

Detection of a structural break in intraday volatility pattern

Piotr Kokoszka^{a,*}, Tim Kutta^a, Neda Mohammadi^a, Haonan Wang^a, Shixuan Wang^b

^a Colorado State University, USA

^b University of Reading, United Kingdom

ARTICLE INFO

MSC:

62R10

62G10

62M10

Keywords:

Change point

Functional data

Intraday volatility

Itô integral

ABSTRACT

We develop theory leading to testing procedures for the presence of a change point in the intraday volatility pattern. The new theory is developed in the framework of Functional Data Analysis. It is based on a model akin to the stochastic volatility model for scalar point-to-point returns. In our context, we study intraday curves, one curve per trading day. After postulating a suitable model for such functional data, we present three tests focusing, respectively, on changes in the shape, the magnitude and arbitrary changes in the sequences of the curves of interest. We justify the respective procedures by showing that they have asymptotically correct size and by deriving consistency rates for all tests. These rates involve the sample size (the number of trading days) and the grid size (the number of observations per day). We also derive the corresponding change point estimators and their consistency rates. All procedures are additionally validated by a simulation study and an application to US stocks.

1. Introduction

Consider a sample of intraday price curves $\{P_i(t), t \in [0, 1]\}$, $1 \leq i \leq N$, where i indexes the trading day and t is intraday time normalized to the standard unit interval. For each i , we study the limits, as $\Delta \rightarrow 0$, of cumulative intraday realized volatility curves

$$RV_i(\Delta)(t) = \sum_{1 \leq k \leq Kt} |\log[P_i(k\Delta)] - \log[P_i((k-1)\Delta)]|^2, \quad t \in [0, 1]. \quad (1.1)$$

To illustrate, five consecutive curves $RV_i(\Delta)(\cdot)$ are shown in Fig. 1.

Under suitable assumptions, see Section 2, for each $t \in [0, 1]$,

$$RV_i(\Delta)(t) \xrightarrow{P} \int_0^t v_i^2(u) du, \quad \text{as } \Delta \rightarrow 0.$$

The object of our study are basically the curves v_i , but a more precise problem statement is needed. We represent the curves v_i as

$$v_i(u) = h_i \sigma_i(u), \quad u \in [0, 1], \quad i \in \mathbb{Z}, \quad (1.2)$$

where the $h_i > 0$ describe the evolution of the curves from day to day (between-day volatility), while the functions σ_i quantify the residual volatility after the between days volatility has been accounted for by the sequence $\{h_i\}$. The identifiability of the components in decomposition (1.2) is addressed in Lemma 3.1. We develop a statistical framework to test if the functions σ_i change over a time period of many days. The volatilities h_i typically exhibit persistent magnitude clusters. We propose methodology, and supporting theory, that allows us to test the constancy of the functions σ_i in index i . We emphasize that we do not test if each σ_i

* Correspondence to: Department of Statistics, Colorado State University, Fort Collins, CO 80521-1877, USA.

E-mail address: Piotr.Kokoszka@colostate.edu (P. Kokoszka).




Fig. 1. Five consecutive realized volatility curves computed with $K = 78$ for Tesla Inc. from Dec 27 to Dec 31, 2021.

is a constant function because this is well-known not to be true. The functions σ_i are approximately the derivatives of functions like those in Fig. 1, divided by h_i . The shape of the functions σ_i can change, and such a change will be undetectable if one focuses only on realized daily volatilities $RV_i(\Delta)$ because they are dominated by the h_i . In the following, we refer to the $\sigma_i(\cdot)$ as volatility or diffusion functions, while keeping in mind that the stochastic volatility functions are the products $h_i\sigma_i(\cdot)$.

Model (1.2) is a useful approximation that allows us to construct an effective test and justify it. Using a related perspective, [12] assume that $v_i(u) = h_i(u)\sigma_i(u)$ and propose a test of $H_0 : h_i(u) = h_i, \forall u \in [0, 1]$. They apply it to over thousand days i and to 30 stocks, about 30,000 tests in total. There are overall more rejections than acceptances of their H_0 , the results depend on the implementation of the test. The rejections dominate if the test is implemented with a good estimate of the function σ^2 , under the assumption that it does not change, i.e.

$$\sigma_i^2(\cdot) = \sigma^2(\cdot), \quad \forall i = 1, 2, \dots, N. \quad (1.3)$$

If condition (1.3) does not hold, any estimate of the function $\sigma^2(\cdot)$ may be meaningless. Our test is thus complementary to that of [12] and impacts its implementation; we test assumption (1.3) with an unknown $\sigma^2(\cdot)$. We evaluate and apply the test to five minute intraday returns, so we do not need to be concerned with microstructure noise and price jumps. Our objective is to provide an effective principled way of testing condition (1.3) within a broad framework of statistical change point detection and functional data analysis. Our tests will be useful in any context that requires the verification that an intraday volatility pattern remains constant over the period of many trading days.

Research on the detection and estimation of a change point in various statistical models is over 70 years old and forms a well-established subfield of statistics. Its importance stems from the fact that most statistical models assume a single data generating mechanism, so if this mechanism changes over the observational period, their application will be meaningless. There are several monographs and thousands of research papers; the paper of [21] can serve as a concise and modern introduction to the general framework of this paper. It includes change point detection for different data structures, among them functional data. A broad review of inference and estimation techniques is given in [11], who also consider a range of different applications. The related literature on sequential (online) tests for structural changes is reviewed in [5]. In the following two paragraphs, we briefly review the most closely related research. An important point to note is that in the framework of functional data analysis (FDA), the tests are $\sqrt{N}$ -consistent (N is the sample size), whereas in our framework of replications of a diffusion process, we can obtain $\sqrt{NK}$ -consistency, where K is the number of sampling points for each replication. Moreover, these rates depend on the type of local alternative, shape change vs. size change, and arise because we explicitly use modeling through Itô integrals. A key starting point is new concentration results, Proposition A.1 and its corollaries, that can be used in other contexts that require information about the rates, in terms of grid size, at which population volatility functions can be approximated by realized volatility curves. As far as we know, the framework we study has been considered neither in change point research nor in intraday volatility research. We hope that the theoretical advances we make together with a comprehensive application will motivate further research at the nexus of FDA and SDEs.

Related change point research in the framework of FDA In FDA, the observations are random elements in some function space, such the space L^2 equipped with the canonical L^2 -norm, or C (continuous functions) equipped with the supremum norm. The space L^2 has played a particularly important role in FDA since, under weak assumptions, it is a separable Hilbert space. [8] proposed a test for a change in the mean function in an L^2 setting; extensions were considered by [3,6,17,20], among others. Structural breaks of time series in the space $C([0, 1])$ were studied in [14], see also [33] for a more abstract context. Changes in the covariance of

a functional time series were considered in [37] and in the cross-covariance operator in [34]. [13] considered inference for the covariance kernel of continuous data. More recently, [22] employed a weighted CUSUM statistic for the detection and localization of changes in the covariance operator. Tests for the stability of eigenvalues and principal components were presented in [7,16]. A self-normalization approach, see [35,36], was used in the context of change point detection in functional time series by [15]. The cited works comprise only a small fraction of the relevant literature.

Change point detection in Itô semimartingales We are not aware of any research that considers a change point problem in a sample of trajectories, each of which is an Itô semimartingale. Research to date has focused on the detection of a change point in a single continuous-time realization. [1] modified the commonly used CUSUM approach to detect jumps in Itô semimartingales. In particular, in order to detect jumps in asset returns, they proposed a test statistic based on multiplicative difference of realized truncated p -th variation. [10] used a similar approach to detect structural changes in the volatility of Itô semimartingales. They addressed detection of jumps, the so called local changes, as well as changes in the roughness of sample path, the so called global changes. Other related papers are [18,19], and [9] who provide further references to the general area of change point detection in Itô semimartingales.

The remainder of the paper is organized as follows. In Section 2, we collect the minimum required background on Itô integrals and the FDA. Section 3 is dedicated to the precise formulation of the problem outlined above. Testing and estimation approaches are developed in Section 4. Their finite sample properties are examined in Section 5. Section 6 contains an application to sequences of intraday returns on US stocks. The E-component contains proofs of all results stated in Section 4, details of practical implementation of all procedures, and some additional information.

