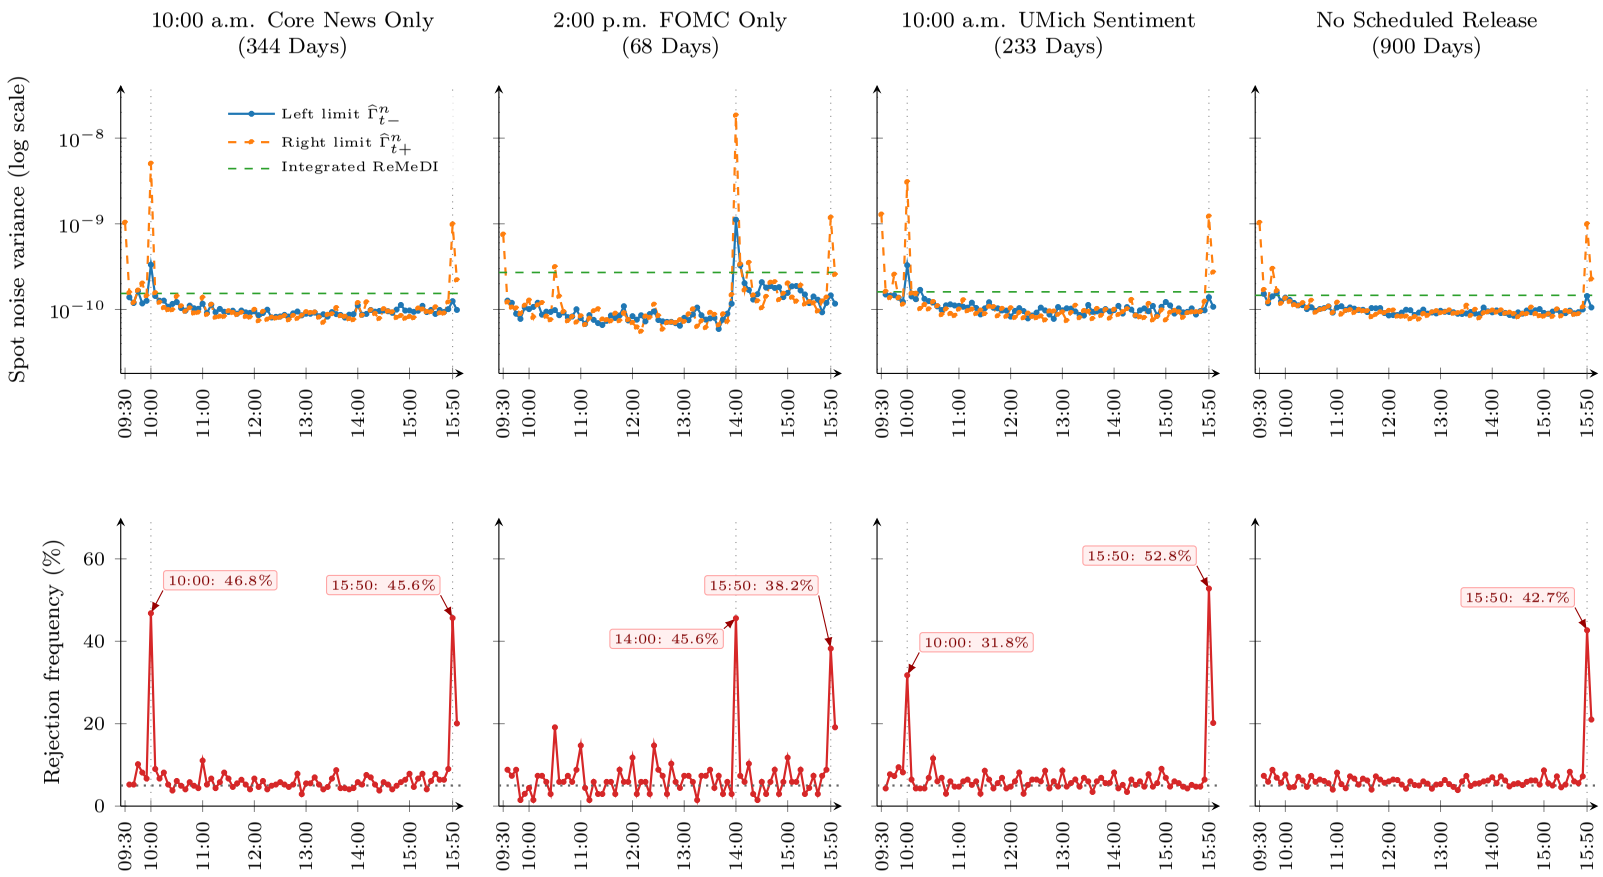


Figure 4: Intraday spot noise variance (top) and two-sided 5% rejection frequencies (bottom) on a five-minute grid for the day groups of Section 5.1, SPY 2016–2025. The top row shows cross-day means of $\widehat{\Gamma}^n_{t-}$ and $\widehat{\Gamma}^n_{t+}$ on a logarithmic scale and the group-average full-session single-horizon benchmark. Dotted lines mark the release time, 15:50, and the 5% level.

discussed in Section 5.2. Under the last rule regime, noise variance spikes in the deadline minute, which coincides with a peak in rejection frequencies.

5.4 Empirical Implications

Our findings have several implications for the intraday analysis of high-frequency data. First, the noise variance changes abruptly at times known in advance, and these changes coincide with scheduled events. They cluster at announcements and are absent at the same times on comparison days. Both facts support adding a prescheduled jump component to the noise variance model, especially when studying intraday noise dynamics. Existing frontier models of noise omit this component.

Second, the findings have implications for empirical practice. Researchers typically estimate intraday patterns by pooling observations across days at each time of day. Pooling announcement days, which are frequent, with other days combines event-driven changes with the regular intraday pattern. Researchers should therefore estimate intraday profiles conditional on the announcement calendar, for example by estimating separate profiles for announcement and non-announcement days.

