

MULTIPLE CHANGE-POINT DETECTION FOR NON-STATIONARY TIME SERIES USING WILD BINARY SEGMENTATION

Karolos K. Korkas and Piotr Fryzlewicz

London School of Economics

Abstract: We propose a new technique for consistent estimation of the number and locations of the change-points in the second-order structure of a time series. The core of the segmentation procedure is the Wild Binary Segmentation method (WBS), a technique which involves a certain randomised mechanism. The advantage of WBS over the standard Binary Segmentation lies in its localisation feature, thanks to which it works in cases where the spacings between change-points are short. In addition, we do not restrict the total number of change-points a time series can have. We also ameliorate the performance of our method by combining the CUSUM statistics obtained at different scales of the wavelet periodogram, our main change-point detection statistic, which allows a rigorous estimation of the local autocovariance of a piecewise-stationary process. We provide a simulation study to examine the performance of our method for different types of scenarios. A proof of consistency is also provided. Our methodology is implemented in the R package `wbsts`, available from CRAN.

Key words and phrases: Binary segmentation, change-points, locally stationary wavelet processes, non-stationarity.

1. Introduction

The assumption of stationarity has been the dominant framework for the analysis of many real data. However, in practice, time series entail changes in their dependence structure and therefore modelling non-stationary processes using stationary methods to capture their time-evolving dependence aspects will most likely result in a crude approximation. As pointed out by Mercurio and Spokoiny (2004) the risk of fitting a stationary model to non-stationary data can be high in terms of prediction and forecasting. Many examples of non-stationary data exist; for example, in biomedical signal processing of electroencephalograms (EEG) see Ombao et al. (2001); in audio signal processing see Davies and Bland (2010); in finance see Stărică and Granger (2005); in oceanography see Killick, Eckley and Jonathan (2013). In this paper we deal with piecewise stationarity, arguably the simplest type of deviation from stationarity. This implies a time-varying process where its parameters evolve through time but remain constant for a specific period of time.

The problem of change-point estimation has attracted significant attention. A branch of the literature deals with the estimation of a single change-point (for a change in mean see e.g. Sen and Srivastava (1975); for time series see Davis, Huang and Yao (1995), Gombay (2008), Gombay and Serban (2009), and references therein) while another extends it to multiple change-points with many changing parameters such as Ombao et al. (2001) who divide a time series into dyadic segments and choose the one with the minimum cost. The latter branch can be further categorised. The multiple change-point estimation can be formulated as minimising a multivariate cost function (or criterion). When the number of change-points N is unknown then a penalty is typically added e.g. the Schwarz criterion (see Yao (1988)). In addition, the user can adopt certain cost functions to deal with the estimation of specific models: the least-squares for change in the mean of a series (Yao and Au (1989) or Lavielle and Moulines (2000)), the Minimum Description Length criterion (MDL) for non-stationary time series (Davis, Lee and Rodriguez-Yam (2006)), the Gaussian log-likelihood function for changes in the volatility (Lavielle and Teyssiere (2007)) or the covariance structure of a multivariate time series (Lavielle and Teyssiere (2006)).

Several algorithms for minimising a cost function are based on dynamic programming (Bellman and Dreyfus (1966) and Kay (1998)) and they are often used in solving change-point problems, see e.g. Perron (2006) and references therein. Auger and Lawrence (1989) propose the Segment Neighbourhood method with complexity $\mathcal{O}(QT^2)$, where Q is the maximum number of change-points. An alternative method is the exact method of Optimal Partitioning by Jackson et al. (2005), but its complexity of $\mathcal{O}(T^2)$ still makes it suitable mostly for smaller samples.

Change-point estimators that adopt a multivariate cost function often come with a high computational cost. An attempt to reduce the computational burden is found in Killick, Fearnhead and Eckley (2012) who extend the Optimal Partitioning method of Jackson et al. (2005) (termed PELT) and show that the computational cost is $\mathcal{O}(T)$ when the number of change-points increases linearly with T . Another attempt is found in Davis, Lee and Rodriguez-Yam (2006) and Davis, Lee and Rodriguez-Yam (2008) who suggest a genetic algorithm to detect change-points in a piecewise-constant AR model or non-linear processes, respectively, where the MDL criterion is used.

The estimation of change-points can also be formulated as a problem of minimising a series of univariate cost functions, detecting a single change-point and then progressively moving to identify more. The Binary Segmentation method (BS) belongs to this category and uses a test statistic (such as the CUSUM) to reject the null hypothesis of no change-point. The BS has been widely used

and the main reasons are its low computational complexity and the fact that it is conceptually easy to implement: after identifying a change-point the detection of further change-points continues to the left and to the right of the initial change-point until no further changes are found.

