Intraday Periodic Volatility Curves

Torben G. Andersen* Tao Su† Viktor Todorov‡ Zhiyuan Zhang§

January 24, 2023

Abstract

The volatility of financial asset returns displays pronounced variation over the trading day. Our goal is nonparametric inference for the *average* intraday volatility pattern, viewed as a *function* of time-of-day. The functional inference is based on a long span of high-frequency return data. Our setup allows for general forms of volatility dynamics, including time-variation in the intraday pattern. The estimation is based on forming local volatility estimates from the high-frequency returns over overlapping blocks of asymptotically shrinking size, and then averaging these estimates across days in the sample. The block-based estimation of volatility renders the error in the estimation due to the martingale return innovation asymptotically negligible. As a result, the centered and scaled calendar volatility effect estimator converges to a Gaussian process determined by the empirical process error associated with estimating average volatility across the trading day. Feasible inference is obtained by consistently estimating the limiting covariance operator. Simulation results corroborate our theoretical findings. In an application to S&P 500 futures data, we find evidence for a shift in the intraday volatility pattern over time, including a more pronounced role for volatility outside U.S. trading hours in the latter part of the sample.

Keywords: calendar effect; volatility; functional central limit theorem; high-frequency data; nonparametric estimation.
JEL Classification: C14; C22; G12.

*Kellogg School of Management, Northwestern University, NBER, and CREATES; e-mail: t-andersen@northwestern.edu.

†School of Statistics and Mathematics, Zhejiang Gongshang University, e-mail:su.tao@mail.zjgsu.edu.cn.

‡Kellogg School of Management, Northwestern University, e-mail: v-todorov@northwestern.edu.

§School of Statistics and Management, Shanghai University of Finance and Economics, e-mail: zhang.zhiyuan@mail.shufe.edu.cn.

1 Introduction

Time-varying volatility is a ubiquitous feature of financial data and a large body of work has developed parametric and nonparametric techniques for studying volatility dynamics. The primary focus of the earlier literature was on daily and lower frequency volatility variation. However, volatility exhibits systematic and dramatic intraday changes too, with volatility at the market open and close being significant higher, on average, than in the middle of the trading day, see, e.g., [2] and the references therein. Given these pronounced intraday, or diurnal, volatility patterns and their importance for inference from high-frequency return data, our goal is functional inference for average volatility, viewed as a function of time-of-day. Our analysis is conducted in a general setting accommodating, in particular, time-varying intraday periodicity in volatility, evidence for which is provided in [5].

Conducting functional inference for volatility is intricate because volatility is not directly observable. We use high-frequency returns to form estimates of the volatility process which, in turn, can be used to study the calendar effects. In particular, exploiting that volatility remains approximately constant over short intervals, with high probability, we form local volatility estimates from blocks of high-frequency returns of asymptotically shrinking time span. We then take time-series averages of the recovered volatility process for the trading days within the sample to form an estimate of the average intraday volatility function.

We derive the asymptotic behavior of our statistic in the space of L^2 functions over time-of-the-day. This reflects a joint asymptotic setting: the mesh of the observation time grid shrinks jointly with the increase in the time span of the data. We assume initially that, while the intraday volatility pattern may change stochastically over time, it does so in a stationary and ergodic way. In this scenario, our estimate of the average intraday volatility function is consistent and asymptotically normal. The limiting process is Gaussian and

determined by the empirical process error associated with measuring average volatility over the trading day from a sample of daily observations of the intraday volatility process.

We further extend the above analysis to a setting in which the intraday periodic component of volatility shifts gradually over time in a nonstationary way. In this case, our intraday volatility calendar effect estimator recovers a mixture of intraday periodic functions, with the weights assigned to the different periodic functions being determined by the fractions of the sample corresponding to the different intraday periodic volatility regimes. Consequently, the limit distribution of our estimator in this nonstationary setup provides a natural generalization of the one described above in the stationary case.

The error associated with measuring volatility from local blocks of high-frequency increments is asymptotically negligible, provided the time span of the data grows at a slower rate than the square of the sampling frequency (the number of times we sample during a trading day). There are two leading components for the error associated with the recovery of volatility from returns. One is a bias due to the fact that the latent volatility process can change over the local window used for its estimation. The second is a centered error associated with the martingale component of the price. We show that the length of the local window controls the size of these two types of error, with a bias-variance tradeoff determining the optimal window length for forming local volatility estimates. This window size is smaller than the optimal one for estimating spot volatility, see e.g., [21]. This is due to the current long-span asymptotic setting, which reduces the size of the centered error component in measuring volatility in our inference procedures.

For conducting feasible inference, we develop a consistent estimator of the covariance operator using the local volatility proxies. This allows us to quantify the precision in estimating the intraday volatility function and construct associated confidence intervals. Importantly, the newly developed limit theory also enables us to check for nonstationary

trends in the intraday volatility pattern by testing for equality of the intraday volatility functions recovered from non-overlapping subsamples (each of asymptotically increasing time span). Such tests feature a non-standard infinite mixture of independent chi-squared random variables. We develop an easy-to-implement simulation-based method to evaluate corresponding critical test values using our estimate of the covariance operator for the limiting distribution of the statistics. Extensive Monte Carlo simulations corroborate our theoretical results and document good finite-sample size and power of the test.

We go on to estimate the average intraday volatility function and associated pointwise confidence bands using round-the-clock high-frequency return observations for the e-mini S&P 500 futures contract over 2005–2020. In addition, we test for nonstationarity of the intraday calendar effects across the early, middle and later parts of the sample. The average intraday volatility curve is generally obtained with good precision, but the estimation errors tend to rise in concert with the intraday level of volatility. Our test for a stationary intraday volatility calendar effect is rejected, providing strong indication of a shift over time towards a proportionally larger contribution to overall volatility from regions outside the regular U.S. trading day, and especially the Asian zone. Furthermore, there is evidence of enhanced asymmetry during active U.S. trading, with volatility spiking even more dramatically going into the market close in recent years than early in the sample. We reiterate that our null hypothesis accommodates (stationary) variability in the intraday volatility calendar effect. Hence, this represents new formal evidence for a nonstationary shift in the volatility calendar effect that is conceptually distinct from mere stochasticity in the latter.

The current paper is related to several strands of existing work. There is a large body of work documenting calendar volatility effects and testing various features of it. Early empirical studies documenting a U-shaped diurnal pattern in volatility can be found in [34], [17], [28] and [6]. [2] and [32] model the intraday volatility pattern parametrically via

a multiplicative specification of the volatility process as a deterministic function of time-of-day and a standard stationary volatility component. [10] examine various parametric and nonparametric calendar-effect estimation methods in a setting with price jumps and show that their use improves jump detection. [13] test whether a constant deterministic function of time-of-day can explain all variation of volatility across the trading day, while [5] test whether the intraday volatility pattern is time varying but stationary. Contrary to the above-cited work, we develop functional limit results for the average intraday volatility curve without imposing parametric assumptions or assuming constancy of the intraday volatility pattern across time. Thus, unlike existing work, our approach enables us to test for genuine nonstationarity of the intraday calendar effect in volatility.

More broadly, our work relates to the problem of nonparametric measurement of realized return variation. Within an in-fill asymptotic setting, formal inference procedures have been developed for the components of the return process, including diffusive volatility, jump arrival and jump sizes along with the associated microstructure noise process, see, e.g., [8], [23], [1], [26], and [15], as well as [30] for an overview. However, the pronounced (and nonstationary) volatility calendar effects require a different approach, as a non-trivial (asymptotically increasing) time span is necessary for identification of such recurrent patterns in the data. Providing such estimators is important because the economic origins of the periodic and nonperiodic volatility components are distinct, motivating the desire to estimate them separately, see, e.g., [7]. In addition, for problems like identification of jumps from high-frequency returns, accounting for the intraday volatility periodicity can provide nontrivial finite-sample improvements, see e.g., [10]. As such, we fill a notable gap in the literature by developing inference for the intraday volatility curves via formal joint in-fill and long-span asymptotic methods.

Finally, our paper relates to a large and growing literature dealing with functional time

series, see, e.g., [9], [18], [19] for a review of the existing statistical methods and [11], [12], [22], [24], among others, for various applications. The major difference between our paper and this strand of work is that our process of interest, the intraday volatility curve, is latent and has to be inferred from the high-frequency return data. [27] consider functional data analysis of intraday volatility from high-frequency returns, but their inference relies on repeated sampling of the volatility process, while we rely on one realization of the volatility path for estimation. Moreover, [27] require a Lipschitz-continuous volatility process, while we work in the far more general assumption of a jump-diffusive volatility process. The latter is the commonly adopted setting for modern arbitrage-free asset pricing.