Third, prescheduled jumps in Γ can arise from sources other than news releases. The deterministic set $\widetilde{\mathcal{T}}$ in (4) can also include times specified by trading rules, such as the 15:50 and 15:55 cutoffs for submitting closing-auction orders and the times when exchanges publish order imbalances. Exchange rulebooks specify these times ex ante. When the rules changed, the spike shifted to a different time. Researchers should therefore avoid pooling observations across rule regimes.

Finally, our findings have implications for the intraday measurement of other quantities, including spot volatility. Researchers often estimate spot volatility using sparse subsamples of high-frequency data, such as 1-minute returns (Bollerslev, Li, and Xue, 2018; Bollerslev, Li, and Liao, 2021), and evaluate these estimates on a coarser grid, such as every 30 minutes. Because markets react quickly to news releases, estimates evaluated on a coarse grid may miss abrupt changes in spot volatility. Finer grids and ultra-high-frequency data can better capture these changes but require accounting for microstructure noise. Our results show that noise changes abruptly around these events. Researchers should therefore model and estimate noise separately before and after each event. Otherwise, they may mistake a jump in noise for a jump in volatility.

6 Conclusion

We develop estimation and inference for the spot variance of serially dependent microstructure noise, allowing irregular sampling and scheduled variance jumps. Averaging across differencing horizons yields one-sided estimators asymptotically equivalent to an oracle that observes the noise. A joint stable central limit theorem and consistent two-scale variance estimators provide confidence intervals for the left and right limits and a test for scheduled jumps across bandwidth regimes.

