AARHUS UNIVERSITY

Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the accepted manuscript version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Andersen, T. G., Thyrgaard, M., & Todorov, V. (2019). Time-Varying Periodicity in Intraday Volatility. *Journal of the American Statistical Association*, 114(528), 1695-1707.
<https://doi.org/10.1080/01621459.2018.1512864>

Publication metadata

Title:	Time-Varying Periodicity in Intraday Volatility
Author(s):	Andersen, Torben Gustav ; Thyrgaard, Martin ; Todorov, Viktor
Journal:	Journal of the American Statistical Association
DOI/Link:	10.1080/01621459.2018.1512864
Document version:	Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

If the document is published under a Creative Commons license, this applies instead of the general rights.

Time-Varying Periodicity in Intraday Volatility*

Torben G. Andersen[†] Martin Thyrsgaard[‡] Viktor Todorov[§]

June 18, 2018

Abstract

We develop a nonparametric test for whether return volatility exhibits time-varying intraday periodicity using a long time series of high-frequency data. Our null hypothesis, commonly adopted in work on volatility modeling, is that volatility follows a stationary process combined with a constant time-of-day periodic component. We construct time-of-day volatility estimates and studentize the high-frequency returns with these periodic components. If the intraday periodicity is invariant, then the distribution of the studentized returns should be identical across the trading day. Consequently, the test compares the empirical characteristic function of the studentized returns across the trading day. The limit distribution of the test depends on the error in recovering volatility from discrete return data and the empirical process error associated with estimating volatility moments through their sample counterparts. Critical values are computed via easy-to-implement simulation. In an empirical application to S&P 500 index returns, we find strong evidence for variation in the intraday volatility pattern driven in part by the current level of volatility. When volatility is elevated, the period preceding the market close constitutes a significantly higher fraction of the total daily integrated volatility than during low volatility regimes.

JEL classification: C51, C52, G12.

Keywords: high-frequency data, periodicity, semimartingale, specification test, stochastic volatility.

*Andersen's and Todorov's research is partially supported by NSF grant SES-1530748. We would like to thank anonymous referees for many helpful comments and suggestions.

[†]Department of Finance, Kellogg School, Northwestern University, NBER and CREATES.

[‡]CREATES, Department of Economics and Business Economics, Aarhus University.

[§]Department of Finance, Kellogg School, Northwestern University.

1 Introduction

Stock returns have time-varying volatility and this has important theoretical as well as practical ramifications. Most existing work on volatility assumes it is a stationary process. However, there is both theoretical (see, e.g., [1, 27]) and empirical evidence (see, e.g., [2, 3]) for the presence of intraday periodicity in volatility. To illustrate this phenomenon, on Figure 1, we plot the average level of the S&P 500 index return volatility as a function of time-of-day. As seen from the figure, the intraday periodic component of volatility is nontrivial. Specifically, the average volatility at the market close is about three times the average volatility around lunch.

Figure 1: Intraday Volatility Periodicity for the S&P 500 Index. The plot presents smoothed estimates of the average time-of-day volatility, normalized by the trading day volatility. Details regarding the construction of the series are provided in Section 5.

High-frequency data is commonly used, as it offers significant efficiency gains for measuring and forecasting volatility, see, e.g., [5]. The pronounced periodic pattern in Figure 1

has strong implications regarding the appropriate methodology for studying volatility using intraday return data. The usual approach stipulates that the time-of-day component of volatility is constant across days, and then one standardizes the high-frequency returns by the corresponding estimates, see, e.g., [2, 3], [9] and [38] among many others. However, this only annihilates the intraday volatility component from the returns, if the periodicity is time-invariant. The goal of the current paper is to test this (null) hypothesis within a general nonparametric setting. Moreover, if the null hypothesis is rejected, we provide techniques that can help identify the sources of variation in the intraday periodic component. The statistical analysis is conducted using a long span of high-frequency return data.

The major challenges in designing the test stem from the fact that volatility is not directly observed and both the stationary and periodic component of volatility can change over the course of the day. We exploit the long time span as well as the short distance between the intraday observations to circumvent these latency problems. We first estimate the average periodic component of volatility from the high-frequency returns, and then standardize the returns with these estimated time-of-day components. Under the null hypothesis, this studentization of the returns annihilates the periodic volatility component. Therefore, the studentized high-frequency returns should possess an identical distribution, regardless of time-of-day. In contrast, under the alternative hypothesis, this is violated, as the studentized return distribution is given by a convolution of the distributions for the stationary volatility component and the standardized (time-varying) periodic volatility component. Hence, the distribution of the studentized returns depends on the time-of-day, when the periodic volatility component fluctuates over time.

Given the discrepancy in the distributional properties of the studentized returns under the null and alternative hypotheses, our test statistic is designed to measure the distance between the studentized return distribution for different parts of the trading day. In par-

ticular, we rely on a weighted L^2 norm of the difference in the real parts of the empirical characteristic functions of the studentized returns. We use only the real parts of these functions, because the high-frequency returns are approximately conditionally Gaussian, when conditioning on the information at the start of the return interval.

Under the null hypothesis, the limiting distribution of our test statistic depends on the error in recovering volatility from the returns as well as the empirical process error associated with estimating population moments of volatility using their sample counterparts. The contribution of the first of these errors to the limiting distribution is a distinctive feature of our test, setting it apart from other estimation and testing problems involving joint in-fill and long-span asymptotics for high-frequency return data, where this error is negligible asymptotically. The reason for the added complication is that, due to the nature of the testing problem at hand, we only use a limited number of high-frequency returns per day in forming the statistic. Hence, we cannot derive the limiting distribution assuming that volatility, effectively, is observed, which is a convenient simplification in determining the asymptotic properties of existing joint in-fill and long-span inference procedures. As a consequence, the limit distribution of our test statistic is non-standard, but its quantiles are readily evaluated through simulation.

We further extend our theoretical results to cover situations in which the observed prices are contaminated with microstructure noise. We follow [13], [16, 17] and [35] and assume that the noise can be modeled as a (known) function of observables and a finite-dimensional parameter vector. Given an estimate from the data for the latter, we form estimates for the unobservable prices at the observation times and plug them into our no-noise test statistics for the volatility periodicity described above. Provided the noise parameters are estimated at a sufficiently fast rate, we show that the step involved in extracting an estimate of the true price from the noisy observations has no first-order effect on our test.

We implement our new testing procedure on high-frequency return data for the S&P 500 index. Even after excluding trading days comprising scheduled macroeconomic announcements, our test rejects the null hypothesis of a time-invariant intraday periodicity in volatility. Additional analysis shows that a significant driver of the variation in the periodic volatility component is the concurrent level of volatility, as proxied by the VIX volatility index at the market open. Upon separating the trading days into regimes of low, medium, and high volatility, according to the level of the VIX, we find that our test rejects significantly less on these three subsamples. Specifically, when volatility is elevated, the period before the market close contributes a substantially higher fraction of the total integrated daily volatility compared with regimes featuring lower volatility.

Our paper is related to several strands of earlier work. First, there is a large literature on detecting and modeling periodicity in discrete time series. Examples include [14], [22], [24], [29], [34], [37] and [39]. Second, there is a sizable literature that estimates (assumed) constant intraday volatility patterns. This includes empirical work by [2, 3], [23], [25] and [40]. The (constant) intraday periodicity is further explicitly modeled or accounted for in papers estimating volatility models and detecting jumps, such as [9], [12], [19] and [30]. [26] study the commonality in the intraday periodicity across many assets ([41] and [42] recognize the importance of intraday volatility variation for high-dimensional multivariate volatility estimation). Finally, [15] assume that the stationary volatility component is constant within the day and test whether the periodic volatility component can explain the full dynamic evolution across each day in a setting where prices are contaminated by noise. We reiterate that a common feature of the above literature is the assumed invariance for the periodicity of volatility, while the goal of the current paper is to test this underlying assumption of existing work, and to explore potential sources of deviation from this hypothesis. Third, [4] consider testing for changes in the periodic component of

volatility at a specific (known) point in time and in a parametric volatility setting (both for the stationary and periodic components). Unlike that paper, the analysis here is fully nonparametric and we can test for changes in the periodic component, which can happen at unknown times and be stochastic. Fourth, our paper is related to a voluminous statistical literature that tests for the equality of two distributions by means of a weighted L^2 distance between the associated empirical characteristic functions. Applications of this approach for testing dependence between two variables can be found in, e.g., [7, 8], [18], [21] and [36]. The empirical characteristic function has also been used to study serial dependence in time series by [28] and to test for Gaussianity in stationary time series by [20]. The major difference between this strand of the literature and the current paper is that the variables, whose distributions are compared, are not directly observable and need to be “filtered” from the data, and this filtering procedure has a significant impact on the limiting distribution of the test statistic.

The rest of the paper is organized as follows. Section 2 presents the formal setup and introduces the statistics for assessing stochastic time-variation in the periodicity of intraday volatility. In Section 3 we derive the asymptotic limit theory for these statistics, and we then apply it for developing a feasible testing procedure for whether the intraday volatility pattern is time-invariant. Section 4 summarizes the results obtained from a large-scale simulation study (with additional simulation results reported in a Supplementary Appendix), and Section 5 presents an empirical implementation of the testing procedure. Section 6 concludes. All proofs are deferred to a Supplementary Appendix.