The rest of the paper is organized as follows. We introduce the setup and assumptions in Section 2, and define our calendar-effect estimator in Section 3.1. The functional asymptotic theory and a formal test for nonstationarity of the intraday volatility curve are developed in Sections 3.2-3.4. The asymptotic behavior of the estimator in a nonstationary setting is analysed in Section 3.5. Feasible inference methods are provided in Section 4. A finite-sample bias-variance tradeoff analysis is presented in Section 5. Section 6 contains our empirical analysis. Auxiliary results, a simulation study and all proofs are provided in a Supplementary Appendix.

2 Setup and Assumptions

We assume all processes and random variables are defined on a common filtered probability space $(\Omega, \mathscr{F}, (\mathscr{F}(t))_{0\leq t\leq T}, P)$, with the log price process $X(t)$ of the financial asset of interest being governed by an Itô semimartingale of the form,

$$X(t) \;=\; X(0) \;+\; \int_0^t \mu(s)\, ds \;+\; \int_0^t \sigma(s)\, dW(s) \;+\; \int_0^t \int_{\mathbb{R}} x\, \nu(ds, dx), \tag{1}$$

where $\mu(t)$ and $\sigma(t)$ are adapted càdlàg processes, $W(t)$ is a standard Brownian motion, ν is an integer-valued random measure counting the jumps in X with compensator $\chi(t)\,dt \otimes F(dx)$, $\chi(t)$ is an adapted càglàd process and F is a measure on $\mathbb{R}$. The stochastic volatility $\sigma(t)$ follows another Itô semimartingale,

$$\sigma^2(t) \;=\; \sigma^2(0) \;+\; \int_0^t \widetilde{\mu}(s)\,ds \;+\; \int_0^t \breve{\sigma}(s)\,dW(s) \;+\; \int_0^t \widetilde{\sigma}(s)\,d\widetilde{W}(s) \;+\; \int_0^t \int_{\mathbb{R}} x\,\widetilde{\nu}(ds,dx), \quad (2)$$

where $\widetilde{\mu}(t)$, $\widetilde{\sigma}(t)$ and $\breve{\sigma}(t)$ are adapted càdlàg processes, $\widetilde{W}(t)$ is a standard Brownian motion independent from $W(t)$. Moreover, $\widetilde{\nu}$ is the counting jump measure of σ^2 with compensator $\widetilde{\chi}(t)\,dt \otimes \widetilde{F}(dx)$, where $\widetilde{\chi}(t)$ is an adapted càglàd process, and $\widetilde{F}$ is a measure on $\mathbb{R}$.

Most models used in applied work are covered by the volatility specification in (2), with the notable exception of volatility models driven by fractional Brownian motion and/or infinite activity jumps. In the Supplementary Appendix, we extend our analysis to cover such much more general volatility specifications.

Our focus is on the calendar effect in volatility. We assume that,

$$E\left(\sigma^2(t)\right) \;=\; g\left(t - \lfloor t \rfloor\right),$$

for some positive bounded function g on $[0,1]$ with $g(0) = g(1)$. Hence, the volatility process may be nonstationary due to "calendar" effects. Our goal is to estimate,

$$f(\kappa) \;=\; \frac{g(\kappa)}{\int_0^1 g(u)du}, \qquad \textit{for} \;\; \kappa \in [0,1], \quad (3)$$

which we henceforth refer to as the volatility calendar effect. We define $\eta := \int_0^1 g(u)du$ for later use. The standard approach of modeling calendar effects in volatility is through a decomposition, $\sigma^2(t) = c(t - \lfloor t \rfloor)\,\breve{\sigma}^2(t)$, where $c(t - \lfloor t \rfloor)$ is a deterministic function and

$\breve{\sigma}^2(t)$ a stationary process. However, this rules out time variation in the calendar effect, which is documented in [5]. Our general setup accommodates such time variation.

We now turn to our assumptions, starting with one concerning the existence of moments.

Assumption I. *(i) The drift term $\mu(t)$ satisfies $E\left|\mu(t)-\mu(s)\right|^2 \leq C\left|t-s\right|$, for any $s,t\in[0,\infty)$ and some positive constant C that does not depend on s and t.*
(ii) $\sup_{t\in\mathbb{R}_+} Ee^{|\mu(t)|} + \sup_{t\in\mathbb{R}_+} Ee^{|\sigma(t)|} + \sup_{t\in\mathbb{R}_+} Ee^{|\chi(t)|} < \infty$. *Moreover,* $F(\mathbb{R}) < \infty$, $\tilde{F}(\mathbb{R}) < \infty$, $\int_{\mathbb{R}} |x|^2 \tilde{F}(dx) < \infty$, *and,*

$$\sup_{t\in\mathbb{R}_+} E\left|\tilde{\mu}(t)\right|^8 + \sup_{t\in\mathbb{R}_+} E\left|\breve{\sigma}(t)\right|^8 + \sup_{t\in\mathbb{R}_+} E\left|\tilde{\sigma}(t)\right|^8 + \sup_{t\in\mathbb{R}_+} E\left|\tilde{\chi}(t)\right|^8 < \infty.$$

Assumption I(i) is a weak assumption that is satisfied if the drift itself is an Itô semimartingale. In Assumption I(ii), we assume that jumps in price and volatility are of finite activity. This is imposed mainly for simplicity, but still covers a lot of popular models in applied work. This assumption is further relaxed in the Supplementary Appendix. The moment conditions in Assumption I(ii) are stronger than required, but streamline the exposition of our theoretical results. In this regard, note that $X(t)$ is the log-price. Hence, the existence of exponential moments in Assumption I(ii) is a relatively weak condition that typically is satisfied for exponentially-affine models.

Our next assumption concerns stationarity and ergodicity of the volatility process.

Assumption II. *For any positive integer i and $\kappa\in[0,1)$, $\sigma^2(i-1+\kappa)$ is a function (depending on κ) of $Y(i-1+\kappa)$, where $\{Y(t)\}_{t\in\mathbb{R}_+}$ is a (multivariate) Markov process, which is stationary, ergodic and α-mixing with coefficient $\alpha_s = O(s^{-q-\iota})$ for some $q>0$, positive constant ι (which can be arbitrarily close to zero), where for $\mathscr{G}_t = \sigma(Y(u), u\leq t)$,*

$\mathscr{G}^t = \sigma\left(Y(u), u \geq t\right)$ *and* $s > 0$, *we denote,*

$$\alpha_s \;=\; \sup_{t\geq 0} \sup \left\{ |P(A \cap B) - P(A)P(B)| : \; A \in \mathscr{G}_t, B \in \mathscr{G}^{t+s} \right\}.$$

Assumption II covers a wide range of scenarios including, for instance, the case where the volatility takes the following mixture form,

$$\sigma^2(t) \;=\; f_1(t - \lfloor t \rfloor)\, \breve{\sigma}_1^2(t) \;+\; f_2(t - \lfloor t \rfloor)\, \breve{\sigma}_2^2(t),$$

for $\breve{\sigma}_1^2(t)$ and $\breve{\sigma}_2^2(t)$ being stationary processes and some deterministic functions $f_1(\kappa)$ and $f_2(\kappa)$, defined on $[0, 1]$. This setup can accommodate the situation in which the calendar effect in volatility can vary over time, e.g., as a function of the current level of volatility.

