

MULTISCALE AND MULTILEVEL TECHNIQUE FOR CONSISTENT SEGMENTATION OF NONSTATIONARY TIME SERIES

Haeran Cho and Piotr Fryzlewicz

London School of Economics

Abstract: In this paper, we propose a fast, well-performing, and consistent method for segmenting a piecewise-stationary, linear time series with an unknown number of breakpoints. The time series model we use is the nonparametric Locally Stationary Wavelet model, in which a complete description of the piecewise-stationary second-order structure is provided by wavelet periodograms computed at multiple scales and locations. The initial stage of our method is a new binary segmentation procedure, with a theoretically justified and rapidly computable test criterion that detects breakpoints in wavelet periodograms separately at each scale. This is followed by within-scale and across-scales post-processing steps, leading to consistent estimation of the number and locations of breakpoints in the second-order structure of the original process. An extensive simulation study demonstrates good performance of our method.

Key words and phrases: Binary segmentation, breakpoint detection, locally stationary wavelet model, piecewise stationarity, post-processing, wavelet periodogram.

1. Introduction

A stationarity assumption is appealing when analysing short time series. But it is often unrealistic, for example when observing time series evolving in naturally nonstationary environments. One such example can be found in econometrics, where price processes are considered to have time-varying variance in response to events taking place in the market; Mikosch and Stărică (1999), Kokoszka and Leipus (2000), and Stărică and Granger (2005), among others, argued in favour of nonstationary modelling of financial returns. For example, given the explosion of market volatility during the recent financial crisis, it is unlikely that the same stationary time series model can accurately describe the evolution of market prices before and during the crisis.

Piecewise stationarity is a well studied and arguably the simplest form of departure from stationarity, and one task of interest is to detect breakpoints in the dependence structure. Breakpoint detection has received considerable attention and the methods that have been developed can be broadly categorized

into retrospective (a posteriori) methods and on-line methods. In the interest of space, we do not review on-line breakpoint detection approaches here but refer the reader to Lai (2001).

The “a posteriori” approach takes into account the entire set of observations at once and detects breakpoints which occurred in the past. Our interest here lies in the “a posteriori” segmentation category, and we propose a retrospective segmentation procedure that achieves consistency in identifying multiple breakpoints for a class of nonstationary processes. (Note that we use the term “segmentation” interchangeably with “multiple breakpoint detection”.)

Early segmentation literature was mostly devoted to testing the existence of a single breakpoint in the mean or variance of independent observations (Chernoff and Zacks (1964), Sen and Srivastava (1975), Hawkins (1977), Hsu (1977), Worsley (1986)). When the presence of more than one breakpoint is suspected, an algorithm for detecting multiple breakpoints is needed. In Vostrikova (1981), a “binary segmentation” procedure was introduced, a computationally efficient multilevel breakpoint detection procedure that recursively locates and tests for multiple breakpoints, producing consistent breakpoint estimators for a class of random processes with piecewise constant means. However, the critical values of the tests at each stage are difficult to compute in practice due to stochasticity in previously selected breakpoints. Venkatraman (1993) employed the same procedure to find multiple breakpoints in the mean of independent and normally distributed variables and showed the consistency of the detected breakpoints with the tests depending on the sample size only, and thus are easier to compute. The binary segmentation procedure was also adopted to detect multiple shifts in the variance of independent observations (Inclán and Tiao (1994), Chen and Gupta (1997)).

Various multiple breakpoint detection methods have been proposed for time series of dependent observations. In Lavielle and Moulines (2000), least squares estimators of breakpoint locations were developed for linear processes with changing mean, extending the work of Bai and Perron (1998). Adak (1998) and Ombao et al. (2001) proposed methods that divided the time series into dyadic blocks and chose the best segmentation according to suitably tailored criteria. Whitcher, Guttorp, and Percival (2000); Whitcher et al. (2002) and Gabbanini et al. (2004) suggested segmenting long memory processes by applying the iterative cumulative sum of squares (ICSS) algorithm (proposed in Inclán and Tiao (1994)) to discrete wavelet coefficients of the process which were approximately Gaussian and decorrelated. Davis, Lee, and Rodriguez-Yam (2006) developed the Auto-PARM procedure which found the optimal segmentation of piecewise stationary AR processes via the minimum description length principle, later extended to the segmentation of non-linear processes in Davis, Lee, and Rodriguez-Yam (2008).

In Lavielle and Teyssière (2005), a breakpoint detection method was developed for weakly or strongly dependent processes with time-varying volatility that minimises a penalised contrast function based on a Gaussian likelihood. Andreou and Ghysels (2002) studied a heuristic segmentation procedure for the GARCH model with changing parameters, based on the work of Lavielle and Moulines (2000).

The aim of our work is to propose a well-performing, theoretically tractable, and fast procedure for detecting breakpoints in the second-order structure of a piecewise stationary time series that is linear but otherwise does not follow any particular parametric model. The nonparametric model we use for this purpose is the Locally Stationary Wavelet (LSW) model first proposed by Nason, von Sachs, and Kroisandt (2000) and later studied by Fryzlewicz and Nason (2006) and Van Bellegem and von Sachs (2008). Detailed justification of our model choice is given in Section 2. In the LSW model, the piecewise constant second-order structure of the process is completely described by local wavelet periodograms at multiple scales, and it is those basic statistics that we use as a basis of our segmentation procedure.

To achieve the multiple breakpoint detection, we propose a binary segmentation method that is applied to wavelet periodograms separately at each scale, and followed by a within-scale and across-scales post-processing procedure to obtain consistent estimators of breakpoints in the second-order structure of the process. We note that wavelet periodograms follow a multiplicative statistical model, but our binary segmentation procedure is different from previously proposed binary segmentation methods for multiplicative models (Inclán and Tiao (1994), Chen and Gupta (1997)) in that it allows for correlated data, which is essential when working with wavelet periodograms. We note that Kouamo, Moulines, and Roueff (2010) proposed a CUSUM-type test for detecting a *single* change in the wavelet variance at one or several scales that also permits correlation in the data. We emphasise other unique ingredients of our breakpoint detection procedure which lead to its good performance and consistency in probability: the theoretical derivation of our test criterion (which only depends on the length of the time series and is thus fast to compute); the novel across-scales post-processing step, essential in combining the results of the binary segmentation procedures performed for each wavelet periodogram scale separately. We note that our method can simultaneously be termed “multiscale” and “multilevel”, as the basic time series model used for our purpose is wavelet-based and thus is a “multiscale” model, and the core methodology to segment each scale of the wavelet periodogram in the model is based on binary segmentation and is thus a “multilevel” procedure.

The paper is organised as follows. Section 2 explains the LSW model and justifies its choice. Our breakpoint detection methodology (together with the

post-processing steps) is introduced in Section 3, where we also demonstrate its theoretical consistency in estimating the total number and locations of breakpoints. In Section 4, we describe the outcome of an extensive simulation study that demonstrates the good performance of our method. In Section 5, we apply our technique to the segmentation of the Dow Jones index and this results in the discovery of two breakpoints: one coinciding with the initial period of the recent financial crisis, and the other coinciding with the collapse of Lehman Brothers, a major financial services firm. The proofs of our theoretical results are provided in the online supplementary material. Software (an R script) implementing our methodology is available from: http://personal.lse.ac.uk/choh1/msml_technique.html.

2. Locally Stationary Wavelet Time Series

In this section, we define the Locally Stationary Wavelet (LSW) time series model (noting that our definition is a slight modification of that of Fryzlewicz and Nason (2006)), describe its properties, and justify its choice as an attractive framework for developing our time series segmentation methodology.

Definition 1. A triangular stochastic array $\{X_{t,T}\}_{t=0}^{T-1}$ for $T = 1, 2, \dots$, is in a class of Locally Stationary Wavelet (LSW) processes if there exists a mean-square representation

$$X_{t,T} = \sum_{i=-\infty}^{-1} \sum_{k=-\infty}^{\infty} W_i\left(\frac{k}{T}\right) \psi_{i,t-k} \xi_{i,k} \quad (2.1)$$

with $i \in \{-1, -2, \dots\}$ and $k \in \mathbb{Z}$ as scale and location parameters, respectively, the $\psi_i = (\psi_{i,0}, \dots, \psi_{i,\mathcal{L}_i-1})$ are discrete, real-valued, compactly supported, non-decimated wavelet vectors with support lengths $\mathcal{L}_i = O(2^{-i})$, and the $\xi_{i,k}$ are zero-mean, orthonormal, identically distributed random variables. For each $i \leq -1$, $W_i(z) : [0, 1] \rightarrow \mathcal{R}$ is a real-valued, piecewise constant function with a finite (but unknown) number of jumps. If the L_i denote the total magnitude of jumps in $W_i^2(z)$, the variability of functions $W_i(z)$ is controlled so that

- $\sum_{i=-\infty}^{-1} W_i^2(z) < \infty$ uniformly in z ,
- $\sum_{i=-I}^{-1} 2^{-i} L_i = O(\log T)$ where $I = \log_2 T$.

