Monitoring the intraday volatility pattern*

July 1, 2013

Abstract

A functional time series consists of curves, typically one curve per day. The most important parameter of such a series is the mean curve. We propose two methods of detecting a change in the mean function of a functional time series. The change is detected on line, as new functional observations arrive. The general methodology is motivated by, and applied to, the detection of a change in the mean intraday volatility pattern. The methodology is asymptotically justified by applying a new notion of weak dependence for functional time series. It is calibrated and validated by simulations based on real intraday volatility curves.

Keywords: Change point detection; Intraday volatility; Functional data analysis; Sequential analysis.

Abbreviated Title: Monitoring intraday volatility.

*This research was partially supported by Communauté française de Belgique—Actions de Recherche Concertées (2010–2015) and the IAP research network grant nr. P7/06 of the Belgian government (Belgian Science Policy).

1 Motivation and introduction

The presence of a U–shaped intraday volatility pattern has been observed and studied for over two decades, see Admati and Pfeiderer (1988), Becker *et al.* (1993), Andersen and Bollerslev (1997b, 1998), Gwilym *et al.* (1999), Cyree *et al.* (2004), Wang *et al.* (2006), McMillan and Speight (2006), Hughes *et al.* (2007), among others.

As the above references show, various classes of assets have been examined, including stocks, bonds and futures. The data frequency ranges from hourly to tick data, depending on the structure of the asset or availability of the data. The intraday pattern in both volume and/or volatility has been carefully examined. The empirical studies of this pattern for specific asset classes have advanced the understanding and modeling of their intraday dynamics. Changes in intraday volatility have some predictive power for changes in asset levels, and affect volatility estimates used in pricing formulas.

This paper focuses on the volatility of one minute returns on US stocks and indexes, but the statistical methodology we develop is very general, and can be applied to other asset classes, and to volume as well as volatility. In fact, it can be applied to any functional time series with a very general linear or nonlinear dependence structure, but we concentrate on detecting a change in the intraday volatility pattern. If an asset has different volatility patters over two disjoint periods of time, estimating a common pattern by averaging the data from the two periods will generally lead to misleading conclusions, as the estimated pattern will be a mixture of two different patterns. This is analogous to a common statistical setting in which the data come from two populations with different means. If the populations are not separated, the common sample mean will not be close to any of the two means.

To make our work more useful for operational purposes, we consider a sequential setting which allows to detect changes in the intraday volatility pattern, as new data become available. We follow the paradigm developed by Chu *et al.* (1996) for monitoring the parameters of a linear regression. In our setting, the objects we want to monitor are curves defined over a time interval corresponding to the beginning and end of the trading day at the NYSE. The novelty of this paper lies in the development of methods for monitoring the stability of a parameter which is a curve rather than a finite dimensional vector. We use some aspects of the methodology, known as functional data analysis, which was developed for statistical inference for curves. A comprehensive introduction to functional data analysis can be found in Ramsay and Silverman (2005)); Bosq (2000) develops theory for linear functional processes; Ferraty and Vieu (2006) focus on nonparametric approaches; Horváth and Kokoszka (2012) consider inferential procedures based on functional principal components; computational issues are explained in Ramsay *et al.* (2009).

Following Andersen and Bollerslev (1997a), the average intraday MAD-based volatility is defined as follows. Suppose $P_n(t)$ is the price of an asset at time t on the nth day of trading. In this paper, the data are $P_n(t_j), 0 < t_j \le T$, where T is the length of the trading day in minutes, and the length of each interval $[t_{j-1}, t_j]$ is one minute. The return at time t_j of day n is then defined as

$$r_n(t_j) = 100\left(\log P_n(t_j) - \log P_n(t_{j-1})\right),$$

i.e. as the percentage change in the log price over the interval $[t_{j-1}, t_j]$. The average intraday volatility based on N trading days indexed by $n = 1, 2, \ldots, N$ is

$$\widehat{\mu}(t_j) = N^{-1} \sum_{n=1}^{N} |r_n(t_j)|. \tag{1.1}$$

The objective of this paper is to develop a methodology that would allow us to detect in a sequential manner a change in an appropriately defined population parameter, which we call the intraday volatility pattern (IDVP). A practical implication of the detection of such a change is that the days before the change point follow a different IDVP than those after the change point, so the two periods must be treated separately to avoid erroneous inference and conclusions. This simple change point model is clearly only an operational approximation: the detection of a change point means that it is no longer possible to treat the IDVP as constant.

The problem can be approached using the framework introduced by Robbins (1970) and later Chu *et al.* (1996), and developed in Leisch *et al.* (2000), Horváth *et al.* (2004), Aue *et al.* (2006), and Horváth *et al.* (2006). The main difficulty is that the existing research on monitoring econometric parameters is based on scalar valued observations like daily returns, but the IDVP's are essentially curves. It is thus necessary to first develop a suitable statistical framework in which the problem of monitoring can be meaningfully posed.

The functional (curve–valued) observations are the absolute values of intradaily returns, i.e.

$$X_n(t_j) = |r_n(t_j)|, \quad 0 < t_j \le T.$$

Five consecutive curves X_n are shown in Figure 1.1.

We view the curves X_n as functions in $L^2([0,1])$, the space of square integrable functions defined on the unit interval (after a time transformation), whose values at the discrete times t_j are known. We postulate that

$$X_n(t) = \mu_n(t) + \varepsilon_n(t), \tag{1.2}$$

where μ_n is the mean curve, and ε_n is the error curve. If the intraday volatility pattern remains constant over the observed time span, then the μ_n are equal to a single function

Figure 1.1 Five consecutive curves X_n of absolute values of intraday returns from 3/8/2007-3/14/2007. Dashed lines separate trading days

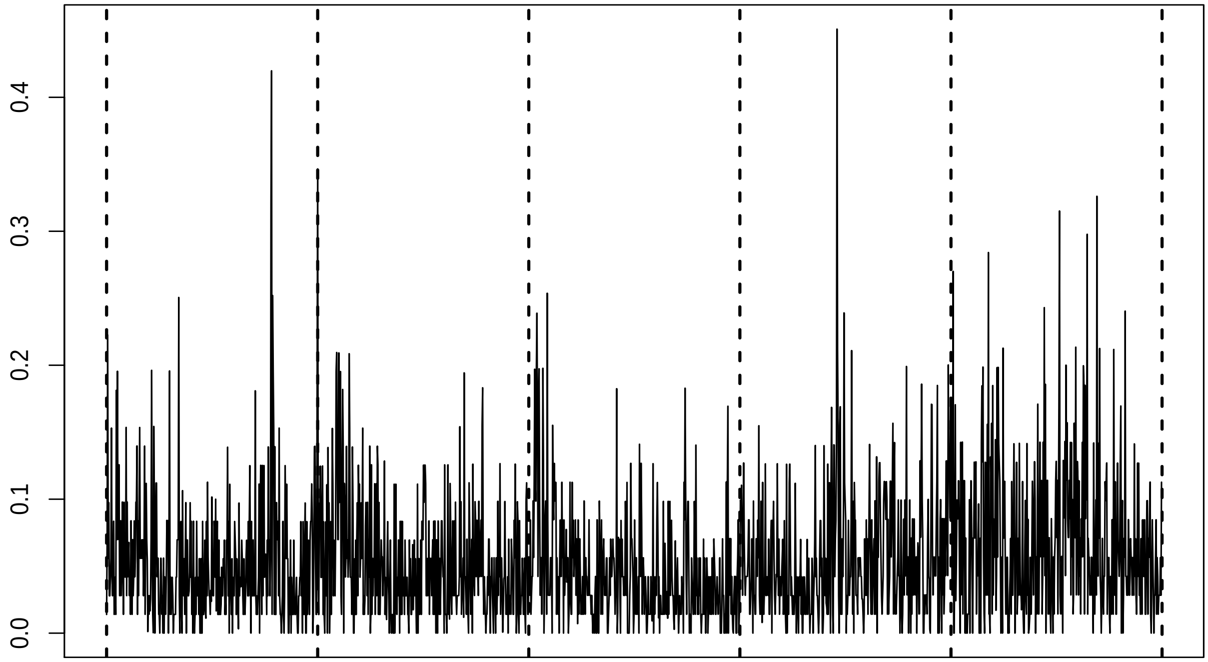

μ, i.e.

$$X_n(t) = \mu(t) + \varepsilon_n(t), \quad n = 1, 2, \ldots, N. \tag{1.3}$$

The average intraday volatility (1.1) is then a consistent estimator of the IDVP μ. The objective of this paper is to develop procedures for monitoring the constancy of the function–valued parameter μ. For example, Figure 1.2 shows the estimators $\hat{\mu}$ computed over two disjoint time periods. We expect that the IDVP μ is different for these two periods.

In order to develop a monitoring procedure we assume that μ_n does not change in m days before the monitoring commenced, i.e.

$$\mu_1 = \mu_2 = \ldots = \mu_m. \tag{1.4}$$

We start monitoring on day $m+1$ and wish to test the null hypothesis

$$H_0: \ \mu_n = \mu_m \quad \text{for } n \geq m+1$$

against the alternative

$$H_A: \ \text{there is } n^* \geq m+1 \text{ such that}$$

$$\mu_n = \mu_m \text{ for } n \leq n^* - 1 \text{ and } \mu_n = \mu_* \neq \mu_m \text{ for } n \geq n^*.$$

FIGURE 1.2 Black lines show the curves $\hat{\mu}$ (left plot is based on the period 8/24/2005-10/5/2005, right on the period 10/6/2005 - 11/16/2005). White solid lines show the smoothed curves $\hat{\mu}$, denoted $\tilde{\mu}$ in Section 4.