2. Mathematical preliminaries

We begin by providing some mathematical background, beginning with stochastic differential equations. We first recall the definition of the quadratic variation of a stochastic process $\{X(t) : t \in [0, 1]\}$. We assume throughout that $0 = t_0^{(K)} < t_1^{(K)} < \dots < t_K^{(K)} = 1$ is a grid on the interval $[0, 1]$ with $\max_k [t_k - t_{k-1}] \rightarrow 0$, as $K \rightarrow \infty$. Then,

$$\sum_{k=1}^K \left| X(t_k^{(K)}) - X(t_{k-1}^{(K)}) \right|^2 \mathbb{I}_{\{t_k^{(K)} \leq t\}} \xrightarrow{P} [X, X]_t \quad t \in [0, 1], \quad (2.1)$$

The limit $[X, X]_t$ exists for any semimartingale, and is called the quadratic variation at time t , see e.g. Theorem 1.14 and relation (3.23) in [2]. In this work, we assume that the process X is given by the Itô integral

$$X(t) := \int_0^t v(u) dW(u),$$

where W is a standard Wiener process and $v : [0, 1] \rightarrow (0, \infty)$ is a continuous function (for a detailed discussion of the existence and properties of this process, see Theorem 5.2.1 in [31]). It is well-known that for Itô integrals the quadratic variation is given by

$$[X, X]_t := \int_0^t v^2(u) du, \quad t \in [0, 1], \quad (2.2)$$

see e.g. equation (2.1) in [25]. Moreover, in this case, (2.1) can be strengthened to

$$\sup_{0 \leq t \leq 1} \left| \sum_k \left| X(t_k^{(K)}) - X(t_{k-1}^{(K)}) \right|^2 \mathbb{I}_{\{t_k^{(K)} \leq t\}} - \int_0^t v^2(u) du \right| \xrightarrow{P} 0, \quad K \rightarrow \infty, \quad (2.3)$$

see again Theorem 1.14 and relation (3.23) in [2]. Next, we present a consequence of the Dambis–Dubins–Schwarz theorem which states that *any* continuous local martingale can be expressed as a time change of a Brownian motion, see e.g. Section 5.3.2 in [28] for a general statement. In our setting,

$$\left\{ \int_0^t v(u) dW(u), \quad t \in [0, 1] \right\} \stackrel{d}{=} \left\{ W \left(\int_0^t v^2(u) du \right), \quad t \in [0, 1] \right\}, \quad (2.4)$$

where the equality in distribution is in the space $C([0, 1])$ of continuous functions, equipped with the topology of uniform convergence. Identity (2.4) entails

$$\mathbb{E} \left[\int_0^t v(u) dW(u) \right]^2 = \int_0^t v^2(u) du, \quad (2.5)$$

which is a special case of the Itô isometry for a deterministic, square integrable integrand $v(\cdot)$.

In identities (2.4) and (2.5), the Itô process is treated as a random function in $C([0, 1])$. However, in the context of FDA, it is often useful to embed the smaller space of continuous functions in the larger Hilbert space of square integrable functions. More precisely, we define $L^2([0, 1])$ as the space of measurable functions $f : [0, 1] \rightarrow \mathbb{R}$ that satisfy $\int_0^1 f^2(x) dx < \infty$. Equipped with the inner product

$$\langle f, g \rangle := \int_0^1 f(t)g(t)dt, \quad f, g \in L^2([0, 1]),$$

and the induced norm $\|\cdot\|_{L^2}$, $L^2([0, 1])$ becomes a separable Hilbert space, where we identify functions equal almost everywhere. A random function X in $L^2([0, 1])$ is then a measurable map $X : (\Omega, \mathcal{A}, \mathbb{P}) \rightarrow L^2([0, 1])$, where $(\Omega, \mathcal{A}, \mathbb{P})$ is a probability space. If the

first moment of X exists in the sense that $\mathbb{E}\|X\| < \infty$, we can define the expectation $\mu \in L^2([0, 1])$ of X , which is characterized by the identity

$$\mathbb{E}\langle X, f \rangle = \langle \mu, f \rangle \quad \forall f \in L^2([0, 1]).$$

Similarly, if the second moment of X exists, $\mathbb{E}\|X\|^2 < \infty$, we can define the covariance operator $C_X : L^2([0, 1]) \rightarrow L^2([0, 1])$ of X by the identity

$$\langle C_X[f], g \rangle := \mathbb{E}[\langle X - \mu, f \rangle \langle X - \mu, g \rangle] \quad \forall f, g \in L^2([0, 1]).$$

It is known that C_X is a self-adjoint, positive semi-definite, Hilbert–Schmidt operator and as such it can be identified with a square integrable kernel function $c_X : [0, 1]^2 \rightarrow \mathbb{R}$ via

$$C_X[f](x) := \int_0^1 c_X(x, y) f(y) dy \quad \forall f \in L^2([0, 1]).$$

Chapters 10 and 11 of [27] provide a concise introduction to the L^2 framework of FDA. For a comprehensive treatment see [23].

3. Statistical model and problem formulation

Suppressing the superscript (K) , consider the grid $0 = t_0 < t_1 < \dots < t_K = 1$ introduced in Section 2, and the cumulative returns

$$R_i(t_k) = \log[P_i(t_k)] - \log[P_i(0)].$$

The realized volatility curves (1.1) for this grid can be written as

$$RV_i(t) = \sum_{k=1}^K |R_i(t_k) - R_i(t_{k-1})|^2 \mathbb{I}_{\{t_k \leq t\}}.$$

Setting $h_i = \exp(g_i)$ in (1.2), we postulate the model

$$R_i(t) = \exp(g_i) \int_0^t \sigma_i(u) dW_i(u), \quad t \in [0, 1], \quad i \in \mathbb{Z}. \quad (3.1)$$

The $W_i(\cdot)$ are independent standard Wiener processes. The sequence g_i is a centered real-valued, weakly stationary time series independent of $(W_i)_{i \in \mathbb{Z}}$. The following lemma shows that for each i the volatility function $\sigma_i(\cdot)$ depends only on $R_i(\cdot)$, so g_i and $\sigma_i(\cdot)$ are identifiable.

Lemma 3.1. *Suppose g satisfies $\mathbb{E}g = 0$ and is independent of the Wiener process $W(\cdot)$. Setting*

$$R(t) = e^g \int_0^t \sigma(u) dW(u), \quad t \in [0, 1],$$

for a continuous function $\sigma(\cdot)$, we have

$$\int_0^t \sigma^2(u) du = \exp \{ \mathbb{E} \log[R, R]_t \}.$$

Proof. By (2.2), $[R, R]_t = \exp(2g) \int_0^t \sigma^2(u) du$, so

$$\log[R, R]_t = 2g + \log \int_0^t \sigma^2(u) du.$$

Since $\mathbb{E}(g) = 0$, $\mathbb{E} \log[R, R]_t = \log \int_0^t \sigma^2(u) du$, and the claim follows. ■

To test for changes in the volatility functions $\sigma_i(\cdot)$, we propose the following change point model. Let $\theta \in (0, 1)$ be a parameter that locates a potential change in the discrete time index i and let $\sigma_{(1)}, \sigma_{(2)} : [0, 1] \rightarrow (0, \infty)$ denote two continuous volatility functions. We postulate that

$$\begin{cases} \sigma_i(\cdot) = \sigma_{(1)}(\cdot), & \text{for } i \leq \lfloor N\theta \rfloor, \\ \sigma_i(\cdot) = \sigma_{(2)}(\cdot), & \text{for } i > \lfloor N\theta \rfloor. \end{cases} \quad (3.2)$$

A change occurs if $\sigma_{(1)}(\cdot) \neq \sigma_{(2)}(\cdot)$. The testing problem is thus

$$H_0 : \sigma_{(1)}(\cdot) = \sigma_{(2)}(\cdot), \quad \text{vs.} \quad H_A : \sigma_{(1)}(\cdot) \neq \sigma_{(2)}(\cdot). \quad (3.3)$$

A cornerstone of our statistical analysis is the translation of changes in volatility to changes of certain features in the quadratic variation process. This allows us to take advantage of regularities of the quadratic variation process compared to the process $R_i(\cdot)$. Indeed, (2.2) directly entails

$$Q_i(t) := [R_i, R_i]_t = \exp(2g_i) \int_0^t \sigma_i^2(u) du, \quad t \in [0, 1], \quad i = 1, 2, \dots, N. \quad (3.4)$$

This, together with the stationarity of the time series $(g_i)_{i \in \mathbb{Z}}$, indicates that a change in volatility corresponds to a change in the distribution of the quadratic variation process $Q_i(\cdot)$, over index i .