The reader unfamiliar with basic concepts in wavelet analysis is referred to the monograph by Vidakovic (1999). By way of example, we recall the simplest discrete, non-decimated wavelet system: the Haar wavelets. Here

$$\psi_{i,k} = 2^{i/2} \mathbf{I}_{\{0, \dots, 2^{-j-1}-1\}}(k) - 2^{i/2} \mathbf{I}_{\{2^{-j-1}, \dots, 2^{-j}-1\}}(k),$$

for $i = -1, -2, \dots, k \in \mathbb{Z}$, where $\mathbf{I}_A(k)$ is 1 if $k \in A$ and 0 otherwise. We note that discrete non-decimated wavelets $\psi_{i,k}$ can be shifted to any location defined by the finest-scale wavelets, and not just to ‘dyadic’ locations (i.e. those with shifts being multiples of 2^{-i}) as in the discrete wavelet transform. Therefore, discrete non-decimated wavelets are no longer an orthogonal, but an overcomplete collection of shifted vectors (Nason, von Sachs, and Kroisandt (2000)).

Throughout, the $\xi_{i,k}$ are assumed to follow the normal distribution; extensions to non-Gaussianity are possible but technically difficult. Comparing the above definition with the Cramér’s representation of stationary processes, $W_i(k/T)$ is a (scale- and location-dependent) transfer function, the wavelet vectors ψ_i are analogous to the Fourier exponentials, and the innovations $\xi_{i,k}$ correspond to the orthonormal increment process. Small negative values of the scale parameter i denote “fine” scales where the wavelet vectors are the most localised and oscillatory; large negative values denote “coarser” scales with longer, less oscillatory wavelet vectors. By assuming that $W_i(z)$ is piecewise constant, we are able to model processes with a piecewise constant second-order structure where, between any two breakpoints in $W_i(z)$, the second-order structure remains constant. The Evolutionary Wavelet Spectrum (EWS) is defined as $S_i(z) = W_i(z)^2$, and it is in a one-to-one correspondence with the time-dependent autocovariance function of the process $c(z, \tau) := \lim_{T \rightarrow \infty} \text{Cov}(X_{[zT],T}, X_{[zT]+\tau,T})$ (Nason, von Sachs, and Kroisandt (2000)). We note that $W_i(z)$ is a valid transfer function; the variance of the resulting time series $X_{t,T}$ is uniformly bounded over t , and the one-to-one correspondence between the autocovariance function and $S_i(z)$ leads to model identifiability. Our objective is to develop a consistent method for detecting breakpoints in the EWS, and consequently to provide a segmentation of the original time series. The following technical assumption is placed on the breakpoints present in the EWS.

Assumption 1. *The set of locations z where (possibly infinitely many) functions $S_i(z)$ contain a jump, is finite; with $\mathcal{B} = \{z; \exists i \lim_{u \rightarrow z-} S_i(u) \neq \lim_{u \rightarrow z+} S_i(u)\}$, then $B = |\mathcal{B}| < \infty$.*

We further define the wavelet periodogram of the LSW time series.

Definition 2. Let $X_{t,T}$ be an LSW process as in (2.1). The triangular stochastic array

$$I_{t,T}^{(i)} = \left| \sum_s X_{s,T} \psi_{i,s-t} \right|^2 \tag{2.2}$$

is called the wavelet periodogram of $X_{t,T}$ at scale i .

With the autocorrelation wavelets $\Psi_i(\tau) := \sum_k \psi_{i,k} \psi_{i,k-\tau}$, the wavelet operator matrix is defined as $\mathbf{A} = (A_{i,k})_{i,k < 0}$ with $A_{i,k} := \langle \Psi_i, \Psi_k \rangle = \sum_\tau \Psi_i(\tau) \Psi_k(\tau)$. Fryzlewicz and Nason (2006) showed that the expectation of a wavelet periodogram $\mathbb{E}I_{t,T}^{(i)}$ is “close” (in the sense of the integrated squared bias converging to zero) to the function $\beta_i(z) = \sum_{j=-\infty}^{-1} S_j(z) A_{i,j}$, a piecewise constant function with at most B jumps, all of which occur in the set \mathcal{B} . Thus, there exists a one-to-one correspondence between EWS, the time-dependent autocovariance function, and the function $\beta_i(z)$ (being the asymptotic expectation of the wavelet periodogram). Every breakpoint in the autocovariance structure then results in a breakpoint in at least one of the $\beta_i(z)$ ’s, and is thus detectable, at least with $T \rightarrow \infty$, by analysing the wavelet periodogram sequences. We note that $\mathbb{E}I_{t,T}^{(i)}$ itself is piecewise constant by definition, except on the intervals of length $C2^{-i}$ around the discontinuities occurring in \mathcal{B} (C denotes an arbitrary positive constant throughout the paper); given a breakpoint $\nu \in \mathcal{B}$, the computation of $I_{t,T}^{(i)}$ for $t \in [\nu - C2^{-i}, \nu + C2^{-i}]$ involves observations from two stationary segments, which results in $\mathbb{E}I_{t,T}^{(i)}$ being “almost” piecewise constant yet not completely so.

The finiteness of \mathcal{B} implies that there exists a fixed index $I^* < \lfloor \log_2 T \rfloor$ such that each breakpoint in \mathcal{B} can be found in at least one of the functions $S_i(z)$ for $i = -1, \dots, -I^*$. Thus, from the invertibility of \mathbf{A} and the closeness of $\beta_i(z)$ and $\mathbb{E}I_{t,T}^{(i)}$, as noted above, we conclude that every breakpoint is detectable from the wavelet periodogram sequences at scales $i = -1, \dots, -I^*$. Since I^* is fixed but unknown, in our theoretical considerations we permit it to increase slowly to infinity with T , see the supplementary material for further details. A further reason for disregarding the coarse scales $i < -I^*$ is that the autocorrelation within each wavelet periodogram sequence becomes stronger at coarser scales; similarly, the intervals on which $\mathbb{E}I_{t,T}^{(i)}$ is not piecewise constant become longer. Thus, for coarse scales, wavelet periodograms provide little useful information about breakpoints and can safely be omitted.

We end this section by briefly summarising our reasons behind the choice of the LSW model as a suitable framework for developing our methodology:

- (i) The entire piecewise constant second-order structure of the process is encoded in the (asymptotically) piecewise constant sequences $\mathbb{E}I_{t,T}^{(i)}$.
- (ii) Due to the “whitening” property of wavelets, the wavelet periodogram sequences are often much less autocorrelated than the original process. In Section 9.2.2 of Vidakovic (1999), the “whitening” property of wavelets is formalised for a second-order stationary time series $X_{t,T}$ with a sufficiently smooth spectral density; defining the wavelet coefficient as $r_{i,k} := \sum_s X_{s,T} \psi_{i,s-k}$, the across-scale covariance of the wavelet coefficients

$\mathbb{E}(r_{i,k}r_{i',k'})$ vanishes for $|i - i'| > 1$, is arbitrarily small for $|i - i'| = 1$, and decays as $o(|k - k'|^{-1})$ within each scale, provided the wavelet used is also sufficiently smooth. However, we emphasise that our segmentation method permits autocorrelation in the wavelet periodogram sequences, as described later in Section 3.1.

- (iii) The entire array of the wavelet periodograms at all scales is easily and rapidly computable via the non-decimated wavelet transform.
- (iv) The use of the “rescaled time” $z = k/T$ in (2.1) and the associated regularity assumptions on the transfer functions $W_i(z)$ permit us to establish rigorous asymptotic properties of our procedure.

3. Binary Segmentation Algorithm

Noting that each wavelet periodogram sequence follows a multiplicative model, as described in Section 3.1, we introduce a binary segmentation algorithm for such class of sequences. Binary segmentation is a computationally efficient tool that searches for multiple breakpoints in a recursive manner (and can be classed as a “greedy” and “multilevel” algorithm). Venkatraman (1993) applied the procedure to a sequence of independent normal variables with multiple breakpoints in its mean and showed that the detected breakpoints were consistent in terms of their number and locations. In the following, we aim at extending these consistency results to the multiplicative model where dependence between observations is permitted.