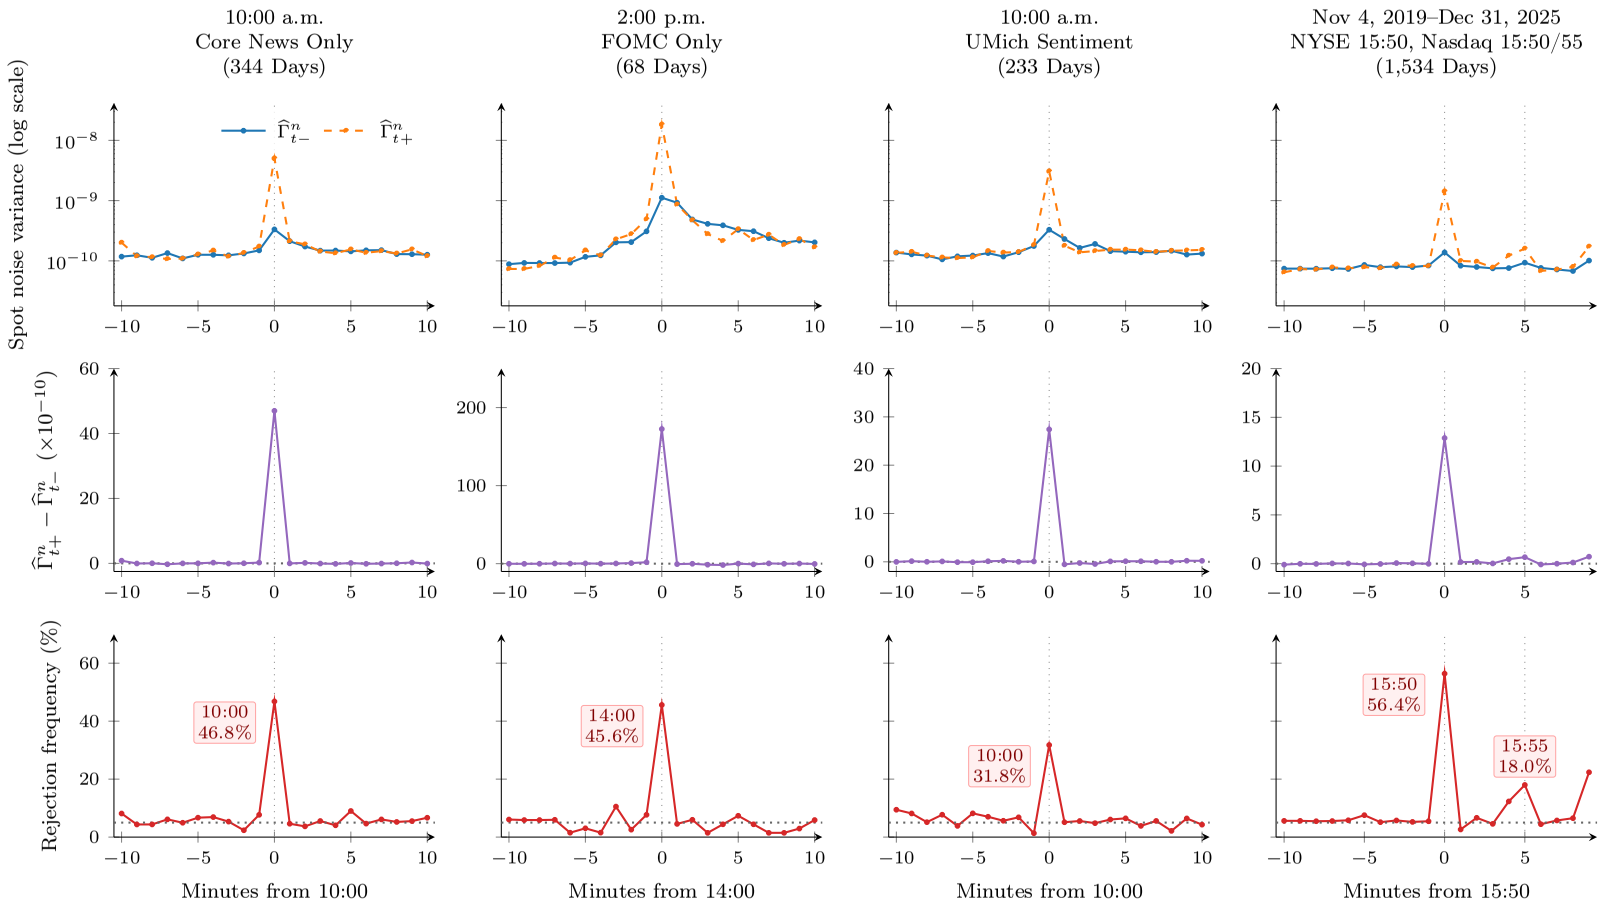


Figure 5: One-minute event-time profiles of spot noise variance, SPY 2016–2025. The rows show means of paired spot estimates on a logarithmic scale, the mean right-minus-left difference in units of 10^{-10}, and two-sided 5% rejection frequencies. The first three columns are the announcement groups of Figure 4; the last covers the closing auction from November 4, 2019. Dotted lines mark the event time, 15:55, and the 5% level.