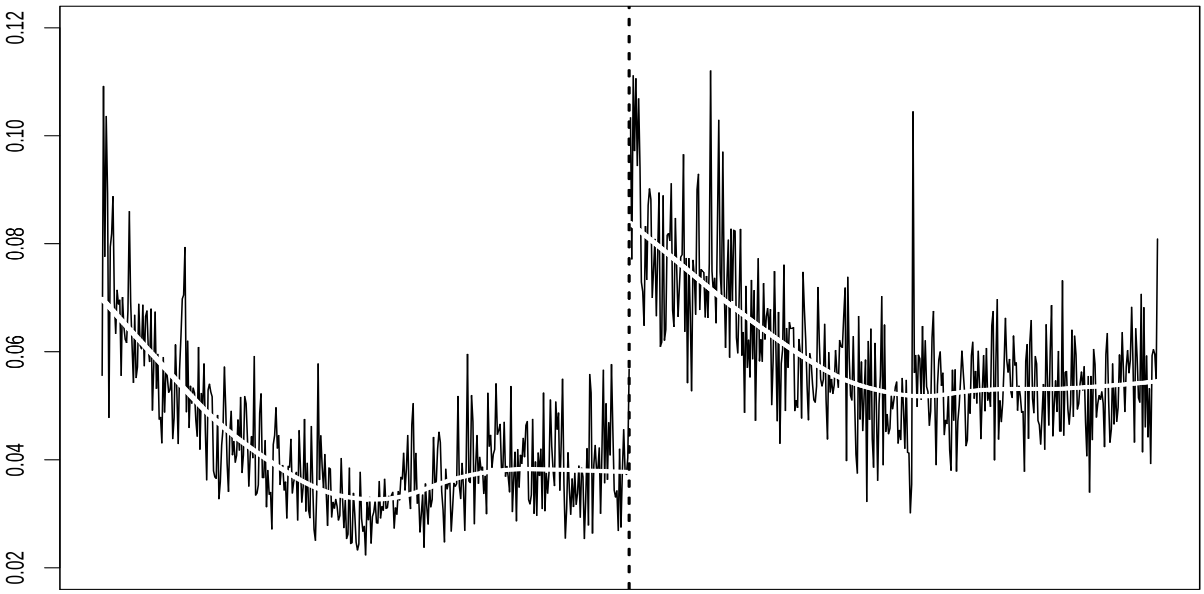

As will be explained in Section 2, for stock returns, the curves ε_n are approximately uncorrelated, in an appropriate sense. They are however not assumed independent, because even the daily returns are dependent (but uncorrelated). However, the methodology we develop does not require that the error curves are uncorrelated, and such assumptions are not made in the general theoretical framework described in Section 3. This methodology can thus be applied to other econometric data in functional form where the daily curves exhibit serial correlation, for example to the Euro–Dollar futures studied by Kargin and Onatski (2008).

The literature on detecting change points, typically called structural breaks in econometrics, is too extensive to even attempt listing important references. A recent review can be found in Andreou and Ghysels (2009). The work of Andersen *et al.* (2001) is particularly related, but these authors focus on a possible change in the IDVP of an FX asset following a specific event. They use a different statistical tool, Fourier flexible form regression, which is suitable for the different problem they approach.

The paper is organized as follows. In Section 2, we describe the detection procedures. Next, in Section 3, we study their asymptotic properties. Their finite sample performance is assessed in Section 4. The proofs are presented in Section 5.

2 Detection procedures

In this section, we provide a description of the monitoring procedures in a form suitable for monitoring the intraday volatility of stock or stock index returns. An assumption made in this section is that the error curves ε_n in (1.2) are uncorrelated, in a sense that for any curves $x, y \in L^2$ $(= L^2([0,1])$,

$$E[\langle \varepsilon_i, x\rangle \langle \varepsilon_j, y\rangle] = 0, \quad \text{if } i \neq j. \tag{2.1}$$

In (2.1), and throughout the paper, we use the inner product notation

$$\langle x, y\rangle = \int_0^1 x(t)y(t)dt, \quad x, y \in L^2.$$

Relation (2.1) involves an infinite number of conditions and cannot be verified exactly. Gabrys and Kokoszka (2007) developed an approximate test based on functional principal components, which suggests that there is no correlation in the data treated here. We also studied acf plots and applied the test of Ljung and Box (1978) to sequences of the form $\langle \varepsilon_n, x\rangle$, for trigonometric functions x, and found that they are uncorrelated in most cases, especially for shorter periods. For longer periods, possible changes in the IDVP may confound these tests. This indicates that the curves ε_n can be treated as uncorrelated, and we shall therefore restrict ourselves to this setup in the empirical part of the paper. Assumption (2.1) is however not needed in our large sample arguments in Section 3.

We consider a basis $\{e_i, \ i \geq 1\}$ in the space L^2 which is used to smooth and represent the volatility curves X_n. In the application to intraday returns, we use the B–spline basis (see for example de Boor (1993)), but other bases could be used as well, for example the polynomial or trigonometric functions, similarly as in Andersen *et al.* (2001). We focus on B-splines because they represent any function in a piecewise manner, and by attaching weights to them, we can monitor for volatility changes at the beginning, middle or end of the trading day. We thus also introduce the weights $\{w_i, \ 1 \leq i \leq d\}$ which satisfy

$$\sum_{i=1}^{d} w_i = 1, \quad w_i \geq 0.$$

The number d is the count of the spline functions used.

To describe the detection procedures, we define the following column vectors:

$$\boldsymbol{w} = [w_1, w_2, \ldots, w_d]^T,$$

$$\mathbf{X}_k = [\langle X_k, e_1\rangle, \langle X_k, e_2\rangle, \ldots, \langle X_k, e_d\rangle]^T,$$

$$\boldsymbol{\varepsilon}_k = [\langle \varepsilon_k, e_1\rangle, \langle \varepsilon_k, e_2\rangle, \ldots, \langle \varepsilon_k, e_d\rangle]^T.$$

The idea is to compare the sample mean of the observations $X_1, X_2, \dots, X_n, n > m$, to the sample mean of the initial data $X_1, X_2, \dots, X_m$, for which the population mean is assumed to be constant.

Define

$$D_n = n(\boldsymbol{w}^T \boldsymbol{\Sigma} \boldsymbol{w})^{-1/2} \sum_{i=1}^{d} w_i \left\langle \overline{X}_n - \overline{X}_m, e_i \right\rangle ,$$

where

$$\overline{X}_n = \frac{1}{n} \sum_{k=1}^{n} X_k$$

is the usual sample mean and

$$\boldsymbol{\Sigma} = [\sigma_{ij},\ 1 \le i, j \le d], \quad \sigma_{ij} = E[\langle \varepsilon_k, e_i \rangle \langle \varepsilon_k, e_j \rangle],$$

is the covariance matrix of each vector $\boldsymbol{\varepsilon}_k$ (and $\mathbf{X}_k$).

To calculate D_n, we must estimate $\sigma_y^2 := \boldsymbol{w}^T \boldsymbol{\Sigma} \boldsymbol{w}$. It is the variance of the scalar sequence $y_k = \boldsymbol{w}^T \mathbf{X}_k$, and can be estimated by the sample variance $\hat{\sigma}_y^2$. (If the $\mathbf{X}_k$ are correlated, $\boldsymbol{\Sigma}$ must be replaced by the long run covariance matrix, and must be estimated differently, see Section 3.) Note that this sample variance can be computed using only $y_k, k \le m$, or all available $y_k, k \le n$. Only the first estimate is consistent under the alternative, but it is possible that the second estimate gives better results in finite samples, even under the alternative.

Setting

$$\hat{D}_n = n(\hat{\sigma}_y^2)^{-1/2} \sum_{i=1}^{d} w_i \left\langle \overline{X}_n - \overline{X}_m, e_i \right\rangle ,$$

the alarm time is the first $n > m$ for which

$$|\hat{D}_n| > m^{1/2} \frac{n-m}{m} g\left(\frac{n}{n-m} \right), \tag{2.2}$$

where $g(\cdot)$ is a function such that

$$P\left(\sup_{1<s<\infty} \frac{|W(s)|}{g(s)} \ge 1 \right) = \alpha. \tag{2.3}$$

In (2.3), $W(\cdot)$ is the standard Wiener process, and α is the confidence level. We work with $\alpha = 5\%$ and $\alpha = 10\%$. Under H_0, the asymptotic probability, as $m \to \infty$, that (2.2) occurs for some $n > m$ is α.

We are now ready to describe two detection procedures, which we refer to in the following as Method 1 and Method 2. Method 1 is based on the approach of Chu *et al.* (1996). Method 2 uses a time inversion to construct a parametric class of functions $g(\cdot)$ in (2.3).

Method 1: Use the function

$$g_a(s) = [s(a^2 + \ln s)]^{1/2},$$

where ln denotes the natural logarithm; $a^2 = 7.78$ gives $\alpha = 5\%$, and $a^2 = 6.25$ gives $\alpha = 10\%$ (in (2.3)).

Method 2: Set

$$g_\gamma(s) = c_\alpha(\gamma)s^{1-\gamma}, \quad \gamma \in [0, 1/2).$$

Then using the well known distributional identity

$$\{W(t),\, t \in [0,1]\} \stackrel{d}{=} \{tW(1/t),\, t \in [0,1]\}$$

we have

$$\sup_{1<t<\infty} \frac{|W(t)|}{c_\alpha(\gamma)t^{1-\gamma}} = \sup_{0<t<1} \frac{|tW(1/t)|}{c_\alpha(\gamma)t^\gamma} \stackrel{d}{=} \sup_{0<t<1} \frac{|W(t)|}{c_\alpha(\gamma)t^\gamma}.$$

The tuning parameter γ can be any number in the interval $[0, 1/2)$. The critical value $c_\alpha(\gamma)$ is defined by

$$P\left\{\sup_{0<t<1} t^{-\gamma}|W(t)| > c_\alpha(\gamma)\right\} = \alpha.$$

For selected values of γ and α, the critical values $c_\alpha(\gamma)$ are tabulated in Horváth *et al.* (2004).

3 Asymptotic theory

In order to derive asymptotic results we must cast our procedure into some appropriate mathematical framework. The easiest approach would be to impose assumptions on the scalar sequence (y_k) on which our detector is based. However, we prefer to impose technical assumptions on the functional data X_k, and not on certain projections. This is because the objects of interests in our context are curves: we model the (functional) volatility process rather than the projected (vector) process. The problem then is to find a reasonable class of functional data for which our approach works. One strategy would be to impose certain mixing conditions, but to date we are not aware of any results showing that commonly studied functional time series models are mixing. Consequently, we will impose below a notion of weak dependence that is known to cover a wide range of functional time series, including functional linear and ARCH processes. We will show in Section 5 that the projected variables y_k are then weakly dependent in the sense of Berkes *et al.* (2010), giving access to our key technical ingredient for the proofs. We notice that the theory will not make use of condition (2.1), postulated in Section 2.