3 Estimation and Inference

3.1 Estimating the Intraday Volatility Calendar Effect

We now present our estimator for the volatility calendar effect function $f(\kappa)$ defined in equation (3). We assume that log price process $X(t)$ is discretely observed and observation times are equally spaced over $[0, T]$. In each period $[i-1, i]$, $i = 1, 2, ..., T$, there are $n+1$ observation times $(t_{i,j})_{0\leq j\leq n}$ with $i - 1 \equiv t_{i,0} < t_{i,1} < \ldots < t_{i,j} < \ldots < t_{i,n} \equiv i$, and,

$$t_{i,j} \;=\; i - 1 + j/n \quad \text{with } \Delta = t_{i,j} - t_{i,j-1} \;=\; 1/n, \qquad \text{for} \quad j = 1, 2, ..., n.$$

Note that in the notation above, $t_{i,n} = t_{i+1,0}$. We further adopt the convention that $t_{i,k} = t_{i-1,n+k}$ for $i \geq 2$ and $-n \leq k \leq 0$. To approximate spot volatilities, using observations within a local window of some time-of-period $\kappa \in [0, 1]$, we consider intervals of the form

$[t_{i,j_\kappa-\ell}, t_{i,j_\kappa}]$, where ℓ is an integer and,

$$j_\kappa \;=\; \lfloor \kappa n \rfloor \,.$$

We define the high-frequency log returns as follows,

$$\Delta_{i,j}^n X \;=\; X\left(t_{i,j}\right) - X\left(t_{i,j-1}\right), \text{ for } i = 1, 2, \ldots, T, \;\; j = -n+1, \ldots, n \;\text{ and }\; t_{i,j-1} \geq 0.$$

In all cases except for $i = 1$ and $\kappa \in [0, \ell\Delta)$, the following quantities are used to approximate the spot volatility $\sigma^2(i - 1 + \kappa)$ and the mean of the integrated volatility, respectively,

$$\widehat{\sigma}_{i,\kappa}^2 \;=\; \frac{1}{\ell\Delta} \sum_{k=j_\kappa-\ell+1}^{j_\kappa} \left(\Delta_{i,k}^n X\right)^2 1_{\left\{\left|\Delta_{i,k}^n X\right| \leq u_n\right\}}, \quad \widehat{\eta} \;=\; \frac{1}{T} \sum_{i=1}^{T} \sum_{j=1}^{n} \left(\Delta_{i,j}^n X\right)^2 1_{\left\{\left|\Delta_{i,j}^n X\right| \leq u_n\right\}}, \quad (4)$$

where $u_n = \beta\Delta^{\varpi}$, for some $\varpi \in (0, 1/2)$ and constant $\beta > 0$. For initialization, we simply set $\widehat{\sigma}_{1,\kappa}^2 \equiv \widehat{\sigma}_{1,\ell\Delta}^2$ for $\kappa \in [0, \ell\Delta)$. Based on these quantities, we propose a general estimator of the volatility calendar effect as follows,

$$\widehat{f}(\kappa) \;=\; \frac{1}{T} \sum_{i=1}^{T} \widehat{\sigma}_{i,\kappa}^2 \Bigg/ \widehat{\eta} \;. \quad (5)$$

We shall establish asymptotic theory that takes place in the Hilbert space $\mathcal{L}^2$,

$$\mathcal{L}^2 = \left\{ k : [0,1] \to \mathbb{R} \middle| \int_{[0,1]} k(u)^2 du < \infty \right\}.$$

We denote the inner product and the norm on $\mathcal{L}^2$ by $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$, respectively. Throughout the paper, we adopt the convention that $x_n \asymp y_n$ means $1/C \leq x_n/y_n \leq C$ for some positive constant C.

3.2 Consistency

We start by showing consistency of $\widehat{f}(\kappa)$. Unlike the existing literature, the consistency is established in a functional sense rather than pointwise.

Theorem 1. *Assume that Assumptions I(ii) and II with $q = 1$ hold. Let $\ell \to \infty$ with $\ell\,\Delta \to 0$, and moreover, suppose $T \asymp n^b$ and $\ell \asymp n^c$, for some nonnegative exponents b and c, which satisfy the following condition,*

$$b + c \;>\; 1 - 4\,\varpi\,, \tag{6}$$

where $0 < \varpi < 1/2$. Then, we have, as $n \to \infty$,

$$\parallel \widehat{f}(\kappa) - f(\kappa) \parallel \quad \xrightarrow{P} \quad 0\,.$$

The lower bound condition (6) is imposed on $b + c$ to ensure that the squared difference between the spot volatility approximation with price-jump-truncation and that based on the continuous part of the price process without price-jump-truncation is asymptotically negligible. This condition is very weak and trivially satisfied, when ϖ is greater than $1/4$.

3.3 Convergence in Distribution

We require additional notation to state our formal functional CLT result. First, we define,

$$A_i\,(\kappa) \;=\; \sigma^2\,(i - 1 + \kappa) \;-\; f(\kappa) \int_{i-1}^{i} \sigma^2(t)\,dt, \quad \kappa \in [0,1], \;\; i = 1, 2, \ldots, T. \tag{7}$$

We then set,

$$C(\kappa, \kappa') \;=\; 1/\eta^2 \sum_{h=-\infty}^{\infty} \phi_{\kappa,\kappa'}(h), \quad \kappa, \kappa' \in [0,1]\,, \tag{8}$$

where $\eta = \int_0^1 g(u)du$, $\phi_{\kappa,\kappa'}(h) = Cov\left(A_1(\kappa), A_{1+h}(\kappa')\right)$, if h is a nonnegative integer; and $\phi_{\kappa,\kappa'}(h) = \phi_{\kappa',\kappa}(-h)$, if h is a negative integer.

Theorem 2. *Suppose $\ell \to \infty$ with $\ell\Delta \to 0$ and Assumptions I and II with $q = 3$ hold. Moreover, let $T \asymp n^b$ and $\ell \asymp n^c$, for some nonnegative exponents b and c, which satisfy the following conditions,*

$$0 < b < 4\varpi \qquad \text{and} \qquad 1 - 4\varpi < c < 1 - b/2, \tag{9}$$

where $0 < \varpi < 1/2$. Then, as $n \to \infty$,

$$\sqrt{T}\left(\widehat{f}(\kappa) - f(\kappa)\right) \xrightarrow{d} \mathcal{G}_{\mathcal{K}} \quad \text{in} \quad \mathcal{L}^2,$$

where $\mathcal{G}_{\mathcal{K}}$ is an $\mathcal{L}^2$-valued zero-mean Gaussian process with covariance operator $\mathcal{K}$ defined through the kernel $C(\kappa, \kappa')$ in (8) as follows,

$$\mathcal{K}\, y(\kappa') \;=\; \int_{[0,1]} C(\kappa, \kappa')\, y(\kappa)\, d\kappa, \qquad \forall y \in \mathcal{L}^2\,.$$

The convergence rate for $\widehat{f}(\kappa)$ is naturally determined by the time span of the data. Condition (9) imposes restrictions on the asymptotic order of n relative to T. In evaluating the severity of this condition, we recall that ϖ – determining the truncation threshold for removing price jumps – is optimally set to a value close to $1/2$, see e.g., [16]. In this case, b may take any value in $(0, 2)$, implying Theorem 2 imposes very weak restrictions on the rate at which we sample intraday relative to the time span of the data. For example, n may grow at a rate slower or faster than T. This robustness to the size of n and T is contrary to related joint in-fill and long-span asymptotic results, which require $T/n \to 0$, see e.g., [33]. We discuss the optimal choice of ℓ in Section 5.

The covariance kernel $C(\kappa, \kappa')$ includes autocovariances of all lags in order to accom-

modate the time-series persistence in volatility. We note, however, that the asymptotic variance of $\widehat{f}(\kappa)$ is determined by that of sample averages of $A_i(\kappa)$, which only depend on *changes* in volatility over unit time intervals (trading days). Hence, even if $\sigma^2(t)$ is very persistent, $A_i(\kappa)$ typically displays limited dependence over time. As a result, one can estimate $C(\kappa, \kappa')$ well using only a small number of lags.

The following corollary derives the limiting distribution of the norm $\| \widehat{f}(\kappa) - f(\kappa) \|$.

Corollary 3. *If all Assumptions in Theorem 2 hold, then as* $n \to \infty$,

$$T \| \widehat{f}(\kappa) - f(\kappa) \|^2 \xrightarrow{d} \mathcal{Z},$$

where $\mathcal{Z} = \| \mathcal{G}_{\mathcal{K}} \|^2$ *is a weighted sum of independent* $\chi^2(1)$ *variables, defined on an extension of the original probability space and independent from* $\mathscr{F}$. *The weights are given by the eigenvalues* $(\pi_i)_{i \geq 1}$ *of the covariance operator* $\mathcal{K}$ *in Theorem 2.*

3.4 Testing for Nonstationary Intraday Volatility Periodicity

We now develop a formal test for nonstationarity in the intraday calendar effect. Specifically, it is a test for equality of the estimated intraday volatility curve versus the alternative of a shift in the calendar functional across two separate time periods, using the theory detailed in Section 3.3. Letting $f_{\mathrm{P}}(\kappa)$ and $f_{\mathrm{P}'}(\kappa)$ denote the true intraday calendar effects for the non-overlapping periods P and P$'$, the null and alternative hypotheses take the form,

$$H_0 : \ \| f_{\mathrm{P}}(\kappa) - f_{\mathrm{P}'}(\kappa) \| = 0 \qquad \text{vs.} \qquad H_1 : \ \| f_{\mathrm{P}}(\kappa) - f_{\mathrm{P}'}(\kappa) \| > 0. \tag{10}$$