Suppose we observe a sample $R_1, \dots, R_N$. Reflecting practically available data, we assume that the functions R_i are observed at $K + 1$ equidistant points in $[0, 1]$. This means that inference will be based on the matrix of observations

$$\{R_i(k/K) : i = 1, \dots, N, k = 0, \dots, K\}. \quad (3.5)$$

In our theory, we assume that the number of grid points, $K + 1$, as well as the number of curves, N , tend to infinity. In view of approximations (2.1) and (2.3), we consider the realized quadratic variation processes

$$\begin{aligned} \hat{Q}_i(t) &= \sum_{k=1}^K |R_i(k/K) - R_i((k-1)/K)|^2 \mathbb{I}\{k/K \leq t\} \\ &= \exp(2g_i) \sum_{k=1}^K \left| \int_{(k-1)/K}^{k/K} \sigma(u) dW_i(u) \right|^2 \mathbb{I}\{k/K \leq t\}, \quad t \in [0, 1]. \end{aligned} \quad (3.6)$$

as estimators of the $Q_i(\cdot)$ in (3.4). Observe that $\hat{Q}_i(t)$ is equal to the realized volatility function (1.1), with the second line reflecting the assumed model.

Assuming the g_i have exponential moments, a test could be based on the approximation

$$\mathbb{E}[\hat{Q}_i(t)] \approx \mathbb{E}[Q_i(t)] = \mathbb{E}[\exp(2g_i)] \cdot \int_0^t \sigma_i^2(u) du \quad (3.7)$$

which indicates that volatility function changes translate to mean changes in the realized quadratic variation process. Detecting changes in the mean of a functional time series is a well-studied problem, as discussed in Section 1. However, a test based on (3.7), requires the existence of exponential moments of the g_i and is not robust against distributional changes in g_i , which might be mistaken for changes in the volatility functions $\sigma_i(\cdot)$. Moreover, CUSUM based FDA tests are $\sqrt{N}$ -consistent, but we demonstrate that against large classes of common alternatives a much stronger consistency rate of $\sqrt{NK}$ is attainable by some tests we propose. For these reasons, we present in this paper a different, more effective method to test the hypotheses (3.3). Our approach does not require exponential moments of g_i , is more stable against distributional changes (or spurious changes) in the g_i , and benefits from $\sqrt{NK}$ -consistency under typical alternatives. As a first step, we express the hypothesis H_0 in (3.3) in terms of two null hypotheses, $H_0^{(1)}$ and $H_0^{(2)}$, that are together equivalent to H_0 :

$$H_0^{(1)} : \frac{\int_0^t \sigma_{(1)}^2(u) du}{\int_0^1 \sigma_{(1)}^2(u) du} = \frac{\int_0^t \sigma_{(2)}^2(u) du}{\int_0^1 \sigma_{(2)}^2(u) du}, \quad \forall t \in [0, 1], \quad (3.8)$$

$$H_0^{(2)} : \int_0^1 \sigma_{(1)}^2(u) du = \int_0^1 \sigma_{(2)}^2(u) du. \quad (3.9)$$

Heuristically, $H_0^{(1)}$ states that the volatility function does not change its shape, while $H_0^{(2)}$ states that the total volatility stays the same. In Section 4, we formulate statistical tests of $H_0^{(1)}$ and $H_0^{(2)}$ separately, and then combine them to test the H_0 in (3.3).

4. Change point tests

We begin by stating assumptions for our subsequent analysis.

Assumption 4.1.

1. The volatility function $\sigma_{(1)}, \sigma_{(2)} : [0, 1] \rightarrow (0, \infty)$ are continuous.
2. The standard Wiener processes W_i , $i \in \mathbb{Z}$, are independent.
3. The two sequences $(g_i)_{i \in \mathbb{Z}}$ and $(W_i)_{i \in \mathbb{Z}}$ are independent of each other.
4. The time series $(g_i)_{i \in \mathbb{Z}}$ is centered, weakly stationary and satisfies a weak invariance principle of the form

$$\left\{ \frac{1}{\sqrt{N}} \sum_{i=1}^{\lfloor Nx \rfloor} g_i : x \in [0, 1] \right\} \xrightarrow{d} \{\lambda W(x) : x \in [0, 1]\},$$

where W is a standard Wiener process and $\lambda^2 > 0$ denotes the long-run variance.

Assumption 4.1 is satisfied in many different scenarios. It basically postulates a very general functional stochastic volatility model. Condition 1 (before and after the change) is common in the literature on diffusion processes and intuitive in our setting. Conditions 2 and 3 determine the dependence structure along our functional time series, which is moderated by the scaling factors e^{g_i} . In [26], the g_i follow an AR(p) model, but for our theory the precise dependence structure is immaterial. If the dependence is sufficiently weak, the partial sum process on the left-hand side of condition 4 converges to a Wiener process. This is true under a multitude of dependence conditions, see e.g. [29], so instead of choosing some of them, we postulate the general condition 4.

4.1. Inference for a shape change

Recall the hypothesis $H_0^{(1)}$ in (3.8). We begin with the simple observation that, according to (3.4), the functions in (3.8) can be represented by the *standardized quadratic variation* as follows:

$$F_i(t) := \frac{Q_i(t)}{Q_i(1)} = \frac{\int_0^t \sigma_{(j)}^2(u) du}{\int_0^1 \sigma_{(j)}^2(u) du}, \quad \text{with } j = \begin{cases} 1, & \text{for } i = 1, \dots, \lfloor N\theta \rfloor, \\ 2, & \text{for } i = \lfloor N\theta \rfloor + 1, \dots, N. \end{cases} \quad (4.1)$$

This motivates using for statistical inference the empirical versions:

$$\begin{aligned} \hat{F}_i(t) &:= \frac{\hat{Q}_i(t)}{\hat{Q}_i(1)} := \frac{\sum_{k=1}^K |R_i(k/K) - R_i((k-1)/K)|^2 \mathbb{I}\{k/K \leq t\}}{\sum_{k=1}^K |R_i(k/K) - R_i((k-1)/K)|^2} \\ &= \frac{\sum_{k=1}^K \left| \int_{(k-1)/K}^{k/K} \sigma_i(u) dW_i(u) \right|^2 \mathbb{I}\{k/K \leq t\}}{\sum_{k=1}^K \left| \int_{(k-1)/K}^{k/K} \sigma_i(u) dW_i(u) \right|^2}. \end{aligned} \quad (4.2)$$

Remark 4.1. We highlight two useful properties of $\hat{F}_i$:

- (i) $\hat{F}_i$ is monotonically increasing with $\hat{F}_i(0) = 0$ and $\hat{F}_i(1) = 1$ and in particular it is a random cdf (and thus measurable). It can be interpreted as a random function, mapping into the space $L^2[0, 1]$ of square integrable functions on the unit interval.
- (ii) The functions $\hat{F}_1, \dots, \hat{F}_N$ are independent, and they do not involve the g_i .

Property (ii) implies that any test statistic based on the $\hat{F}_i$ s will be unaffected by the structure of the g_i , or a potentially misspecified model for them.

Lemma 4.1. Suppose that Conditions 1 and 2 of Assumption 4.1 hold. Then, each $\hat{F}_i$ is a consistent estimator of the standardized quadratic variation F_i (defined in (4.1)) and satisfies a functional central limit theorem of the form

$$\sqrt{K} \{ \hat{F}_i(\cdot) - F_i(\cdot) \} \xrightarrow{d} Z(\cdot) \quad (4.3)$$

where Z is a centered, Gaussian process in $L^2([0, 1])$, with distribution depending on the volatility function σ_i .

The proof of Lemma 4.1 follows by an application of Theorem 5.3.6 in [24] together with the continuous mapping theorem. In view of the convergence in (4.3), we expect a test statistic based on $\hat{F}_1, \dots, \hat{F}_N$ to have variance of order $O(1/(NK))$, or a corresponding test for $H_0^{(1)}$ to be $\sqrt{NK}$ -consistent.

To test $H_0^{(1)}$, we use the CUSUM statistic

$$\hat{S}^{(1)} := \frac{K}{N^2} \sum_{n=1}^N \int_0^1 \left(\sum_{i=1}^n \hat{F}_i(u) - \frac{n}{N} \sum_{i=1}^N \hat{F}_i(u) \right)^2 du. \quad (4.4)$$

In the following result, the asymptotics “ $N, K \rightarrow \infty$ ” should be understood in terms of a sequence $K = K_N$ of natural numbers that diverges as $N \rightarrow \infty$. We do not impose any restrictions on the growth rate of K relative to N , making our method valid regardless of the interplay between K and N .