The notion of weak dependence for functional data which we employ has been recently introduced in Hörmann and Kokoszka (2010). In order to give its definition we have to provide some notation first.
For $f \in L^2(= L^2([0,1]))$, we set $\|f\|^2 = \int_0^1 |f(t)|^2 dt$. All random elements are defined on some common probability space $(\Omega, \mathcal{A}, P)$. For $p \geq 1$, we denote by $\mathcal{L}^p = L^p(\Omega, \mathcal{A}, P)$ the space of (classes of) real valued random variables such that $\|X\|_p = (E|X|^p)^{1/p} < \infty$. Further we let $\mathcal{L}_H^p = \mathcal{L}_H^p(\Omega, \mathcal{A}, P)$ be the space of L^2 valued random variables X such that

$$\nu_p(X) = \big(E\|X\|^p\big)^{1/p} < \infty. \tag{3.1}$$

For real sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n \ll b_n$ if $\limsup_{n\to\infty} |a_n/b_n| < \infty$.

Definition 3.1 A sequence $\{X_n\} \in \mathcal{L}_H^p$ is called *$\mathcal{L}^p$–m–approximable* if each X_n admits the representation

$$X_n = f(u_n, u_{n-1}, \ldots), \tag{3.2}$$

where the u_i are iid elements taking values in a measurable space S, and f is a measurable function $f : S^\infty \to H$. Moreover we assume that if $\{u_i'\}$ is an independent copy of $\{u_i\}$ defined on the same probability space, then letting

$$X_n^{(m)} = f(u_n, u_{n-1}, \ldots, u_{n-m+1}, u_{n-m}', u_{n-m-1}', \ldots) \tag{3.3}$$

we have

$$\sum_{m=1}^{\infty} \nu_p\big(X_m - X_m^{(m)}\big) < \infty. \tag{3.4}$$

The sequence is called *$\mathcal{L}^p$–m–approximable with rate* (δ_m), if

$$\nu_p\big(X_m - X_m^{(m)}\big) \leq \delta_m, \tag{3.5}$$

for some $\delta_m \to 0$ as $m \to \infty$.

The u_i in Definition 3.1 should be identified with model errors; typically they are L^2–valued iid random functions. To give a feel for what Definition 3.1 is about, consider a functional linear process $X_n = \sum_{j=0}^{\infty} \Psi_j(u_{n-j})$, where Ψ_j are bounded linear operators acting on functions in L^2. Then

$$X_n^{(m)} = \sum_{j=1}^{m-1} \Psi_j(u_{n-j}) + \sum_{j=m}^{\infty} \Psi_j(u_{n-j}');$$

$\mathcal{L}^p$–m–approximability means that the remainder terms $\sum_{j=m}^{\infty} \Psi_j(u_{n-j})$ are asymptotically negligible. The conditions of Definition 3.1 hold for functional linear time series, including the important ARH(1) model of Bosq (2000), but also for several non-linear functional time series. For more details and applications, as well as a comparison with existing dependence measures like strong mixing or NED we refer to Hörmann and Kokoszka (2010).

We are now ready to state the assumptions needed for our asymptotic results.

Assumption 3.1 *The sequence* $\{\varepsilon_n,\, n \geq 1\}$ *defined in* (1.3) *is* $\mathcal{L}^p$*–*m*–approximable for some* $p > 2$ *with rate* $\delta_m \ll m^{-A}$, $A > 1$.

If the ε_n are correlated, the variance σ_y^2 defined in Section 2 must be replaced by the long run variance of the sequence $y_k = \boldsymbol{w}^T \mathbf{X}_k$; $\hat{\sigma}_y^2$ is then supposed to be a consistent estimator of the long run variance σ_y^2. Conditions on the consistency of the kernel estimators of σ_y^2 are derived in Hörmann and Kokoszka (2012). They hold for all kernel estimators used in practice and for a broad class of $\mathcal{L}^4$–m–approximable time series, which includes a relatively large class of linear processes. These conditions are stated for projection vectors $\mathbf{X}_k$, and are inherited by the linear combinations $\boldsymbol{w}^T \mathbf{X}_k$. For ease of reference, we state the following formal assumption.

Assumption 3.2 *There exists a consistent estimator* $\hat{\sigma}_y^2$ *for* σ_y^2*, i.e.* $\hat{\sigma}_y^2 \to \sigma_y^2$*, in probability.*

For notational convenience we suppress the dependence of any estimator $\hat{\sigma}_y^2$ on the sample size m or n.

If, as it was assumed in Section 2, the curves X_n are uncorrelated, then under H_0 the usual sample variance

$$\frac{1}{N} \sum_{k=1}^{N} (y_k - \overline{y}_N)^2$$

with $N = m$ or $N = n$ is a consistent estimator. This follows from the ergodic theorem as Assumption 3.1 and H_0 imply that the sequence $\{y_k\}$ is strictly stationary and ergodic.

Assumption 3.3 *The critical function* $g(\cdot)$ *is non-decreasing and satisfies for some* $\varsigma > 0$

$$g(s) \geq \varsigma \sqrt{s} \left(\log_+(\log s) + 1 \right), \quad \forall s > 1,$$

where $\log_+(x) = \max\{\log x, 0\}$.

The function g_a of Method 1 satisfies Assumption 3.3 because $(\ln s)^{1/2}$ is larger than $\log_+(\log s)$ for sufficiently large s. The function g_γ of Method 2 also satisfies it because $s^{1-\gamma} > s^{1/2} \log_+(\log s)$, if $\gamma < 1/2$.

Our first theorem states that our detection methods have asymptotically correct size.

Theorem 3.1 (Asymptotics under H_0) *If H_0 and Assumptions 3.1–3.3 hold, then*

$$\hat{T}_m := \sup_{m<n<\infty} \frac{\sqrt{m}\,|\hat{D}_n|}{(n-m)\,g\left(\frac{n}{n-m}\right)} \xrightarrow{d} \sup_{1<t<\infty} \frac{|W(t)|}{g(t)}, \quad (m \to \infty).$$

To obtain consistency of the test, we must make some assumptions on the behavior of the estimator $\hat{\sigma}_n^2$, because only $\hat{\sigma}_m^2$ is consistent under H_A. A simple condition is stated in the following assumption (which covers both $\hat{\sigma}_n^2$ and $\hat{\sigma}_m^2$ for ease of reference).

Assumption 3.4 *The estimator $\hat{\sigma}_y^2$ is bounded in probability.*

It is easy to check that Assumption 3.4 holds under H_A for the standard sample variance. It also holds for kernel estimators whose weights decay sufficiently fast, but we do not want to go into such questions here.

Theorem 3.2 (Asymptotics under H_A) *If H_A and Assumptions 3.1, 3.3 and 3.4 hold, and $Ey_1 \neq Ey_{n^*}$, then $\hat{T}_m \xrightarrow{P} \infty$, as $m \to \infty$.*

4 Finite sample performance and application to intraday volatility of US stocks

In this section, we assess the performance of the monitoring procedures by simulations which use synthetic but realistic data, and an application to stocks of large US corporations.

Size and power. To make our simulation study empirically relevant, we attempt to generate synthetic sequences that resemble real volatility data. The starting point are one minute stock returns from Exxon, $r_n(t_j)$, $n = 1, 2, \ldots, N = 1843$, $j = 1, 2, \ldots, J = 389$, recorded daily from 9:30 AM till 15:59 PM from 12/01/1999 to 14/03/2007. First, we split this data set into consecutive nonoverlapping subsets. On each segment we take the average $\hat{\mu}$ as an estimate of the unobservable IDVP μ, assumed constant over the segment. For stationary data, longer segments imply lower variance of $\hat{\mu}$. On the other hand, if data are non-stationary, the estimate becomes biased. To find a trade-off, we computed the scalar observations y_k defined in Section 2. Our empirical analysis showed that for periods of about ≤ 30 days the y_k seem to have negligible lagged auto-correlations, while this was no longer true for larger intervals. This indicates that the mean function cannot be treated as approximately constant any more. Mikosch and Stărică (2002) and Berkes *et al.* (2006), among many others, study the effect of mean shifts on statistics for measuring dependence. Based on these observation, we chose segments of length $m = 30$, and

Figure 4.1 Average IDV's $\hat{\mu}$ (black line) with LOWESS smooths (white line) over all 61 nonoverlapping 30 day periods from 12/01/1999 to 14/03/2007 (Exxon returns).

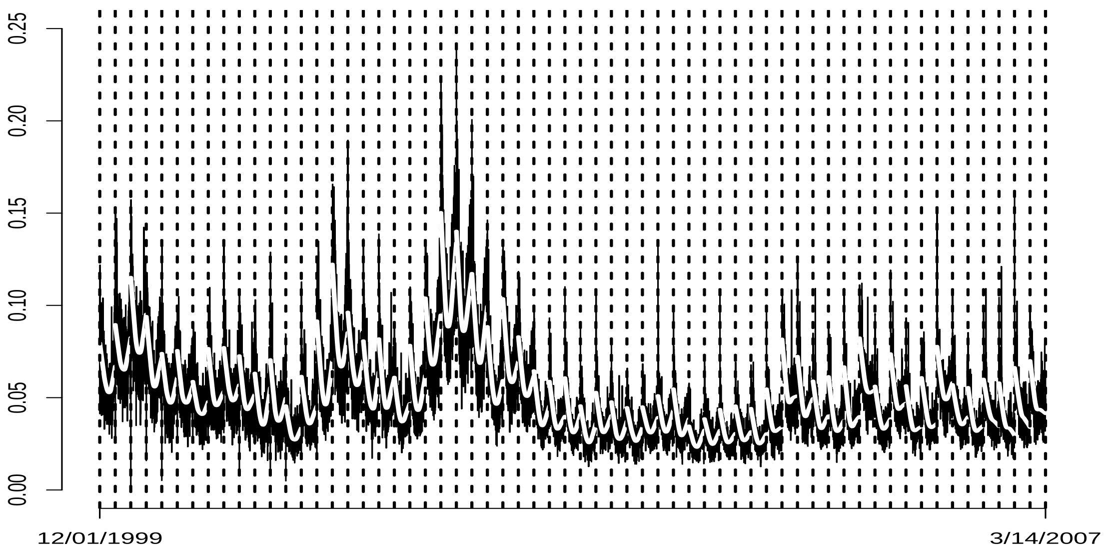

calculate for each of them $\hat{\mu}_i(\mathbf{t}),\ i = 1, 2, \ldots, I$, where $I = 61$ is the number of periods, and $\mathbf{t} = (1, 2, \ldots, 389)^T$ is the vector of minutes in a trading day. A LOWESS smoother is passed through each $\hat{\mu}_i(\mathbf{t})$ to produce $\tilde{\mu}_i(\mathbf{t})$. Figure 4.1 shows both $\hat{\mu}(\mathbf{t})$ and $\tilde{\mu}(\mathbf{t})$. It also illustrates the scope of possible values of μ, which can be approximately identified with the smoothed $\hat{\mu}$.