3.1. Generic multiplicative model

Recall that each wavelet periodogram ordinate is simply a squared wavelet coefficient of a zero-mean Gaussian time series, is distributed as a scaled χ_1^2 variable, and satisfies $I_{t,T}^{(i)} = \mathbb{E}I_{t,T}^{(i)} \cdot Z_{t,T}^2$, where $\{Z_{t,T}\}_{t=0}^{T-1}$ are autocorrelated standard normal variables. Hence we develop a generic breakpoint detection tool for multiplicative sequences

$$Y_{t,T}^2 = \sigma_{t,T}^2 \cdot Z_{t,T}^2, \quad t = 0, \dots, T - 1; \tag{3.1}$$

$I_{t,T}^{(i)}$ and $\mathbb{E}I_{t,T}^{(i)}$ can be viewed as special cases of $Y_{t,T}^2$ and $\sigma_{t,T}^2$, respectively. We assume additional conditions that are, in particular, satisfied for $I_{t,T}^{(i)}$ and $\mathbb{E}I_{t,T}^{(i)}$ by the assumptions of Theorem 2.

- (i) $\sigma_{t,T}^2$ is deterministic and “close” to a piecewise constant function $\sigma^2(t/T)$ in the sense that $\sigma_{t,T}^2$ is piecewise constant apart from intervals of length at most $C2^{J^*}$ around the discontinuities in $\sigma^2(z)$, and $T^{-1} \sum_{t=0}^{T-1} |\sigma_{t,T}^2 - \sigma^2(t/T)|^2 =$

$o(\log^{-1} T)$, where the latter rate comes from the rate of convergence of the integrated squared bias between $\beta_i(t/T)$ and $\mathbb{E}I_{t,T}^{(i)}$ (see Fryzlewicz and Nason (2006) for details) and from the fact that our attention is limited to the I^* finest scales only. Further, $\sigma^2(z)$ is bounded from above and away from zero, with a finite but unknown number of jumps.

- (ii) $\{Z_{t,T}\}_{t=0}^{T-1}$ is a sequence of standard Gaussian variables and the function $\rho(\tau) = \sup_{t,T} |\text{Corr}(Z_{t,T}, Z_{t+\tau,T})|$ satisfies $\rho_\infty^1 < \infty$ where $\rho_\infty^p = \sum_\tau |\rho(\tau)|^p$.

Once the breakpoint detection algorithm for the generic model (3.1) has been established, we apply it to the wavelet periodograms.

3.2. Algorithm

The first step of the binary segmentation procedure is to find the likely location of a breakpoint. We locate such a point in the interval $(0, T-1)$ as the one which maximizes the absolute value of

$$\mathbb{Y}_{0,T-1}^\nu = \sqrt{\frac{T-\nu}{T \cdot \nu}} \sum_{t=0}^{\nu-1} Y_{t,T}^2 - \sqrt{\frac{\nu}{T \cdot (T-\nu)}} \sum_{t=\nu}^{T-1} Y_{t,T}^2. \quad (3.2)$$

Here $\mathbb{Y}_{0,T-1}^\nu$ can be interpreted as a scaled difference between the partial means of two segments $\{Y_{t,T}^2\}_{t=0}^{\nu-1}$ and $\{Y_{t,T}^2\}_{t=\nu}^{T-1}$, where the scaling is chosen so as to keep the variance $\mathbb{Y}_{0,T-1}^\nu$ constant over ν in the idealised case of $Y_{t,T}^2$ being i.i.d.. Once such a ν has been found, we use $\mathbb{Y}_{0,T-1}^\nu$ (but not only this quantity; see below for details) to test the null hypothesis of $\sigma^2(t/T)$ being constant over $[0, T-1]$. The test statistic and its critical value are established such that when a breakpoint is present, the null hypothesis is rejected with probability converging to 1. If the null hypothesis is rejected, we continue the simultaneous locating and testing of breakpoints on the two segments to the left and right of ν in a recursive manner until no further breakpoints are detected. The algorithm is summarised below, where j is the level index and l is the location index of the node at each level. Here the term ‘‘level’’ is used to indicate the progression of the segmentation procedure.

Algorithm

Step 1. Begin with $(j, l) = (1, 1)$. Let $s_{j,l} = 0$ and $e_{j,l} = T-1$.

Step 2. Iteratively compute $\mathbb{Y}_{s_{j,l}, e_{j,l}}^b$ as in (3.2) for $b \in (s_{j,l}, e_{j,l})$. Then, find $b_{j,l}$ which maximizes its absolute value while satisfying

$$\max \left\{ \sqrt{\frac{(e_{j,l} - b_{j,l})}{(b_{j,l} - s_{j,l} + 1)}}, \sqrt{\frac{(b_{j,l} - s_{j,l} + 1)}{(e_{j,l} - b_{j,l})}} \right\} \leq c$$

for a fixed constant $c \in [1/2, 1)$. Let $n_{j,l} = e_{j,l} - s_{j,l} + 1$, $d_{j,l} = \mathbb{Y}_{s_{j,l}, e_{j,l}}^{b_{j,l}}$, and $m_{j,l} = \sum_{t=s_{j,l}}^{e_{j,l}} Y_{t,T}^2 / \sqrt{n_{j,l}}$.

Step 3. Perform hard thresholding on $|d_{j,l}|/m_{j,l}$ with the threshold $t_{j,l} = \tau T^\theta \sqrt{\log T/n_{j,l}}$ so that $\hat{d}_{j,l} = d_{j,l}$ if $|d_{j,l}| > m_{j,l} \cdot t_{j,l}$, and $\hat{d}_{j,l} = 0$ otherwise. The choice of θ and τ is discussed in Section 3.4.

Step 4. If either $\hat{d}_{j,l} = 0$ or $\max\{b_{j,l} - s_{j,l} + 1, e_{j,l} - b_{j,l}\} < \Delta_T$ for l , stop the algorithm on the interval $[s_{j,l}, e_{j,l}]$; if not, let $(s_{j+1,2l-1}, e_{j+1,2l-1}) = (s_{j,l}, b_{j,l})$ and $(s_{j+1,2l}, e_{j+1,2l}) = (b_{j,l} + 1, e_{j,l})$, and update the level j as $j \rightarrow j + 1$. The choice of Δ_T is discussed in Section 3.4.

Step 5. Repeat Steps 2–4.

The condition imposed on $b_{j,l}$ in Step 2 implies that the breakpoints should be sufficiently scattered over time without being too close to each other, and a similar condition is required of the true breakpoints in $\sigma^2(t/T)$, see Assumption 2 in Section 3.3. The set of detected breakpoints is $\{b_{j,l}; \hat{d}_{j,l} \neq 0\}$. The test statistic $|d_{j,l}|/m_{j,l}$ is a scaled version of the test statistics in the ICSS algorithm (Inclán and Tiao (1994)). However, the test criteria in that paper are derived empirically under the assumption of independent observations, and there is no guarantee that their algorithm produces consistent breakpoint estimates.

Fryzlewicz and Nason (2006) and Fryzlewicz, Sapatinas, and Subba Rao (2006) introduced “Haar-Fisz” techniques in different contexts; the former for estimating the time-varying local variance of an LSW time series, and the latter for estimating time-varying volatility in a locally stationary model for financial log-returns. Each Haar-Fisz method has a device (termed the “Fisz transform”) for stabilising the variance of the Haar wavelet coefficients of the data and thereby bringing the distribution of the data close to Gaussianity with constant variance. This is similar to the step in our algorithm where the differential statistic ($d_{j,l}$) is divided by the local mean ($m_{j,l}$), with the convention $0/0 = 0$. However, the Fisz transform was only defined for the case $b = (1/2)(e_{j,l} + s_{j,l} + 1)$ (meaning the segments were split in half) and it was not used for the purposes of breakpoint detection.

3.2.1. Post-processing within a sequence

We equip the procedure with an extra step aimed at reducing the risk of overestimating the number of breakpoints. The ICSS algorithm in Inclán and Tiao (1994) has a “fine-tune” step whereby if more than one breakpoint is found, each breakpoint is checked against the adjacent ones to reduce the risk of over-estimation. We propose a post-processing procedure performing a similar task

within the single-sequence multiplicative model (3.1). At each breakpoint, the test statistic is re-calculated over the interval between two neighbouring breakpoints and compared with the threshold. Denote the breakpoint estimates as $\hat{\eta}_p, p = 1, \dots, \hat{N}$ and $\hat{\eta}_0 = 0, \hat{\eta}_{\hat{N}+1} = T$. For each $\hat{\eta}_p$, we examine whether $\left| \mathbb{Y}_{\hat{\eta}_{p-1}+1, \hat{\eta}_p}^{\hat{\eta}_p} \right| > \tau T^\theta \sqrt{\log T} \cdot \sum_{t=\hat{\eta}_{p-1}+1}^{\hat{\eta}_p} Y_{t,T}^2 / (\hat{\eta}_p - \hat{\eta}_{p-1})$. If this inequality does not hold, $\hat{\eta}_p$ is removed and the same procedure is repeated with the reduced set of breakpoints until the set does not change.