We denote the number of trading days within P and P$'$ as T and T', respectively. In what follows, we further assume, for $r \in (0, \infty)$, that, $\frac{T}{T'} \to r$. Motivated by Corollary

3, we introduce the test statistic, $T \parallel \widehat{f}_{\mathrm{P}}(\kappa) - \widehat{f}_{\mathrm{P}'}(\kappa) \parallel^2$, where $\widehat{f}_{\mathrm{P}}(\kappa)$ and $\widehat{f}_{\mathrm{P}'}(\kappa)$ denote the calendar effect estimators for time periods P and P$'$. Because we split data into segments over non-overlapping time periods, $\sqrt{T}\,(\widehat{f}_{\mathrm{P}}(\kappa) - f_{\mathrm{P}}(\kappa))$ and $\sqrt{T'}\,(\widehat{f}_{\mathrm{P}'}(\kappa) - f_{\mathrm{P}'}(\kappa))$ are asymptotically independent. Therefore, under the null hypothesis, the test statistic $T \parallel \widehat{f}_{\mathrm{P}}(\kappa) - \widehat{f}_{\mathrm{P}'}(\kappa) \parallel^2$ has the following limiting distribution,

$$\mathcal{M} \;=\; \parallel \mathcal{G}_{\mathcal{K}}(\kappa) - \sqrt{r}\,\mathcal{G}'_{\mathcal{K}}(\kappa) \parallel^2,$$

where $\mathcal{G}_{\mathcal{K}}$ and $\mathcal{G}'_{\mathcal{K}}$ are independent Gaussian processes with the identical covariance operator $\mathcal{K}$ defined in Theorem 2. We summarize the above results formally in the following theorem.

Theorem 4. *Suppose that all Assumptions in Theorem 2 hold for non-overlapping time periods* P *and* P$'$ *with numbers of trading days* T *and* T'*, respectively. Moreover, let* $\frac{T}{T'} \to r$*, for some* $r \in (0, \infty)$*. Then, under the null hypothesis that* $\parallel f_{\mathrm{P}}(\kappa) - f_{\mathrm{P}'}(\kappa) \parallel = 0$*, we have*

$$T \parallel \widehat{f}_{\mathrm{P}}(\kappa) - \widehat{f}_{\mathrm{P}'}(\kappa) \parallel^2 \;\xrightarrow{d}\; \mathcal{M},$$

where $\mathcal{M} = \parallel \mathcal{G}_{\mathcal{K}}(\kappa) - \sqrt{r}\,\mathcal{G}'_{\mathcal{K}}(\kappa) \parallel^2$ *and* $\mathcal{G}_{\mathcal{K}}$ *and* $\mathcal{G}'_{\mathcal{K}}$ *are independent Gaussian processes with the identical covariance operator* $\mathcal{K}$ *defined in Theorem 2.*

Now, testing H_0 versus H_1 is easy. Letting $\mathcal{C}_{1-\alpha}$ denote a consistent estimate of the $1-\alpha$ quantile of $\mathcal{M}$, which is asymptotically bounded under the alternative, we define the rejection region as,

$$\mathcal{R} \;=\; \left\{ T \parallel \widehat{f}_{\mathrm{P}}(\kappa) - \widehat{f}_{\mathrm{P}'}(\kappa) \parallel^2 \geq \mathcal{C}_{1-\alpha} \right\}. \tag{11}$$

The construction of $\mathcal{C}_{1-\alpha}$ is provided in Section 4.

Given Theorem 4, we have,

$$P\left(\mathcal{R}|H_0\right) \to \alpha \qquad \text{and} \qquad P\left(\mathcal{R}|H_1\right) \to 1.$$

3.5 Limit Behavior of the Estimator in a Nonstationary Setting

What happens with our estimator $\widehat{f}(\kappa)$ when the intraday volatility calendar effect is nonstationary? In this section, we answer this question for a type of nonstationarity which we think is most relevant from an empirical point of view. Mainly, one in which the intraday periodic component of volatility shifts gradually over time. To this end, we make the following assumption for the latent volatility process.

Assumption I-NS. *Assume that the price process follows* (1) *and satisfies Assumption I. The latent volatility process* $\sigma^2(t)$ *follows*

$$\sigma^2(t) = g_{\lceil t \rceil}(t - \lfloor t \rfloor)\breve{\sigma}^2(t), \tag{12}$$

where $\breve{\sigma}^2(t)$ *satisfies* (2) *and Assumptions I and II (with* $\sigma^2(t)$ *in* (2) *and Assumptions I and II being replaced with* $\breve{\sigma}^2(t)$*) and*

$$g_{\lceil t \rceil}(t - \lfloor t \rfloor) := g(t - \lfloor t \rfloor) + \gamma(t - \lfloor t \rfloor)h_{\tau,n}\left(\lceil t \rceil / T\right), \tag{13}$$

where $h_{\tau,n}(u)$ *is a bounded and continuous function on* $[0,1]$ *satisfying*

$$h_{\tau,n}(u) = \begin{cases} 0 & \text{when } u \in [0, \tau - \epsilon_n], \\ 1 & \text{when } u \in [\tau + \epsilon_n, 1], \end{cases}$$

for some $0 < \tau < 1$, $[\tau - \epsilon_n, \tau + \epsilon_n] \subset [0,1]$ *and* $\sqrt{T}\epsilon_n \to 0$, g *and* γ *are defined on* $[0,1]$ *such that* $\widetilde{g}(t) := g(t - \lfloor t \rfloor)$ *and* $\widetilde{\gamma}(t) := \gamma(t - \lfloor t \rfloor)$ *are differentiable with bounded first derivatives on* $[0, \infty)$*. Without loss of generality, we further assume that* $E\breve{\sigma}^2(t) \equiv 1$.

In this setting the intraday volatility periodic component, $g_{\lceil t \rceil}(t - \lfloor t \rfloor)$ is nonstationary,

with dynamics similar to that of [20]. We note that $h_{\tau,n}(u) \to h_\tau(u) := 1_{\{u>\tau\}}$ everywhere excluding the point $u = \tau$ as $n \to \infty$. While $h_\tau(u)$ corresponds to the change point alternative of [20], $h_{\tau,n}(u)$ allows for a smooth change of the diurnal pattern in volatility.

We shall derive an asymptotic theory for the calendar effect estimator $\widehat{f}(\kappa)$ defined in equation (5) under the nonstationarity setup (12) and (13) for the latent volatility process. Under the above nonstationarity setup, we need to redefine $f(\kappa)$ as follows,

$$f(\kappa) := \frac{g(\kappa) + (1-\tau)\gamma(\kappa)}{\eta}, \qquad for \quad \kappa \in [0,1], \tag{14}$$

where η is redefined as $\eta := \int_0^1 g(u)du + (1-\tau)\int_0^1 \gamma(u)du$. That is, our estimator now estimates a mixture of g and $g+\gamma$.

In order to present the asymptotic theory in the nonstationary setting, we need some additional notation. First, we redefine $A_i(\kappa)$ as,

$$A_i(\kappa) := g(\kappa)\,(\breve{\sigma}^2(i-1+\kappa)-1) - f(\kappa)\int_{i-1}^{i} g(t-\lfloor t \rfloor)(\breve{\sigma}^2(t)-1)\,dt, \tag{15}$$

and define $B_i(\kappa)$ as,

$$B_i(\kappa) := g_s(\kappa)\,(\breve{\sigma}^2(i-1+\kappa)-1) - f(\kappa)\int_{i-1}^{i} g_s(t-\lfloor t \rfloor)(\breve{\sigma}^2(t)-1)\,dt, \tag{16}$$

with the shorthand notation $g_s(\kappa) = g(\kappa) + \gamma(\kappa)$. We then redefine $C(\kappa,\kappa')$ as,

$$C(\kappa,\kappa') := \tau\, C_A(\kappa,\kappa') + (1-\tau)\, C_B(\kappa,\kappa'), \tag{17}$$

where,

$$C_A(\kappa,\kappa') := \frac{1}{\eta^2}\left(EA_1(\kappa)A_1(\kappa') + \sum_{h=1}^{\infty}\left(EA_1(\kappa)A_{1+h}(\kappa') + EA_{1+h}(\kappa)A_1(\kappa')\right)\right),$$

and

$$C_B(\kappa,\kappa') := \frac{1}{\eta^2}\left(EB_1(\kappa)B_1(\kappa') + \sum_{h=1}^{\infty}\left(EB_1(\kappa)B_{1+h}(\kappa') + EB_{1+h}(\kappa)B_1(\kappa')\right)\right).$$

We are now ready to present the asymptotic theory for our calendar effect estimator under Assumption I-NS.