Theorem 4.1. Suppose that Conditions 1 and 2 of Assumption 4.1 hold and that $N, K \rightarrow \infty$. Then, under $H_0^{(1)}$, the weak convergence

$$\hat{S}^{(1)} \xrightarrow{d} S^{(1)} := \sum_{\ell=1}^{\infty} \lambda_{\ell} \int_0^1 \mathbb{B}_{\ell}(u)^2 du \quad (4.5)$$

holds, where $(\mathbb{B}_{\ell})_{\ell \in \mathbb{N}}$ is a sequence of i.i.d. Brownian bridges and $(\lambda_{\ell})_{\ell \in \mathbb{N}}$ the collection of eigenvalues of the asymptotic covariance kernel

$$c_F(u, v) := \lim_{K \rightarrow \infty} K \cdot \mathbb{E} \left[\{ \hat{F}_1(u) - \mathbb{E}[\hat{F}_1(u)] \} \{ \hat{F}_1(v) - \mathbb{E}[\hat{F}_1(v)] \} \right]. \quad (4.6)$$

Moreover, if $H_0^{(1)}$ is violated, $\hat{S}^{(1)} \xrightarrow{\mathbb{P}} \infty$.

Denoting for any $\alpha \in (0, 1)$ the upper α -quantile of $S^{(1)}$ by $q_{1-\alpha}^{(1)}$, the decision

$$\text{reject if } \hat{S}^{(1)} > q_{1-\alpha}^{(1)}$$

yields a consistent asymptotic level α test of $H_0^{(1)}$. While in practice, we do not know the distribution of $S^{(1)}$, it is uniquely determined by the eigenvalues of c_F , which can be estimated by off-the-shelf methods (we provide details in Appendix C in the E-component). An explicit formula for the kernel $c_F(u, v)$ is given in Theorem B.1, but it is not needed to estimate the λ_{ℓ} because (4.6) is a limit of covariance kernels, and many FDA packages output their eigenvalues. The distribution of the integral in (4.5) is easy to simulate, and it is fairly well-known how to compute the approximate quantiles of the right-hand side of (4.5).

In the next theorem, we demonstrate the consistency of our test procedure against local alternatives.

Theorem 4.2. Suppose Conditions 1 and 2 of Assumption 4.1 hold. Let $\tilde{\sigma} : [0, 1] \rightarrow (0, \infty)$ be a continuous function such that $\tilde{\sigma}(\cdot)/\sigma_{(1)}(\cdot)$ is not constant and let $(a_N)_{N \in \mathbb{N}}$ be a bounded sequence of positive numbers. Then, defining $\sigma_{(2)} := \sigma_{(1)} + a_N \tilde{\sigma}$ and imposing the growth conditions $a_N \sqrt{NK} \rightarrow \infty$ and $a_N K \rightarrow \infty$, it follows that

$$\lim_{N, K \rightarrow \infty} \mathbb{P}(\hat{S}^{(1)} > c) = 1 \quad \forall c \geq 0.$$

Finally, we define the change point estimator

$$\hat{\theta}^{(1)} := \frac{1}{N} \operatorname{argmax}_{n \in \{1, \dots, N\}} \int_0^1 \left(\sum_{i=1}^n \hat{F}_i(u) - \frac{n}{N} \sum_{i=1}^N \hat{F}_i(u) \right)^2 du. \quad (4.7)$$

If the hypothesis of no change in the shape of volatility is violated, i.e.

$$\begin{cases} F_i(\cdot) = F_{(1)}(\cdot), & \text{for } i \leq \lfloor N\theta \rfloor \\ F_i(\cdot) = F_{(2)}(\cdot), & \text{for } i > \lfloor N\theta \rfloor, \end{cases} \quad F_{(1)}(\cdot) \neq F_{(2)}(\cdot) \quad (4.8)$$

for some $\theta \in (0, 1)$, we can show that the estimator $\hat{\theta}^{(1)}$ in (4.7) is consistent under local alternatives.

Theorem 4.3. Under the assumptions of Theorem 4.2,

$$\hat{\theta}^{(1)} - \theta = \mathcal{O}_P \left(\max \left\{ \frac{a_N^{-2}}{NK}, \frac{1}{N} \right\} \right),$$

where $\theta \in (0, 1)$ is the rescaled time of the change in (4.8).

The rate in Theorem 4.3 can be explained as follows: For a change of size a_N (potentially tending to 0), it is well-known that an optimal approximation rate is given by

$$\hat{\theta}^{(1)} - \theta = \mathcal{O}_P \left(\frac{a_N^{-2}}{\text{sample size}} \right).$$

In our case the “sample size” is NK , yielding a rate of $\mathcal{O}_P(a_N^{-2}/(NK))$. However, since the number of curves in discrete time is N , it is also clear that a convergence rate cannot be faster than $\mathcal{O}_P(1/N)$. This limitation is simply due to the discretization of time in N steps. As a consequence, the best attainable rate is as specified in Theorem 4.3. Notice that in the special case of $a_N = O(1/\sqrt{K})$, we obtain the optimal rate $\mathcal{O}_P(1/N)$ on the right-hand side, the same rate as for fully observed functions, see e.g. [4].

4.2. Inference for a change in total volatility

Recall the hypothesis $H_0^{(2)}$ in (3.9). The integrals in (3.9) are closely related to the total quadratic variation $Q_i(1)$ and taking its logarithm, we obtain

$$\log(Q_i(1)) = 2g_i + \log \left(\int_0^1 \sigma_{(j)}^2(u) du \right), \quad \text{for } \begin{cases} i = 1, \dots, \lfloor N\theta \rfloor, & j = 1, \\ i = \lfloor N\theta \rfloor + 1, \dots, N, & j = 2. \end{cases} \quad (4.9)$$

Since the g_i are centered, any change in total volatility translates into a mean change of the real-valued time series $\{\log(Q_i(1))\}$. An empirical analogue of (4.9) is

$$\log(\hat{Q}_i(1)) = 2g_i + w_i, \quad (4.10)$$

where

$$w_i := \log \left(\sum_{k=1}^K \left| \int_{(k-1)/K}^{k/K} \sigma_i(u) dW_i(u) \right|^2 \right). \quad (4.11)$$

This decomposition shows that the observations $\log(\hat{Q}_1(1)), \dots, \log(\hat{Q}_N(1))$ form (for any fixed K) a dependent time series that is stationary before and after a potential change. For the purpose of statistical inference, we use the following CUSUM statistics:

$$\hat{S}^{(2)} := \frac{1}{N^2} \sum_{n=1}^N \left(\sum_{i=1}^n \log(\hat{Q}_i(1)) - \frac{n}{N} \sum_{i=1}^N \log(\hat{Q}_i(1)) \right)^2. \quad (4.12)$$

Theorem 4.4. If Assumption 4.1 holds and $N, K \rightarrow \infty$, then, under $H_0^{(2)}$,

$$\hat{S}^{(2)} \xrightarrow{d} S^{(2)} := (4\lambda) \cdot \int_0^1 \mathbb{B}(u)^2 du, \quad (4.13)$$

where $\mathbb{B}$ is a standard Brownian bridge and the long-run variance λ is defined as

$$\lambda := \sum_{i \in \mathbb{Z}} \operatorname{Cov}(g_0, g_i). \quad (4.14)$$

Moreover, if $H_0^{(2)}$ is violated, $\hat{S}^{(2)} \xrightarrow{\mathbb{P}} \infty$.

The fact that the long-run variance λ only depends on the g_i s is not an accident. As we will see in the next section, the statistic $\hat{S}^{(2)}$ is asymptotically only dependent on the g_i s and independent of the W_i s. This implies that $\hat{S}^{(1)}$ (which does not depend on the g_i s) and $\hat{S}^{(2)}$ are asymptotically independent.

Theorem 4.4 implies that if we denote by $q_{1-\alpha}^{(2)}$ the upper α -quantile of the distribution $S^{(2)}$, then the decision to

$$\text{reject if } \hat{S}^{(2)} > q_{1-\alpha}^{(2)}$$

yields a consistent asymptotic level α test of the hypothesis $H_0^{(2)}$. Again, $q_{1-\alpha}^{(2)}$ cannot be directly computed, but it can be approximated, if a consistent estimator for the long-run variance is given (see Appendix C in the E-component).

We now show consistency of the test against local alternatives.

Theorem 4.5. Suppose [Assumption 4.1](#) holds and $(a_N)_{N \in \mathbb{N}}$ is a bounded sequence of positive numbers. Then, defining $\sigma_{(2)} := (1 + a_N)\sigma_{(1)}$ and imposing the growth conditions $a_N \sqrt{N} \rightarrow \infty$ and $a_N K \rightarrow \infty$, it follows that

$$\lim_{N, K \rightarrow \infty} \mathbb{P}(\hat{S}^{(2)} > c) = 1, \quad \forall c \geq 0.$$

Finally, with the change point estimator

$$\hat{\theta}^{(2)} := \frac{1}{N} \operatorname{argmax}_{n \in \{1, \dots, N\}} \left(\sum_{i=1}^n \log(\hat{Q}_i(1)) - \frac{n}{N} \sum_{i=1}^N \log(\hat{Q}_i(1)) \right)^2, \quad (4.15)$$

we can localize a change in total volatility. If the hypothesis of no change in the total volatility is violated, i.e.

$$\begin{cases} \log(Q_i(1)) = \log(Q_{(1)}(1)), & \text{for } i \leq \lfloor N\theta \rfloor \\ \log(Q_i(1)) = \log(Q_{(2)}(1)), & \text{for } i > \lfloor N\theta \rfloor, \end{cases} \quad \log(Q_{(1)}(1)) \neq \log(Q_{(2)}(1)), \quad (4.16)$$

for some $\theta \in (0, 1)$, we obtain the following result.