We emphasise that our within-scale post-processing step is in line with the theoretical derivation of breakpoint detection consistency as (a) the extra checks are of the same form as those done in the original algorithm, (b) the locations of the breakpoints that survive the post-processing are unchanged. The next section provides details of our consistency result.

3.3. Consistency of detected breakpoints

In this section, we first show the consistency of our algorithm for a multiplicative sequence as in (3.1), which corresponds to the wavelet periodogram sequence at a single scale. Later, Theorem 2 shows how this consistency result carries over to the consistency of our procedure in detecting breakpoints in the entire second-order structure of the input LSW process $X_{t,T}$.

Denote the number of breakpoints in $\sigma^2(t/T)$ by N and the breakpoints themselves by $0 < \eta_1 < \dots < \eta_N < T - 1$, with $\eta_0 = 0, \eta_{N+1} = T - 1$.

Assumption 2. For $\theta \in (1/4, 1/2)$ and $\Theta \in (\theta + 1/2, 1)$, the length of each segment in $\sigma^2(t/T)$ is bounded from below by $\delta_T = CT^\Theta$. Further, there exists some constant $c \in [1/2, 1)$ such that,

$$\max_{1 \leq p \leq N} \left\{ \sqrt{\frac{\eta_p - \eta_{p-1}}{\eta_{p+1} - \eta_p}}, \sqrt{\frac{\eta_{p+1} - \eta_p}{\eta_p - \eta_{p-1}}} \right\} \leq c.$$

Theorem 1. Suppose that $\{Y_{t,T}\}_{t=0}^{T-1}$ follows model (3.1). Assume there exist $M, m > 0$ such that $\sup_t |\sigma^2(t/T)| \leq M$ and $\inf_{1 \leq i \leq N} |\sigma^2((\eta_i + 1)/T) - \sigma^2(\eta_i/T)| \geq m$. Under Assumption 2, the number and locations of the detected breakpoints are consistent. That is, $\mathbf{P}\{\hat{N} = N; |\hat{\eta}_p - \eta_p| \leq C\epsilon_T, 1 \leq p \leq N\} \rightarrow 1$ as $T \rightarrow \infty$, where $\hat{\eta}_p, p = 1, \dots, \hat{N}$ are detected breakpoints and $\epsilon_T = T^{1/2} \log T$. (Interpreting this in the rescaled time interval $[0, 1]$, $\epsilon_T/T = T^{-1/2} \log T$.)

3.3.1. Post-processing across the scales

We only consider wavelet periodograms $I_{t,T}^{(i)}$ at scales $i = -1, \dots, -I^*$, choosing I^* to satisfy $2^{I^*} \ll \epsilon_T = T^{1/2} \log T$, so that the bias between $\sigma_{t,T}^2$ and

$\sigma^2(t/T)$ does not preclude the results of Theorem 1. Recall that any breakpoint in the second-order structure of the original process $X_{t,T}$ must be reflected in a breakpoint in at least one of the asymptotic wavelet periodogram expectations $\beta_i(z)$, $i = -1, \dots, -I^*$, and vice versa: a breakpoint in one of the $\beta_i(z)$'s implies a breakpoint in the second-order structure of $X_{t,T}$. Thus, it is sensible to combine the estimated breakpoints across the periodogram scales by, roughly speaking, selecting a breakpoint as significant if it appears in *any* of the wavelet periodogram sequences. This section provides a precise algorithm for doing this, and states a consistency result for the final set of breakpoints arising from combining them across scales.

The complete across-scales post-processing algorithm follows. Denote the set of detected breakpoints from the sequence $I_{t,T}^{(i)}$ as $\hat{\mathcal{B}}_i = \{\hat{\eta}_p^{(i)}, p = 1, \dots, \hat{N}_i\}$. Then the post-processing finds a subset of $\cup_{i=-1}^{-I^*} \hat{\mathcal{B}}_i$, say $\hat{\mathcal{B}}$, as follows.

Step 1. Arrange all breakpoints into groups so that those from different sequences and within the distance of Λ_T from each other are classified to the same group; denote the groups by $\mathcal{G}_1, \dots, \mathcal{G}_{\hat{B}}$.

Step 2. Find $i_0 = \max\{\arg \max_{-I^* \leq k \leq -1} \hat{N}_k\}$, the finest scale with the most breakpoints.

Step 3. Check whether there exists $\hat{\eta}_{p_0}^{(i_0)}$ for every $\hat{\eta}_p^{(i)}$, $i \neq i_0$, $1 \leq p \leq \hat{N}_i$, which satisfies $|\hat{\eta}_p^{(i)} - \hat{\eta}_{p_0}^{(i_0)}| < \Lambda_T$. If so, let $\hat{\mathcal{B}} = \hat{\mathcal{B}}_{i_0}$ and stop the post-processing.

Step 4. Otherwise let $\hat{\mathcal{B}} = \{\hat{\nu}_p, p = 1, \dots, \hat{B}\}$ where each $\hat{\nu}_p \in \mathcal{G}_p$ with the maximum i .

We set $\Lambda_T = \lfloor \epsilon_T/2 \rfloor$ in order to take into account the bias arising in deriving the results of Theorem 1. Breakpoints detected at coarser scales are likely to be less accurate than those detected at finer scales; therefore, our algorithm prefers the latter. The across-scales post-processing procedure preserves the number of “distinct” breakpoints and also their locations as determined by the algorithm. Hence the breakpoints in set $\hat{\mathcal{B}}$ are still consistent estimates of true breakpoints in the second-order structure of the original nonstationary process $X_{t,T}$. Although this is not the only way to combine the breakpoints across scales consistent with our theory, we advocate it due to its good performance.

Denote the set of the true breakpoints in the second-order structure of $X_{t,T}$ by $\mathcal{B} = \{\nu_p, p = 1, \dots, B\}$, and the estimated breakpoints by $\hat{\mathcal{B}} = \{\hat{\nu}_p, p = 1, \dots, \hat{B}\}$.

Theorem 2. *Suppose that $X_{t,T}$ satisfies Assumption 1 and that $\nu_p, 1 \leq p \leq B$ satisfy the condition required of the η_p 's in Assumption 2. Further assume that the conditions on $\sigma^2(z)$ in Theorem 1 hold for each $\beta_i(z)$. Then $\mathbf{P}\{\hat{B} = B; |\hat{\nu}_p - \nu_p| \leq C\epsilon_T, 1 \leq p \leq B\} \rightarrow 1$ as $T \rightarrow \infty$.*

3.4. Choice of Δ_T , θ , τ and I^*

To ensure that each estimated segment is of sufficiently large length so as not to distort our theoretical results, Δ_T is chosen so that $\Delta_T \geq C\epsilon_T$. In practice our method works well for smaller values of Δ_T as well, and in the simulation experiments, $\Delta_T = C\sqrt{T}$ is used. As $\theta \in (1/4, 1/2)$, we use $\theta = 0.251$ (we have found that the method works best when θ is close to the lower end of its permitted range) and elaborate on the choice of τ below.

The selection of τ is not a straightforward task and to get some insight into the issue, a set of numerical experiments was conducted. A vector of random variables $\mathbf{X} \sim \mathcal{N}_T(0, \Sigma)$ was generated, $\mathbf{X} = (X_1, \dots, X_T)^T$, then transformed into sequences of wavelet periodograms $I_{t,T}^{(i)}$. The covariance matrix satisfied $\Sigma = (\sigma_{i,j})_{i,j=1}^T$ where $\sigma_{i,j} = \rho^{|i-j|}$. Then we found $b \in (1, T)$ that maximised

$$\mathbb{I}_i^b = \left| \sqrt{\frac{T-b}{T \cdot b}} \sum_{t=1}^b I_{t,T}^{(i)} - \sqrt{\frac{b}{T(T-b)}} \sum_{t=b+1}^T I_{t,T}^{(i)} \right|,$$

and computed $\mathbb{U}_{i,\rho,T} = \mathbb{I}_i^b \cdot \{T^{-1} \sum_{t=1}^T I_{t,T}^{(i)} \cdot T^\theta \sqrt{\log T}\}^{-1}$. This was repeated with a varying covariance matrix ($\rho = 0, 0.3, 0.6, 0.9$) and sample size ($T = 512, 1,024, 2,048$), 100 times for each combination.