Next, we calculate the residuals in each period:

$$\tilde{\varepsilon}_{ik}(\mathbf{t}) = X_{ik}(\mathbf{t}) - \tilde{\mu}_i(\mathbf{t}),\ i = 1, 2, \ldots, I,\ k = 1, 2, \ldots, 30,$$

and denote

$$\tilde{\boldsymbol{\varepsilon}} = \{\tilde{\varepsilon}_{ik} \mid i = 1, 2, \ldots, I,\ k = 1, 2, \ldots, 30\}.$$

Our objective is to resample the whole residual curves because we do not want to model the microstructure of the intradaily returns.

The two methods introduced in Section 2 are referred to as M1 and M2. As explained in Section 2, the variance σ_y^2 can be estimated using either the initial m days, or all days up to the current day n. We refer to these two methods as Sm and Sn. Both methods are based on the variables y_k which are weighted projections of X_k onto the cubic splines $\{e_i \mid i = 1, 2, \ldots, d\}$. Two values of d were considered, $d = 5$ and $d = 10$. The results are reported for $d = 5$, as those for $d = 10$ are very similar. The weights are constructed

FIGURE 4.2 Weights, w_i, $i = 1, 2, \ldots, d$, for $d = 5$ (left to right): (1) uniform, (2) concave, (3) convex, (4) increasing, and (5) decreasing.

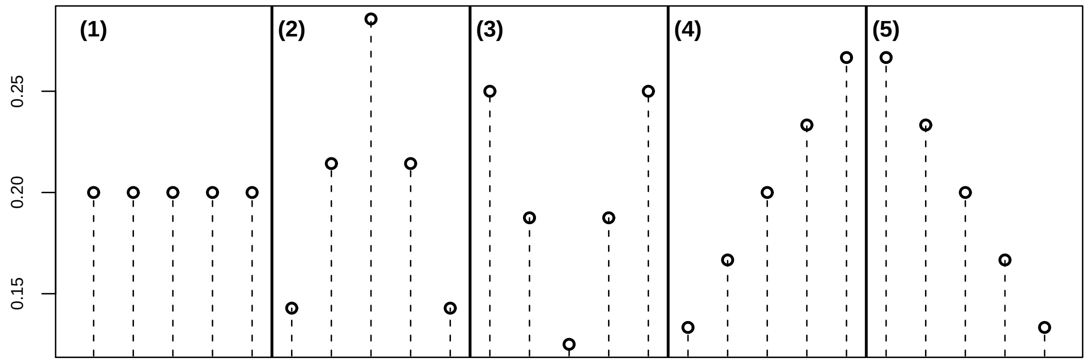

as $w_i = \frac{f(i)}{\sum_{j=1}^{d} f(j)}$, $i = 1, 2, \ldots, d$, $f \geq 0$, where for $i = 1, 2, \ldots, d$, and where $f(i)$ is: uniform, concave, convex, increasing or decreasing. Figure 4.2 shows the weights for each scenario if $d = 5$. The simulations results are reported for the uniform weight case. The results for other weight choices are fairly similar. Procedures on how we generate observations under H_0 and H_A are described in detail below, in subsections *Empirical Size* and *Empirical power*.

The rejection rates up to time n, RR_n, are calculated as

$$RR_n = \frac{1}{R} \sum_{r=1}^{R} I \left\{ \max_{m \leq k \leq n} K_i(k) > 1 \right\}, \quad i = 1, 2,$$

where $I\{\cdot\}$ is the indicator function, and

$$K_1(n) = |\hat{D}_n| / \left[m^{1/2}(n-m)/m g_a \left(n/(n-m) \right) \right], \quad (M1),$$
$$K_2(n) = |\hat{D}_n| / \left[c_\alpha(\gamma) m^{-1/2} n^{1-\gamma} (n-m)^\gamma \right], \quad (M2).$$

All rejection rates are based on $R = 10,000$ replications.

If we put $n = m + k$, then it is not difficult to show that under H_0

$$K_2(n) = O_P \left(\left(\frac{k}{m} \right)^{1/2-\gamma} \right)$$

and under H_A

$$K_2(n) = O_P \left(\sqrt{k} \left(\frac{k}{m} \right)^{1/2-\gamma} \right).$$

Hence, the smaller γ, the less sensitive is the monitoring to early changes. This is why for Method 2, we report the results for $\gamma = 0.49$.

From the practical point of view it is important for a test to perform well for short monitoring periods, 1 or 2 weeks long. On the other hand, we want to see if the simulations agree with the asymptotic properties. For these reasons, the rejection rates are calculated for the following values of n :

$$\mathbf{n} = 30 + \{1, 2, \ldots, 29, 30, 60, 90, \ldots, 270, 300\}.$$

Empirical size; independent errors. The observations without a change point are simulated as

$$X_i^*(\mathbf{t}) = \tilde{\mu}(\mathbf{t}) + \tilde{\varepsilon}_i^*(\mathbf{t}),\ i = 1, 2, \ldots, 330, \tag{4.1}$$

where $\tilde{\mu}(\mathbf{t})$ is the same for all i; it is one of the 61 smoothed means shown in Figure 4.1, the one in the left panel of Figure 1.2. The bootstrap errors $\tilde{\varepsilon}_i^*(\mathbf{t})$ are randomly drawn with replacement from the set of errors $\tilde{\boldsymbol{\varepsilon}}$.

The empirical sizes based on the asympotic theory are shown in Figure 4.3. It shows that the empirical size is highly affected by the choice of the approach used to estimate σ_y^2. While for method Sn, the empirical rejection rates are too conservative, method Sm seems slightly too sensitive. It should be noted, that a certain bias is well expected. Our asymptotic results are established for $m \to \infty$, while we have used here only $m = 30$. Overall, methods M1 and M2 perform relatively similar. Depending on the targeted application, it is left to the practitioner to choose a more conservative or more sensitive setting. If a reduction of the bias is desired, we recommend a calibration of the constants a^2 and $c_\gamma(\alpha)$ defined in methods M1 and M2. For example, one may approximate the finite sample distribution of

$$\max_{m<k\le n} \frac{\sqrt{m}\,|\hat{D}_k|}{(k-m)\, g\left(\frac{k}{k-m}\right)}$$

for given values of m and n via the bootstrap. We performed this for both methods M1 and M2 by bootstrapping from $\tilde{\boldsymbol{\varepsilon}}$. With a bootstrap sample of 10,000 we determine $c_{0.49}^{\text{calibrated}}$ and $a^2_{\text{calibrated}}$ for methods M1 and M2 for both Sm and Sn and compare them to $c_{0.49}(\alpha)$ and a^2 respectively. In accordance with the previous setting we choose $m = 30$, $n = 330$ and $\alpha \in \{0.05, 0.1\}$. See Table 1.

To conserve space, we shall now focus for the rest of the paper on method M2.

Empirical size; dependent (uncorrelated) errors. We want to check how the methods perform if the X_n are dependent but uncorrelated. To generate such functional observations, we replace the $\tilde{\varepsilon}_i^*(\mathbf{t})$ in the data generating process by innovations defined as:

$$\eta_i^*(\mathbf{t}) = \tilde{\varepsilon}_i^*(\mathbf{t}) G_i,$$

where the scalar process G_i is weakly dependent in i, but independent of the process $\tilde{\varepsilon}_i^*(\mathbf{t})$. We model the G_i as an AR(1) process:

$$G_i = \phi G_{i-1} + U_i,$$

	M1		M2	
	$a^2_{\text{calibrated}}(a^2)$		$c^{\text{calibrated}}_{0.49}(c_{0.49})$	
	$\alpha = 5\%$	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 10\%$
Sm	3.33 (3.07)	2.83 (2.83)	10.34 (7.78)	7.05 (6.25)
Sn	2.79 (3.07)	2.51 (2.83)	6.79 (7.78)	5.33 (6.25)

Table 1: Calibrated values $a^2_{\text{calibrated}}$ and $c^{\text{calibrated}}_{0.49}$ for M1 and M2 respectively compared with asymptotic ones (in brackets).

FIGURE 4.3 Empirical size (independent errors, $d = 5$). Empirical rejection rates in percent versus monitoring days; $\circ$ corresponds to Sn and $\bullet$ to Sm. Top row: rejection rates at $\alpha = 5\%$. Bottom row: rejection rates at $\alpha = 10\%$. The vertical dashed line is drawn at $n = 30$.