2 Setup and Estimation of Periodic Volatility

The (log) price process X is defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. Consistent with the absence of arbitrage, it follows an Itô semimartingale of the form,

$$dX_t = a_t dt + \tilde{\sigma}_t dW_t + \int_{\mathbb{R}} x \mu(dt, dx), \quad (1)$$

where a_t is the drift, W_t a Brownian motion, $\tilde{\sigma}_t$ the (diffusive) stochastic volatility, μ the counting measure for jumps in X with compensator $b_t dt \otimes F(dx)$, where b_t is a càdlàg process and $F : \mathbb{R} \rightarrow \mathbb{R}_+$. Our main focus is the stochastic volatility component. Beyond the customary stationary part, we assume it contains a periodic component with a cycle spanning one unit of time. Specifically,

$$\tilde{\sigma}_t^2 = \sigma_t^2 f_{[t], t - [t]}, \quad t \in \mathbb{R}_+, \quad (2)$$

for some stationary process σ_t and time-of-day function $f : \mathbb{N}_+ \times [0, 1] \rightarrow \mathbb{R}_+$ with $f_{[t], 0} = f_{[t], 1}$, where the time unit is one day.

In the standard setting, adopted in most current work, f is deterministic and depends only on the time-of-day, i.e., its second argument. However, it is plausible that the periodic component might vary with the concurrent level of (the stationary component of) volatility, as well as the occurrence of events such as prescheduled macroeconomic announcements and, more generally, any shifts in the organization and operation of the financial markets. The goal of the current paper is to test whether the time-of-day periodic component of volatility changes over time.

The formal assumptions for the process X needed for our analysis are stated below. In what follows we use the shorthand notation $V_t = \sigma_t^2$.

Assumption 1. *We have $\sup_{t \in \mathbb{R}_+} \mathbb{E}|a_t|^8 + \sup_{t \in \mathbb{R}_+} \mathbb{E}|b_t|^4 + \sup_{t \in \mathbb{R}_+} \mathbb{E}|V_t|^8 < \infty$ as well as*

$F(\mathbb{R}) < \infty$ and $\int_{\mathbb{R}} |x|^4 F(dx) < \infty$, and further

$$\inf_{t \in \mathbb{N}_+} \inf_{\kappa \in (0,1]} f_{t,\kappa} > \underline{\epsilon} \quad \text{and} \quad \sup_{t \in \mathbb{N}_+} \sup_{\kappa \in (0,1]} f_{t,\kappa} < \bar{\epsilon}, \quad (3)$$

for some non-random $0 < \underline{\epsilon} < \bar{\epsilon} < \infty$.

Assumption 2. The following smoothness in expectation conditions hold for $0 < s \leq t$:

$$\mathbb{E}|a_t - a_s|^2 + \mathbb{E}|b_t - b_s|^2 + \mathbb{E}|\tilde{\sigma}_t - \tilde{\sigma}_s|^2 + \mathbb{E}|V_t - V_s|^2 \leq C|t - s|, \quad (4)$$

for some positive and finite constant C that does not depend on t and s .

Assumption 3. For a given finite set $\mathcal{K}$ of numbers in $(0, 1]$, denote with $\mathbf{Y}_t = \{f_{t,\kappa}\sigma_{t+\kappa}^2\}_{\kappa \in \mathcal{K}}$ and suppose that it is a function of a Markov process $\{\tilde{\mathbf{Y}}_t\}_{t \in \mathbb{N}_+}$. We then assume that $\{\tilde{\mathbf{Y}}_t\}_{t \in \mathbb{N}_+}$ is stationary, ergodic and α -mixing with $\alpha_t = o(t^{-16/3})$ when $t \rightarrow \infty$, and where

$$\alpha_t = \sup\{|\mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)| : A \in \mathcal{G}_0, B \in \mathcal{G}^t\},$$

for $\mathcal{G}_0 = \sigma(\tilde{\mathbf{Y}}_s, s \leq 0)$ and $\mathcal{G}^t = \sigma(\tilde{\mathbf{Y}}_s, s \geq t)$.

Assumption 1 is our moment condition. This assumption also imposes finite activity of the jumps in X . Assumption 2 is a “smoothness in expectation” for the various processes involved in defining the dynamics of X . We note that this is a rather weak assumption which is satisfied if the processes are modeled as Itô semimartingales. In particular, the volatility process is allowed to contain a jump component. Finally, Assumption 3 imposes stationarity and mixing conditions for the process $\mathbf{Y}_t$. The specific order for the mixing coefficient α is needed for deriving the CLT for our statistic (and in particular to control the decay in the dependence between W_t and past V_{t-k} , for large k) as well as for establishing the consistency of the estimator of the asymptotic covariance operator. As usual for the analysis of the limiting behavior of processes with dependence, see e.g., Theorem VIII.3.79

in [32], there is a tradeoff between the requirement for the existence of moments and the tail decay of the α -mixing coefficient, with stronger moment conditions requiring slower tail decay of α .

We continue next with the construction of our statistics. The inference will be based on discrete observations of the process X at equidistant times $0, \frac{1}{n}, \frac{2}{n}, \dots, T$, where the integer T represents the time span, and the integer n indicates the number of times we sample within a unit interval. We denote the length of the sampling interval by $\Delta_n = 1/n$ and the high-frequency increments of X by

$$\Delta_{t,\kappa}^n X = X_{((t-1)n + \lfloor \kappa n \rfloor)/n} - X_{((t-1)n + \lfloor \kappa n \rfloor - 1)/n}, \quad t \in \mathbb{N}_+ \text{ and } \kappa \in (0, 1]. \quad (5)$$

The asymptotic setting involves $n \rightarrow \infty$ and $T \rightarrow \infty$, where, intuitively, the increasing sampling frequency assists in the nonparametric identification of the level of stochastic volatility from discrete observations of X , and the long time span allows us to separate the stationary and periodic components of volatility.

Our estimate for the time-of-day component of volatility is given by,

$$\hat{f}_\kappa = \frac{n}{T} \frac{\pi}{2} \sum_{t=1}^T |\Delta_{t,\kappa}^n X| |\Delta_{t,\kappa-\Delta}^n X| 1_{\{\mathcal{A}_{t,\kappa}^n\}}, \quad \mathcal{A}_{t,\kappa}^n = \{|\Delta_{t,\kappa}^n X| \leq v_n \cap |\Delta_{t,\kappa-\Delta}^n X| \leq v_n\}, \quad (6)$$

for $v_n = \alpha \Delta_n^\varpi$ with $\varpi \in (0, 1/2)$ and $\alpha > 0$. Under appropriate conditions, $\hat{f}_\kappa$ converges in probability to $\mathbb{E}(f_{t,\kappa} \sigma_{t+\kappa}^2)$, for $t \in \mathbb{N}_+$. Therefore, up to the constant $\mathbb{E}(\sigma_t^2)$, $\hat{f}_\kappa$ provides an estimate for the periodic component of volatility, when the latter is time-invariant.

We test for invariance of the intraday component of volatility by comparing the distribution of estimates for volatility deseasonalized by $\hat{f}_\kappa$ over different parts of the trading day. Under the null hypothesis, these distributions are identical, while they differ under the alternative. The inference for the distribution of volatility at different parts of the day

will be based on the studentized returns

$$\Delta_{t,\kappa}^n X / \sqrt{\widehat{f}_\kappa}, \quad t \in \mathbb{N}_+ \text{ and } \kappa \in (0, 1], \quad (7)$$

and the result in [38] that (the real part of) the empirical characteristic function of the high-frequency increments in X is an estimate for the Laplace transform of stochastic volatility. Therefore, we introduce,

$$\widehat{L}_\kappa^n(u) = \frac{1}{T} \sum_{t=1}^T \cos \left(\sqrt{2un} \Delta_{t,\kappa}^n X / \sqrt{\widehat{f}_\kappa} \right), \quad u \in \mathbb{R}_+ \text{ and } \kappa \in (0, 1], \quad (8)$$

and, as shown in the next section, $\widehat{L}_\kappa^n$ converges in probability (in a functional sense) to,

$$\mathcal{L}_\kappa(u) = \mathbb{E} \left[e^{-u f_{t,\kappa} \sigma_{t+\kappa}^2 / \mathbb{E}[f_{t,\kappa} \sigma_{t+\kappa}^2]} \right], \quad \text{for } t \in \mathbb{N}_+ \text{ and } \kappa \in (0, 1]. \quad (9)$$

3 Testing for Time-Invariant Periodicity of Volatility

We proceed with the formal asymptotic results for $\widehat{L}_\kappa^n$, which in turn will allow us to construct a feasible test for detecting time-variation in the intraday periodicity of volatility.

3.1 Infeasible Limit Theory

Our results will be based on the function $\widehat{L}_\kappa^n(u)$ in u , and the functional convergence results below take place in the Hilbert space $\mathcal{L}^2(w)$,

$$\mathcal{L}^2(w) = \left\{ f : \mathbb{R}_+ \rightarrow \mathbb{R} \left| \int_{\mathbb{R}_+} |f(u)|^2 w(u) du < \infty \right. \right\}, \quad (10)$$

for some positive-valued continuous weight function w with exponential tail decay. As usual, we denote the inner product and the norm on $\mathcal{L}^2(w)$ by $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$, respectively. Convergence in probability for $\widehat{L}_\kappa^n$ is established in the following theorem.

Theorem 1. *Suppose Assumptions 1-3 hold with $\mathcal{K} = \{\kappa\}$, for some $\kappa \in (0, 1]$, and $\varpi \in (0, \frac{1}{2})$. Then, as $n \rightarrow \infty$ and $T \rightarrow \infty$, we have,*

$$\widehat{L}_\kappa^n \xrightarrow{\mathbb{P}} \mathcal{L}_\kappa. \quad (11)$$

The intuition behind the above result is the following. First, $\widehat{f}_\kappa$ is an estimate of $\mathbb{E}(f_{t,\kappa}\sigma_{t+\kappa}^2)$. Second, over small time intervals, we have $\Delta_{t,\kappa}^n X \approx \widetilde{\sigma}_{t-1+\frac{[\kappa n]}{n}} \Delta_{t,\kappa}^n W$. From here, the result in Theorem 1 follows by a Law of Large Numbers.