The quantity $\mathbb{U}_{i,\rho,T}$ is the ratio between our test statistic and the time-dependent factor $T^\theta \sqrt{\log T}$ appearing in the threshold defined in the algorithm of Section 3.2. $\mathbb{U}_{i,\rho,T}$ is computed under the ‘‘null hypothesis’’ of no breakpoints being present in the covariance structure of X_t , and its magnitude serves as a guideline as to how to select the value of τ , for each scale i , to prevent spurious breakpoint detection in the null hypothesis case. The results showed that the values of $\mathbb{U}_{i,\rho,T}$ and their range tended to increase for coarser scales, this due to the increasing dependence in the wavelet periodogram sequences. In comparison to the scale factor i , the parameters ρ or T had relatively little impact on $\mathbb{U}_{i,\rho,T}$.

We thus propose to use different values of τ in Step 3 of Algorithm of Section 3.2 and in the within-scale post-processing procedure of Section 3.2.1. Denoting the former by $\tau_{i,1}$ and the latter by $\tau_{i,2}$, we chose $\tau_{i,1}$ differently for each i as the 95% quantile, and $\tau_{i,2}$ as the 97.5% quantile of $\mathbb{U}_{i,\rho,T}$ for given i and T and ρ chosen from the set $\{0, 0.3, 0.6, 0.9\}$ with equal probability. The numerical values of $\mathbb{U}_{i,0,T}$ when $T = 1,024$ are summarised in Table 1.

Finally, we discuss the choice of I^* . We first detect breakpoints in wavelet periodograms at scales $i = -1, \dots, -\lfloor \log_2 T/3 \rfloor$ and perform the across-scale post-processing as described in Section 3.3.1, obtaining the set of breakpoints $\hat{\mathcal{B}} = \{\hat{\nu}_p, p = 1, \dots, \hat{B}\}$. Subsequently, for the wavelet periodogram at the next

Table 1. Values of τ for each scale $i = -1, \dots, -4$.

scale i	-1	-2	-3	-4
$\tau_{i,1}$	0.39	0.46	0.67	0.83
$\tau_{i,2}$	0.48	0.52	0.75	0.96

finest scale, we compute the quantity \mathbb{V}_p , $p = 1, \dots, \hat{B} + 1$ as

$$\mathbb{V}_p = \max_{\nu \in (\hat{\nu}_{p-1}, \hat{\nu}_p)} \left| \frac{\sqrt{\frac{\hat{\nu}_p - \nu}{(\hat{\nu}_p - \hat{\nu}_{p-1}) \cdot (\nu - \hat{\nu}_{p-1})}} \sum_{t=\hat{\nu}_{p-1}+1}^{\nu} I_{t,T}^{(i)} - \sqrt{\frac{\nu - \hat{\nu}_{p-1}}{(\hat{\nu}_p - \hat{\nu}_{p-1}) \cdot (\hat{\nu}_p - \nu)}} \sum_{t=\nu+1}^{\hat{\nu}_p} I_{t,T}^{(i)}}{\sum_{t=\hat{\nu}_{p-1}+1}^{\hat{\nu}_p} I_{t,T}^{(i)} / (\hat{\nu}_p - \hat{\nu}_{p-1})} \right|,$$

where $\hat{\nu}_0 = -1$ and $\hat{\nu}_{\hat{B}+1} = T - 1$. Then \mathbb{V}_p is compared to $\tau_{i,1} \cdot T^\theta \sqrt{\log T}$ to see whether there are any further breakpoints yet to be detected in $I_{t,T}^{(i)}$ that have not been included in $\hat{\mathcal{B}}$. (This step is similar to our within-scale post-processing.) If there is an interval $[\hat{\nu}_{p-1} + 1, \hat{\nu}_p]$ where \mathbb{V}_p exceeds the threshold, I^* is updated as $I^* := I^* + 1$ and the above procedure is repeated to update $\hat{\mathcal{B}}$ until either no further changes are made, or $I^* \geq \lfloor \log_2 T/2 \rfloor$.

We note that this approach is in line with the theoretical consistency of our breakpoint detection procedure; \mathbb{V}_p is of the same form as the test statistic and Lemma 6 in the supplementary material implies that, if there are no more breakpoints to be detected from $I_{t,T}^{(i)}$ for $i < -I^*$ other than those already chosen ($\hat{\mathcal{B}}$), then \mathbb{V}_p does not exceed the threshold, and vice versa by Lemma 5.

4. Simulation Study

In Davis, Lee, and Rodriguez-Yam (2006), the performance of the Auto-PARM was assessed and compared with the Auto-SLEX (Ombao et al. (2001)) through simulation in various settings. The Auto-PARM was shown to be superior to Auto-SLEX in identifying both dyadic and non-dyadic breakpoints in piecewise stationary time series. Some examples in the following are adopted from Davis, Lee, and Rodriguez-Yam (2006) for the comparative study between our method and the Auto-PARM, alongside some other new examples. We also applied the breakpoint detection method proposed in Lavielle and Teyssière (2005) to the same simulated processes and, while the performance was found to be good, it was inferior to both Auto-PARM and our method for these particular examples, so we do not report these results. In the simulations, wavelet periodograms were computed using Haar wavelets and both post-processing procedures (Section 3.2.1 and Section 3.3.1) followed the application of the segmentation algorithm.

Table 2. Summary of breakpoint detection from Simulation (A): our method (CF) and Auto-PARM (AP). Results over 100 simulations.

a	number of breakpoints											
	0.7		0.4		0.1		-0.1		-0.4		-0.7	
	CF	AP	CF	AP	CF	AP	CF	AP	CF	AP	CF	AP
0	100	100	100	100	100	100	99	100	99	100	94	100
1	0	0	0	0	0	0	1	0	1	0	5	0
≥ 2	0	0	0	0	0	0	0	0	0	0	1	0
total	100	100	100	100	100	100	100	100	100	100	100	100

Table 3. Summary of breakpoint detection from simulations: our method (CF) and Auto-PARM (AP). Results over 100 simulations.

	number of breakpoints											
	model (B)		model (C)		model (D)		model (E)		model (F)		model (G)	
	CF	AP	CF	AP	CF	AP	CF	AP	CF	AP	CF	AP
0	0	0	0	0	0	0	1	42	1	20	0	0
1	0	0	0	0	97	100	0	31	14	68	1	16
2	93	99	96	100	3	0	97	16	84	7	6	55
3	4	1	3	0	0	0	2	9	1	3	76	29
4	3	0	1	0	0	0	0	0	0	1	17	0
5	0	0	0	0	0	0	0	2	0	1	0	0
total	100	100	100	100	100	100	100	100	100	100	100	100

In our examples, $T = 1,024$ and therefore I^* was set as 3 at the start of each application of the algorithm, then updated automatically if necessary, as described in Section 3.4. Simulation outcomes are given in Tables 2–3, where the total number of detected breakpoints are summarised over 100 simulations.

(A) Stationary AR(1) process with no breakpoints

We consider a stationary AR(1) process,

$$X_t = aX_{t-1} + \epsilon_t \quad \text{for } 1 \leq t \leq 1,024, \quad (4.1)$$

where $\epsilon_t \sim \text{i.i.d. } \mathcal{N}(0, 1)$ (as in all subsequent examples unless specified otherwise). For a range of values of a , we summarise the breakpoint detection outcome in Table 2.

(B) Piecewise stationary AR process with clearly observable changes

This example is taken from Davis, Lee, and Rodriguez-Yam (2006). The target nonstationary process was generated as

$$X_t = \begin{cases} 0.9X_{t-1} + \epsilon_t & \text{for } 1 \leq t \leq 512, \\ 1.68X_{t-1} - 0.81X_{t-2} + \epsilon_t & \text{for } 513 \leq t \leq 768, \\ 1.32X_{t-1} - 0.81X_{t-2} + \epsilon_t & \text{for } 769 \leq t \leq 1,024. \end{cases} \quad (4.2)$$

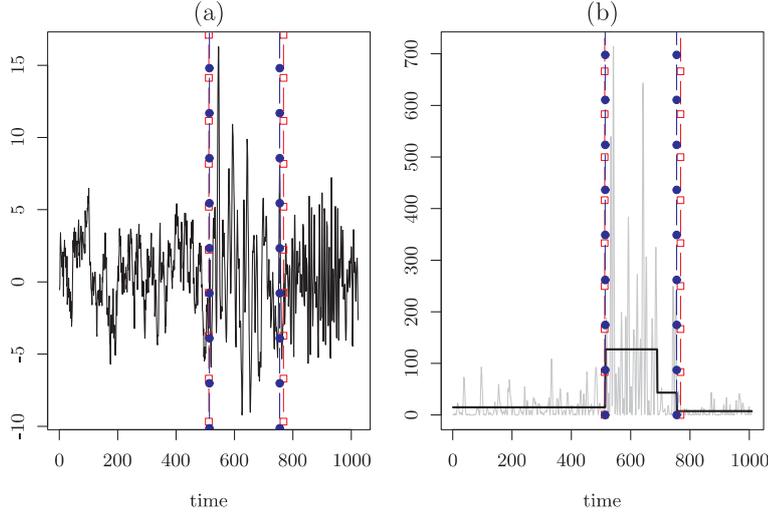


Figure 1. (a) A realisation of model (4.2) with true (empty square) and detected (filled circle) breakpoints; (b) $I_{t,T}^{(i)}$ at $i = -4$ and the breakpoint detection outcome.