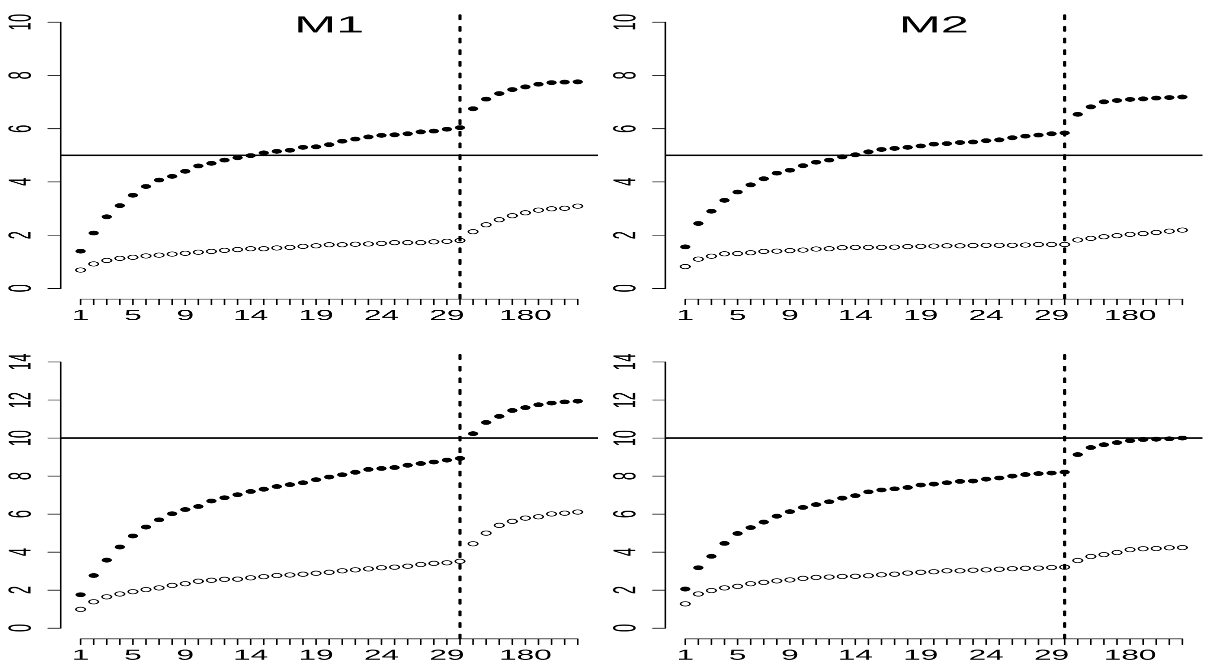

where $0 < \phi < 1$ and U_i are uniformly distributed on (a,b), for some $a,\ b > 0$. Since $G_i = \sum_{j=1}^{\infty} \phi^j U_{i-j}$ we have

$$\frac{a}{1-\phi} \leq G_i \leq \frac{b}{1-\phi}. \tag{4.2}$$

The lower and upper bounds in (4.2) indicate the degree of the modification of the amplitude of the original residuals, the autoregressive coefficient ϕ defines the degree of dependence. We consider two cases:

1. moderate dependence, strong amplitude modification:
   $\phi = 0.5,\ b/(1-\phi) = 1.1,\ a/(1-\phi) = 0.5;$
2. strong dependence, moderate amplitude modification:
   $\phi = 0.9,\ b/(1-\phi) = 1.1,\ a/(1-\phi) = 0.8.$

The size graphs for both scenarios are very similar to those shown in Figure 4.3.

Empirical power. To evaluate the power of the test we simulate observations with change points. We consider two types of changes: abrupt and gradual. Denote $\tilde{\mu}_1(\mathbf{t})$ and $\tilde{\mu}_2(\mathbf{t})$ smoothed mean functions shown, respectively, in the left and right panels of Figure 1.2. Suppose the change occurs at some point $m+k^*$. To simulate the data under the abrupt change point scenario we use the following model:

$$X_i^*(\mathbf{t}) = \begin{cases} \tilde{\mu}_1(\mathbf{t}) + \tilde{\varepsilon}_i^*(\mathbf{t}), & 1 \leq i < m+k^*, \\ \tilde{\mu}_2(\mathbf{t}) + \tilde{\varepsilon}_i^*(\mathbf{t}), & m+k^* \leq i \leq 330. \end{cases} \tag{4.3}$$

Under the gradual change assumption the data is simulated in the following manner:

$$X_i^*(\mathbf{t}) = \begin{cases} \tilde{\mu}_1(\mathbf{t}) + \tilde{\varepsilon}_i^*(\mathbf{t}), & 1 \leq i < m+k^*, \\ (1-\theta)\tilde{\mu}_1(\mathbf{t}) + \theta\tilde{\mu}_2(\mathbf{t}) + \tilde{\varepsilon}_i^*(\mathbf{t}), & m+k^* \leq i < m+\ell^*, \quad \theta = \frac{i-k^*-m}{\ell^*-k^*} \\ \tilde{\mu}_2(\mathbf{t}) + \tilde{\varepsilon}_i^*(\mathbf{t}), & m+\ell^* \leq i \leq 330, \end{cases} \tag{4.4}$$

where $m+k^*$ and $m+\ell^*$ indicate the beginning and the end of the gradual change. Figure 4.4 displays the empirical power for Method 2. As expected, the power is lower for the gradual change scenario. Asymptotically, the power approaches 100% for all methods and all data generating processes. Introducing dependence in the amplitude through the process G_i increases the power by a few percent.

The change from $\tilde{\mu}_1(\mathbf{t})$ to $\tilde{\mu}_2(\mathbf{t})$ shown in Figure 1.2 is relatively small. The power increases much faster after the change point if the change is larger, in level or in shape. Also,

FIGURE 4.4 Empirical power ($k^* = 5, \ell^* = 10$, independent errors, $d = 5$). Empirical rejection rates in percent versus monitoring days; ∘ corresponds to Sn and • to Sm; A to abrupt, G to gradual change. Top row: rejection rates at $\alpha = 5\%$. Bottom row: rejection rates at $\alpha = 10\%$. The vertical dashed line is drawn at $n = 30$.

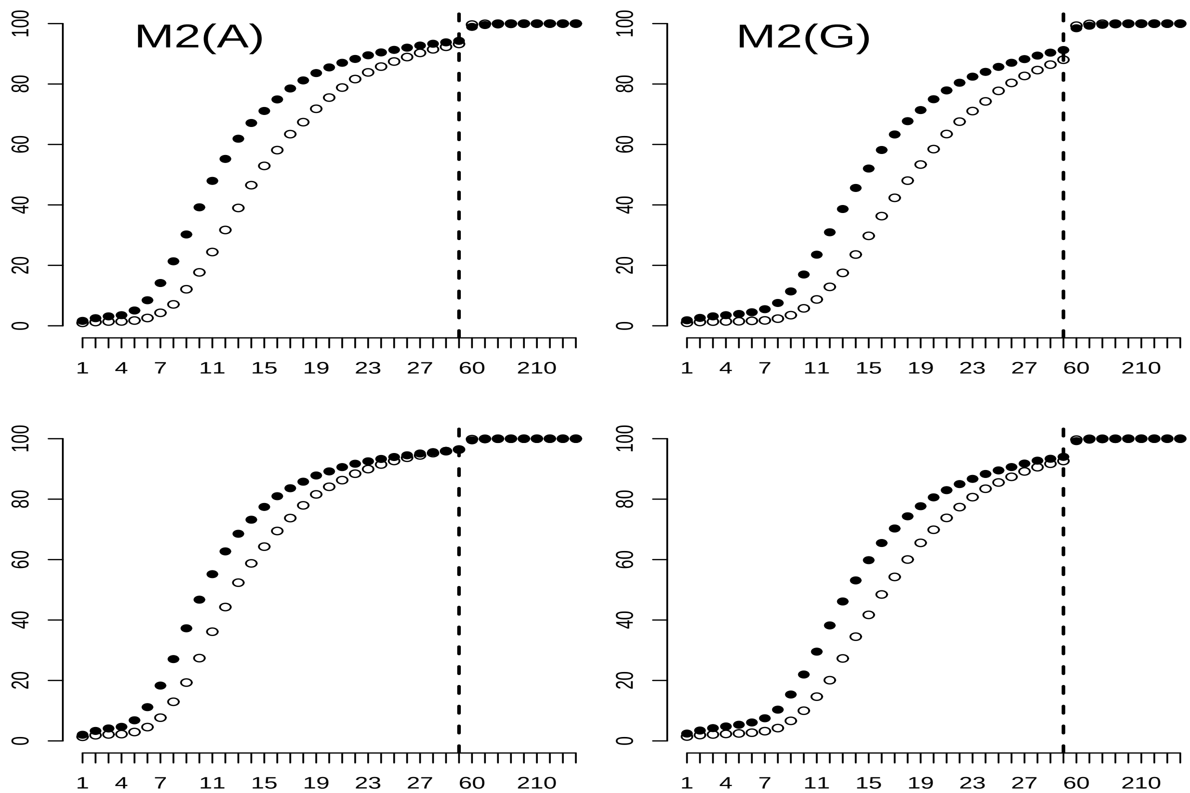

FIGURE 4.5 Empirical power ($k^* = 1, \ell^* = 5$, independent errors, $d = 5$). Empirical rejection rates in percent versus monitoring days; $\circ$ corresponds to Sn and $\bullet$ to Sm; A to abrupt, G to gradual change. Top row: rejection rates at $\alpha = 5\%$. Bottom row: rejection rates at $\alpha = 10\%$. The vertical dashed line is drawn at $n = 30$.