Theorem 4.6. Under the assumptions of [Theorem 4.5](#), $\hat{\theta}^{(2)} - \theta = O_p(a_N^{-2}/N)$.

Taking $a_N = 1$, we obtain the optimal rate.

4.3. Inference for an arbitrary change

In the previous subsections, we have developed test statistics $\hat{S}^{(1)}, \hat{S}^{(2)}$ for the null hypotheses $H_0^{(1)}, H_0^{(2)}$ in (3.8) and (3.9), respectively. We now want to combine these two tests to yield a test for the global null hypothesis H_0 in (3.3). As a first step, we show that as $N, K \rightarrow \infty$, the two test statistics (4.4) and (4.12) become independent of each other.

Proposition 4.1. If [Assumption 4.1](#) and H_0 in (3.3) hold, then, as $N, K \rightarrow \infty$,

$$(\hat{S}^{(1)}, \hat{S}^{(2)}) \xrightarrow{d} (S^{(1)}, S^{(2)}),$$

where $S^{(1)}, S^{(2)}$ are independent and defined in (4.5), (4.13), respectively.

In order to combine the results from both test statistics, we employ their asymptotic p -values. To be precise, if $\Lambda^{(j)}$ is the (continuous) cumulative distribution function of $S^{(j)}$, we define the p -values

$$p^{(j)} = 1 - \Lambda^{(j)}(\hat{S}^{(j)}), \quad j = 1, 2. \quad (4.17)$$

In practice, the $\Lambda^{(j)}$ are not known, but can uniformly approximated, yielding empirical p -values. We discuss this issue in Appendix C in the E-component. To combine our test statistics, we recall that under H_0 , both p -values $p^{(1)}, p^{(2)}$ are asymptotically uniformly distributed on $[0, 1]$ and according to [Proposition 4.1](#) asymptotically independent. Hence, using Fisher's method, see e.g. [32], we can combine them to

$$\hat{S} := -2\{\log(p^{(1)}) + \log(p^{(2)})\}, \quad (4.18)$$

which then converges under H_0 to a chi-squared distribution with four degrees of freedom. Denoting the upper α -quantile of this distribution by $q_{1-\alpha}$, gives us the test decision

$$\text{reject if } \hat{S} > q_{1-\alpha}. \quad (4.19)$$

We make this result precise in the following proposition.

Proposition 4.2. Under the assumptions of [Proposition 4.1](#),

$$\hat{S} \xrightarrow{d} \chi_4^2,$$

where χ_4^2 is a chi-squared distribution with four degrees of freedom. If H_0 is violated, $\hat{S} \xrightarrow{P} \infty$.

It is a simple consequence of [Theorems 4.2](#) and [4.5](#) that the test [\(4.19\)](#) is consistent against local alternatives of shape changes and changes in total volatility, with the rates discussed in those theorems.

Remark 4.2. In view of [Proposition 4.1](#), there are different ways of combining the test statistics $\hat{S}^{(1)}, \hat{S}^{(2)}$ for a joint test, apart from our choice of $\hat{S}$. Such combinations correspond to different rejection regions in $\mathbb{R}_{\geq 0}^2$ for $(\hat{S}^{(1)}, \hat{S}^{(2)})$. Generically, we can define for a function $f : \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{\geq 0}$ the combined statistic $\hat{S}_f = f(\hat{S}^{(1)}, \hat{S}^{(2)})$. A simple choice might be a sum $f_{\text{sum}}(x, y) = x + y$, which has linear, downward sloping contour lines and thus triangular rejection regions. Our choice

$$f_{\text{Fisher}}(x, y) = -2 \left\{ \log(1 - \Lambda^{(1)}(x)) + \log(1 - \Lambda^{(2)}(y)) \right\}$$

has astroid shaped contour lines (like a p -norm with $0 < p < 1$). Accordingly rejection regions are shaped like ellipsoids. The precise shape of the contour lines depends on the asymptotic distributions $\Lambda^{(1)}, \Lambda^{(2)}$. The function f implies how evidence against the null hypothesis is interpreted in different scenarios. Roughly speaking, f_{sum} is indifferent between large x , large y or large $x + y$. This means that more evidence against the null might come just as well from one statistic, or the other, or their sum. In contrast f_{Fisher} is largest if both x and y are large, treating evidence against the null hypothesis as strongest, when it comes from both statistics and weaker if it only comes from one.

Finally, we discuss the problem of change point localization. For this purpose, we introduce the pooled change point estimator

$$\hat{\theta} := \frac{p^{(1)}}{p^{(1)} + p^{(2)}} \hat{\theta}^{(2)} + \frac{p^{(2)}}{p^{(1)} + p^{(2)}} \hat{\theta}^{(1)}. \quad (4.20)$$

Intuitively, $\hat{\theta}$ combines information from both estimators $\hat{\theta}^{(1)}, \hat{\theta}^{(2)}$, putting priority on the one where the change is more pronounced (smaller p -value). Our proof rests on a careful investigation of the tail behavior of the distributions $\Lambda^{(1)}, \Lambda^{(2)}$, see [Theorem B.3](#) in the E-component. The tail behavior of these distributions determines the relative size of the p -values $p^{(1)}, p^{(2)}$ in the above weights.

Proposition 4.3. Suppose [Assumption 4.1](#) holds, $K \rightarrow \infty$, $K/N \rightarrow 0$, and the continuous function $\bar{\sigma} : [0, 1] \rightarrow (0, \infty)$ is such that $\bar{\sigma}(\cdot)/\sigma_{(1)}(\cdot)$ is not constant.

(i) If only $H_0^{(2)}$ is violated with $\sigma_{(2)} = (1 + 1/\sqrt{K})\sigma_{(1)}$, then

$$|\hat{\theta} - \theta| = \mathcal{O}_P\left(\frac{K}{N}\right).$$

(ii) If, in addition, $H^{(1)}$ is violated with $\sigma_{(2)} = (1 + 1/\sqrt{K})\sigma_{(1)} + \bar{\sigma}/\sqrt{K}$, then

$$|\hat{\theta} - \theta| = \mathcal{O}_P\left(\frac{1}{N}\right).$$

5. Finite sample properties

5.1. Empirical size

We generate data under the null hypothesis according to the Functional Stochastic Volatility Model of [\[26\]](#):

$$R_i(t) = \exp(g_i) \int_0^t \sigma(u) dW_i(u), \quad t \in [0, 1], \quad i = 1, \dots, N,$$

$$g_i = \varphi g_{i-1} + \varepsilon_i, \quad \varepsilon_i \sim i.i.d. \mathcal{N}(0, \sigma_\varepsilon^2).$$

There are a number of settings to be carefully chosen:

- Following [\[26\]](#), we set $\varphi = 0.55$ and $\sigma_\varepsilon^2 = 0.25$ in order to reflect real-world data.
- We have four settings of $\sigma(\cdot)$
 - **Flat:** $\sigma(u) = 0.2$. This is a simple case that we have the same intraday volatility throughout the day.
 - **Slope:** $\sigma(u) = 0.1 + 0.2u$. This is a case that the intraday volatility is increasing in a linear manner.
 - **Sine:** $\sigma(u) = 0.1 \sin(2\pi u) + 0.2$. This is the case we have higher volatility in the morning, but lower volatility in the afternoon.
 - **U-shape:** $\sigma(u) = (u - 0.5)^2 + 0.1145299$. This choice is the most relevant one because it reflects the stylized fact that volatility is typically higher at the beginning and the end of a trading day.

The coefficients in $\sigma(\cdot)$ are set to ensure that the above four $\sigma(\cdot)$ have a similar scale.

- The continuous time t in $[0, 1]$ is discretized as $[t_0, t_1, \dots, t_K]$, where $t_k = k\Delta$ and $k = 1, \dots, K$. This is the same for all random curves.
- The number of intraday observations is $K = 26, 39, 78$, which corresponds to 15-min, 10-min 5-min sampling intervals in our data analysis respectively. Their corresponding stepsizes are $\Delta = 1/26, 1/39, 1/78$.
- The sample size is $N = 100, 200, 500$.

Table 1
Empirical size.