Theorem 5. *Assume that Assumption I-NS with $q = 3$ holds. Moreover, let $T \asymp n^b$ and $\ell \asymp n^c$, for some nonnegative exponents b and c, which satisfy the following conditions,*

$$0 < b < 4\varpi \qquad \text{and} \qquad 1 - 4\varpi < c < 1 - b/2,$$

where $0 < \varpi < 1/2$. Then, as $n \to \infty$,

$$\sqrt{T}\left(\widehat{f}(\kappa) - f(\kappa)\right) \xrightarrow{d} \mathcal{G}_{\mathcal{K}} \quad \text{in} \quad \mathcal{L}^2,$$

where f is given in (14), $\mathcal{G}_{\mathcal{K}}$ is an $\mathcal{L}^2$-valued zero-mean Gaussian process with covariance operator $\mathcal{K}$ defined through the kernel $C(\kappa,\kappa')$ in (17) as follows,

$$\mathcal{K}\,y(\kappa') = \int_{[0,1]} C(\kappa,\kappa')\,y(\kappa)\,d\kappa, \qquad \forall y \in \mathcal{L}^2\,.$$

4 Feasible Inference

To render Theorem 2, and hence Corollary 3 as well as the test in Section 3.4 feasible, we need to estimate the covariance operator $\mathcal{K}$. Towards this end, for each $\kappa \in (0,1]$ and

$i \in \{1, 2, ..., T\}$, we denote,

$$\widehat{A}_i(\kappa) \;=\; \widehat{\sigma}^2_{i,\kappa} \;-\; \widehat{f}(\kappa) \sum_{j=1}^{n} \left(\Delta^n_{i,j} X\right)^2 1_{\left\{|\Delta^n_{i,j} X| \leq u_n\right\}}.$$

Notice that the covariance operator $\mathcal{K}$ is defined through the covariance kernel $C(\kappa, \kappa')$ in equation (8), which is the autocovariance function of the limiting Gaussian process $\mathcal{G}_{\mathcal{K}}$ in Theorem 2. Hence, we estimate $\mathcal{K}$ by estimating $C(\kappa, \kappa')$. To do this, notice further that from equation (8), we can rewrite $C(\kappa, \kappa')$ as follows,

$$C(\kappa, \kappa') \;=\; \frac{1}{\eta^2} \sum_{h=-\infty}^{\infty} Cov\left(A_1(\kappa), A_{1+h}(\kappa')\right).$$

We have already provided a consistent estimator of η in equation (4). It remains to find consistent estimators of the covariances $\{Cov\left(A_1(\kappa), A_{1+h}(\kappa')\right)\}_{h\in\mathbb{Z}}$ for the sequences $(A_i(\kappa))_{i\geq 1}$ and $(A_i(\kappa'))_{i\geq 1}$, that are stationary under Assumption II.

It is natural to estimate $\{Cov\left(A_1(\kappa), A_{1+h}(\kappa')\right)\}_{h\in\mathbb{Z}}$ via their sample counterparts. Since $A_i(\kappa)$ and $A_i(\kappa')$ are unobserved, we proceed by substituting $\widehat{A}_i(\kappa)$ and $\widehat{A}_i(\kappa')$ for $A_i(\kappa)$ and $A_i(\kappa')$ in the sample covariance estimators. Using this approach, we obtain the following estimator of the kernel $C(\kappa, \kappa')$,

$$\widehat{C}(\kappa, \kappa') = \frac{1}{\widehat{\eta}^2} \left\{ \frac{1}{T} \sum_{i=1}^{T} \widehat{A}_i(\kappa)\widehat{A}_i(\kappa') + \sum_{h=1}^{L_n} \frac{1}{T-h} \sum_{i=1}^{T} \left[\widehat{A}_i(\kappa) \left(\widehat{A}_{i+h}(\kappa') + \widehat{A}_{i-h}(\kappa') \right) \right] \right\} \quad (18)$$

where $\widehat{A}_i(\kappa) = \widehat{A}_i(\kappa') = 0$, if $i \leq 0$ or $i > T$, and the integer L_n diverges as specified in Theorem 6. Thus, an estimator $\widehat{\mathcal{K}}$ of the covariance operator $\mathcal{K}$ is naturally given by,

$$\widehat{\mathcal{K}}\, y(\kappa') \;=\; \int_{[0,1]} \widehat{C}(\kappa, \kappa')\, y(\kappa)\, d\kappa, \qquad \forall y \in \mathcal{L}^2. \quad (19)$$

Now let $\mathcal{G}_{\widehat{\mathcal{K}}}$ be an $\mathscr{F}$-conditional $\mathcal{L}^2$-valued zero-mean Gaussian process with the covariance operator $\widehat{\mathcal{K}}$ defined in equation (19). The next theorem shows that $\mathcal{G}_{\widehat{\mathcal{K}}}$ converges in distribution to $\mathcal{G}_{\mathcal{K}}$ in $\mathcal{L}^2$.

Theorem 6. *Suppose that all Assumptions in Theorem 2, Assumption II with* $q = 4$, $b + c > 1 - 16\varpi/7$ *and* $c > (3 - 8\varpi)/3$ *hold. Moreover,* $\int_{\mathbb{R}} |x|^8 \tilde{F}(dx) < \infty$ *and* $L_n \asymp n^{\varrho}$ *for some strictly positive exponent* ϱ*, which satisfies the following condition,*

$$\varrho < \min\Big\{c/2, (1-c)/4, 2\varpi - 3(1-c)/4, b/2, 2\varpi - 7/8 + 7(b+c)/8\Big\}. \tag{20}$$

Then,

$$\mathcal{G}_{\widehat{\mathcal{K}}} \xrightarrow{d} \mathcal{G}_{\mathcal{K}} \quad \text{in} \quad \mathcal{L}^2.$$

Compared with Theorem 2, Theorem 6 imposes two additional constraints on the rate parameters, $c > (3-8\varpi)/3$ and $b+c > 1-16\varpi/7$, but they remain weak. The requirement of $c > (3-8\varpi)/3$ is stronger than $c > 1-4\varpi$ in Theorem 2, because $(3-8\varpi)/3 > 1-4\varpi$. Nonetheless, if one chooses $\varpi \geq 3/8$, the condition $c > (3-8\varpi)/3$ is trivially satisfied. A similar comment applies to the requirement $b + c > 1 - 16\varpi/7$.

Armed with Theorem 6, we can now approximate the distribution of the limiting variable $\mathcal{Z} = \| \mathcal{G}_{\mathcal{K}} \|^2$, and hence of $T \| \widehat{f}(\kappa) - f(\kappa) \|^2$, in Corollary 3 by that of $\| \mathcal{G}_{\widehat{\mathcal{K}}} \|^2$. The distribution of $\| \mathcal{G}_{\widehat{\mathcal{K}}} \|^2$, in turn, is readily obtained by simulation. We first partition the interval [0,1] into 100 equal subintervals. Define $\kappa_i = \frac{i}{100}$ for $i = 1, 2, \ldots, 100$, and,

$$\widehat{\mathcal{Z}} \;=\; \frac{1}{100} \sum_{i=1}^{100} \mathcal{G}_{\widehat{\mathcal{K}}}(\kappa_i)^2 \;\stackrel{d}{=}\; \frac{1}{100} \sum_{i=1}^{100} \widehat{\pi}_i \, \mathcal{X}_i^2 \,, \tag{21}$$

where $(\mathcal{G}_{\widehat{\mathcal{K}}}(\kappa_1), \ldots, \mathcal{G}_{\widehat{\mathcal{K}}}(\kappa_{100}))$ is, conditionally on $\mathscr{F}$, a multivariate normal random vector with conditional covariance matrix $(\widehat{C}(\kappa_i, \kappa_j))_{1 \leq i,j \leq 100}$, the $\mathcal{X}_i^2$ are independent and $\chi^2(1)$

distributed, defined on an extension of the original probability space and independent from $\mathscr{F}$, and $\widehat{\pi}_i$'s are eigenvalues of the matrix $(\widehat{C}(\kappa_i, \kappa_j))_{1\leq i,j\leq 100}$. In practice, the matrix $(\widehat{C}(\kappa_i, \kappa_j))_{1\leq i,j\leq 100}$ may occasionally produce negative eigenvalues. If this occurs, we retain only the terms with positive eigenvalues in equation (21). Our simulation evidence confirms that this approach generates negligible approximation errors. Finally, we note that $\widehat{\mathcal{Z}}$ is just a Riemann sum associated with $\| \mathcal{G}_{\widehat{\mathcal{K}}} \|^2 = \int_0^1 \mathcal{G}_{\widehat{\mathcal{K}}}(\kappa)^2 \, d\kappa$.