Theorem 1 requires both $n \rightarrow \infty$ and $T \rightarrow \infty$, but imposes no restriction on their relative rate of growth. We emphasize that the above result is functional, i.e., we recover the Laplace transform $\mathcal{L}_\kappa$ as a function of u . As is well known, the Laplace transform of a positive-valued random variable uniquely identifies its distribution. Therefore, any differences in $\mathcal{L}_\kappa$ for different times-of-day (different values of κ) must stem from time variation in the periodic component of volatility. In this case, studentizing the high-frequency increments by the time-of-day estimate $\sqrt{\widehat{f}_\kappa}$, obtained from sample return moment as in equation (6), will not suffice to eliminate the intraday periodic component.

We next derive a Central Limit Theorem (CLT) for the difference in $\widehat{L}_\kappa^n$, for two different values of κ , under the null hypothesis.

Theorem 2. *Suppose Assumptions 1-3 hold with $\mathcal{K} = \{\kappa, \kappa'\}$ and $f_{t,\kappa} \equiv f_\kappa$ (constant time-of-day periodicity) for $t \in \mathbb{N}_+$. Let $\varpi \in (0, \frac{5}{12}]$. Then, for any $\kappa, \kappa' \in (0, 1]$, as $n \rightarrow \infty$ and $T \rightarrow \infty$ with $T \Delta_n \rightarrow 0$, we have,*

$$\sqrt{T} \left(\widehat{L}_\kappa^n - \widehat{L}_{\kappa'}^n \right) \xrightarrow{\mathcal{L}} N(0, K), \quad (12)$$

where K is a covariance operator with integral representation,

$$Kh(z) = \int_{\mathbb{R}^+} k(z, u) h(u) w(u) du, \quad \forall h \in \mathcal{L}^2(w), \quad (13)$$

with kernel $k(z, u) = \sum_{j=-\infty}^{\infty} \mathbb{E}[d_1(z)d_{j+1}(u)]$ and,

$$\begin{aligned} d_t(u) = & \cos\left(\sqrt{2u\sigma_{t-1+\kappa}^2}Z_t\right) - \cos\left(\sqrt{2u\sigma_{t-1+\kappa'}^2}Z'_t\right) \\ & - u\mathcal{L}'(u)\frac{\pi}{2}\left(\sigma_{t-1+\kappa}^2|Z_t||\tilde{Z}_t| - \sigma_{t-1+\kappa'}^2|Z'_t||\tilde{Z}'_t|\right), \end{aligned} \quad (14)$$

for $\{Z_t\}$, $\{\tilde{Z}_t\}$, $\{Z'_t\}$ and $\{\tilde{Z}'_t\}$ being sequences of independent standard normal random variables defined on an extension of the original probability space and independent of $\mathcal{F}$.

The limit result has some notable features. First, the rate of convergence is controlled by the time span of the data. The limit result requires $T\Delta_n \rightarrow 0$, that is, we sample slightly faster than we increase the time span of the data (note that this is strictly a rate condition, i.e., it is not scale free but depends on whether time is measured in years or days, so it is not meaningful to assess it through the size of this product). This is a standard condition in joint in-fill and long-span asymptotic settings. It ensures that certain biases due to the time variation of volatility over the discrete observation intervals are negligible. These biases are very small in empirically relevant situations as we document in the Monte Carlo study. We further note that, if we were to replace $\hat{f}_\kappa$ with the truncated volatility estimator $\frac{n}{T} \sum_{t=1}^T |\Delta_{t,\kappa}^n X|^2 1_{\{|\Delta_{t,\kappa}^n X| \leq v_n\}}$, then, under somewhat stronger moment conditions than in Assumption 1, a result corresponding to Theorem 2 can be shown to hold for the modified statistic, under the weaker condition $T\Delta_n^{2-\iota} \rightarrow 0$, for some arbitrarily small $\iota > 0$ (the stronger moment conditions are needed for showing the analogue of Lemma 1 in the Supplementary Appendix for the truncated volatility).

Second, $\hat{L}_\kappa^n - \hat{L}_{\kappa'}^n$ is based on the difference of functions of increments over different parts of the day. Consequently, the asymptotic covariance operator K depends only on the autocovariance of the differential between transforms of volatility at different times-of-day. As a result, the persistence in $d_t(u)$ is typically small, even if σ_t^2 contains a very

persistent stationary component. To illustrate, suppose σ_s^2 is constant during the day, i.e., for $s \in [t-1, t]$ and $t \in \mathbb{N}_+$. Then we have $\mathbb{E}(d_1(z)d_{j+1}(u)) = 0$ for $j \neq 0$. The implication is that, even if volatility is highly persistent (which is true empirically), we do not require a large time span for reliable recovery of $\mathcal{L}_\kappa - \mathcal{L}_{\kappa'}$. This is unlike the situation, where one seeks to recover $\mathcal{L}_\kappa$ and $\mathcal{L}_{\kappa'}$ separately, as the precision of those estimates will be compromised by strong volatility persistence.

Third, the asymptotic limit in Theorem 2 reflects two sources of error. The first is associated with uncovering the latent stochastic variance σ_t^2 from high-frequency data. The second is the empirical process error capturing the deviation of sample averages for transforms of volatility from their unconditional means. This is unlike most existing joint in-fill and long span asymptotic limit results, in which the error from recovering the latent volatility is asymptotically negligible. The reason is that here, unlike in previous work, we do not integrate functions of the high-frequency data over the full trading day, but rather rely on only a fixed number of high-frequency increments each day. The main error in measuring volatility from high-frequency returns of X stems from the increments of the Brownian motion over the small sampling intervals. We allow for these increments to be correlated with the innovations of σ^2 , that is, the so-called leverage effect is accommodated. Nevertheless, since the length of the high-frequency intervals shrinks asymptotically, this dependence has an asymptotically negligible effect on the limit result in Theorem 2. Hence, for the purposes of the CLT of $\hat{L}_\kappa^n - \hat{L}_{\kappa'}^n$, our asymptotic setting becomes equivalent to conducting inference from observations of $\left(Z_t \sqrt{f_\kappa \sigma_{t-1+\kappa}^2}, \tilde{Z}_t \sqrt{f_\kappa \sigma_{t-1+\kappa}^2}\right)_{t \in \mathbb{N}_+}$ and $\left(Z'_t \sqrt{f_{\kappa'} \sigma_{t-1+\kappa'}^2}, \tilde{Z}'_t \sqrt{f_{\kappa'} \sigma_{t-1+\kappa'}^2}\right)_{t \in \mathbb{N}_+}$, where $\{Z_t\}$, $\{\tilde{Z}_t\}$, $\{Z'_t\}$ and $\{\tilde{Z}'_t\}$ are i.i.d. sequences of standard normals independent from the volatility process. This situation mirrors some features of the CLT for measuring quantities associated with the jump part of X , such as their quadratic variation, see, e.g., [31]. In that case, only the increments of the Brownian

motion over the intervals containing the jumps drive the asymptotics. In our case, because we study time-of-day volatility patterns, we similarly rely on only a finite number of high-frequency increments per day. Unlike the high-frequency analysis of jumps, however, we also have $T \rightarrow \infty$ and, consequently, we have an additional source of error driving the CLT, namely the empirical process error associated with the recovery of unconditional moments of volatility from the corresponding sample averages.

Finally, we note that, under the conditions in Theorem 2, but with f_κ possibly time-varying (given Assumptions 1-3 hold), we can show $\|\hat{f}_\kappa - f_\kappa\| = O_p(1/\sqrt{T})$. We are not going to make use of this result henceforth, however.

Given Theorems 1 and 2, our test statistic is quite intuitive. It is given by the weighted squared difference of the estimates for the volatility Laplace transforms over the two distinct periods across the trading day,

$$TS_{n,T}(\kappa, \kappa') = T\|\hat{L}_\kappa^n - \hat{L}_{\kappa'}^n\|^2 \equiv T \int_{\mathbb{R}_+} \left(\hat{L}_\kappa^n(u) - \hat{L}_{\kappa'}^n(u) \right)^2 w(u) du, \quad \kappa, \kappa' \in (0, 1]. \quad (15)$$

The asymptotic behavior of $TS_{n,T}(\kappa, \kappa')$ under the null hypothesis follows directly from the CLT in Theorem 2. It is stated formally in Corollary 1.

Corollary 1. *Under the conditions of Theorem 2, we have,*

$$TS_{n,T}(\kappa, \kappa') \xrightarrow{\mathcal{L}} Z(\kappa, \kappa'),$$

where $Z(\kappa, \kappa')$ is a weighted sum of independent chi-squared distributions with one degree of freedom, defined on an extension of the original probability space and independent from $\mathcal{F}$. The weights are given by the eigenvalues of the covariance operator K , defined in Theorem 2.

When the alternative hypothesis is true, i.e., when the time-of-day periodic component of volatility varies over time, then $\mathcal{L}_\kappa$ and $\mathcal{L}_{\kappa'}$ differ and, from Theorem 1, we conclude that $TS_{n,T}(\kappa, \kappa')$ diverges to infinity.