Simulations support the method's finite-sample performance, including when noise variance jumps coincide with jumps in efficient prices and volatility. In SPY transactions over 2016–2025, noise variance rises sharply at macroeconomic releases, with recovery dynamics varying across announcement types. Near the close, bursts align with auction deadlines and imbalance disclosures, and their timing shifts with trading-rule changes. These findings highlight the importance of accounting for announcement calendars and trading rules when measuring intraday market quality, and of estimating noise variance separately on each side of events.

Appendix A Asymptotic Conditions for Theorems

In the sequel, we let $k_n \asymp \Delta_n^{-\kappa}$, $p_n \asymp \Delta_n^{-\varrho}$, $l_n \asymp \Delta_n^{-\lambda}$; and for $\ell \in \{1,2\}$, we let $d_{n,\ell} \asymp \Delta_n^{-\iota_\ell}$, $q_{n,\ell} \asymp \Delta_n^{-\theta_\ell}$.

Condition A.1. *Assume that $v > 1$, and let $\lambda \in (0,1)$.*

Condition A.2. *Assume that $v > 1$, and let $\lambda \in (0,1)$ with $c_n := l_n^2 \Delta_n \to c \in [0,\infty]$.*

Condition A.3. *Assume that $v > 1$, and let $\lambda \in (0,1)$, $\kappa \in (0,\lambda)$ and $\varrho \in [0,\lambda)$.*

Condition A.4. *Assume that $v > 2$ and $c_n := l_n^2 \Delta_n \to c \in [0,\infty]$, and let*

$$\lambda \in (0,1), \quad \kappa \in \left(\frac{\lambda \wedge (1-\lambda)}{v}, \lambda - \frac{\lambda \wedge (1-\lambda)}{2} \right), \quad \varrho \in \left(0, \lambda - \frac{\lambda \wedge (1-\lambda)}{2} \right).$$

Condition A.5. *Define $\bar{\iota}_\ell := \iota_\ell \wedge (1-\iota_\ell)$ for $\ell \in \{1,2\}$. Assume that $v > 2$, and let*

$$\kappa \in \left(\frac{\bar{\iota}_1 \vee \bar{\iota}_2}{v}, \bar{\iota}_1 \wedge \bar{\iota}_2 \right), \quad \varrho \in (0, \bar{\iota}_1 \wedge \bar{\iota}_2), \quad \bar{\iota}_\ell \in (\kappa \vee \varrho, v\kappa), \quad \theta_\ell \in (0, 1-\iota_\ell),$$

and require $\left(\iota_1 - \frac{1}{2}\right)\left(\iota_2 - \frac{1}{2}\right) \leq 0$. In addition, let $\delta_{n,\ell} := d_{n,\ell}^2 \Delta_n \to \delta_\ell \in [0,\infty]$; if $\iota_1 = \iota_2 = \frac{1}{2}$, we require $\delta_1 \neq \delta_2$.

Appendix B Asymptotic Variance Estimators

Inspired by Mykland and Zhang (2017), we propose the following estimator for the asymptotic variance. For any two sequences of integers $\{d_n\}_n$ and $\{L_n\}_n$, let $q_n := \lfloor L_n/d_n \rfloor$, $I_{t+}^n := n_t$, $I_{t-}^n := n_t - L_n$, and define $\mathbb{I}(q)_{t\pm}^n := \{I_{t\pm}^n + qd_n + 2\overline{k}_n, \ldots, I_{t\pm}^n + (q+1)d_n - 2\overline{k}_n\}$. Let $\widehat{\Gamma}(q;d_n)_{t\pm}^n := \frac{1}{d_n p_n} \sum_{p=0}^{p_n-1} \sum_{i \in \mathbb{I}(q)_{t\pm}^n} f(k_{n,p})_i^n$, $\Theta(q;d_n)_{t\pm}^n := \widehat{\Gamma}(q;d_n)_{t\pm}^n - \widehat{\Gamma}(q-1;d_n)_{t\pm}^n$, and $\Sigma(L_n,d_n)_{t\pm}^n := \frac{d_n}{2q_n} \sum_{q=1}^{q_n-1} \Theta(q;d_n)_{t\pm}^{n\,2}$ for any $t \in \mathcal{T}$. Then, let $(L_{n,1},d_{n,1})$ and $(L_{n,2},d_{n,2})$ be two pairs of parameters of different scales. We denote $\Sigma(L_{n,1},d_{n,1})_{t\pm}^n$ and $\Sigma(L_{n,2},d_{n,2})_{t\pm}^n$ by $\Sigma(1)_{t\pm}^n$ and $\Sigma(2)_{t\pm}^n$, respectively. Let $\delta_{n,1} := d_{n,1}^2 \Delta_n$ and $\delta_{n,2} := d_{n,2}^2 \Delta_n$. We define