We can similarly approximate the limiting distribution of the test in Section 3.4. Let $\mathcal{G}_{\widehat{\mathcal{K}}}$ and $\mathcal{G}'_{\widehat{\mathcal{K}}}$, conditionally on $\mathscr{F}$, be independent Gaussian processes with common covariance operator $\widehat{\mathcal{K}}$ defined in equation (19). As above, we approximate the limiting distribution of $T \| \widehat{f}_{\mathrm{P}}(\kappa) - \widehat{f}_{\mathrm{P}'}(\kappa) \|^2$ under the null hypothesis by that of,

$$\widehat{\mathcal{M}} \;=\; \frac{1}{100} \sum_{i=1}^{100} \left(\mathcal{G}_{\widehat{\mathcal{K}}}(\kappa_i) - \sqrt{T/T'}\, \mathcal{G}'_{\widehat{\mathcal{K}}}(\kappa_i) \right)^2 \;\stackrel{d}{=}\; \frac{1}{100} \sum_{i=1}^{100} \widehat{\pi}_i \, \mathcal{X}_i^2 \tag{22}$$

where $\kappa_i = i/100$ for $i = 1, 2, \ldots, 100$, $\widehat{\pi}_i$'s are eigenvalues of the estimated limiting covariance matrix $\{(1 + T/T')\, \widehat{C}(\kappa_i, \kappa_j)\}_{1\leq i,j\leq 100}$, whose entries are defined in equation (18) and computed from pooled data across the periods P and P′, while the $\mathcal{X}_i^2$ are iid $\chi^2(1)$, defined on an extension of the original probability space independent of $\mathscr{F}$. We retain only the terms with positive eigenvalues, noting again that the simulation evidence in Section A.3 confirms the finite-sample bias associated with this approach is negligible. Hence, the distribution of $\widehat{\mathcal{M}}$ may be approximated by the empirical distribution of a large number of iid copies generated from $(\mathcal{X}_1^2, \ldots, \mathcal{X}_{100}^2)$. As a result, $\mathcal{C}_{1-\alpha}$ in equation (11) is the $(1-\alpha)$th quantile of the distribution of $\widehat{\mathcal{M}}$.

We conclude this section with a corollary enabling pointwise feasible inference.

Corollary 7. *In the setting of Theorem 6, we have*

$$\frac{\sqrt{T}\left(\widehat{f}(\kappa)-f(\kappa)\right)}{\sqrt{\widehat{C}(\kappa,\kappa)}} \quad \xrightarrow{d} \quad \mathrm{N}(0,1), \qquad as \quad n\to\infty.$$

We may also construct uniform confidence bounds based on our functional limit theory, but the analysis is fairly involved and it is relegated to the Supplementary Appendix.

5 A Finite-Sample Bias-Variance Tradeoff for ℓ

There is a bias-variance tradeoff in picking ℓ, which we review in this section. The following analysis adopts a pointwise perspective. From equations (C.3) and (C.6) in Section C.1 of the Supplementary Appendix, for a fixed κ, the error in the calendar-effect estimation has the following decomposition,

$$\begin{aligned}
\widehat{f}(\kappa)-f(\kappa) \;=\; & \underbrace{\frac{1}{\widehat{\eta}}\left(\frac{1}{T}\sum_{i=1}^{T}\sigma^2(i-1+\kappa)-\frac{f(\kappa)}{T}\sum_{i=1}^{T}\int_{i-1}^{i}\sigma^2(t)\,dt\right)}_{\text{term I}=O_P\left(\frac{1}{\sqrt{T}}\right)} + \\
& \underbrace{\frac{1}{\widehat{\eta}\,T\,\ell\,\Delta}\sum_{i=1}^{T}\sum_{k=j_\kappa-\ell+1}^{j_\kappa}\left\{\left[\int_{t_{i,k-1}}^{t_{i,k}}\sigma(t)dW(t)\right]^2-\int_{t_{i,k-1}}^{t_{i,k}}\sigma^2(t)\,dt\right\}}_{\text{term II}=O_P\left(\frac{1}{\sqrt{T\ell}}\right)} + \\
& \underbrace{\frac{1}{\widehat{\eta}\,T\,\ell\,\Delta}\sum_{i=1}^{T}\sum_{k=j_\kappa-\ell+1}^{j_\kappa}\int_{t_{i,k-1}}^{t_{i,k}}\left[\sigma^2(t)-\sigma^2(i-1+\kappa)\right]dt}_{\text{term III}=O_P\left(\frac{\ell}{n}\vee\sqrt{\frac{\ell}{nT}}\right)} \\
& + \;\ell\text{-independent terms of order } o_P\left(\frac{1}{\sqrt{T}}\right) \;+\; \text{higher-order terms.}
\end{aligned}$$

Under the assumptions of Theorem 2, terms II and III are of smaller order than term I, so only the latter contributes to the central limit theorem. However, in finite samples, terms II and III do matter. In fact, they are the only relevant terms involving ℓ, so they govern the associated tradeoff. Close inspection reveals that term II constitutes a martingale innovation that generates finite-sample variation for the calendar-effect estimator, while term III is a bias term arising from treating volatility as constant over the local window. This type of decomposition arises also in other high-frequency asymptotic settings, see e.g., Theorem 1 in [29] and also Theorem 2 in [31].

Given that $T \asymp n^b$ and $\ell \asymp n^c$, we may now provide guidance on a suitable choice for the order of ℓ, for predetermined values of ϖ and $b \in (0, 4\varpi)$ (see equation (9)) in terms of minimizing the order of the sum of the terms II and III. We define,

$$c_L \;=\; \max\left\{1 - 4\varpi, 0\right\}, \qquad \text{and} \qquad c_U \;=\; 1 - \frac{b}{2}\,.$$

The optimal choice of c, denoted c_{opt}, depends on the specific configuration of ϖ and b. It may be shown to take the following form (see Section C.8 of the Supplementary Appendix),

$$c_{\text{opt}} \;=\; \begin{cases} \frac{2-b}{3}, & \text{when } b \geq 1/2\ , \\ \frac{1}{2}, & \text{when } b < 1/2 \quad \text{and} \quad c_L < 1/2, \\ c_L+, & \text{when } b < 1/2 \quad \text{and} \quad c_L \geq 1/2, \end{cases}$$

where c_L+ denotes a value exceeding c_L by an arbitrary small amount. In the standard case, $\varpi > 1/4$ and $n \asymp T$ (i.e., $b = 1$), the optimal choice implies $\ell \asymp n^{1/3}$. This order is smaller than the optimal window for spot volatility estimation, which is $\sqrt{n}$, see, e.g., [21]. The reason is that our current setting involves T asymptotically increasing. This reduces

the impact of the variance component associated with term II, rendering it beneficial to employ a smaller sized local window.

6 Empirical Analysis

In this section, we apply the general calendar-effect estimator and the associated inference techniques to high-frequency transaction prices of e-mini S&P 500 futures (e-mini for short), a stock market index futures contract traded on the CME Group's electronic trading platform Globex. We investigate the diurnal pattern across different time periods.

6.1 Data

Our e-mini transaction price data are provided by Tick Data. We first construct a fifteen-year record of thirty-second transaction prices from November 30, 2005, through November 30, 2020, based on the most active futures contracts. We note that the calendar-effect estimator (5) is affected by microstructure noise in the same way that standard estimators of integrated volatility are impacted. For the average daily realized volatility calculated from thirty-second returns, the impact of microstructure noise on volatility/calendar-effect estimation is near negligible, consistent with the evidence for only a limited autocorrelation in the noise process, see [14] and [25]. The e-mini contract trades from Sunday to Friday 5:00pm–4:00pm Chicago Time (CT) with a 15-minute trading halt from 3:15–3:30pm CT, and a daily maintenance period from 4:00–5:00pm CT. Hence, one regular trading week consists of five trading days and each complete trading day consists of 22.75 trading hours. We eliminate the price jumps arising from the roll-over between futures contracts via the *forward ratio adjustment* method, as detailed in the TickWrite manual developed by Tick Data. Likewise, we discard all returns that cover time gaps in the price series stemming

from faulty price dissemination, trading halts, daily maintenance periods, weekends and holidays. This leaves us with a sample of 3,379 trading days for a total of 9,224,671 thirty-second trade price observations.