Flat		Shape of volatility			Total volatility			Global		
		10%	5%	1%	10%	5%	1%	10%	5%	1%
$N = 100$	$K = 26$	11.4%	5.9%	1.4%	10.5%	5.1%	0.5%	11.4%	5.4%	1.2%
	$K = 39$	10.9%	5.8%	1.2%	10.5%	4.7%	0.4%	10.6%	5.0%	1.0%
	$K = 78$	11.9%	5.9%	0.9%	9.8%	4.3%	0.4%	11.0%	5.0%	0.9%
$N = 200$	$K = 26$	10.6%	5.3%	1.1%	10.5%	5.2%	0.9%	10.8%	5.6%	1.2%
	$K = 39$	10.8%	5.5%	1.3%	10.6%	5.1%	0.7%	11.3%	5.3%	1.0%
	$K = 78$	11.8%	5.4%	1.2%	10.1%	5.0%	0.9%	11.2%	5.4%	0.9%
$N = 500$	$K = 26$	11.0%	5.6%	1.0%	10.4%	5.6%	1.1%	11.0%	5.9%	1.1%
	$K = 39$	11.2%	5.5%	1.2%	11.0%	5.1%	0.8%	11.5%	5.6%	1.0%
	$K = 78$	11.1%	5.6%	1.3%	10.2%	4.7%	0.9%	10.7%	5.5%	1.2%
Slope										
$N = 100$	$K = 26$	11.3%	5.9%	1.4%	10.4%	4.6%	0.4%	11.1%	5.3%	0.9%
	$K = 39$	10.9%	5.3%	1.2%	10.0%	4.3%	0.4%	10.7%	5.2%	1.0%
	$K = 78$	10.5%	5.6%	1.4%	9.5%	4.0%	0.6%	10.5%	5.2%	0.8%
$N = 200$	$K = 26$	10.3%	5.4%	1.1%	10.9%	5.4%	1.1%	11.3%	5.9%	1.2%
	$K = 39$	11.7%	5.9%	1.3%	9.8%	4.8%	0.7%	11.1%	5.4%	1.0%
	$K = 78$	11.2%	6.1%	1.2%	10.7%	5.1%	0.5%	11.2%	5.8%	1.1%
$N = 500$	$K = 26$	10.7%	5.1%	0.9%	10.7%	5.5%	1.2%	11.1%	5.6%	0.8%
	$K = 39$	11.3%	5.6%	1.2%	10.3%	5.0%	1.0%	11.3%	5.6%	1.1%
	$K = 78$	11.2%	5.6%	1.1%	10.2%	5.1%	0.9%	10.9%	5.4%	1.1%
Sine										
$N = 100$	$K = 26$	11.3%	5.6%	1.1%	11.0%	5.5%	0.7%	11.6%	5.3%	0.8%
	$K = 39$	11.3%	5.9%	1.2%	10.8%	5.1%	0.7%	11.6%	5.7%	1.0%
	$K = 78$	11.8%	6.4%	1.5%	9.6%	4.6%	0.6%	11.5%	5.3%	0.9%
$N = 200$	$K = 26$	10.7%	5.1%	1.3%	11.6%	6.2%	1.2%	11.5%	6.2%	1.4%
	$K = 39$	11.1%	5.8%	1.3%	10.3%	4.8%	0.8%	11.5%	5.5%	1.1%
	$K = 78$	10.8%	5.5%	1.0%	10.3%	5.0%	0.8%	11.3%	5.4%	0.8%
$N = 500$	$K = 26$	10.6%	5.2%	1.0%	11.0%	5.5%	1.0%	11.2%	5.8%	0.9%
	$K = 39$	11.4%	5.3%	1.1%	9.8%	4.9%	0.8%	11.1%	5.7%	1.2%
	$K = 78$	11.1%	5.8%	1.0%	10.0%	5.0%	0.6%	11.1%	4.9%	0.8%
U-shape										
$N = 100$	$K = 26$	11.0%	5.7%	1.2%	10.6%	4.7%	0.5%	11.0%	5.8%	1.0%
	$K = 39$	11.3%	6.1%	1.3%	11.0%	4.9%	0.4%	11.4%	5.2%	1.1%
	$K = 78$	10.9%	5.8%	1.4%	10.1%	4.3%	0.4%	10.7%	5.4%	1.1%
$N = 200$	$K = 26$	11.0%	5.9%	1.3%	11.0%	5.4%	0.9%	11.4%	6.1%	1.2%
	$K = 39$	11.2%	6.0%	1.1%	10.5%	5.4%	1.0%	11.2%	5.8%	1.1%
	$K = 78$	11.4%	6.0%	1.4%	10.4%	5.0%	0.8%	11.2%	6.3%	1.2%
$N = 500$	$K = 26$	11.0%	5.5%	1.1%	11.3%	5.2%	0.8%	11.0%	5.8%	1.0%
	$K = 39$	9.7%	4.9%	0.8%	10.4%	5.4%	1.0%	10.3%	4.9%	1.0%
	$K = 78$	10.6%	5.2%	1.1%	10.6%	5.2%	1.1%	11.0%	5.7%	1.0%

Details on the computation of $\int_0^t \sigma(u) dW(u)$ and $\int_0^1 \mathbb{B}^2(u) du$, both use special approaches, are presented in Section D, which also contains step-by-step formulas for the computation of the three test statistics. The long-run variance of the $\log \hat{Q}_t(1)$ was computed using the Bartlett kernel with bandwidth selected by the procedure of [30] with prewhitening.

Table 1 provides the empirical sizes of the three tests under four different shapes of $\sigma(\cdot)$. We see that the test performs very well, even for fairly small sample sizes N and low resolution K .

One advantage of using our tests is that it is robust against changes in g_i , which should not be mistaken as changes in the volatility function $\sigma_i(\cdot)$. To verify this property, we consider

$$g_i = \begin{cases} 0.45g_{i-1} + \varepsilon_i, & \varepsilon_i \sim i.i.d. \mathcal{N}(0, \sigma_\varepsilon^2), & i = 1, \dots, \lfloor N/20 \rfloor, \\ 0.65g_{i-1} + \varepsilon_i, & \varepsilon_i \sim i.i.d. \mathcal{N}(0, \sigma_\varepsilon^2), & i = \lfloor N/20 \rfloor + 1, \dots, N, \end{cases}$$

and all other settings are the same as before. Table 2 presents the empirical sizes of the three tests under the U-Shaped $\sigma_i(\cdot)$. The other three shapes yield similar results. As can be seen, the empirical sizes of our three tests are not affected by the change in g_i and match their theoretical levels reasonably well.

5.2. Empirical power

We set the time of the change at $\theta = 0.25, 0.5, 0.75$ and consider $N = 250$ and $N = 500$. All other settings are the same as under the null.

Table 2
Empirical size under a change in g_t .

		Shape of volatility			Total volatility			Global		
		10%	5%	1%	10%	5%	1%	10%	5%	1%
$N = 100$	$K = 26$	12.0%	6.0%	1.4%	11.9%	5.8%	0.8%	12.2%	6.0%	1.3%
	$K = 39$	11.4%	6.0%	1.2%	10.9%	5.3%	0.5%	11.9%	5.4%	1.0%
	$K = 78$	10.7%	5.6%	1.1%	11.2%	4.8%	0.5%	11.0%	5.1%	0.8%
$N = 200$	$K = 26$	11.6%	6.3%	1.2%	12.8%	6.7%	1.2%	13.3%	6.7%	1.2%
	$K = 39$	10.4%	5.1%	1.0%	12.0%	6.0%	1.2%	11.6%	5.7%	1.3%
	$K = 78$	10.6%	5.0%	0.9%	11.2%	5.3%	0.9%	10.7%	5.5%	1.1%
$N = 500$	$K = 26$	10.8%	5.3%	1.0%	12.0%	6.2%	1.2%	11.5%	5.8%	1.3%
	$K = 39$	11.6%	5.9%	1.2%	11.6%	6.5%	1.4%	12.9%	7.2%	1.3%
	$K = 78$	11.3%	5.6%	1.3%	11.6%	6.3%	1.1%	12.3%	6.3%	1.3%

There are unlimited possibilities for a change in $\sigma_t(\cdot)$. To focus on the scenarios emphasized in this paper, we consider the following three alternative hypotheses:

1. $H_{A,1}$: a shape change in volatility, but no change in total volatility,
2. $H_{A,2}$: a change in total volatility, but no change in the shape of volatility,
3. $H_{A,3}$: a simultaneous change in the shape of volatility and total volatility.