3.2 Feasible Inference and Construction of the Test

The feasible version of our test statistic will be based on the limit results in Theorem 2 and Corollary 1. For implementation, we need to obtain an estimate of the covariance operator K from the data. To this end, we first construct the feasible counterpart of $d_t(u)$ given by $\widehat{d}_{t,n}(u) = \widehat{d}_{t,n}^\kappa(u) - \widehat{d}_{t,n}^{\kappa'}(u)$ with,

$$\widehat{d}_{t,n}^\kappa(u) = \cos\left(\sqrt{2un}\Delta_{t,\kappa}^n X / \sqrt{\widehat{f}_\kappa}\right) + (|u\widehat{\mathcal{L}}'_\kappa(u)| \wedge e^{-0.5}) \frac{\pi}{2} \frac{n|\Delta_{t,\kappa}^n X| |\Delta_{t,\kappa}^n X - \Delta_n X|}{\widehat{f}_\kappa} 1_{\{\mathcal{A}_{t,\kappa}^n\}}, \quad (16)$$

for $u \in \mathbb{R}_+$ and $\kappa \in (0, 1]$ and with,

$$\widehat{\mathcal{L}}'_\kappa(u) = -\frac{1}{T} \sum_{t=1}^T \sin\left(\sqrt{2un}\Delta_{t,\kappa}^n X / \sqrt{\widehat{f}_\kappa}\right) \frac{\sqrt{n}\Delta_{t,\kappa}^n X}{\sqrt{2u\widehat{f}_\kappa}} 1_{\{|\Delta_{t,\kappa}^n X| \leq v_n\}}. \quad (17)$$

In defining $\widehat{d}_{t,n}^\kappa(u)$, we impose a small sample correction by using $|u\widehat{\mathcal{L}}'_\kappa(u)| \wedge e^{-0.5}$ instead of $u\widehat{\mathcal{L}}'_\kappa(u)$. This is because we have $\sup_{u \in \mathbb{R}} |u\mathcal{L}'_\kappa(u)| \leq e^{-1}$, so it follows that the above correction has no asymptotic effect (the same will hold if we replace $e^{-0.5}$ in the truncation of $|u\widehat{\mathcal{L}}'_\kappa(u)|$ with any number above e^{-1}).

Given $\widehat{d}_{t,n}^\kappa(u)$, the feasible kernel-type estimator of the covariance operator is,

$$K_T f(s) = \int_{\mathbb{R}^+} k_T(s, u) f(u) w(u) du, \quad (18)$$

where k_T is given by,

$$k_T(u, s) = \frac{1}{T} \sum_{t=1}^T \widehat{d}_{t,n}(u) \Gamma \widehat{d}_{t,n}(s), \quad (19)$$

and Γ is the linear operator defined by,

$$\Gamma \widehat{d}_{t,n}(s) = h(0) \widehat{d}_{t,n}(s) + \sum_{j=1}^T h\left(\frac{j}{B_T}\right) \left(\widehat{d}_{t-j,n}(s) + \widehat{d}_{t+j,n}(s)\right), \quad (20)$$

with the convention that $\widehat{d}_{t,n}(s) = 0$ if $t \leq 0$ or $t > T$, and h is a kernel which satisfies the following regularity condition.

Assumption 4. *The kernel function h used for constructing K_T satisfies the following conditions: $h : \mathbb{R} \rightarrow [-1, 1]$, $h(0) = 1$, $h(x) = h(-x)$, h is continuous and is further continuously differentiable in a neighborhood of zero with the potential exception at zero where $h'_\pm(0)$ exist and are bounded.*

To conduct feasible inference for $TS_{n,T}(\kappa, \kappa')$, we require estimates for the eigenvalues of the operator K . Since K is a Hilbert-Schmidt operator, it follows that the eigenvalues of K converge to zero. The limiting distribution of $TS_{n,T}(\kappa, \kappa')$ depends on all the eigenvalues of the covariance operator K . However, since the eigenvalues converge to zero, it is natural to approximate the distribution of $Z(\kappa, \kappa')$ through estimates for only the p_T largest ones, where p_T is a sequence of positive integers that asymptotically diverge to infinity. This is what we do below.

Our estimates for the eigenvalues of K will be based on its estimate K_T . By construction, k_T is a degenerate kernel and, thus, it has only finitely many non-zero eigenvalues. Furthermore, the range of K_T is spanned by $\hat{d}_{1,n}(u), \dots, \hat{d}_{T,n}(u)$, so the eigenfunctions are of the form $\hat{\psi}_j(u) = \frac{1}{T} \sum_{t=1}^T \beta_{j,t} \hat{d}_{t,n}(u)$, for a sequence of coefficients $\{\beta_{j,t}\}$ and $j = 1, \dots, T$. The estimated eigenvalues are then obtained by solving the following equation,

$$K_T \hat{\psi}_j(u) = \hat{\lambda}_j \hat{\psi}_j(u), \quad j = 1, \dots, T. \quad (21)$$

Solving for these eigenvalues is equivalent to finding the eigenvalues of the matrix C , whose (i, j) 'th element equals $c_{ij} = \frac{1}{T} \int_{\mathbb{R}_+} \hat{d}_{j,n}(u) \Gamma \hat{d}_{i,n}(u) w(u) du$, for $i, j = 1, \dots, T$. Then the eigenvalues of C , denoted $\hat{\lambda}_1, \dots, \hat{\lambda}_T$, are natural estimators of $\lambda_1, \dots, \lambda_T$. Based on these estimated eigenvalues, we construct the following approximation of the limiting distribution in Corollary 1,

$$\hat{Z}_T(\kappa, \kappa') = \sum_{i=1}^{p_T} \hat{\lambda}_i \chi_i^2, \quad (22)$$

where $\{\chi_i^2\}_{i \geq 1}$ denotes the sequence of $\chi^2(1)$ distributed random variables from Corollary 1. The following theorem shows that the limiting distribution of our test statistic can be approximated by $\widehat{Z}_T(\kappa, \kappa')$.

Theorem 3. *Suppose Assumptions 1-4 hold with $\mathcal{K} = \{\kappa, \kappa'\}$ and $f_{t,\kappa} \equiv f_\kappa$ (constant time-of-day periodicity) for $t \in \mathbb{N}_+$. Let $\varpi \in [\frac{1}{4}, \frac{3}{8}]$ and $n \rightarrow \infty$, $T \rightarrow \infty$ with $T \Delta_n \rightarrow 0$. Suppose $B_T \rightarrow \infty$ and $p_T \rightarrow \infty$ such that,*

$$\frac{B_T^2}{T} \rightarrow 0 \quad \text{and} \quad p_T \left(B_T^{-6} \bigvee \frac{B_T^2}{T} \right) \rightarrow 0. \quad (23)$$

We then have,

$$\widehat{Z}_T(\kappa, \kappa') - Z(\kappa, \kappa') \xrightarrow{\mathbb{P}} 0. \quad (24)$$

The sequence B_T controls the number of lags of $\widehat{d}_{t,n}(u)$ we use in the construction of our estimator for the covariance operator K_T . The choice of B_T naturally depends on the persistence of the underlying series. In standard time series applications, see, e.g., [6], it typically takes values like $B_T = O(T^{1/3})$ or $B_T = O(T^{1/5})$. Given our earlier discussion, $d_t(u)$ will typically display limited persistence, and we can therefore reliably estimate K with only a relatively small number of lags included in the construction of K_T .

The second condition in equation (23) puts an upper bound on the rate of growth of p_T which, we recall, controls the number of the largest eigenvalues of K_T included in constructing $\widehat{Z}_T(\kappa, \kappa')$. We determine the upper bound on p_T using the connection between the error in recovering the eigenvalues of K and the one associated with the estimation of the covariance operator (measured in Hilbert-Schmidt norm). This error, in turn, stems from the sampling error in inferring K as well as the bias due to using only B_T autocovariances of $\widehat{d}_{t,n}(u)$ (and their smoothing with the kernel h) in the construction of K_T . As we later document, the eigenvalues of K typically die out very fast and, hence, the test has very limited sensitivity with respect to the choice of p_T .

With the feasible approximation $\widehat{Z}_T(\kappa, \kappa')$ of $Z(\kappa, \kappa')$, we are now ready to formally define our test. For some $\kappa, \kappa' \in (0, 1]$ with $\kappa \neq \kappa'$, the null and alternative hypotheses are given by,

$$H_0 : \{\mathcal{L}_\kappa = \mathcal{L}_{\kappa'}\} \quad \text{and} \quad H_A : \{\mathcal{L}_\kappa \neq \mathcal{L}_{\kappa'}\}, \quad (25)$$

where the equality and inequality are to be understood in the $\mathcal{L}^2(w)$ sense. Define next,

$$cv_{n,T}^\alpha(\kappa, \kappa') = Q_{1-\alpha}(\widehat{Z}_T(\kappa, \kappa')|\mathcal{F}), \quad (26)$$

where $Q_\alpha(Z)$ denotes the α -quantile of a generic random variable Z . $cv_{n,T}^\alpha(\kappa, \kappa')$ is computed numerically using the estimated eigenvalues $\{\widehat{\lambda}_i\}_{i=1,\dots,p_T}$ and the simulation of a sequence of i.i.d. $\chi^2(1)$ distributed random variables. We then have the following result.