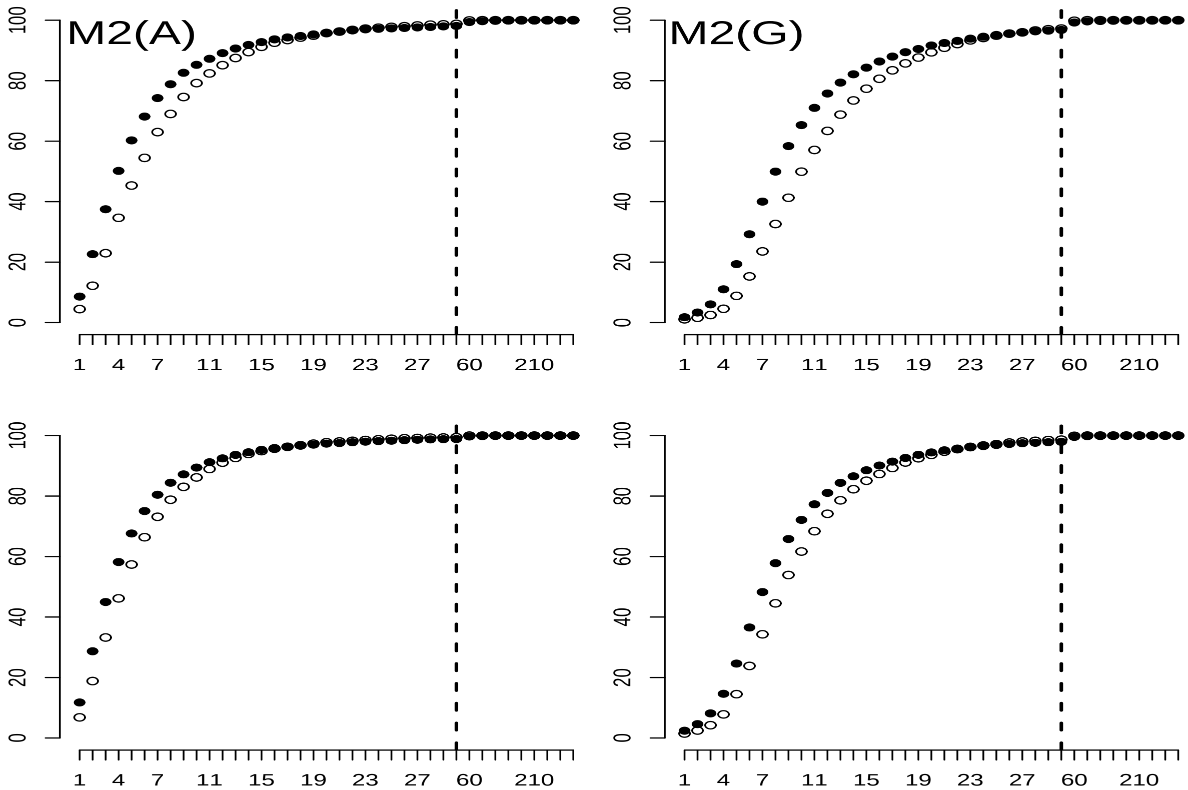

$k^* = 5$, as in Figure 4.4, means that the change begins at the end of a week, if monitoring commenced at the beginning of the week.

The earlier a change occurs, the more powerful is our test. For example, in the abrupt change point scenario this intuitively follows by observing that $E(\overline{X}_m - \overline{X}_{m+k^*+j}) = (\mu_1 - \mu_2)\frac{j}{m+k^*+j}$. Once for some $j \geq 1$ the difference $\overline{X}_m - \overline{X}_{m+k^*+j}$ is large enough, our detector, which is based on this quantity, will reject H_0. The smaller k^*, the larger the factor $\frac{j}{m+k^*+j}$ and hence the threshold is exceeded more easily. This observation is confirmed by Figure 4.5 where we chose $k^* = 1$. In practical applications the initial m days might be updated on a regular basis, so that a small k^* is a reasonable change point.

Application to US stocks. We apply our tests to monitor the IDVP for a number of financial indices and individual stocks. We considered nine periods in 2005-7 consisting of 60 consecutive trading days. The first m =30 days were considered as the training period, and the next 30 days as the monitoring period. Table 2 displays the results for the period 1/4/2007 to 3/30/2007. The training period is from 1/4/2007 to 2/15/2007. Starting 2/16/2007 the monitoring commences. If a change is detected, the day of the change is recorded. If the tests fail to reject the null over the monitoring period, we denote this by ∞. Table 2 shows that changes take place on similar days for the stocks from corresponding sectors. For both methods, Sm detects a change point earlier than Sn; M2 and M1 perform fairly similar. The results do not change much when using the calibrated parameters determined by the bootstrap in Table 1. These findings are in agreement with the simulation results, and remain valid for other monitoring periods. For the specific period considered in Table 2, we conclude that the IDVP of the energy and IT stocks does not change, whereas it does change for financials and consumer staples. This pattern is reflected in the indices, the IT-heavy NASDAQ exhibits no change, whereas the SP100 and DJ, which contain large financial and consumer stables companies, exhibit a change at times at which the change in these sectors is detected. In other periods, the patterns are different, but generally stocks form the same sectors share a similar behavior.

Figure 4.6 illustrates what kind of change in the IDVP our methods can detect: the panels show the averages $\hat{\mu}$ computed, respectively, using the training and the monitoring period (the same as in Table 2). In the top row, COP and MSFT, there is no statistically significant difference between the IDVP over the two periods. In the bottom row, BAC and WMT, the two patterns are significantly different. The differences between training and monitoring periods are visualized in Figure 4.7.

Conclusions. The simulations and the data example lead to the following conclusions:

1. Methods M1 and M2 have similar performance. Method M2/Sm has the smallest bias at the 5% and 10% level and also has good power.

2. As applied, both methods are too conservative under Sn and too sensitive under Sm. This bias is due to the very small size ($m = 30$) of the training sample.

3. One may consider further calibrating the testing procedures for small sample sizes. We have done this, but it didn't lead to substantially different findings in our application. (See Table 2.)

4. If a fast detection is important, and a risk of a false detection is acceptable, Sm should be used rather than Sn.

5. Power depends strongly on how soon the change point occurs. If it occurs 5 or more days after the monitoring has commenced, it may take 10 or more days to detect it, for a moderate change in the IDVP. If no change is detected within several days, the initial period of $m = 30$ days should be updated.

6. The properties of the tests do not change if the IDV curves are dependent, but uncorrelated.

Table 2: Detected days when change in the IDVP takes place. M1 refers to Method 1, M2 refers to Method 2. Sm refers to sample standard deviation estimated using only training data, while Sn refers to standard deviation using pooled data from training and monitored sets. B refers to bootstrap distribution, while T refers to theoretical distribution. B and T were used to determine a^2 for M1 and $c_\alpha(\gamma)$ for M2. Training period: 1/4/2007 - 2/15/2007, monitoring period: 2/16/2007 - 3/30/2007. Sign ∞ indicates no change was detected through the monitoring period.

Stock/Index	M1				M2			
	Sm		Sn		Sm		Sn	
	B	T	B	T	B	T	B	T
Financial Indices								
S & P 100	8	8	29	18	9	8	18	26
COMPX[1]	10	10	∞	∞	11	11	∞	∞
DINDU[2]	8	8	15	11	8	8	10	11
Energy Sector								
XOM[3]	∞	∞	∞	∞	∞	∞	∞	∞
CVX[4]	9	9	∞	∞	9	9	∞	∞
COP[5]	∞	∞	∞	∞	∞	∞	∞	∞
Information Technology Sector								
IBM	9	9	∞	∞	9	9	∞	∞
MSFT[6]	∞	∞	∞	∞	∞	∞	∞	∞
GOOGLE	∞	∞	∞	∞	∞	∞	∞	∞
Financial Sector								
JPM[7]	8	8	29	18	9	8	18	26
WFS[8]	8	8	25	19	9	9	18	23
BAC[9]	7	8	∞	18	8	8	13	19
Consumer Staples Sector								
PG[10]	8	8	18	11	8	8	10	11
KO[11]	7	7	18	11	7	7	9	11
WMT[12]	8	8	20	11	8	8	10	11
Consumer Discretionary Sector								
MCD[13]	10	10	∞	19	11	10	18	∞
DIS[14]	8	8	∞	∞	8	8	11	∞
CMCSA[15]	∞	∞	∞	∞	∞	∞	∞	∞

[1]NASDAQ Composite Index, [2]Dow Jones Industrial Index, [3]Exxon Mobil Corporation, [4]Chevron Corporation, [5]Conocophillips, [6]Microsoft Corporation, [7]JPMmorgan Chase, [8]Wells Fargo [9]Bank of America Corporation, [10]Procter & Gamble, [11]Coca-cola, [12]Wal-Mart Stores, [13]McDonald's Corporation, [14]The Walt Disney Corporation, [15]Comcast.

FIGURE 4.6 Estimated IDVP $\hat{\mu}$ (black solid line) and its LOWESS smooth (white solid line) for both training, 1/4/2007 - 2/15/2007, and monitoring, 2/16/2007 - 3/30/2007, periods separated by a dashed vertical line. Top row: COP (left) and MSFT (right). Bottom row: BAC (left) and WMT (right).

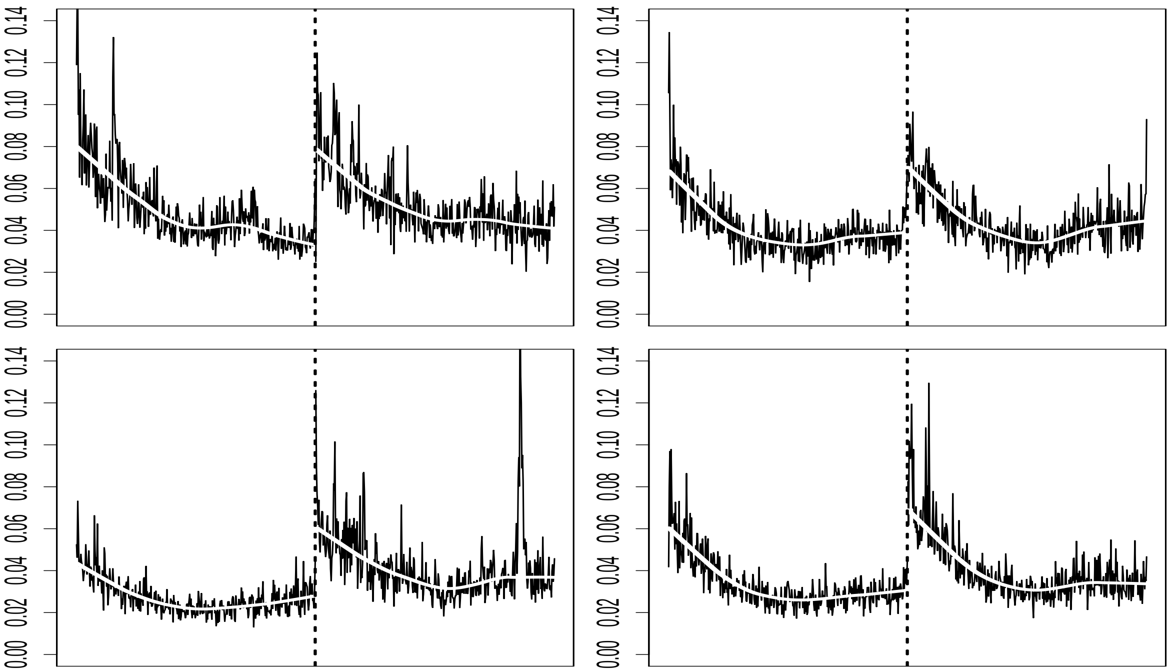

To keep this contribution within a reasonable length, and to focus on a specific class of assets, in our empirical work we did not consider the case of serially correlated X_n. This case is however carefully studied in Section 3. We emphasize that there is no need to estimate a high dimensional (long run) covariance matrix. Once the weights are chosen, it is enough to estimate the (long run) variance of the scalar sequence $y_k = \boldsymbol{w}_k^T \mathbf{X}_k$ defined in Section 2.