$$\widehat{\mathfrak{S}}(1)_{t\pm}^n := \frac{\delta_{n,1}^{-1}\Sigma(1)_{t\pm}^n - \delta_{n,2}^{-1}\Sigma(2)_{t\pm}^n}{\delta_{n,1}^{-1} - \delta_{n,2}^{-1}}, \ \widehat{\mathfrak{S}}(2)_{t\pm}^n := \frac{\Sigma(1)_{t\pm}^n - \Sigma(2)_{t\pm}^n}{\delta_{n,1} - \delta_{n,2}}, \ \widehat{\mathfrak{S}}_{t\pm}^n := \frac{\widehat{\mathfrak{S}}(1)_{t\pm}^n + c_n \widehat{\mathfrak{S}}(2)_{t\pm}^n}{1 + c_n}.$$

We have the following theorem for the consistency of this feasible asymptotic variance estimator.

Theorem B.1. *Let Condition A.5 hold. For any $l_n \to \infty$, $c_n := l_n^2 \Delta_n \to c \in [0, \infty]$, and $t \in \mathcal{T}$, we have*

$$\widehat{\mathfrak{S}}_{t\pm}^n \xrightarrow{\mathbb{P}} \mathfrak{S}_{t\pm},$$

where $\mathfrak{S}_{t\pm}$ is defined in Theorem 3.2.

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 J. JACOD (2014): *High-Frequency Financial Econometrics*, Princeton University Press.

AÏT-SAHALIA, Y. AND J. YU (2009): "High frequency market microstructure noise estimates and liquidity measures," *Annals of Applied Statistics*, 3, 422–457.

ANDERSEN, T. G., R. RIVA, M. THYRSGAARD, AND V. TODOROV (2023): "Intraday Cross-Sectional Distributions of Systematic Risk," *Journal of Econometrics*, 235, 1394–1418.

ANDERSEN, T. G., T. SU, V. TODOROV, AND Z. ZHANG (2024): "Intraday periodic volatility curves," *Journal of the American Statistical Association*, 119, 1181–1191.

ANDERSEN, T. G., M. THYRSGAARD, AND V. TODOROV (2021): "Recalcitrant Betas: Intraday Variation in the Cross-Sectional Dispersion of Systematic Risk," *Quantitative Economics*, 12, 647–682.

BANDI, F. M. AND J. R. RUSSELL (2006): "Separating microstructure noise from volatility," *Journal of Financial Economics*, 79, 655–692.

——— (2008): "Microstructure noise, realized variance, and optimal sampling," *Review of Economic Studies*, 75, 339–369.

BOLLERSLEV, T., J. LI, AND Z. LIAO (2021): "Fixed-k inference for volatility," *Quantitative Economics*, 12, 1053–1084.

BOLLERSLEV, T., J. LI, AND Y. REN (2024): "Optimal inference for spot regressions," *American Economic Review*, 114, 678–708.

BOLLERSLEV, T., J. LI, AND Y. XUE (2018): "Volume, Volatility, and Public News Announcements," *Review of Economic Studies*, 85, 2005–2041.

BROWNLEES, C. T. AND G. M. GALLO (2006): "Financial econometric analysis at ultra-high frequency: Data handling concerns," *Computational Statistics & Data Analysis*, 51, 2232–2245.