Corollary 2. *Suppose Assumptions 1-4 hold with $\mathcal{K} = \{\kappa, \kappa'\}$ and the sequences B_T and p_T satisfy condition (23). The test defined by the critical region $\{TS_{n,T}(\kappa, \kappa') > cv_{n,T}^\alpha(\kappa, \kappa')\}$ has asymptotic size α under the null and asymptotic power one under the alternative, i.e.,*

$$\mathbb{P}(TS_{n,T}(\kappa, \kappa') > cv_{n,T}^\alpha(\kappa, \kappa')|H_0) \longrightarrow \alpha, \quad \mathbb{P}(TS_{n,T}(\kappa, \kappa') > cv_{n,T}^\alpha(\kappa, \kappa')|H_A) \longrightarrow 1. \quad (27)$$

3.3 Extensions

3.3.1 Averaging Multiple Time-of-Day Intervals

One natural extension is to compare the average Laplace transforms of volatility over two distinct sets of time-of-day intervals. This has the benefit of reducing the measurement error, and hence increases the power of the test. Of course, the averaging ignores the potential differences in the Laplace transforms that we average. Therefore, this procedure is most advantageous for intervals in which the periodic volatility component is similar, even if it is time-varying. This is naturally satisfied for adjacent intervals during the trading day,

e.g., neighboring five-minute intervals within one hour. In fact, given the smoothness of the periodic component in Assumption 2, the distance between the Laplace transforms of the deseasonalized volatilities over high-frequency intervals within neighborhoods of asymptotic size of order $o(1/\sqrt{T})$ of the time-of-day κ and κ' will be $o(1/\sqrt{T})$. Therefore, the averaging of the Laplace transforms over these blocks of high-frequency data will continue to provide a valid test for the null hypothesis of equality between $\mathcal{L}_\kappa$ and $\mathcal{L}_{\kappa'}$, but with more power relative to the original test in Corollary 2.

We formalize this extension of our test without a formal proof, as it follows straightforwardly from our earlier results. For simplicity, we restrict attention to the case where the number of high-frequency intervals, over which averaging is performed, remains fixed, that is, it does not increase with the sampling frequency. The location of the elements in the two sets on $(0, 1]$ may change with the sampling frequency (but deviates from fixed points on $(0, 1]$ by terms which are $o(1/\sqrt{T})$). Denote two disjoint finite sets of numbers in $(0, 1]$ by $\mathcal{K}_n$ and $\mathcal{K}'_n$. The typical example of such a set, $\mathcal{K}_n$, takes the form $\mathcal{K}_n = \left\{ \frac{\lfloor \kappa n \rfloor}{n}, \frac{\lfloor \kappa n \rfloor + 1}{n}, \dots, \frac{\lfloor \kappa n \rfloor + kn}{n} \right\}$, for some fixed integer $k \geq 1$. This corresponds to using several high-frequency increments located in the vicinity of κ during the trading day. We then define,

$$\widehat{L}_{\mathcal{K}}^n(u) = \frac{1}{|\mathcal{K}_n|} \sum_{\kappa \in \mathcal{K}_n} \widehat{L}_{\kappa}^n(u), \quad (28)$$

with $|\mathcal{K}_n|$ denoting the cardinality of the set $\mathcal{K}_n$. The test statistic is now generalized to,

$$T S_{n,T}(\mathcal{K}_n, \mathcal{K}'_n) = \|\widehat{L}_{\mathcal{K}}^n - \widehat{L}_{\mathcal{K}'}^n\|^2. \quad (29)$$

We define the counterpart of $d_{t,n}^\kappa(u)$ by,

$$d_{t,n}^{\mathcal{K}}(u) = \frac{1}{|\mathcal{K}_n|} \sum_{\kappa \in \mathcal{K}_n} d_{t,n}^\kappa(u). \quad (30)$$

The extension of the test is then based on the critical region $\{TS_{n,T}(\mathcal{K}_n, \mathcal{K}'_n) > cv_{n,T}^\alpha(\mathcal{K}_n, \mathcal{K}'_n)\}$, where $cv_{n,T}^\alpha(\mathcal{K}_n, \mathcal{K}'_n) = Q_{1-\alpha}(\widehat{Z}_T(\mathcal{K}_n, \mathcal{K}'_n)|\mathcal{F})$ with $\widehat{Z}_T(\mathcal{K}_n, \mathcal{K}'_n)$ constructed from $d_{t,n}^\mathcal{K}(u)$ and $d_{t,n}^{\mathcal{K}'}(u)$, exactly as $\widehat{Z}_T(\kappa, \kappa')$ is constructed from $d_{t,n}^\kappa(u)$ and $d_{t,n}^{\kappa'}(u)$.

3.3.2 Incorporating Additional Information

We can further extend the analysis by considering conditioning information,

$$\widehat{L}_{\kappa, \mathcal{B}}^n(u) = \frac{1}{T} \sum_{t=1}^T 1_{\{\mathcal{B}_{t-1}\}} \cos \left(\sqrt{2un} \Delta_{t,\kappa}^n X / \sqrt{\widehat{f}_{\kappa, \mathcal{B}}} \right), \quad (31)$$

for

$$\widehat{f}_{\kappa, \mathcal{B}} = \frac{n}{T} \frac{\pi}{2} \sum_{t=1}^T |\Delta_{t,\kappa}^n X| |\Delta_{t,\kappa-\Delta}^n X| 1_{\{\mathcal{A}_{t,\kappa}^n \cap \mathcal{B}_{t-1}\}}, \quad (32)$$

where $\{\mathcal{B}_t\}_{t \in \mathbb{N}_+}$ is a sequence of $\mathcal{F}_t$ -adapted random sets. Provided appropriate ergodicity and mixing conditions hold, $\widehat{L}_{\kappa, \mathcal{B}}^n(u)$ converges in probability to,

$$\mathcal{L}_{\kappa, \mathcal{B}} = \mathbb{E} \left[e^{-uf_{t,\kappa} \sigma_{t+\kappa}^2 / \mathbb{E}[f_{t,\kappa} \sigma_{t+\kappa}^2 1_{\{\mathcal{B}_{t-1}\}}]} 1_{\mathcal{B}_{t-1}} \right], \quad \text{for } t \in \mathbb{N}_+, \quad (33)$$

and the CLT of Theorem 2 continues to apply with $d_t(u)$ replaced by,

$$\begin{aligned} d_t^\mathcal{B}(u) &= 1_{\{\mathcal{B}_{t-1}\}} \left[\cos \left(\sqrt{2u\sigma_{t-1+\kappa}^2} Z_t \right) - \cos \left(\sqrt{2u\sigma_{t-1+\kappa'}^2} Z'_t \right) \right] \\ &\quad - 1_{\{\mathcal{B}_{t-1}\}} u \mathcal{L}'(u) \frac{\pi}{2} \left(\sigma_{t-1+\kappa}^2 |Z_t| |\widetilde{Z}_t| - \sigma_{t-1+\kappa'}^2 |Z'_t| |\widetilde{Z}'_t| \right). \end{aligned} \quad (34)$$

We can similarly define $TS_{n,T}^\mathcal{B}(\kappa, \kappa')$ and $\widehat{d}_{t,n}^\mathcal{B}(u)$ from $TS_{n,T}(\kappa, \kappa')$ and $\widehat{d}_{t,n}(u)$, and then conduct tests on the basis of $TS_{n,T}^\mathcal{B}(\kappa, \kappa')$ and critical regions constructed exactly as in Corollary 2. We omit formal proofs of these extensions, as they follow directly from our results in Sections 3.1 and 3.2.

The above generalization may be used to estimate Laplace transforms conditional on specific events, e.g., level of volatility or the occurrence of a prescheduled announcement.

This enables us to investigate potential sources of variation in the periodic component of volatility.

3.3.3 Accounting for Microstructure Noise

The final extension of our theoretical analysis we consider here is to allow for presence of microstructure noise in the observations. This extension is of particular importance if one is to apply the above analysis to very high frequencies, e.g., when sampling every second. A common approach to dealing with noise is to do local averaging of the prices, and use the pre-averaged prices in the construction of the statistics, see e.g., [33]. The presence of noise and the pre-averaging slow down the rate of recovery of volatility from high-frequency data. On the other hand, [13], [16, 17] and [35] consider the situation where the microstructure noise is a (known) function of observable variables and a finite-dimensional parameter. In this parametric setting for the noise, volatility can still be recovered at the fast no-noise rate. We follow this approach below in our treatment of the noise. Specifically, we assume that, instead of observing directly X , we observe,

$$Y_{\frac{i}{n}} = X_{\frac{i}{n}} + g\left(Z_{\frac{i}{n}}, \theta\right), \quad i = 1, \dots, nT, \quad (35)$$

where Z_t is $p \times 1$ vector of observable variables, θ is a $k \times 1$ vector of parameters which are not known, with the true value denoted by θ_0 , and $g : \mathbb{R}^p \times \mathbb{R}^k \rightarrow \mathbb{R}$ is a known function. We will further assume that we have an estimator $\hat{\theta}$ from the data of θ_0 for which

$$\hat{\theta} - \theta_0 = O_p(1/n). \quad (36)$$

Examples of estimators $\hat{\theta}$ satisfying this condition are given in [16, 17] and [35]. We note that to satisfy equation (36), one must impose assumptions on the process Z and the function g . Finally, for our analysis we assume that the function g satisfies the following

Lipschitz type condition,

$$\|g(Z, \theta_2) - g(Z, \theta_1)\| \leq C\|\theta_2 - \theta_1\|, \quad \forall Z \in \mathbb{R}^p, \quad \theta_1, \theta_2 \in \mathbb{R}^k, \quad (37)$$

for some finite constant $C > 0$. Given the estimator $\hat{\theta}$, we can recover the true price via

$$\hat{X}_{\frac{i}{n}} = Y_{\frac{i}{n}} - g\left(Z_{\frac{i}{n}}, \hat{\theta}\right), \quad i = 1, \dots, nT. \quad (38)$$

Using $\hat{X}_{\frac{i}{n}}$, we can construct the counterparts of $\hat{L}_{\kappa}^n(u)$, $\hat{f}_{\kappa}$ and $\mathcal{A}_{t,\kappa}^n$. We denote these quantities by $\hat{L}_{\kappa}^{r,n}(u)$, $\hat{f}_{\kappa}^r$ and $\mathcal{A}_{t,\kappa}^{r,n}$, respectively. We then have the following result,

Theorem 4. *Suppose the observations are given by equation (35), and that the conditions (36)-(37) apply.*

(a) *Under the conditions of Theorem 1, we have*

$$\hat{L}_{\kappa}^{r,n} \xrightarrow{\mathbb{P}} \mathcal{L}_{\kappa}. \quad (39)$$

(b) *Under the conditions of Theorem 2, we have*

$$\sqrt{T} \left(\hat{L}_{\kappa}^{r,n} - \hat{L}_{\kappa'}^{r,n} \right) \xrightarrow{\mathcal{L}} N(0, K), \quad (40)$$

where K is the covariance integral operator defined in Theorem 2.