As seen in Figure 1(a), there is a clear difference between the three segments in the model. Figure 1(b) shows the wavelet periodogram at scale -4 and the estimation results, where the lines with empty squares indicate the true breakpoints ($\eta_1 = 512, \eta_2 = 768$) while the lines with filled circles indicate the detected ones ($\hat{\eta}_1 = 519, \hat{\eta}_2 = 764$). Note that although initially the procedure returned three breakpoints, the within-sequence post-processing successfully removed the false one.

(C) Piecewise stationary AR process with less clearly observable changes

In this example, the piecewise stationary AR model is revisited, but its breakpoints are less clear-cut, as seen in Figure 2.

$$X_t = \begin{cases} 0.4X_{t-1} + \epsilon_t & \text{for } 1 \leq t \leq 400, \\ -0.6X_{t-1} + \epsilon_t & \text{for } 401 \leq t \leq 612, \\ 0.5X_{t-1} + \epsilon_t & \text{for } 613 \leq t \leq 1,024 \end{cases} \quad (4.3)$$

Figure 2 (b) shows the wavelet periodogram at scale -1 for the realisation in the left panel with its breakpoint estimates ($\hat{\eta}_1 = 403, \hat{\eta}_2 = 622$). Both procedures achieved good performance.

(D) Piecewise stationary AR process with a short segment

This example is again from Davis, Lee, and Rodriguez-Yam (2006). A single

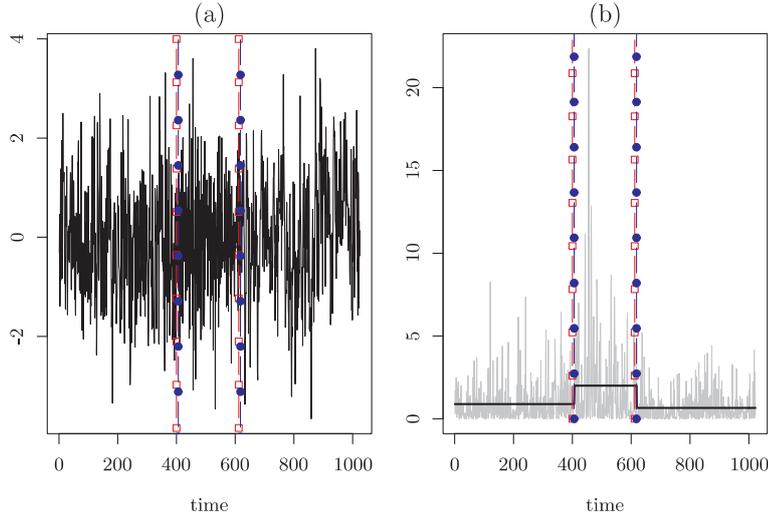


Figure 2. (a) A realisation of model (4.3); (b) $I_{t,T}^{(i)}$ at $i = -1$ and the breakpoint detection outcome.

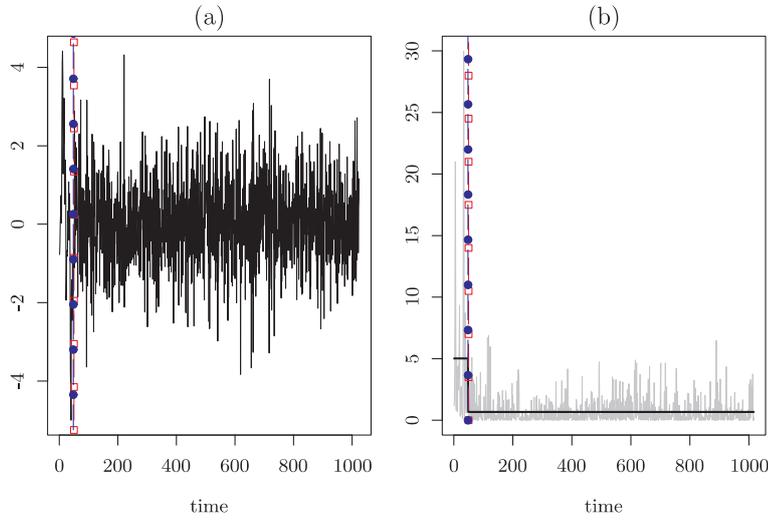


Figure 3. (a) A realisation of model (4.4); (b) $I_{t,T}^{(i)}$ at $i = -3$ and the breakpoint detection outcome.

breakpoint occurs and one segment is much shorter than the other.

$$X_t = \begin{cases} 0.75X_{t-1} + \epsilon_t & \text{for } 1 \leq t \leq 50, \\ -0.5X_{t-1} + \epsilon_t & \text{for } 51 \leq t \leq 1,024. \end{cases} \quad (4.4)$$

A typical realisation of (4.4), its wavelet periodogram at scale -3 , and the estimation outcome are shown in Figure 3, where the jump at $\eta_1 = 50$ was identified

as $\hat{\eta}_1 = 49$. Even though one segment is substantially shorter than the other, our procedure was able to detect exactly one breakpoint in 97% of the cases and underestimation did not occur even when it failed to detect exactly one.

(E) Piecewise stationary near-unit-root process with changing variance

Financial time series, such as stock indices, individual share or commodity prices, or currency exchange rates are, for such purposes as pricing of derivative instruments, often modelled by a random walk with a time-varying variance. We generated a piecewise stationary, near-unit-root example following (4.5), where the variance has two breakpoints over time and the AR parameter remains constant and very close to 1; a typical realisation is given in Figure 4(a). Note that, within each stationary segment, the process can be seen as a special case of the near-unit-root process of Phillips and Perron (1988).

$$X_t = \begin{cases} 0.999X_{t-1} + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, 1) & \text{for } 1 \leq t \leq 400, \\ 0.999X_{t-1} + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, 1.5^2) & \text{for } 401 \leq t \leq 750, \\ 0.999X_{t-1} + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, 1) & \text{for } 751 \leq t \leq 1,024. \end{cases} \quad (4.5)$$

Recall that the Auto-PARM is designed to find the “best” combination of the total number and locations of breakpoints, and adopts a genetic algorithm to traverse the vast parameter space. However, due to the stochastic nature of the algorithm, it occasionally fails to return consistent estimates. This instability was emphasised here, with each run often returning different breakpoints. For one typical realisation, it detected $t = 21$ and 797 as breakpoints, and then only $t = 741$ in the next run on the same sample path. Overall, the performance of Auto-PARM leaves much to be desired for this particular example, whereas our method performed well, though this is not a criticism of Auto-PARM in general, as it performed well in other examples. Note that it was at scale -1 of the wavelet periodogram that both breakpoints were consistently identified the most frequently. The computation of the wavelet periodogram at scale -1 with Haar wavelets is a differencing operation and naturally “whitens” the near-unit-root process (4.5), clearly revealing any changes of variance in the sequence.

(F) Piecewise stationary AR process with high autocorrelation

The features of this AR model are a high degree of autocorrelation and less obvious breakpoints compared to previous examples. A typical realisation is shown in Figure 5(a).

$$X_t = \begin{cases} 1.399X_{t-1} - 0.4X_{t-1} + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, 0.8^2) & \text{for } 1 \leq t \leq 400, \\ 0.999X_{t-1} + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, 1.2^2) & \text{for } 401 \leq t \leq 750, \\ 0.699X_{t-1} + 0.3X_{t-1} + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, 1) & \text{for } 751 \leq t \leq 1,024. \end{cases} \quad (4.6)$$

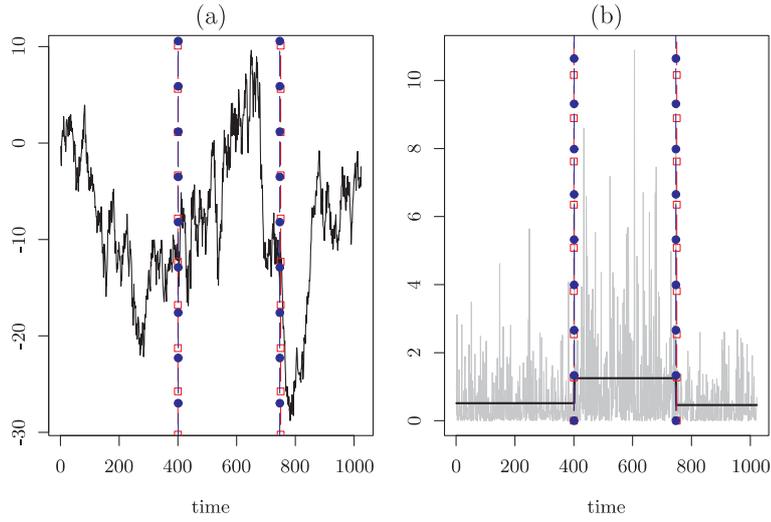


Figure 4. (a) A realisation of model (4.5); (b) $I_{t,T}^{(i)}$ at $i = -1$ and the breakpoint detection outcome.