CHONG, C. H., T. DELERUE, AND G. LI (2025): "When frictions are fractional: Rough noise in high-frequency data," *Journal of the American Statistical Association*, 120, 1531–1544.

CHORDIA, T. AND A. SUBRAHMANYAM (2004): "Order imbalance and individual stock returns: Theory and evidence," *Journal of Financial Economics*, 72, 485–518.

CREGO, J. A. (2020): "Why does public news augment information asymmetries?" *Journal of Financial Economics*, 137, 72–89.

CUI, W. (2026): "The Explicative Market Microstructure Noise," *Journal of the American Statistical Association*, forthcoming.

DA, R. AND D. XIU (2021): "When Moving-Average Models Meet High-Frequency Data: Uniform Inference on Volatility," *Econometrica*, 89, 2787–2825.

——— (2026): "Disentangling Autocorrelated Intraday Returns," Working paper. Available at SSRN: https://ssrn.com/abstract=3857785.

DELBAEN, F. AND W. SCHACHERMAYER (1994): "A general version of the fundamental theorem of asset pricing," *Mathematische Annalen*, 300, 463–520.

FLEMING, M. J. AND E. M. REMOLONA (1999): "Price formation and liquidity in the U.S. Treasury market: The response to public information," *Journal of Finance*, 54, 1901–1915.

FOSTER, D. P. AND D. B. NELSON (1996): "Continuous Record Asymptotics for Rolling Sample Variance Estimators," *Econometrica*, 64, 139–174.

HASBROUCK, J. (1993): "Assessing the quality of a security market: A new approach to transaction-cost measurement," *Review of Financial Studies*, 6, 191–212.

——— (2007): *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*, Oxford University Press.

HENDERSHOTT, T. AND A. J. MENKVELD (2014): "Price pressures," *Journal of Financial Economics*, 114, 405–423.

JACOD, J., Y. LI, AND X. ZHENG (2017): "Statistical properties of microstructure noise," *Econometrica*, 85, 1133–1174.

——— (2019): "Estimating the integrated volatility with tick observations," *Journal of Econometrics*, 208, 80–100.

JACOD, J. AND P. E. PROTTER (2012): *Discretization of Processes*, vol. 67, Springer Science & Business Media.

KALNINA, I. AND O. LINTON (2008): "Estimating quadratic variation consistently in the presence of endogenous and diurnal measurement error," *Journal of Econometrics*, 147, 47–59.

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

LEE, C. M. C., B. MUCKLOW, AND M. J. READY (1993): "Spreads, depths, and the impact of earnings information: An intraday analysis," *Review of Financial Studies*, 6, 345–374.

LEHMANN, B. N. (1990): "Fads, martingales, and market efficiency," *Quarterly Journal of Economics*, 105, 1–28.

LI, Z. M. AND O. LINTON (2022): "A ReMeDI for microstructure noise," *Econometrica*, 90, 367–389.

——— (2026): "Robust estimation of integrated and spot volatility," *Journal of Econometrics*, 254, 105614.

LI, Z. M. AND X. YANG (2026): "Multi-horizon test for market frictions," *Journal of Econometrics*, 253, 106171.

LI, Z. M. AND H. ZHANG (2026): "Supplementary Materials for 'Estimation and Inference for Spot Variance of Microstructure Noise'," Working Paper.

MYKLAND, P. A. AND L. ZHANG (2017): "Assessment of uncertainty in high frequency data: The observed asymptotic variance," *Econometrica*, 85, 197–231.

XIU, D. (2010): "Quasi-maximum likelihood estimation of volatility with high frequency data," *Journal of Econometrics*, 159, 235–250.

ZHANG, C., J. LI, AND T. BOLLERSLEV (2022): "Occupation Density Estimation for Noisy High-Frequency Data," *Journal of Econometrics*, 227, 189–211.

ZHANG, L., P. A. MYKLAND, AND Y. AÏT-SAHALIA (2005): "A tale of two time scales: determining integrated volatility with noisy high-frequency data," *Journal of the American Statistical Association*, 100, 1394–1411.

ZU, Y. AND H. P. BOSWIJK (2014): "Estimating Spot Volatility with High-Frequency Financial Data," *Journal of Econometrics*, 181, 117–135.