4 Simulation Study

In this section, we assess the finite sample properties of the proposed test through a Monte Carlo study. We rely on the following two-factor affine jump-diffusion with an intraday periodic volatility component,

$$\begin{aligned} X_t &= X_0 + \int_0^t \sqrt{\tilde{V}_s} dW_s + \sum_{s=1}^{N_t} Z_s, \quad \tilde{V}_t = f_{[t], t-[t]} V_t, \quad V_t = \left(V_t^{(1)} + V_t^{(2)} \right), \\ dV_t^{(i)} &= \kappa_i (\theta - V_t^{(i)}) dt + \xi_i \sqrt{V_t^{(i)}} dB_t^{(i)}, \quad i = 1, 2, \end{aligned} \quad (41)$$

where W , $B^{(1)}$ and $B^{(2)}$ are independent standard Brownian motions, N_t is a Poisson process with intensity λ_J , and $\{Z_s\}_{s \geq 1}$ is an i.i.d. sequence of normally distributed random variables with mean zero and variance σ_j^2 . This representation captures the main features of the U.S. equity market index. In accordance with [11], we fix the model parameters as follows, $(\kappa_1, \kappa_2, \theta, \xi_1, \xi_2, \lambda_J, \sigma_j^2) = (0.0128, 0.6930, 0.4068, 0.0954, 0.7023, 0.2, 0.19, 0.932)$.

To explore the size of the test under the null hypothesis, $f_{t,\kappa} \equiv f_\kappa$ (constant time-of-day periodicity), we set f_κ equal to the average time-of-day effect obtained in our empirical application, displayed in Figure 1. Under the alternative hypothesis, $f_{t,\kappa}$ varies with t . Consistent with our empirical findings in Section 5, we let $f_{t,\kappa}$ be a function of the stationary component of volatility, V_t , for investigating the power of the test. Thus, we stipulate,

$$f_{t,\kappa} = \begin{cases} f_\kappa^l, & \text{if } V_t \leq Q_{0.25}(V_t), \\ f_\kappa^m, & \text{if } V_t \in (Q_{0.25}(V_t), Q_{0.75}(V_t)), \\ f_\kappa^h, & \text{if } V_t \geq Q_{0.75}(V_t), \end{cases} \quad t \in \mathbb{N}_+, \quad (42)$$

where f_κ^l , f_κ^m , and f_κ^h equal our empirical estimates for the time-of-day periodic volatility component after conditioning on whether, at the start of the trading day, the VIX volatility index – an option-based indicator of future volatility – belongs to the respective empirical VIX quantile, namely $(0, Q_{0.25}(VIX)]$, $(Q_{0.25}(VIX), Q_{0.75}(VIX))$ and $[Q_{0.75}(VIX), \infty)$. Specifically, we compute $n\hat{f}_{i\Delta_n, \mathcal{B}} / \sum_{i=1}^n \hat{f}_{i\Delta_n, \mathcal{B}}$ on the real data for each of the three regions $\mathcal{B}$ above, and then apply a Nadaraya-Watson kernel regression with a Gaussian kernel and bandwidth corresponding to a five-minute window, to obtain estimates for the standardized periodic component of volatility conditional on the value of the VIX index. The resulting periodic volatility components are displayed on Figure 2.

In the Monte Carlo we set $n = 77$, corresponding to sampling every five minutes across a 6.5 hour trading day and discarding the first 5-minute interval. Given the imprecision associated with evaluation of our test statistic for high values of u , we truncate the integral

Figure 2: Periodic Volatility Components used in the Monte Carlo. The dashed line corresponds to f_{κ}^l , the dotted line to f_{κ}^m , and the solid line to f_{κ}^h .

in equation (15) at u_{max} , which is set to satisfy $\frac{1}{n} \sum_{i=1}^n \widehat{L}_{i\Delta_n}(u_{max}) = 0.05$. The weight function, w , corresponds to the density of a normal distribution with mean zero and variance such that $\int_{-\infty}^{u_{max}} w(u)du = 0.995$. Next, we use the following data-driven truncation for the jumps, $v_n = 3 \sqrt{BV_{[t]-1} \wedge RV_{[t]-1}} \Delta_n^{3/8}$, where BV and RV are the bipower variation (see, e.g., [10]) and realized volatility estimators defined as,

$$BV_t = \frac{\pi}{2} \sum_{i=2}^n |\Delta_{t,i-1}^n X| |\Delta_{t,i}^n X|, \quad RV_t = \sum_{i=1}^n |\Delta_{t,i}^n X|^2, \quad t \in \mathbb{N}_+. \quad (43)$$

RV_t is a measure of the total quadratic variation of X over $[t-1, t]$, while BV_t is a jump-robust counterpart, estimating the diffusive component of the return variation, $\int_{t-1}^t \widetilde{\sigma}_s^2 ds$. For estimation of the covariance operator, we set $B_T = \lfloor T^{1/5} \rfloor$ and we use the Bartlett kernel for h . In the Supplementary Appendix, we show that the testing procedure is robust (in terms of its size properties) to alternative choices of u_{max} and B_T .

Finally, the critical values of the test are calculated on the basis of 10,000,000 simula-

tions for the $\chi^2(1)$ random variables appearing in $\hat{Z}_T(\kappa, \kappa')$. Exactly as in the empirical application, we perform the test over intervals of 30 minutes, i.e., $\mathcal{K}$ in equation (28) equals the fraction of the trading day represented by half an hour.

The Monte Carlo results under the null hypothesis are given in Table 1. We notice the marginal sensitivity with respect to the number of eigenvalues included in the computation of the critical values beyond $p_T = 2$. Similarly, the performance of the test is remarkably similar for different sample sizes, T , and empirical rejection rates are very close to the nominal level of the test.

T	p_T					
	1	2	3	4	5	6
250	0.054	0.046	0.045	0.044	0.044	0.044
500	0.067	0.057	0.055	0.055	0.055	0.055
1000	0.055	0.050	0.049	0.048	0.047	0.047
1500	0.059	0.053	0.048	0.047	0.047	0.047
2000	0.055	0.048	0.047	0.047	0.047	0.047
2500	0.069	0.062	0.061	0.060	0.060	0.060

Table 1: Monte Carlo Results under the Null Hypothesis, $f_{t,\kappa} \equiv f_\kappa$. The table reports empirical rejection rates of the test of nominal size 0.05 using 1,000 simulations. $\mathcal{K}_n$ and $\mathcal{K}'_n$ correspond to 8:40-9:10 and 12:30-13:00, respectively.

Turning to the power of the test, we provide simulation results for the alternative hypothesis in Table 2. We note that the power of the test depends on which time intervals are compared. This is not surprising given that the time variation in the periodic volatility component differs substantially across time-of-day, as depicted in Figure 2. The largest

discrepancies in the periodic component across volatility regimes occur towards the end of the trading day, and our test picks this up, even for moderate sample sizes. The test struggles more with identifying time variation in the periodic component of volatility in the morning versus the middle of the day, because the marginal distribution of volatility is less distinct across those intraday intervals. Furthermore, while power declines slightly as the number of eigenvalues included in calculating the critical values increases, the discrepancies in power, looking beyond the second eigenvalue, are small. Finally, as expected, the power increases as the sample size grows.

T	8:40 - 9:10 vs 12:30- 13:00						8:40 - 9:10 vs 14:30 - 15:00					
	p_T						p_T					
	1	2	3	4	5	6	1	2	3	4	5	6
250	0.093	0.075	0.064	0.062	0.061	0.061	0.322	0.292	0.278	0.269	0.267	0.266
500	0.139	0.100	0.086	0.076	0.075	0.074	0.623	0.588	0.571	0.562	0.555	0.554
1000	0.168	0.126	0.106	0.099	0.096	0.094	0.918	0.903	0.901	0.893	0.890	0.887
1500	0.229	0.173	0.149	0.140	0.135	0.133	0.993	0.990	0.989	0.989	0.989	0.988
2000	0.339	0.285	0.266	0.251	0.240	0.232	1.000	0.998	0.998	0.998	0.998	0.998
2500	0.410	0.347	0.308	0.292	0.286	0.279	0.999	0.999	0.999	0.999	0.999	0.999

Table 2: Monte Carlo Results under the Alternative Hypothesis, $f_{t,\kappa} \neq f_\kappa$. The table reports empirical rejection rates for the test at nominal size 0.05 using 1000 simulations.

5 Empirical Application

Our empirical analysis is based on high-frequency data for the E-mini S&P 500 futures contract, spanning the period January 1, 2005, till January 30, 2015. After removing

partial trading days from the sample, we end up with a total of 2,516 days. Each day, we sample every five minutes over the period 8.35-15.00 CST, which generates 77 returns per day. For part of the analysis, we also make use of the VIX volatility index, recorded at the start of each trading day.