5 Proofs of the results of Section 3

Before developing proofs of Theorems 3.1 and 3.2, we state and prove Lemma 5.1 which shows that under Assumption 3.1 the partial sums of the sequence $\{y_k\}$ can be approximated by the trajectories of a Wiener process.

LEMMA 5.1 *Let $\{X_n\}$ be a sequence in L^2, satisfying Assumption 3.1 for some $p > 2$. Let $\{y_k\}$ be given as in Section 2 and set $S_n = y_1 + \cdots + y_n$. Then the sequence $\{y_k\}$ can be defined together with a standard Wiener process $\{W(t),\, t \geq 0\}$ on a common probability space, such that for some $\epsilon > 0$ we have $S_n - nEy_1 = W(\sigma_y^2 n) + o\left(n^{1/2-\epsilon}\right)$.*

Figure 4.7 Differences between smoothed IDVP of training (1/4/2007 - 2/15/2007) and monitoring (2/16/2007 - 3/30/2007) periods. Top row: COP (left) and MSFT (right). Bottom row: BAC (left) and WMT (right).

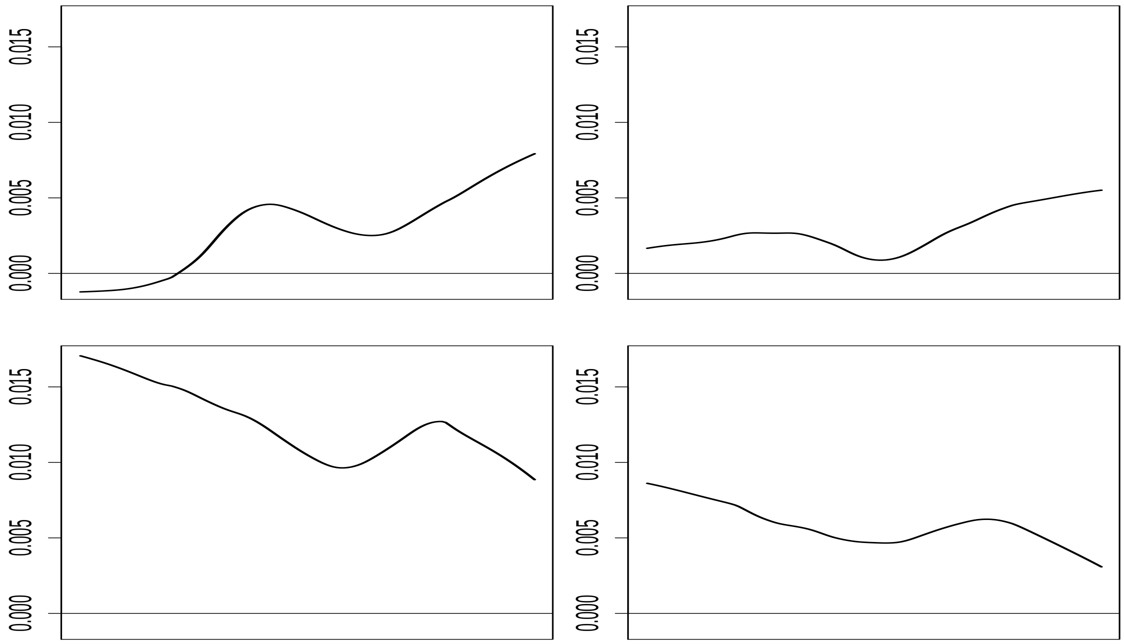

Proof. We define

$$y_k^{(m)} = \sum_{i=1}^{d} w_i \left\langle X_k^{(m)}, e_i \right\rangle .$$

Then we have

$$\begin{aligned}\left\| y_m - y_m^{(m)} \right\|_p &\leq \sum_{i=1}^{d} w_i \left\| \left\langle X_m - X_m^{(m)}, e_i \right\rangle \right\|_p \\ &\leq \sum_{i=1}^{d} w_i \nu_p \left(X_m - X_m^{(m)} \right) \\ &= \nu_p \left(X_m - X_m^{(m)} \right) \ll m^{-A},\end{aligned}$$

showing that the dependence structure (or approximability with m–dependent variables) for the scalar sequence $\{y_m\}$ is inherited from $\{X_m\}$. Hence, by Theorem 1 in (Berkes *et al.*, 2010), we can define $\{y_k\}$ on common space with two Wiener processes $\{W_1(t),\, t \geq 0\}$ and $\{W_2(t),\, t \geq 0\}$ such that for some small enough $\varepsilon_1 > 0$

$$S_n = W_1(s_n^2) + W_2(t_n^2) + o\left(n^{1/2-\epsilon_1}\right). \tag{5.1}$$

The sequences (s_n^2) and (t_n^2) satisfy

$$s_n^2 = n\sigma_y^2 + O\left(n^{1-\epsilon_2}\right) \quad \text{and} \quad t_n^2 = O\left(n^{1-\epsilon_3}\right),$$

with some ϵ_2, $\epsilon_3 > 0$. We have, by an application of Theorem 1.2.1 in (Csörgő and Révész, 1981), that for some large enough constant C

$$|W(s_n^2) - W(\sigma_y^2 n)| \leq \sup_{0 \leq t \leq n} \sup_{0 \leq s \leq Cn^{1-\epsilon_2}} |W(t+s) - W(s)| = o\left(n^{1/2-\epsilon_5}\right) \quad \text{a.s.} \tag{5.2}$$

Furthermore, by the law of the iterated logarithm, it follows immediately that for some small enough $\epsilon_4 > 0$

$$W_2(t_n^2) = o\left(n^{1/2-\epsilon_4}\right) \quad \text{a.s.} \tag{5.3}$$

Combining (5.1), (5.2) and (5.3) yields

$$S_n - nE(y_1) = W(\sigma_y^2 n) + o\left(n^{1/2-\epsilon}\right),$$

with $\epsilon = \min\{\epsilon_1, \epsilon_4, \epsilon_5\}$.

∎

Proof of Theorem 3.1. We have

$$\frac{\hat{\sigma}_y}{\sigma_y} \hat{T}_m = \sup_{m<n<\infty} \frac{\sqrt{m}\left|S_n - \frac{n}{m}S_m\right|}{\sigma_y (n-m)\, g\left(\frac{n}{n-m}\right)} =: T_m.$$

As by Assumption 3.2 $\hat{\sigma}_y/\sigma_y \xrightarrow{P} 1$, it is enough to study the asymptotics of T_m. We assume without loss of generality that $\sigma_y = 1$. Furthermore, since

$$S_n - nEy_1 - \frac{n}{m}(S_m - mEy_1) = S_n - \frac{n}{m}S_m,$$

we can assume in addition that $Ey_1 = 0$. Next we use

$$\begin{aligned} T_m &= \max\left\{\max_{m<n<m+a_m} \frac{\sqrt{m}\left|S_n - \frac{n}{m}S_m\right|}{(n-m)\, g\left(\frac{n}{n-m}\right)}, \sup_{n>m+a_m} \frac{\sqrt{m}\left|S_n - \frac{n}{m}S_m\right|}{(n-m)\, g\left(\frac{n}{n-m}\right)}\right\} \\ &=: \max\{T_m^{(1)}, T_m^{(2)}\}, \end{aligned}$$

where $a_m = m^{1-\epsilon_6}$, for some small $\epsilon_6 \in (0,1)$, which will be chosen later. We show that

$$T_m^{(1)} \xrightarrow{P} 0. \tag{5.4}$$

To this end, we notice that the triangular inequality and basic algebraic manipulations yield

$$\begin{aligned} T_m^{(1)} &\leq \max_{m<n<m+a_m} \frac{\sqrt{m}\,|S_n - S_m|}{(n-m)\, g\left(\frac{n}{n-m}\right)} + \max_{m<n<m+a_m} \frac{|S_m|}{\sqrt{m}} \frac{1}{g\left(\frac{n}{n-m}\right)} \\ &=: T_m^{(1,1)} + T_m^{(1,2)}. \end{aligned}$$

Lemma 5.1 implies that the law of the iterated logarithm applies to S_k, and thus with Assumption 3.3 we can conclude that

$$\begin{aligned}
T_m^{(1,1)} &\stackrel{d}{=} \max_{1\le k\le a_m} \frac{|S_k|}{\sqrt{k\log\log k}} \frac{\sqrt{m\log\log k}}{g\left(\frac{m+k}{k}\right)\sqrt{k}} \\
&\le \max_{1\le k\le a_m} \frac{|S_k|}{\sqrt{k\log\log k}} \frac{\sqrt{m\log\log m}}{\varsigma\sqrt{\frac{m}{k}}\log\log(m^{\epsilon_6})\sqrt{k}} \\
&= O_p(1)\frac{1}{\sqrt{\log\log m}} = o_P(1).
\end{aligned}$$

Under Assumption 3.3 the function $g(\cdot)$ is non-decreasing, and thus again in connection with Lemma 5.1, we conclude that

$$T_m^{(1,2)} \le \frac{|S_m|}{\sqrt{m}} \frac{1}{g\left(\frac{m}{a_m}\right)} = o_P(1).$$

Hence, (5.4) is established.