6.2 The S&P 500 Futures Intraday Volatility Curve

In this section, we provide our estimate for the intraday volatility curve based on the full sample, and then examine subsamples to gauge whether the volatility calendar effect may be changing over time. To this end, we split the sample into five-year segments, November 30, 2005–November 29, 2010, November 30, 2010–November 29, 2015, and November 30, 2015–November 30, 2020. For simplicity, we refer to these periods as 2005–2010, 2010–2015 and 2015–2020. Accounting for trading halts and holidays, this leaves us with a sample of $T = 991$, 1,140 and 1,248 days for a total of 2,705,431, 3,112,201 and 3,407,041 thirty-second transaction prices, respectively, for 2005–2010, 2010–2015 and 2015–2020.

We estimate $f(\kappa)$ using $\ell = 10$ over a grid of 100 equidistant points within $[0, 1]$. In the top panel of Figure 1, we depict the estimated intraday volatility curve for the full sample. The first, second and fourth dashed vertical lines separate regimes associated with active trading in Asia (left), Europe (middle) and American (right), see [4]. It is evident that U-shape style patterns exist for all three trading regimes, while the magnitude of the diurnal effect varies dramatically across regions, increasing manyfold as we move from the Asian to the American regime. Moreover, there is a notable difference between the present e-mini volatility pattern and the one for exchange rates documented in [3]. The latter finds the Deutsche mark–U.S. dollar volatility to decline monotonically in the hours leading up to the close of regular trading in the U.S. This suggests that the intraday pattern for exchange rates is determined, in large part, by the overlap in trading activity across global financial centers while, in contrast, the calendar effects for the e-mini are aligned with the regular

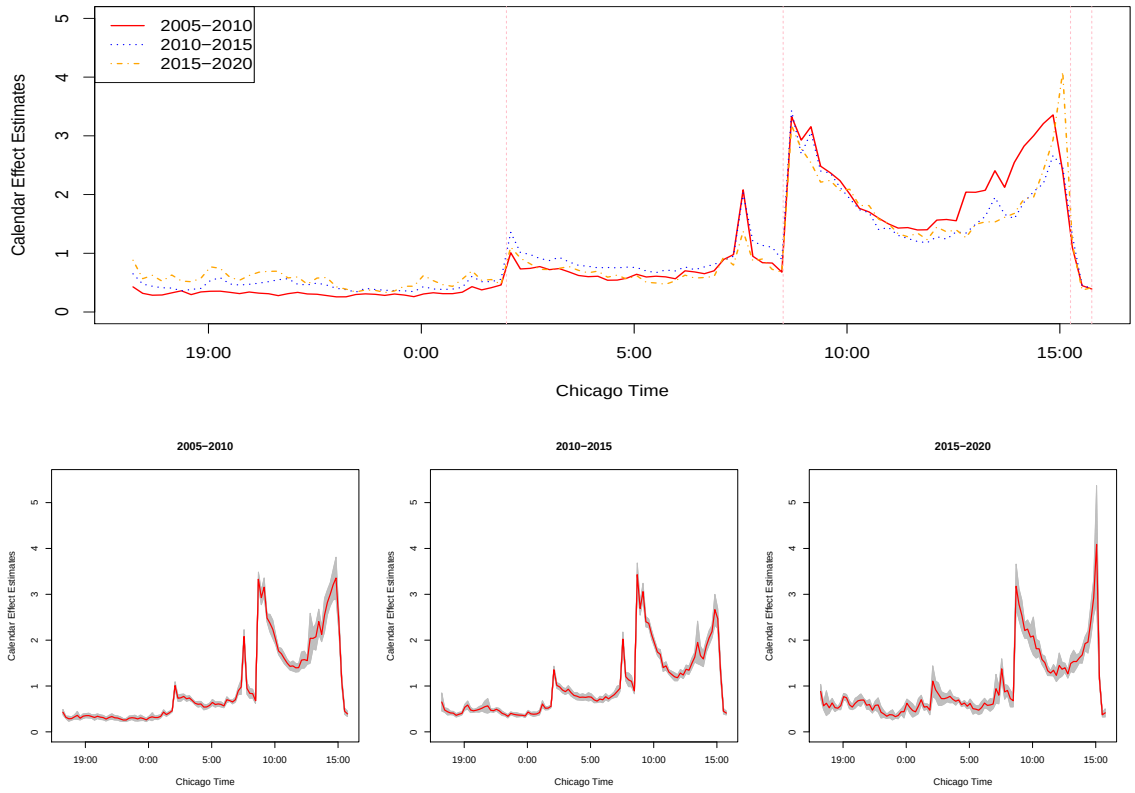

Figure 1. Estimates of the intraday calendar effect in volatility. **Top panel**: The volatility curve estimated from thirty-second e-mini transaction prices over the full sample. The dashed lines at 2:00 CT, 8:30 CT, 15:15 CT and 16:00 CT, separate regular trading hours in Asia (left), Europe (middle), and America (right), and the trading halt from 15:15–15:30 CT. **Bottom panel**: Intraday volatility curves for the e-mini over subsamples. The shaded areas indicate 95% confidence intervals constructed using Corollary 7 with $L_n = 7$.

trading hours of the individual regions. Specifically, the volatility rises up towards the U.S. equity market close at 15:00 CT, but drops off rapidly prior to the halt in futures trading at 15:15 CT, indicated by the third vertical dashed line. Finally, the volatility keeps tailing off during the final short 15:30–16:00 CT trading interval.

The bottom panel of Figure 1 plots the estimated volatility curve for each subsample. In addition, we superimpose 95% pointwise confidence intervals (shaded), based on the feasible CLT in Corollary 7. Evidently, the estimation uncertainty increases with the general level of intraday volatility. In the first two subsamples, the confidence bands are clearly visible only in parts of the American region, but they become notable across the entire trading day over the last period, indicating increased variability in the diurnal pattern. Comparing across subsamples, one notes a relative increase in volatility during the Asian regime and, possibly, some distortions to the U-shape in the American zone. However, the associated estimation uncertainty is also elevated, rendering firm conclusions premature.

6.3 Testing for a Shift in the Volatility Curve over Time

Motivated by these findings, we turn to the formal test for equality in the intraday calendar effect between two separate time periods introduced in Section 3.4. Table 1 reports p-values for the test of identical volatility curves across all possible pairings of subsamples. The test firmly rejects the null that the intraday patterns are equal in the L^2 sense for the pairs (2005-2010, 2010-2015) and (2005-2010, 2015-2020) at the 1% level. However, the evidence against the null is only significant at the 10% level for (2010-2015, 2015-2020). Hence, we have strong indication of a shift in the diurnal pattern after 2010, but only weak evidence that it also changed between 2010-2015 and 2015-2020.

Table 1. Test for changes in the intraday volatility curve. The table provides p-values for the test of the null hypothesis in equation (10), where the non-overlapping subsamples are 2005–2010, 2010–2015, and 2015–2020. The reported p-values equal the percentage of copies from $\widehat{\mathcal{M}}$ that exceeds the realization of the test statistic. The p-values significant at the 1% level are in bold font.

P\P′	2005–2010	2010–2015	2015-2020
2005–2010		**0.0001**	**0.0025**
2010–2015	**0.0001**		0.0992
2015–2020	**0.0025**	0.0992	

6.4 Testing for a Shift in the Regional Volatility Curve

In the preceding section, we found evidence in favor of a shift in the diurnal pattern across time. We now explore whether such a change may be attributed, at least in part, to a shift in the configuration of the diurnal pattern within any particular region of the trading day. To test this hypothesis, we must eliminate the impact of any shift in the relative level of volatility across the regional trading zones. We do this by estimating the volatility calendar effect for each time period separately over each region, thus effectively treating each region, in turn, as comprising the full trading day. That is, we simply perform the test from the preceding section for the Asian, European and American trading regimes separately.

Table 2 reports p-values of the test for pairings of our five-year periods within each trading zone. For the Asian regime, the difference in the diurnal effect is only significant for 2005-2010 versus 2010-2015 at the 5% level, and for 2005-2010 versus 2015-2020 at the 10% level. Likewise, the evidence for the European regime is relatively weak, with the test only being significant for 2005-2010 versus 2010-2015 at the 10% level, and for 2005-2010 versus 2015-2020 at the 5% level. In contrast, for the American regime, the test firmly

Table 2. Test for a shift in the regional intraday volatility pattern. The table provides p-values for the test of the null hypothesis in equation (10) with non-overlapping subsamples 2005–2010, 2010–2015, and 2015–2020. The test is identical to the one described in Section 6.2 and Table 1, but is now applied only within the individual regional zones, Asia, Europe and America. The values that are significant at the 1% level are in bold font.