The BS method has been adopted to solve different types of problems. Inclan and Tiao (1994) detect breaks in the variance of a sequence of independent observations; Berkes, Gombay and Horváth (2009) use a weighted CUSUM to reveal changes in the mean or the covariance structure of a linear process; Lee, Na and Na (2003) apply the test in the residuals obtained from a least squares estimator; and Kim, Cho and Lee (2000) and Lee and Park (2001) extend the Inclan and Tiao (1994) method to a GARCH(1,1) model and linear processes, respectively. A common factor of most of these methods is the estimation of the long-term variance or autocovariance, a rather difficult task when the observations are dependent. Cho and Fryzlewicz (2012) apply the binary segmentation method on the wavelet periodograms to detect change-points in the second-order structure of a non-stationary process. Using the wavelet periodogram, Killick, Eckley and Jonathan (2013) propose a likelihood ratio test under the null and alternative hypotheses. The authors apply the binary segmentation algorithm but assume an upper bound for the number of change-points. Fryzlewicz and Subba Rao (2014) adopt the binary segmentation search to test for multiple change-points in a piecewise constant ARCH model. BS is also used for multivariate (possibly high-dimensional) time series segmentation in Cho and Fryzlewicz (2015) and in Schröder and Fryzlewicz (2013) in the context of trend detection for financial time series.

In this paper we develop a detection method to estimate the number and locations of change-points in the second-order structure of a piecewise stationary time series model using the non-parametric Locally Stationary Wavelet (LSW) process of Nason, Von Sachs and Kroisandt (2000). The LSW model provides a complete description of the second-order structure of a stochastic process and, hence, it permits a fast estimation of the local autocovariance through the evolutionary wavelet spectrum. This choice, however, should not be seen as a restriction, potentially other models can form the basis for our algorithm.

In order to implement the change-point detection we adopt the Wild Binary Segmentation (WBS) method, proposed in the signal+iid Gaussian noise setup by Fryzlewicz (2014), which attempts to overcome the limitations of the BS method. Our motivation for doing so is the good practical performance of the WBS method in this setting. Under specific models in which many change-points are present the BS search may be inefficient in detecting them. This stems from the fact that the BS starts its search assuming a single change-point. To correct this limitation, Fryzlewicz (2014) proposes the WBS algorithm that involves a

“certain random localisation mechanism”. His method can be summarised as follows. At the beginning of the algorithm the CUSUM statistic is over M local segments $[s, e]$. The starting s and ending e points are randomly drawn from a uniform distribution $U(0, T - 1)$ and the hope is that for a large enough M , at least some of the intervals drawn will only contain single change-points, and therefore be particularly suitable for CUSUM-based detection. The location where the largest maximum CUSUM over all intervals drawn is achieved serves as the first change-point candidate. The method then proceeds similarly to BS: if the obtained CUSUM statistic exceeds a threshold then it is deemed to be a change-point and the same procedure continues to its left and right.

In order to adapt the WBS technique to our aim of detecting change-points in the second-order structure of a time series, we first adapt WBS for use in the multiplicative model setting, where the input sequence is a typically autocorrelated random scaled χ_1^2 -distributed sequence with a piecewise constant variance. This is more challenging to achieve than in the standard BS setting (Cho and Fryzlewicz (2012)) due to the fact that many of the intervals considered are short, which typically causes spurious behaviour of the corresponding CUSUM statistics. This phenomenon does not arise in the signal+iid Gaussian noise setting (Fryzlewicz (2014)) and is entirely due to the distributional features of the above multiplicative setting. This challenge requires a number of new solutions that include introducing the smallest possible interval length, and limiting the permitted “unbalancedness” of the CUSUM statistics, which is achieved without a detrimental effect on their operability thanks to the suitably large number of intervals of differing lengths considered at each stage by WBS.

Change-point detection for the second-order structure of a time series is achieved by combining information from local wavelet periodograms (each of which can be viewed as coming from the multiplicative model described above) across the resolution scales at which they were computed. We introduce a new way of combining this information across the scales, with the aim of further improving the practical performance of the methodology.

We attempt here to convey a new *modus operandi* in time series analysis whereby, in order to solve a specific problem, a large number of simple problems are solved on sub-samples of the data of differing lengths, and then the results combined to create an overall answer. Here, this is done via the WBS technique with the aim of detecting change-points, but related techniques could be envisaged e.g. for trend and seasonality detection, stationarity testing, or forecasting. We hope that our work stimulates further work in this direction.

The paper is structured as follows: In Section 2 we present and review the WBS algorithm in the context of time series. The reasons for selecting the LSW model as the core of our detection algorithm are given in Section 3. The main

algorithm is presented in Section 4, along with its theoretical consistency in estimating the number and locations of change-points. In addition, we conduct a simulation study to examine the performance of the algorithm; the results are given in Section 5. In Section 6 we apply our method to two data sets. Proofs of our results are in Section 5 of the supplementary material. Our methodology is implemented in the R package `wbsts`, available from CRAN.