In the implementation of the test, the truncation, the weight function and the tuning parameters of K_T are set exactly as in the Monte Carlo study. In addition, as in the simulation study, we implement the test over half hour intervals with the exception of the first period, which spans an interval of 20 minutes (8:40-9:00 CST). Table 3 reveals that the null hypothesis is rejected, except when the test involves intervals which are very close within the trading day. The failure of the test to reject for adjacent periods is not mechanical, as the respective estimates for the empirical Laplace transforms, $\widehat{L}_{\mathcal{K}}^n$ and $\widehat{L}_{\mathcal{K}'}^n$, have only a minor overlap in terms of the underlying high-frequency data (mainly a five-minute interval which is due to the staggering of returns in the computation of $\widehat{f}_{\kappa}$ for $\kappa \in \mathcal{K}_n$ and $\kappa \in \mathcal{K}'_n$). Instead, this empirical finding is a manifestation of the fact that, although there is a time-variation in the periodic component of volatility, it is quite similar for adjacent intervals. Overall, these results provide strong evidence that the intraday periodicity in volatility is time varying.

One possible explanation for the overwhelming rejection of the null hypothesis is that the intraday volatility pattern is different for days with scheduled release of macroeconomic news. There are numerous such announcements during the trading hours. We focus on the release of news from the Federal Open Market Committee (FOMC), which are regularly scheduled for 1pm CST every six weeks. Other noteworthy announcements during trading hours include the ISM Manufacturing and Non-Manufacturing Indices as well as the Consumer Sentiment report. Unreported results show that these releases have a much smaller impact on the intraday volatility pattern than the FOMC announcement. Hence,

	9:00	9:30	10:00	10:30	11:00	11:30	12:00	12:30	13:00	13:30	14:00	14:30
8:40 - 9:00	0.149	0.417	0.428	0.129	0.004	0.004	0.004	0.000	0.000	0.000	0.000	0.000
9:00 - 9:30		0.386	0.061	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
9:30 - 10:00			0.148	0.006	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
10:00 - 10:30				0.483	0.007	0.000	0.003	0.000	0.000	0.000	0.000	0.000
10:30 - 11:00					0.074	0.018	0.018	0.000	0.001	0.000	0.000	0.000
11:00 - 11:30						0.279	0.504	0.032	0.035	0.001	0.000	0.000
11:30 - 12:00							0.687	0.253	0.274	0.043	0.000	0.000
12:00 - 12:30								0.084	0.226	0.009	0.000	0.000
12:30 - 13:00									0.417	0.315	0.000	0.003
13:00 - 13:30										0.070	0.000	0.000
13:30 - 14:00											0.023	0.119
14:00 - 14:30												0.498

Table 3: Unconditional Test Results. The table reports the results (p-values) from the unconditional test of Section 3.2 over the period 1 January, 2005 to 30 January, 2015. Critical values are computed using $p_T = 3$. Top row indicates the beginning of each half-hour interval.

for brevity, we only analyze the latter here. In total, we have 96 FOMC announcements in our sample, and we label these “FOMC days.” Figure 3 depicts estimates for the periodic volatility component on FOMC and non-FOMC days. The estimates for the periodic component on non-FOMC days are almost identical to those for the full sample displayed on Figure 1, while the corresponding estimates on FOMC days display a sharp increase immediately after the announcement. This elevation in volatility is accompanied by heightened trading volume, as diverse groups of investors assess the impact of the news for asset prices, and the economy more generally.

Given this evidence, we conducted our test for constant periodicity in volatility exclud-

Figure 3: Intraday Volatility Periodicity with and without FOMC Announcements. The figure plots smoothed $n\hat{f}_{i\Delta_n, \mathcal{B}} / \sum_{i=1}^n \hat{f}_{i\Delta_n, \mathcal{B}}$ with Gaussian kernel and bandwidth of 5-minute interval for $\mathcal{B}$ being FOMC days (dashed line) and non-FOMC days (solid line).

ing the FOMC days. The test results are very similar to those for the whole sample, reported in Table 3, and, importantly, the strong rejection of the null hypothesis is preserved (results not reported to conserve space). In summary, scheduled macro announcements cannot explain the variation in the periodic component of volatility.

Additional in-depth analysis of the sources of variation in the intraday periodic volatility component is outside the scope of the current paper. Nonetheless, we illustrate how our approach facilitates direct exploration of this question. The basic rationale from economic theory concerning the observed intraday U-shape in volatility is as follows. The opening hours represent a price discovery phase where overnight news arrivals and large customer orders submitted to different dealers need to be analyzed and processed by the agents in the market. Heterogeneous asset positions and beliefs, asymmetric information, and diverse orders interact to generate elevated volatility, but often only moderately high volume. The

latter is due to the fact that large orders tend to be broken up and processed throughout the trading day to avoid excessive price pressure. The typical incentive scheme for order execution relies on the average trade price achieved for the order relative to some metric like the volume-weighted average price (VWAP) across the trading day. As a consequence, risk-averse dealers will prefer to trade later in the day, when the initial bulk of news and the direction of the order flow have been absorbed into the price, and the price impact typically is lower. Risk aversion will induce agents to postpone some trades until later in the day, unless they are based on short-lived information that must be acted on quickly before others do so or before the information becomes public and prices adjust. Since the tendency to postpone a fraction of the non-informational trades will be common across dealers there, naturally, will be a concentration of uninformed order flow towards the end of the day. Recognizing this feature of the market dynamic, the price impact per trade will be low towards the end of the trading day. As a result, we expect to see highly elevated trading accompanied by some increase in volatility in the final hour of regular trading. This line of reasoning further implies that a period of elevation in the stationary component of volatility should push the intensity of trading and the return volatility further back towards the end of the trading day.

Consequently, we now explore whether the level of volatility affects the shape of the intraday volatility pattern. As a proxy for the latent return volatility at the start of the trading day, we rely on the value of the VIX volatility index. In the left panel of Figure 4, we plot the estimated intraday volatility pattern in high- and low-volatility regimes. Specifically, we identify the high volatility regime as the set of days in the sample in which the VIX index at market open is between its 75th and the 95th quantiles across the sample period. Similarly, the low volatility regime is the set of days in which the VIX index at market open is between the 5th and the 25th quantiles. We exclude days of very low

and very high VIX values (below the 5th and above the 95th quantiles) to guard against the effect of extreme outliers. From Figure 4 we see that the two intraday patterns are roughly identical around noon, differ substantially around the opening and close, with the periodic component in the low volatility regime being almost flat towards the market close as opposed to its steep counterpart in the high volatility regime. The above evidence of elevated contribution of the hour before close to the total daily volatility during days of high volatility is robust to performing the kernel smoothing of $n\hat{f}_{i\Delta_n, \mathcal{B}} / \sum_{i=1}^n \hat{f}_{i\Delta_n, \mathcal{B}}$ with larger bandwidth than the one used for generating Figure 4.

On the right panel of Figure 4, we plot the ratio of the estimated time-of-day effects in the two volatility regimes relative to the one based on the whole sample. As seen from the plot, the periodic component in the high volatility regime is very close to the average one. This is because the high volatility regime contributes, in relative terms, more than the low volatility regime to the estimation of the average time-of-day periodic component. On the other hand, the difference in the average estimate for the periodic component of volatility and the one recovered in the low volatility regime is substantial, particularly in the period before market close. This implies that the periodic component of volatility generally will be severely overstated during periods of low volatility, when relying on the usual procedure of standardizing returns by the average estimates for the time-of-day effect.

The exploratory analysis above does suggest that the level of volatility is an important source of variation in the intraday volatility pattern. We now test formally whether the dependence of the intraday volatility pattern on the volatility regime can explain the high rejection rates of our test (even on non-FOMC days) by incorporating the additional information for the VIX index and following the procedure in Section 3.3. To account for the fact that the Laplace transforms have been shifted downward, we adjust the choice of u_{max} so that it accurately reflects the “effective” sample size (i.e., $T_{obs} = \sum_{t=1}^T 1_{\{\mathcal{B}_{t-1}\}}$).

Figure 4: Intraday Volatility Periodicity and Volatility. The left panel plots the smoothed values for $n\hat{f}_{i\Delta_n, \mathcal{B}} / \sum_{i=1}^n \hat{f}_{i\Delta_n, \mathcal{B}}$ using a Gaussian kernel and bandwidth of five minutes, for $\mathcal{B}$ indicating high VIX (solid line) or low VIX (dashed line). The right panel displays the smoothed ratio $5\hat{f}_{i\Delta_n, \mathcal{B}} / \hat{f}_{i\Delta_n}$ using a Gaussian kernel and bandwidth of ten minutes, for $\mathcal{B}$ indicating high VIX (solid line) or low VIX (dashed line). The low (high) VIX state corresponds to the interval between the 5th and 25th (75th and 95th) empirical quantiles of the VIX index. FOMC days are excluded from the computation.

Formally, we set $u_{max}^* = u_{max} / T_{adj}^2$, where $T_{adj}^* = T / T_{obs}$ reflects how much larger the full sample is relative to the one based on the conditioning information. The results from the tests for the high and low volatility regimes are reported in Table 4 (similar results hold also for a median volatility regime). From Table 4, we conclude that accounting for the level of volatility captures a nontrivial part of the time variation in the intraday volatility periodicity. However, controlling for the volatility level alone is clearly not sufficient to capture the behavior of volatility during the first 90 and the last 30 minutes of the trading day in the low volatility regime. Determining what drives the periodicity during these periods is an important question that we leave for future research.