Trading regimes	Asia			Europe			America		
P\P′	05-10	10-15	15-20	05-10	10-15	15-20	05-10	10-15	15-20
05-10		0.0358	0.0856		0.0846	0.0402		**0.0004**	**0.0027**
10-15	0.0358		0.3735	0.0846		0.3310	**0.0004**		0.0845
15-20	0.0856	0.3735		0.0402	0.3310		**0.0027**	0.0845	

rejects the null of no difference in the diurnal pattern for the pairs (2005-2010, 2010-2015) and (2005-2010, 2015-2020) at the 1% significance level. Meanwhile, the test can only reject the null hypothesis for 2010-2015 versus 2015-2020 at the 10% level.

In summary, the empirical results point to two separate origins behind the shift in the diurnal volatility pattern over our sample period. The first is a shift in the level of volatility across regions, with the relative share of volatility rising during the Asian trading hours, as suggested by the evidence in Section 6.2. The second is the change in the diurnal volatility pattern during the American region, which seems to reflect a shift away from volatility in the middle and towards the close of the regular U.S. trading hours. Interestingly, our findings suggest that the rise in importance of Asian trading – and presumably also for its role in global price discovery for the U.S. equity index – is not associated with any major shift in the qualitative shape of the diurnal pattern within that region.

7 Conclusion

We propose a general estimator of intraday calendar effects in volatility using high-frequency data. Within a general semimartingale setup, we establish the associated functional asymptotic theory in the L^2 metric. To the best of our knowledge, this is the first functional inference theory built for volatility calendar effect estimation using high-frequency data in such a general setup. Both feasible pointwise and functional inference techniques are provided through consistent estimation of the asymptotic covariance operator. Additionally, we provide a finite-sample bias-variance tradeoff analysis for the size of the local window used for volatility estimation. Simulation results corroborate our theoretical findings.

We apply our general estimation procedure to transaction data for the e-mini S&P 500 futures. We find the volatility of the e-mini to exhibit pronounced U-shaped patterns within the Asian, European and American trading regimes. Based on the feasible inference theory, we then test for shifts in the intraday calendar effects across subsamples. The results firmly reject the null of invariance over time, with an elevation in the relative volatility over the Asian regime and a steepening of the U-shape towards the close of regular trading in the U.S. over the recent years being particularly striking.

Acknowledgements

Zhang acknowledges support from the National Natural Science Foundation of China (71871132 and 91546202) and Innovative Research Team of Shanghai University of Finance and Economics (2020110930). The authors report there are no competing interests to declare. We would like to thank an Associate Editor and a referee for many constructive comments and suggestions.

References

[1] Yacine Aït-Sahalia, Jean Jacod, and Jia Li. Testing for jumps in noisy high frequency data. *Journal of Econometrics*, 168:207–222, 2012.

[2] Torben G Andersen and Tim Bollerslev. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4(2-3):115–158, 1997.

[3] Torben G Andersen and Tim Bollerslev. Deutsche mark–dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance*, 53(1):219–265, 1998.

[4] Torben G Andersen, O Bondarenko, Albert S Kyle, and Anna A Obizhaeva. Intraday trading invariance in the e-mini s&p 500 futures market. *Working paper*, 2018.

[5] Torben G Andersen, Martin Thyrsgaard, and Viktor Todorov. Time-varying periodicity in intraday volatility. *Journal of the American Statistical Association*, 114(528):1695–1707, 2019.

[6] Richard T Baillie and Tim Bollerslev. Intra-day and inter-market volatility in foreign exchange rates. *The Review of Economic Studies*, 58(3):565–585, 1991.

[7] Ole E Barndorff-Nielsen and Neil 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.

[8] Ole E Barndorff-Nielsen and Neil Shephard. Power and bipolar variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2:1–37, 2004.

[9] Denis Bosq. *Linear processes in function spaces: theory and applications*, volume 149. Springer Science & Business Media, 2000.

[10] Kris Boudt, Christophe Croux, and Sébastien Laurent. Robust estimation of intraweek periodicity in volatility and jump detection. *Journal of Empirical Finance*, 18(2):353–367, 2011.

[11] Clive G Bowsher and Roland Meeks. The dynamics of economic functions: modeling and forecasting the yield curve. *Journal of the American Statistical Association*, 103(484):1419–1437, 2008.

[12] Brenda López Cabrera and Franziska Schulz. Forecasting generalized quantiles of electricity demand: A functional data approach. *Journal of the American Statistical Association*, 112(517):127–136, 2017.

[13] Kim Christensen, Ulrich Hounyo, and Mark Podolskij. Is the diurnal pattern sufficient to explain intraday variation in volatility? a nonparametric assessment. *Journal of Econometrics*, 205(2):336–362, 2018.

[14] Rui Da and Dacheng Xiu. When moving-average models meet high-frequency data: Uniform inference on volatility. *Econometrica*, 89(6):2787–2825, 2021.

[15] Prosper Dovonon, Silvia Gonçalves, Ulrich Hounyu, and Nour Meddahi. Bootstrapping high-frequency jump tests. *Journal of the American Statistical Association*, 114:793–803, 2019.

[16] José E Figueroa-López and Cecilia Mancini. Optimum thresholding using mean and conditional mean squared error. *Journal of Econometrics*, 208(1):179–210, 2019.

[17] Lawrence Harris. A transaction data study of weekly and intradaily patterns in stock returns. *Journal of Financial Economics*, 16(1):99–117, 1986.

[18] Siegfried Hörmann and Piotr Kokoszka. Functional time series. In *Handbook of statistics*, volume 30, pages 157–186. Elsevier, 2012.

[19] Lajos Horváth and Piotr Kokoszka. *Inference for functional data with applications*, volume 200. Springer Science & Business Media, 2012.

[20] Lajos Horváth, Piotr Kokoszka, and Gregory Rice. Testing stationarity of functional time series. *Journal of Econometrics*, 179(1):66–82, 2014.

[21] Jean Jacod and Philip Protter. *Discretization of processes*. Springer-Verlag, Berlin, Heidelberg, 2011.

[22] Shuhao Jiao, Alexander Aue, and Hernando Ombao. Functional time series prediction under partial observation of the future curve. *Journal of the American Statistical Association*, forthcoming, 2021.

[23] Suzanne S Lee and Per A Mykland. Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies*, 21:2535–2563, 2008.

[24] Degui Li, Peter M Robinson, and Han Lin Shang. Long-range dependent curve time series. *Journal of the American Statistical Association*, 115(530):957–971, 2020.

[25] Z Merrick Li and Oliver B Linton. A remedi for microstructure noise. *Econometrica*, 90(1):367–389, 2022.

[26] Cecilia Mancini. Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics*, 36:270–296, 2009.

[27] Hans-Georg Müller, Rituparna Sen, and Ulrich Stadtmüller. Functional data analysis for volatility. *Journal of Econometrics*, 165(2):233–245, 2011.

[28] Ulrich A Müller, Michel M Dacorogna, Richard B Olsen, Olivier V Pictet, Matthias Schwarz, and Claude Morgenegg. Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. *Journal of Banking & Finance*, 14(6):1189–1208, 1990.

[29] Per A Mykland and Lan Zhang. Inference for volatility-type objects and implications for hedging. *Statistics and its Interface*, 1(2):255–278, 2008.

[30] Per A Mykland and Lan Zhang. The econometrics of high frequency data. In *Kessler, M., Lindner, A., Sørensen, M.(Eds.), Statistical Methods for Stochastic Differential Equations*, pages 109–190. Chapman and Hall/CRC Press New York, 2012.

[31] Per A Mykland and Lan Zhang. Assessment of uncertainty in high frequency data: The observed asymptotic variance. *Econometrica*, 85(1):197–231, 2017.

[32] Stephen J Taylor and Xinzhong Xu. The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance*, 4(4):317–340, 1997.

[33] Viktor Todorov. Estimation of continuous-time stochastic volatility models with jumps using high-frequency data. *Journal of Econometrics*, 148(2):131–148, 2009.

[34] Robert A Wood, Thomas H McInish, and J Keith Ord. An investigation of transactions data for nyse stocks. *The Journal of Finance*, 40(3):723–739, 1985.