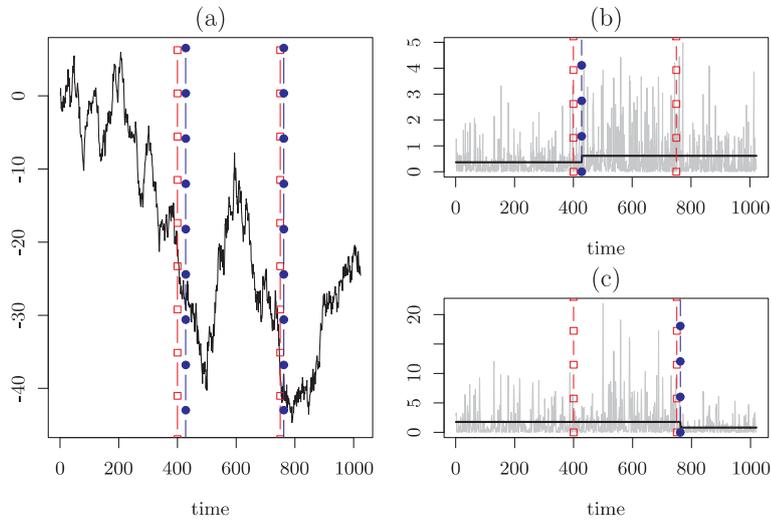


Figure 5. (a) A realisation of model (4.6); (b) $I_{t,T}^{(i)}$ at $i = -1$ and the breakpoint detection outcome; (c) $I_{t,T}^{(i)}$ at $i = -2$ and the breakpoint detection outcome.

Again, the instability of Auto-PARM was notable here, with the second breakpoint at $t = 750$ often left undetected. Our procedure correctly identified both breakpoints in 84% of the cases.

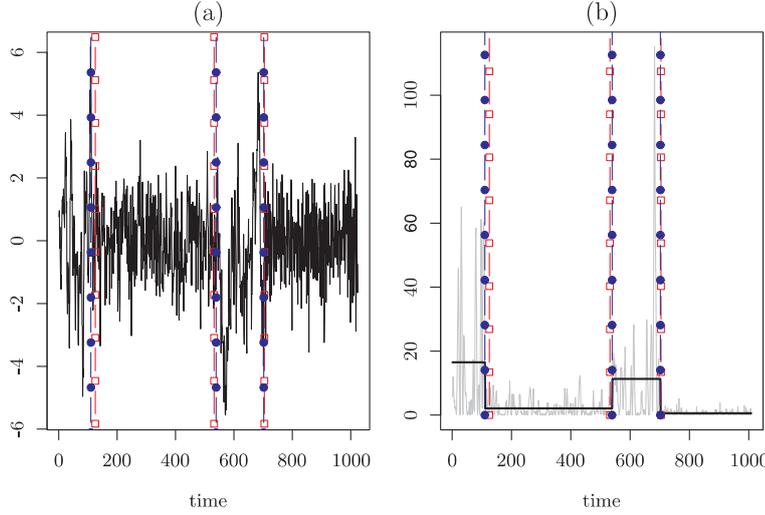


Figure 6. (a) A realisation of model (4.7); (b) $I_{t,T}^{(i)}$ at $i = -4$ and the breakpoint detection outcome.

(G) Piecewise stationary ARMA(1, 1) process

We generated piecewise stationary ARMA processes as

$$X_t = \begin{cases} 0.7X_{t-1} + \epsilon_t + 0.6\epsilon_{t-1} & \text{for } 1 \leq t \leq 125, \\ 0.3X_{t-1} + \epsilon_t + 0.3\epsilon_{t-1} & \text{for } 126 \leq t \leq 532, \\ 0.9X_{t-1} + \epsilon_t & \text{for } 533 \leq t \leq 704, \\ 0.1X_{t-1} + \epsilon_t - 0.5\epsilon_{t-1} & \text{for } 705 \leq t \leq 1,024. \end{cases} \quad (4.7)$$

As illustrated in Figure 6(a), the first breakpoint $t = 125$ is less apparent than the other two. Auto-PARM often left this breakpoint undetected, while our procedure found all three in 76% of cases. We note that it was scale $i = -4$ at which $t = 125$ was detected most frequently by our procedure. With a time series of length $T = 1,024$, default scales provided by our algorithm are $i = -1, -2, -3$, and this example demonstrates the effectiveness of the updating procedure for I^* described in Section 3.4. That is, after completing the examination of $I_{t,T}^{(i)}$ for $i = -1, -2, -3$, our procedure checked if there were more breakpoints to be detected from $I_{t,T}^{(i)}$ for the next scale $i = -4$ and, as it was the case, updated I^* to 4. Figure 6(b) shows the wavelet periodogram at scale -4 for the time series example in the left panel.

5. U.S Stock Market Data Analysis

Many authors, including Stărică and Granger (2005), argue in favour of nonstationary modelling of financial returns. In this analysis, we consider the

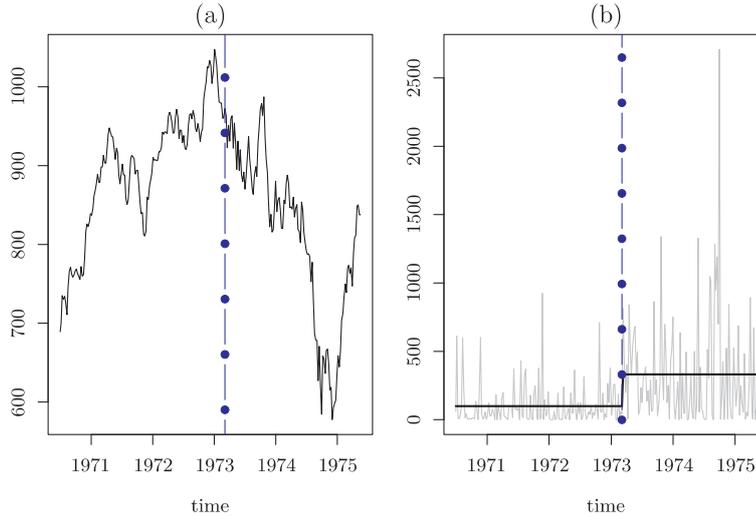


Figure 7. (a) Weekly average values of the Dow Jones IA index (July 1970–May 1975); (b) Wavelet periodogram at scale -1 and the breakpoint detection outcome.

Dow Jones Industrial Average index and regard it as a process with an extremely high degree of autocorrelation (such as in the near-unit-root model of Phillips and Perron (1988)) and a time-varying variance, similar to the simulation model in Section 4 (E).

(A) Dow Jones weekly closing values 1970–1975

The time series of weekly closing values of the Dow Jones Industrial Average index between July 1971 and August 1974 was studied in Hsu (1979) and revisited in Chen and Gupta (1997). Historical data are available on www.google.com/finance/historical?q=INDEXDJK:.DJI, where daily and weekly prices can be extracted for any time period. Both papers concluded that there was a change in the variance of the index around the third week of March 1973. For ease of computation of the wavelet periodogram, we chose the same weekly index between 1 July 1970 and 19 May 1975 so that the data size was $T = 256$ with the above-mentioned time period was contained in this interval. The third week of March 1973 corresponds to $t = 141$ and our procedure detected $\hat{\eta} = 142$ as a breakpoint, as illustrated in Figure 7. The Auto-PARM did not return any breakpoint, while the segmentation procedure proposed in Lavielle and Teysnière (2005), when applied to the log-returns ($\log(X_t/X_{t-1})$) of the data rather than the data X_t themselves, returned $t = 141$ as a breakpoint, which is very close to $\hat{\eta}$.