	9:00	9:30	10:00	10:30	11:00	11:30	12:00	12:30	13:00	13:30	14:00	14:30
<i>Low volatility</i>												
8:40 - 9:00	0.071	0.017	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.114
9:00 - 9:30		0.800	0.050	0.001	0.040	0.020	0.001	0.001	0.001	0.001	0.007	0.738
9:30 - 10:00			0.128	0.006	0.073	0.046	0.005	0.003	0.001	0.007	0.016	0.376
10:00 - 10:30				0.156	0.808	0.703	0.183	0.113	0.080	0.293	0.225	0.013
10:30 - 11:00					0.222	0.562	0.955	0.613	0.591	0.723	0.438	0.000
11:00 - 11:30						0.741	0.291	0.243	0.203	0.491	0.487	0.010
11:30 - 12:00							0.713	0.473	0.377	0.836	0.529	0.005
12:00 - 12:30								0.677	0.657	0.863	0.590	0.000
12:30 - 13:00									0.956	0.850	0.822	0.000
13:00 - 13:30										0.734	0.813	0.000
13:30 - 14:00											0.762	0.000
14:00 - 14:30												0.002
<i>High volatility</i>												
8:40 - 9:00	0.778	0.478	0.684	0.096	0.028	0.004	0.101	0.010	0.008	0.511	0.002	0.229
9:00 - 9:30		0.437	0.649	0.151	0.034	0.008	0.308	0.031	0.019	0.808	0.003	0.442
9:30 - 10:00			0.805	0.424	0.123	0.035	0.534	0.107	0.064	0.868	0.018	0.456
10:00 - 10:30				0.129	0.018	0.003	0.230	0.017	0.004	0.712	0.002	0.177
10:30 - 11:00					0.413	0.075	0.544	0.232	0.230	0.425	0.056	0.242
11:00 - 11:30						0.541	0.242	0.588	0.942	0.119	0.447	0.065
11:30 - 12:00							0.026	0.498	0.733	0.005	0.838	0.007
12:00 - 12:30								0.183	0.095	0.584	0.020	0.883
12:30 - 13:00									0.663	0.037	0.273	0.111
13:00 - 13:30										0.038	0.468	0.013
13:30 - 14:00											0.003	0.746
14:00 - 14:30												0.009

Table 4: Conditional Test Results. The table reports results (p-values) from the test of Section 3.3 over the period 1 January, 2005 to 30 January, 2015. The conditioning set is non-FOMC days and VIX belonging to one of two states: low (between 5th and 25th quantile of its empirical distribution) and high (between the 75th and 95th quantile). Critical values are computed using $p_T = 3$. Top row indicates the beginning of each half-hour interval.

6 Conclusion

In this paper we develop a novel test for deciding whether the intraday periodic component of volatility is time-invariant, using a long span of high-frequency data. The test is based on forming estimates of the average value of the periodic component of volatility and then standardizing by it the high-frequency returns. We exploit a weighted L^2 norm of the distance between the empirical characteristic functions of the studentized high-frequency returns at different times of the day to separate the null and alternative hypotheses. The analysis is extended to allow for testing the hypothesis on conditioning sets which can aid identifying the sources of time variation in the volatility periodicity. Our empirical application reveals that intraday volatility periodicity changes (stochastically) over time, with the shifts partially explained by the level of volatility.

References

- [1] A. Admati and P. Pfleiderer. A theory of intraday patterns: volume and price variability. *Review of Financial Studies*, 1:3–40, 1988.
- [2] T. G. Andersen and T. Bollerslev. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4:115–158, 1997.
- [3] T. G. Andersen and T. Bollerslev. Deutsche mark-dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance*, 53(1):219–265, 1998.
- [4] T. G. Andersen, T. Bollerslev, and A. Das. Variance-ratio statistics and high-frequency

- data: testing for changes in intraday volatility patterns. *Journal of Finance*, 56(1):305–327, 2001.
- [5] T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys. Modeling and forecasting realized volatility. *Econometrica*, 71:579–625, 2003.
 - [6] D. Andrews. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59(3):817–858, 1991.
 - [7] N. K. Bakirov, M. L. Rizzo, and G. J. Székely. A multivariate nonparametric test of independence. *Journal of Multivariate Analysis*, 97:1742–1756, 2006.
 - [8] N. K. Bakirov, M. L. Rizzo, and G. J. Székely. Measuring and testing dependence by correlation of distances. *Annals of Statistics*, 35(6):2769–2794, 2007.
 - [9] O.E. Barndorff-Nielsen and N. Shephard. Non-gaussian ornstein-uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society Series B*, 63(2):167–241, 2001.
 - [10] O.E. Barndorff-Nielsen and N. Shephard. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2:1–37, 2004.
 - [11] T. Bollerslev and V. Todorov. Estimation of jump tails. *Econometrica*, 79(6):1727–1783, 2011.
 - [12] K. Boudt, C. Croux, and S. Laurent. Robust estimation of intraweek periodicity in volatility and jump detection. *Journal of Empirical Finance*, 18(2):353–367, 2011.
 - [13] S. Chaker. On high frequency estimation of the frictionless price: the use of observed liquidity variables. *Journal of Econometrics*, 201(1):127–143, 2017.

- [14] S. Chiu. Detecting periodic components in a white gaussian time series. *Journal of the Royal Statistical Society, Series B*, 51(2):249–259, 1989.
- [15] K. Christensen, U. Hounyo, and M. Podolskij. Is the diurnal pattern sufficient to explain the intraday variation in volatility? A nonparametric assessment. *Journal of Econometrics*, 205(2):336–362, 2018.
- [16] S. Clinet and Y. Potiron. Estimation for high-frequency data under parametric market microstructure noise. *Working paper available at arXiv:1712.01479*, 2017.
- [17] S. Clinet and Y. Potiron. Testing if the market microstructure noise is a function of the limit order book. *Working paper available at arXiv:1709.02502*, 2017.
- [18] Sándor Csörgö. Testing for independence by the empirical characteristic function. *Journal of Multivariate Analysis*, 16:290–299, 1985.
- [19] R. F. Engle and M. E. Sokalska. Forecasting intraday volatility in the us equity market. multiplicative component garch. *Journal of Financial Econometrics*, 10(1):54–83, 2012.
- [20] T. W. Epps. Testing that a stationary time series is gaussian. *Annals of Statistics*, 15(4):1683–1698, 1987.
- [21] Y. Fan, P. L. de Micheaux, S. Penev, and D. Salopek. Multivariate nonparametric test of independence. *Journal of Multivariate Analysis*, 153:189–210, 2017.
- [22] E. Ghysels, A. Hall, and H. S. Lee. On periodic structures and testing for seasonal unit roots. *Journal of the American Statistical Association*, 91(436):1551–1559, 1996.

- [23] D. M. Guillaume, M. M. Dacorogna, R. R. Davé, U. A. Müller, R. B. Olsen, and O. V. Pictet. From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and Stochastics*, 1997.
- [24] W. K. Härdle, B. L. Cabrera, O. Okhrin, and W. Wang. Localizing temperature risk. *Journal of the American Statistical Association*, 111(516):1491–1508, 2016.
- [25] L. Harris. A transaction data study of weekly and intradaily patterns in stock returns. *Journal of Financial Economics*, 16(1):99–117, 1986.
- [26] A. Hecq, S. Laurent, and F. C. Palm. Common intraday periodicity. *Journal of Financial Econometrics*, 10(2):325–353, 2012.
- [27] H. Hong and J. Wang. Trading and returns under periodic market closures. *Journal of Finance*, 55(1):297–354, 2000.
- [28] Y. Hong. Hypothesis testing in time series via the empirical characteristic function: a generalized spectral density approach. *Journal of the American Statistical Association*, 94(448):1201–1220, 1999.
- [29] S. Hylleberg, editor. *Modelling seasonality*. Oxford University Press, 1992.
- [30] R. Ito. Modeling dynamic diurnal patterns in high-frequency financial data. *Cambridge Working Papers in Economics*, 2013.
- [31] J. Jacod and P. Protter. *Discretization of processes*. Springer-Verlag, Berlin, 2012.
- [32] J. Jacod and A.N. Shiryaev. *Limit theorems for stochastic processes*. Springer-Verlag, Berlin, 2nd edition, 2003.

- [33] Jean Jacod, Yingying Li, Per A. Mykland, Mark Podolskij, and Mathias Vetter. Microstructure noise in the continuous case: the pre-averaging approach. *Stochastic Processes and Their Applications*, 119:2249–2276, 2009.
- [34] R. H. Jones and W. M. Brelsford. Time series with periodic structure. *Biometrika*, 54(3-4):403–408, 1967.
- [35] Y. Li, S. Xie, and X. Zheng. Efficient estimation of integrated volatility incorporating trading information. *Journal of Econometrics*, 195(1):33–50, 2016.
- [36] M. L. Rizzo and G. J. Székely. Brownian distance covariance. *Annals of Applied Statistics*, 3(4):1236–1265, 2009.
- [37] A. F. Siegel. Testing for periodicity in a time series. *Journal of the American Statistical Association*, 75(370):345–348, 1980.
- [38] V. Todorov and G. Tauchen. The realized laplace transform of volatility. *Econometrica*, 80(3):1105–1127, 2012.
- [39] B. M. Troutman. Some results in periodic autoregression. *Biometrika*, 66(2):219–228, 1978.
- [40] R. A. Wood, T. H. McInish, and J. K. Ord. An investigation of transactions data for nyse stocks. *Journal of Finance*, 40(3):723–739, 1985.
- [41] N. Xia and X. Zheng. On the inference about the spectral distribution of high-dimensional covariance matrix based on high-frequency noisy observations. *Annals of Statistics*, 46:500–525, 2018.

- [42] X. Zheng and Y. Li. On the estimation of integrated covariance matrices of high dimensional diffusion processes. *Annals of Statistics*, 39:3121 – 3151, 2011.