Under $H_{A,1}$, we have a change in $\sigma_t(\cdot)$ from the flat shape to a sine shape. We have noticed that our test is very effective in detecting changes in the shape of volatility, and it can easily get empirical power of 100%. That is why we deliberately choose a very small change in the shape in order to show the convergence of the empirical power with respect to N and K . Specifically, we set

$$\sigma_t(u) = \begin{cases} 0.2, & \text{for } i = 1, \dots, \lfloor N\theta \rfloor, \\ 0.02 \sin(2\pi u) + \sqrt{199/5000}, & \text{for } i = \lfloor N\theta \rfloor + 1, \dots, N. \end{cases}$$

The constant in the sine function is to ensure that the total volatility before and after the change is the same, i.e. $\int_0^1 0.2^2 du = \int_0^1 [0.02 \sin(2\pi u) + \sqrt{199/5000}]^2 du = 0.04$. Thus, there is a change in the shape of volatility, but no change in the total volatility.

Under $H_{A,2}$, we introduce an upward parallel shift of the flat shape:

$$\sigma_t(u) = \begin{cases} 0.2, & \text{for } i = 1, \dots, \lfloor N\theta \rfloor, \\ 0.4, & \text{for } i = \lfloor N\theta \rfloor + 1, \dots, N. \end{cases}$$

Note that an upward parallel shift in the other three shapes (slope, sine, U-shape) will cause a change in total volatility as well as in the shape of volatility. This is because the other three shapes are actually “compressed” due to a higher total volatility.

Under $H_{A,3}$, we have a simultaneous change in shape and total volatility:

$$\sigma_t(u) = \begin{cases} 0.2, & \text{for } i = 1, \dots, \lfloor N\theta \rfloor, \\ (u - 0.5)^2 + 0.4, & \text{for } i = \lfloor N\theta \rfloor + 1, \dots, N. \end{cases}$$

The shape of $\sigma_t(\cdot)$ is changed from flat to U-shape, and total volatility is changed from $\int_0^1 0.2^2 du = 0.04$ to $\int_0^1 [(u - 0.5)^2 + 0.3]^2 du = 0.1525$.

Table 3 reports the empirical power. The conclusions can be summarized as follows.

1. Under $H_{A,1}$, the empirical power of the shape test and the global test increases with of N and K , in agreement with the $\sqrt{NK}$ -consistency we established theoretically. The rejection rate of total volatility test is always around 5%, as expected since there is no change in total volatility in $H_{A,1}$.
2. Under $H_{A,2}$, the empirical power of the total volatility test increases with the growth of N , not with K . This is exactly what we expected because the test on total volatility is $\sqrt{N}$ -consistent. Additionally, the rejection rate of testing the shape is typically around 5%, since there is no change in the shape in $H_{A,2}$.
3. Under $H_{A,3}$, the empirical power of the shape test and the global test increases with the growth of N and K , and empirical power of the volatility test increases with the growth of N , but not with K , again as predicted by our theory.

In Subsection D.2, we show that the change point estimators under the three alternatives inherit the properties of the corresponding tests: the performance of θ_1 and θ improves with increasing N and K , θ_2 improves with increasing N .

6. Application to US stocks

We begin with an individual stock as a prototype analysis to showcase our developed tests. Then, there are two ways to use the developed tests on a larger scale. First, since there could be multiple changes, we use the binary segmentation to explore all changes

Table 3
Empirical power.

$H_{A,1}$		$\theta = 0.25$			$\theta = 0.5$			$\theta = 0.75$		
		Shape	Total	Global	Shape	Total	Global	Shape	Total	Global
$N = 250$	$K = 26$	62.6%	5.2%	51.1%	86.8%	5.3%	78.2%	62.9%	5.6%	51.3%
	$K = 39$	82.6%	5.0%	72.0%	96.9%	5.1%	93.3%	83.5%	5.0%	72.9%
	$K = 78$	98.7%	5.1%	96.8%	100.0%	4.7%	99.9%	99.2%	4.5%	97.1%
$N = 500$	$K = 26$	92.0%	5.5%	84.5%	99.2%	5.1%	97.6%	91.7%	5.5%	83.4%
	$K = 39$	99.1%	4.6%	97.2%	100.0%	5.1%	99.9%	98.8%	5.2%	96.2%
	$K = 78$	100.0%	5.5%	100.0%	100.0%	4.9%	100.0%	100.0%	5.1%	100.0%
$H_{A,2}$										
$N = 250$	$K = 26$	5.8%	86.8%	74.5%	5.5%	98.6%	95.2%	5.7%	86.8%	73.2%
	$K = 39$	5.9%	86.1%	72.7%	6.0%	98.5%	95.1%	6.0%	86.6%	73.0%
	$K = 78$	5.9%	86.0%	72.3%	5.3%	98.3%	94.9%	5.5%	85.7%	71.6%
$N = 500$	$K = 26$	5.6%	99.6%	97.9%	6.1%	100.0%	100.0%	5.2%	99.5%	98.1%
	$K = 39$	5.4%	99.7%	98.4%	5.9%	100.0%	100.0%	5.3%	99.7%	98.3%
	$K = 78$	5.7%	99.6%	98.3%	5.9%	100.0%	100.0%	5.3%	99.6%	98.0%
$H_{A,3}$										
$N = 250$	$K = 26$	89.2%	97.4%	99.9%	99.9%	100.0%	100.0%	95.2%	97.4%	99.9%
	$K = 39$	99.3%	97.3%	100.0%	100.0%	100.0%	100.0%	99.9%	97.7%	100.0%
	$K = 78$	100.0%	97.1%	100.0%	100.0%	99.9%	100.0%	100.0%	97.3%	100.0%
$N = 500$	$K = 26$	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
	$K = 39$	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
	$K = 78$	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

Table 4
Test results of Tesla (note that the p -values are in percent).

	p -value	Change point estimator	Date of change
Shape of volatility (H_0^1)	0.02%	0.26	Jul 1, 2013
Total volatility (H_0^2)	0.96%	0.34	May 12, 2014
Global (H_0)	0.00%	0.26	Jul 8, 2013

for one stock during a sample period. Second, we apply our test procedure to a large number of stocks and present the summary of first detected changes (without the binary segmentation).

For the purpose of demonstration, we focus on Tesla Inc. (Permno: 93436) for our prototype analysis. We consider 5-min intraday prices, the sample period is from Jun 29, 2010 (the IPO date) to Dec 31, 2021, corresponding to $N = 2891$ trading days. In each trading day i , we have the opening price $P_i(t_0)$ and the subsequent 78 5-min intraday prices $P_i(t_k)$, $k = 1, \dots, 78$, with the last trading price in every 5-min time interval. Thus, the equidistant grid on the unit interval is $t_k = k\Delta$, $k = 0, 1, \dots, K$, where $K = 78$ and the step size $\Delta = 1/78$.

Based on the intraday price data, we calculate the cumulative intraday return (CIDR) curves as

$$R_i(t_k) = \log(P_i(t_k)) - \log(P_i(t_0)), \quad k = 1, \dots, K, \quad i = 1, \dots, N.$$

By definition, the CIDR curves always start from zero, i.e. $R_i(t_0) = 0$, and are scale invariant. We also compute the cumulative intraday realized volatility (CIDRV) curves as

$$RV_i(t_k) = \sum_{k=1}^K |R_i(t_k) - R_i(t_{k-1})|^2 \mathbb{I}\{t_k < t\}, \quad k = 1, \dots, K, \quad i = 1, \dots, N.$$

In order to visualize the important functional objects, Fig. 2 plots the intraday Price $P_i(t_k)$, CIDRs $R_i(t_k)$, and CIDRVs $RV_i(t_k)$ in the upper, middle, and lower panels, respectively.

We apply the tests for the whole sample period, in order to detect (1) a shape change, (2) a change in total volatility, and (3) an arbitrary change. Table 4 presents the test results. The p -value of testing H_0^1 is 0.02%, providing strong evidence of a shape change. The change point estimator $\hat{\theta}_1$ is 0.26, indicating the shape change occurred on Jul 1, 2013. As for testing H_0^2 , we find strong evidence of a change in total volatility with p -value of 0.96%. The date of change in total volatility is May 12, 2014, as suggested by the $\hat{\theta}_2 = 0.34$. Combining the two tests, we have the p -value of 0.00% for the global null hypothesis H_0 , with the pooled change point estimator $\hat{\theta} = 0.26$, implying that the date of arbitrary change in intraday volatility pattern is July 8, 2013.

As the sample period of the Tesla analysis covers more than a decade, there could be multiple changes in the intraday volatility pattern. Thus, we use the standard binary segmentation based on the global test at the 5% significance level and the pooled change point estimator ($\hat{\theta}$). Table 5 presents the result with some associated events that could be used to validate the identified change points.