Next we show that

$$T_m^{(2)} = \sup_{n>m} \frac{\sqrt{m}\left|W(n)-\frac{n}{m}W(m)\right|}{(n-m)\,g\left(\frac{n}{n-m}\right)} + o_P(1). \tag{5.5}$$

Using the same arguments as above, with S_n replaced by the Wiener process $W(n)$, one can show that (5.5) follows from

$$T_m^{(2)} = \sup_{n>m+a_m} \frac{\sqrt{m}\left|W(n)-\frac{n}{m}W(m)\right|}{(n-m)\,g\left(\frac{n}{n-m}\right)} + o_P(1). \tag{5.6}$$

By Lemma 5.1

$$\begin{aligned}
T_m^{(2)} &= \sup_{n>m+a_m} \left\{ \frac{\sqrt{m}\left|W(n)-\frac{n}{m}W(m)\right|}{(n-m)\,g\left(\frac{n}{n-m}\right)} + \sqrt{m}\,\frac{o_P(n^{1/2-\epsilon})+\frac{n}{m}o_P(m^{1/2-\epsilon})}{(n-m)\,g\left(\frac{n}{n-m}\right)} \right\} \\
&=: \sup_{n>m+a_m} \left\{ T_m^{(2,1)}(n) + T_m^{(2,2)}(n) \right\}.
\end{aligned}$$

Thus, (5.6) and consequently (5.5) will follow if

$$T_m^{(2,2,1)} := \sup_{n>m+a_m} \sqrt{m}\,\frac{o_P(n^{1/2-\epsilon})}{(n-m)\,g\left(\frac{n}{n-m}\right)} = o_P(1) \tag{5.7}$$

and

$$T_m^{(2,2,2)} := \sup_{n>m+a_m} \frac{n\,o_P(m^{-\epsilon})}{(n-m)\,g\left(\frac{n}{n-m}\right)} = o_P(1). \tag{5.8}$$

We assume now that ϵ_6 in the definition of $a_m = m^{1-\epsilon_6}$ satisfies $0 < \epsilon_6 < 2\epsilon$. Then for $n > m + a_m$ we have by Assumption 3.3

$$\frac{\sqrt{m}\, n^{1/2-\epsilon}}{(n-m)\, g\left(\frac{n}{n-m}\right)} \leq \frac{\sqrt{m}\, n^{-\epsilon}}{\varsigma\sqrt{n-m}} \leq \frac{m^{1/2-\epsilon}}{m^{1/2(1-\epsilon_6)}} = o(1),$$

showing (5.7). The argument for (5.8) is similar.

Relations (5.4) and (5.5) imply (still assuming $\sigma_y^2 = 1$) that

$$\hat{T}_m = \frac{1}{\hat{\sigma}_y}\left\{\sup_{n>m} \frac{\sqrt{m}\left|W(n) - \frac{n}{m}W(m)\right|}{(n-m)\, g\left(\frac{n}{n-m}\right)} + o_P(1)\right\}.$$

And thus, what remains to show now is that

$$\sup_{n>m} \frac{\sqrt{m}\left|W(n) - \frac{n}{m}W(m)\right|}{(n-m)\, g\left(\frac{n}{n-m}\right)} \xrightarrow{d} \sup_{t>1} \frac{|W(t)|}{g(t)} \quad (m \to \infty). \tag{5.9}$$

This follows from routine arguments for Wiener processes, showing (5.9) and completing the proof of Theorem 3.1.

■

Proof of Theorem 3.2. We use the same notation as for the proof of Theorem 3.1. By Assumption 3.4 it suffices to show that

$$\sup_{m<n<\infty} \frac{\sqrt{m}\,|S_n - \frac{n}{m}S_m|}{(n-m)g\left(\frac{n}{n-m}\right)} \xrightarrow{P} \infty, \quad (m \to \infty). \tag{5.10}$$

Define $S_n^c = S_n - ES_n$. Then we have

$$\begin{aligned} \left|S_n - \frac{n}{m}S_m\right| &= \left|S_n^c - \frac{n}{m}S_m^c + (n^* - n + 1)(Ey_1 - Ey_{n^*})\right| \\ &\geq |(n - n^* + 1)(Ey_1 - Ey_{n^*})| - \left|S_n^c - \frac{n}{m}S_m^c\right|. \end{aligned}$$

From Theorem 3.1 it follows that

$$\sup_{m<n<\infty} \frac{\sqrt{m}\,|S_n^c - \frac{n}{m}S_m^c|}{(n-m)g\left(\frac{n}{n-m}\right)} = O_p(1), \quad (m \to \infty),$$

while

$$\sup_{m<n<\infty} \frac{\sqrt{m}\ |(n - n^* + 1)(Ey_1 - Ey_{n^*})|}{(n-m)g\left(\frac{n}{n-m}\right)} \xrightarrow{P} \infty, \quad (m \to \infty).$$

■

Acknowledgements: The first version of this paper was read by two referees who provided a large number of substantive and precise comments. We acknowledge the time and care these experts dedicated to our work, and thank them for the valuable advice.

References

Admati, A. R. and Pfeiderer, P. (1988). A theory of intraday patters: volume and price variability. *The Review of Financial Studies*, **1,** 3–40.

Andersen, T. G. and Bollerslev, T. (1997a). Heterogeneous information arrivals and return volatility dynamics: uncovering the long run in high frequency data. *Journal of Finance*, **52,** 975–1005.

Andersen, T. G. and Bollerslev, T. (1997b). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, **2–3,** 115–158.

Andersen, T. G. and Bollerslev, T. (1998). DM–dollar volatility: intraday activity patters, microeconomic announcements, and longer run dependencies. *Journal of Finance*, **53,** 219–265.

Andersen, T. G., Bollerslev, T. and Das, A. (2001). Variance-ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns. *Journal of Finance*, **56,** 305–327.

Andreou, E. and Ghysels, E. (2009). Structural breaks in financial time series. In *Handbook of Financial Time Series* (eds T. G. Andersen, R. A. Davis, J-P. Kreiss and T. Mikosch), pp. 839–870. Springer.

Aue, A., Horváth, L., Hušková, M. and Kokoszka, P. (2006). Change–point monitoring in linear models with conditionally heteroskedastic errors. *Econometrics Journal*, **9,** 373–403.

Becker, K. G., Finnerty, J. E. and Kopecky, K. J. (1993). Economic news and intradaily volatility in international bond markets. *Financial Analyst Journal*, **49,** 81–86.

Berkes, I., Hörmann, S. and Schauer, J. (2010). Split invariance principles for stationary processes. *The Annals of Probability*, **00,** 0000–0000; Forthcoming.

Berkes, I., Horváth, L., Kokoszka, P. and Shao, Q-M. (2006). On discriminating between long-range dependence and changes in mean. *The Annals of Statistics*, **34,** 1140–1165.

de Boor, C. (1993). “B(asic)-Spline Basics”. In *Fundamental Developments of Computer-Aided Geometric Modeling, London* (ed. Les Piegl), pp. 27–49. Academic Press.

Bosq, D. (2000). *Linear Processes in Function Spaces.* Springer, New York.

Chu, C.-S. J., Stinchcombe, M. and White, H. (1996). Monitoring structural change. *Econometrica*, **64,** 1045–1065.

Csörgő, M. and Révész, P. (1981). *Strong Approximations in Probability and Statistics.* Academic Press, New York.

Cyree, K. B., Griffiths, M. D. and Winters, D. B. (2004). An empirical examination of the intraday volatility in euro–dollar rates. *The Quarterly Review of Economics and Finance*, **44,** 44–57.

Ferraty, F. and Vieu, P. (2006). *Nonparametric Functional Data Analysis: Theory and Practice.* Springer.

Gabrys, R. and Kokoszka, P. (2007). Portmanteau test of independence for functional observations. *Journal of the American Statistical Association*, **102,** 1338–1348.

Gwilym, O., McMillan, D. G. and Speight, A. E. (1999). The intra–day relationship between volatility and volume in LIFFE futures markets. *Applied Financial Economics*, **9,** 593–604.

Hörmann, S. and Kokoszka, P. (2010). Weakly dependent functional data. *The Annals of Statistics*, **38,** 1845–1884.

Hörmann, S. and Kokoszka, P. (2012). Functional time series. *Handbook of Statistics*, **30,** 157–186.

Horváth, L., Hušková, M., Kokoszka, P. and Steinebach, J. (2004). Monitoring changes in linear models. *Journal of Statistical Planning and Inference*, **126,** 225–251.

Horváth, L. and Kokoszka, P. (2012). *Inference for Functional Data with Applications*. Springer.

Horváth, L., Kokoszka, P. and Zhang, A. (2006). Monitoring constancy of variance in conditionally heteroskedastic time series. *Econometric Theory*, **22,** 373–402.

Hughes, M. P., Smith, S. D. and Winters, D. B. (2007). An empirical examination of intraday volatility in on–the–run U.S. treasury bills. *Journal of Economics and Bussiness*, **59,** 487–499.

Kargin, V. and Onatski, A. (2008). Curve forecasting by functional autoregression. *Journal of Multivariate Analysis*, **99,** 2508–2526.

Leisch, F., Hornik, K. and Kuan, C-M. (2000). Monitoring structural changes with the generalized fluctuation test. *Econometric Theory*, **16,** 835–845.

Ljung, G. and Box, G. (1978). On a measure of lack of fit in time series models. *Biometrika*, **66,** 67–72.

McMillan, D. G. and Speight, E. H. (2006). heterogeneous information flows and intra–day volatility dynamic: Evidence from the UK FTSE–100 stock index futures market. *Applied Financial Economics*, **16,** 959–972.

Mikosch, T. and Stărică, C. (2002). Long-range dependence effects and ARCH modeling. In *Theory and Applications of Long-range Dependence*, *Boston* (eds P. Doukhan, G. Oppenheim and M. S. Taqqu), pp. 439–459. Birkhäuser.

Ramsay, J., Hooker, G. and Graves, S. (2009). *Functional Data Analysis with R and MATLAB*. Springer.

Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. Springer.

Robbins, H. (1970). Statistical methods related to the law of the iterated logarithm. *The Annals of Mathematical Statistics*, **41,** 1397–1409.

Wang, F., Yamazaki, K., Havlin, S. and Stanley, E. (2006). Scaling and memory in intraday volatility return interval in stock market. *Physical Review* E, **73,** 026117:1–8.