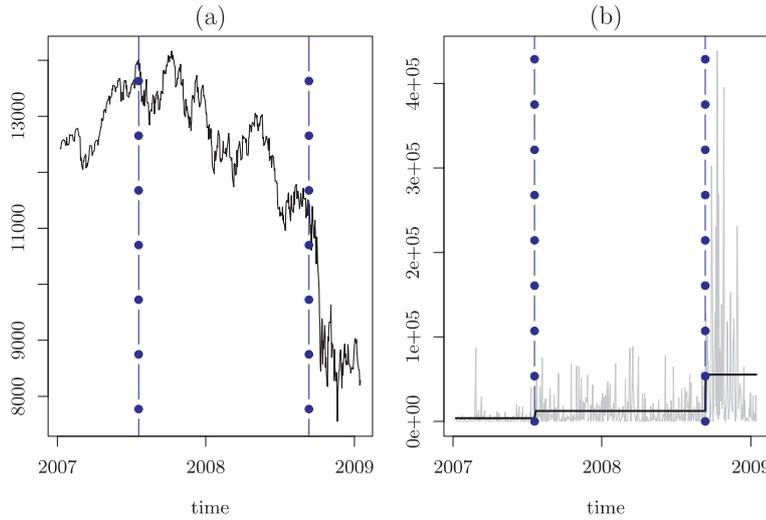


Figure 8. (a) Daily average values of the Dow Jones IA index (Jan 2007–Jan 2009); (b) Wavelet periodogram at scale -1 and the breakpoint detection outcome.

(B) Dow Jones daily closing values 2007–2009

We further investigated more recent *daily* data from the same source, between 8 January 2007 and 16 January 2009. Over this period, the global financial market experienced one of the worst crises in history. Our breakpoint detection algorithm found two breakpoints (see Figure 8), one in the last week of July 2007 ($\hat{\eta}_1 = 135$), and the other in mid-September 2008 ($\hat{\eta}_2 = 424$). The Auto-PARM returned three breakpoints on average, although the estimated breakpoints were unstable as in Section 4 (E): $t = 35, 426$ and 488 were detected most often as breakpoints, while $t = 100$ and $t = 140$ were detected in place of $t = 35$ on other occasions. The segmentation procedure from Lavielle and Teyssière (2005), when applied to the log-returns ($\log(X_t/X_{t-1})$) of the data rather than the data X_t themselves, detected $t = 127$ and 424 as breakpoints, which are very close to $\hat{\eta}_1$ and $\hat{\eta}_2$. The first breakpoint coincided with the outbreak of the worldwide “credit crunch” as subprime mortgage-backed securities were discovered in portfolios of banks and hedge funds around the world. The second breakpoint coincided with the bankruptcy of Lehman Brothers, a major financial services firm, an event that brought even more volatility to the market. Evidence supporting our breakpoint detection outcome is the TED spread (available from <http://www.bloomberg.com/apps/quote?ticker=.tedsp:ind>), an indicator of perceived credit risk in the general economy; it spiked up in late July 2007, remained volatile for a year, then spiked even higher in September 2008. These movements coincide almost exactly with our detected breakpoints.

Acknowledgements

The authors would like to thank Rainer von Sachs for his interesting comments on this work. Also we are grateful to the Editor, an associate editor and two referees for their stimulating reports that led to a significant improvement of the paper.

References

- Adak, S. (1998). Time-dependent spectral analysis of nonstationary time series. *J. Amer. Statist. Assoc.* **93**, 1488-1501.
- Andreou, E. and Ghysels, E. (2002). Detecting multiple breaks in financial market volatility dynamics. *J. Appl. Economics.* **17**, 579-600.
- Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica.* **66**, 47-78.
- Chen, J. and Gupta, A. K. (1997). Testing and locating variance change-points with application to stock prices. *J. Amer. Statist. Assoc.* **92**, 739-747.
- Chernoff, H. and Zacks, S. (1964). Estimating the current mean of a normal distribution which is subject to changes in time. *Ann. Math. Statist.* **35**, 999-1028.
- Davis, R. A., Lee, T. C. M., and Rodriguez-Yam, G. A. (2006). Structural break estimation for non-stationary time series. *J. Amer. Statist. Assoc.* **101**, 223-239.
- Davis, R. A., Lee, T. C. M. and Rodriguez-Yam, G. A. (2008). Break detection for a class of nonlinear time series models. *J. Time Anal.* **29**, 834-867.
- Fryzlewicz, P. and Nason, G. (2006). Haar-Fisz estimation of evolutionary wavelet spectra. *J. Roy. Statist. Soc. Ser. B* **68**, 611-634.
- Fryzlewicz, P., Sapatinas, T. and Subba Rao, S. (2006). A Haar-Fisz technique for locally stationary volatility estimation. *Biometrika.* **93**, 687-704.
- Gabbanini, F., Vannucci, M., Bartoli, G. and Moro, A. (2004). Wavelet packet methods for the analysis of variance of time series with application to crack widths on the Brunelleschi Dome. *J. Comput. Graph. Statist.* **13**, 639-658.
- Hawkins, D. M. (1977). Testing a sequence of observations for a shift in location. *J. Amer. Statist. Assoc.* **72**, 180-186.
- Hsu, D. A. (1977). Tests for variance shifts at an unknown time point. *Journal of Applied Statistics.* **26**, 179-184.
- Hsu, D. A. (1979). Detecting shifts of parameters in Gamma sequences with applications to stock price and air traffic flow analysis. *J. Amer. Statist. Assoc.* **74**, 31-40.
- Inclán, C. and Tiao, G. C. (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *J. Amer. Statist. Assoc.* **89**, 913-923.
- Kokoszka, P. and Leipus, R. (2000). Change-point estimation in ARCH models. *Bernoulli.* **6**, 513-539.
- Kouamo, O., Moulines, E. and Roueff, F. (2010). Testing for homogeneity of variance in the wavelet domain. In *Dependence, with Applications in Statistics and Econometrics* (eds P. Doukhan, G. Lang, D. Surgailis and G. Teyssière). 175-205.
- Lai, T. (2001). Sequential analysis: some classical problems and new challenges. *Statist. Sinica* **11**, 303-350.

- Lavielle, M. and Moulines, E. (2000). Least-squares estimation of an unknown number of shifts in a time series. *J. Time Ser. Anal.* **21**, 33-59.
- Lavielle, M. and Teyssière, G. (2005). Adaptive detection of multiple change-points in asset price volatility. In *Long Memory in Economics* (Edited by G. Teyssière and A. Kirman). 129-156.
- Mikosch, T. and Stărică, C. (1999). Change of structure in financial time series, long range dependence and the GARCH model. Technical Report, University of Groningen.
- Nason, G. P., von Sachs, R. and Kroisandt, G. (2000). Wavelet processes and adaptive estimation of the evolutionary wavelet spectrum. *J. Roy. Statist. Soc. Ser. B* **62**, 271-292.
- Ombao, H. C., Raz, J. A., von Sachs, R. and Malow, B. A. (2001). Automatic statistical analysis of bivariate nonstationary time series. *J. Amer. Statist. Assoc.* **96**, 543-560.
- Phillips, P. and Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika.* **75**, 335-346.
- Sen, A. and Srivastava, M. S. (1975). On tests for detecting change in mean. *Ann. Statist.* **3**, 98-108.
- Stărică, C. and Granger, C. (2005). Nonstationarities in stock returns. *Rev. Economics Statist.* **87**, 503-522.
- Van Bellegem, S. and von Sachs, R. (2008). Locally adaptive estimation of evolutionary wavelet spectra. *Ann. Statist.* **36**, 1879-1924.
- Venkatraman, E. S. (1993). Consistency results in multiple change-point problems. Ph.D. Thesis, Stanford University.
- Vidakovic, B. (1999). *Statistical Modeling by Wavelets*. Wiley, New York.
- Vostrikova, L. J. (1981). Detecting 'disorder' in multidimensional random processes. *Soviet Doklady Math.* **24**, 55-59.
- Whitcher, B., Byers, S. D., Guttorp, P. and Percival, D. B. (2002). Testing for homogeneity of variance in time series: Long memory, wavelets, and the Nile River. *Water Resources Res.* **38**, 10-1029.
- Whitcher, B., Guttorp, P. and Percival, D. B. (2000). Multiscale detection and location of multiple variance changes in the presence of long memory. *J. Statist. Comput. Simulation* **68**, 65-87.
- Worsley, K. J. (1986). Confidence regions and tests for a change-point in a sequence of exponential family random variables. *Biometrika.* **73**, 91-104.

Department of Statistics, London School of Economics, Columbia House, LSE, Houghton Street, London, WC2A 2AE, UK.

E-mail: h.cho1@lse.ac.uk

Department of Statistics, London School of Economics, Columbia House, LSE, Houghton Street, London, WC2A 2AE, UK.

E-mail: p.fryzlewicz@lse.ac.uk

(Received October 2009; accepted October 2010)