It is also interesting to examine change in the intraday volatility pattern for other stocks. Thus, we apply our test procedure to 7293 stocks in the US stock markets. To preserve space, we focus on the first change detected by the global test at 5% significance




Fig. 2. Time series of functional objects derived from intraday Tesla prices. Upper panel: Intraday price $P_t(t_k)$; Middle panel: CIDRs $R_t(t_k)$; Lower panel: CIDRVs $RV_t(t_k)$.

Table 5
Result of binary segmentation to test multiple changes for Tesla.

<i>p</i> -value	Date of change	Related news
0.00%	Jul 8, 2013	Tesla joined the Nasdaq 100 index on Jul 15, 2013
0.00%	Jul 16, 2014	Tesla announced new smaller electric vehicle named Model 3 on Jul 16, 2014
0.08%	Feb 6, 2018	Elon Musk made history launching a car into space on Feb 6, 2018
0.09%	Jan 23, 2019	Tesla posted back-to-back profits for the first time
4.79%	Dec 20, 2019	Tesla's Chinese factory delivered its first cars
0.48%	Mar 31, 2021	NHTSA confirmed no violation of Tesla's touchscreen drive selector




Fig. 3. Dates of first change in the intraday volatility pattern for 7168 US stocks.

level, without using binary segmentation to find additional changes. The stocks used and the data cleaning procedure are the same as in [26]. Their sample period varies in length from 2 to 25 years. Shorter sample periods could be due to IPO dates later than Jan 3, 2006 or stocks delisted before Dec 31, 2021.

Our test indicates that 7168 out of 7293 stocks (98.3%) underwent at least one change in the intraday volatility pattern. This provides the evidence that change in the intraday volatility pattern is a common issue in the US stocks. To provide further insights, we plot the histogram of the first detected changes in Fig. 3. We can clearly see that (1) the highest frequent changes happen during the subprime mortgage crisis in 2008, (2) the second highest frequent changes occur around the European debt crisis in the 2010s, (3) the third highest frequent changes appear after COVID in 2020. These results show that our test is able to detect change points that are consistent with well-known market events, providing additional validation on a very large data set.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We are grateful for Dr. Jian Chen's help with collecting and cleaning the United States stock data. We acknowledge helpful comments from seminar participants at University of Reading, University of Liverpool, and Shanghai Maritime University, a workshop on Central Bank Digital Currencies and Financial Economics at City University of London, and Financial Econometrics Conference at University of Cambridge. We thank two referees for insightful comments that helped us improve the paper. Piotr Kokoszka and Haonan Wang were partially supported by National Science Foundation, United States grant DMS-2123761.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.spa.2024.104426>.

References

- [1] Y. Aït-Sahalia, J. Jacod, Testing for jumps in a discretely observed process, *Ann. Statist.* 37 (1) (2009) 184–222.
- [2] Yacine Aït-Sahalia, Jean Jacod, *High-Frequency Financial Econometrics*, illustrated ed., Princeton University Press, Princeton, 2014.
- [3] M. Aston, C. Kirch, Detecting and estimating changes in dependent functional data, *J. Multivariate Anal.* 109 (2012) 204–220.
- [4] A. Aue, R. Gabrys, L. Horváth, P. Kokoszka, Estimation of a change-point in the mean function of functional data, *J. Multivariate Anal.* 100 (2009) 2254–2269.
- [5] A. Aue, C. Kirch, The state of cumulative sum sequential change point testing seventy years after page, *Biometrika* (2023).
- [6] A. Aue, G. Rice, O. Sonmez, Detecting and dating structural breaks in functional data without dimension reduction, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 80 (2018) 509–529.
- [7] A. Aue, G. Rice, G. Sonmez, Structural break analysis for spectrum and trace of covariance operators, *Environmetrics* 31 (2020) e2617.
- [8] I. Berkes, R. Gabrys, L. Horváth, P. Kokoszka, Detecting changes in the mean of functional observations, *J. R. Statist. Soc. Ser. B* 71 (2009) 927–946.
- [9] M. Bibinger, M. Madensoy, Change-point inference on volatility in noisy Itô semimartingales, *Stochastic Process. Appl.* 129 (2019) 4878–4925.
- [10] M. Bibinger, J. Moritz, M. Vetter, Nonparametric change-point analysis of volatility, *Ann. Statist.* 45 (2017) 1542–1578.
- [11] A. Casini, P. Perron, Structural breaks in time series, in: *Oxford Research Encyclopedia of Economics and Finance*, 2018.
- [12] K. Christensen, U. Hounyo, M. Podolskij, Is the diurnal pattern sufficient to explain intraday variation in volatility? A nonparametric assessment, *J. Econometr.* 205 (2018) 336–362.
- [13] H. Dette, K. Kokot, Detecting relevant differences in the covariance operators of functional time series: a sup-norm approach, *Ann. Inst. Statist. Math.* 74 (2) (2022) 195–231.
- [14] H. Dette, K. Kokot, A. Aue, Functional data analysis in the Banach space of continuous functions, *Ann. Statist.* 48 (2020) 1168–1192.
- [15] H. Dette, K. Kokot, S. Volgushev, Testing relevant hypotheses in functional time series via self-normalization, *J. R. Statist. Soc. Ser. B* 82 (2020) 629–660.
- [16] H. Dette, T. Kutta, Detecting structural breaks in eigensystems of functional time series, *Electron. J. Stat.* 15 (2021) 944–983.
- [17] O. Gromenko, P. Kokoszka, M. Reimherr, Detection of change in the spatiotemporal mean function, *J. R. Statist. Soc. Ser. B* 79 (2017) 29–50.
- [18] M. Hoffmann, H. Dette, On detecting changes in the jumps of arbitrary size of a time-continuous stochastic process, *Electron. J. Stat.* 13 (2) (2019) 3654–3709.
- [19] M. Hoffmann, M. Vetter, H. Dette, Nonparametric inference of gradual changes in the jump behaviour of time-continuous processes, *Stochastic Process. Appl.* 128 (11) (2018) 3679–3723.
- [20] L. Horváth, P. Kokoszka, G. Rice, Testing stationarity of functional time series, *J. Econometrics* 179 (2014) 66–82.
- [21] L. Horváth, G. Rice, Extensions of some classical methods in change point analysis, *Test* 23 (2014) 219–255.
- [22] L. Horváth, G. Rice, L. Zhao, Change point analysis of covariance functions: a weighted cumulative sum approach, *J. Multivariate Anal.* 189 (2022) 104877.
- [23] T. Hsing, R. Eubank, *Theoretical Foundations of Functional Data Analysis, with an Introduction to Linear Operators*, Wiley, 2015.
- [24] J. Jacod, P. Protter, *Discretization of Processes, Stochastic Modelling and Applied Probability*, Springer, Berlin Heidelberg, 2011.
- [25] Ioannis Karatzas, Steven Shreve, *Brownian Motion and Stochastic Calculus*, second ed., Springer, New York, 1991.
- [26] P. Kokoszka, N. Mohammadi, H. Wang, S. Wang, Functional diffusion driven stochastic volatility model, 2023.
- [27] P. Kokoszka, M. Reimherr, *Introduction to Functional Data Analysis*, in: *Texts in Statistical Science*, CRC Press, 2017.
- [28] J.-F. Le Gall, *Brownian Motion, Martingales, and Stochastic Calculus*, Springer, 2016.
- [29] F. Merlevède, M. Peligrad, S. Utev, Recent advances in invariance principles for stationary sequences, *Probab. Surv.* 3 (2006) 1–36.
- [30] Whitney K. Newey, Kenneth D. West, Automatic lag selection in covariance matrix estimation, *Rev. Econ. Stud.* 61 (4) (1994) 631–653.
- [31] B. Øksendal, *Stochastic Differential Equations*, in: *Universitext*, Springer, Berlin Heidelberg, 2003.
- [32] G.P. Quinn, M.J. Keough, *Experimental Design and Data Analysis for Biologists*, Cambridge University Press, 2002.
- [33] A. Račkauskas, C. Suquet, Testing epidemic changes of infinite dimensional parameters, *Stat. Inference Stoch. Process.* 9 (2006) 111–134.
- [34] G. Rice, M. Shum, Inference for the lagged cross-covariance operator between functional time series, *J. Time Series Anal.* 40 (2019) 665–692.
- [35] X. Shao, Self-normalization for time series: A review of recent developments, *J. Amer. Statist. Assoc.* 110 (2015) 1797–1817.
- [36] X. Shao, X. Zhang, Testing for change points in time series, *J. Amer. Statist. Assoc.* 105 (2010) 1228–1240.
- [37] C. Stöhr, J. Aston, C. Kirch, Detecting changes in the covariance structure of functional time series with application to fMRI data, *Econometr. Statist.* 18 (2021) 44–62.