2. The Wild Binary Segmentation Algorithm

The BS algorithm for a stochastic process was first introduced by Vostrikova (1981) who showed its consistency for the number and locations of change-points for a fixed N . A proof of its consistency is also given by Venkatraman (1992) for the Gaussian function+noise model, though the rates for the locations of the change-points are suboptimal. Improved rates of convergence of the locations of the change-points for the BS method are given by Fryzlewicz (2014).

Before considering segmentation in the model at (3.1) we first examine a multiplicative model

$$Y_{t,T}^2 = \sigma_{t,T}^2 Z_{t,T}^2, \quad t = 0, \dots, T-1, \quad (2.1)$$

where $\sigma_{t,T}^2$ is a piecewise constant function and the series $Z_{t,T}$ are possibly auto-correlated standard normal variables. This generic set-up is of interest because the wavelet periodogram, used later in the segmentation of (3.1), follows (2.1), up to a small amount of bias which we show can provably be neglected.

A potential change-point b_0 on a segment $[s, e]$ is given by

$$b_0 = \arg \max_b \left| \frac{\tilde{Y}_{s,e}^b}{q_{s,e}} \right|,$$

where $\tilde{Y}_{s,e}^b$ is the CUSUM statistic

$$\tilde{Y}_{s,e}^b = \sqrt{\frac{e-b}{n(b-s+1)}} \sum_{t=s}^b Y_t^2 - \sqrt{\frac{b-s+1}{n(e-b)}} \sum_{t=b+1}^e Y_t^2, \quad (2.2)$$

$q_{s,e} = \sum_{t=s}^e Y_t^2/n$, and $n = e - s + 1$. It can be shown that b_0 is the least-squares estimator of the change-point location in the case of $[s, e]$ containing exactly one change-point.

The value $|\tilde{Y}_{s,e}^{b_0}/q_{s,e}| = \max_b |\tilde{Y}_{s,e}^b/q_{s,e}|$ is tested against a threshold ω_T in order to decide whether the null hypothesis of no change-point is rejected or not. The BS proceeds by recursively applying the above CUSUM on $[s, b_0]$ and $[b_0 + 1, e]$. The algorithm stops in each current interval when no further change-points are detected, the obtained CUSUM values fall below threshold ω_T .

The BS method has the disadvantage of possibly fitting the wrong model when multiple change-points are present as it searches the whole series. The application of the CUSUM statistic (2.2) can result in spurious change-point detection when e.g. the true change-points occur close to each other. Especially, the BS method can fail to detect a small change in the middle of a large segment (Olshen et al. (2004)) which is illustrated in Fryzlewicz (2014).

Fryzlewicz (2014) proposes a randomised binary segmentation (termed Wild Binary Segmentation – WBS) where the search for change-points proceeds by calculating the CUSUM statistic in smaller segments whose length is random. By doing so, the user is guaranteed, with probability tending to one with the sample size, to draw favourable intervals containing at most a single change-point, which means the CUSUM statistic is an appropriate one to use over those intervals from the point of view of model choice. The maximum of the CUSUM statistics in absolute value, taken over a large collection of random intervals (see Figure 1 for an illustration), is considered to be the first change-point candidate, and is tested for significance. The binary segmentation procedure is not altered, after identifying a change-point the problem is divided into two sub-problems where for each segment we again test for further change-points in the same way. The computational complexity of the method can be reduced by noticing that the randomly drawn intervals and their corresponding CUSUM statistics can be calculated once at the start of the algorithm. Then, as the algorithm proceeds at a generic segment $[s, e]$, the obtained statistics can be reused making sure the random starting and end points fall within $[s, e]$.

The main steps of the WBS algorithm, modified for the model (2.1), are outlined as follows.

- Calculate the CUSUM statistics over a collection of random intervals $[s_m, e_m]$. The starting and ending points are sampled from a uniform distribution making sure that

$$e_m \geq s_m + \Delta_T, \quad (2.3)$$

where $\Delta_T > 0$ is the minimum size of the interval drawn.

$\mathcal{M}_{s,e}$ is the set of indices m of all random intervals $[s_m, e_m]$, $m = 1, \dots, M$, such that $[s_m, e_m] \subseteq [s, e]$; then the likely location of a change-point is

$$(m_0, b_0) = \arg \max_{(m \in \mathcal{M}_{s,e}, b \in s_m, \dots, e_m - 1)} \left| \frac{\tilde{Y}_{s_m, e_m}^b}{q_{s_m, e_m}} \right| \quad (2.4)$$

such that

$$\max \left(\frac{e_{m_0} - b_0}{e_{m_0} - s_{m_0} + 1}, \frac{b_0 - s_{m_0} + 1}{e_{m_0} - s_{m_0} + 1} \right) \leq c_\star, \quad (2.5)$$

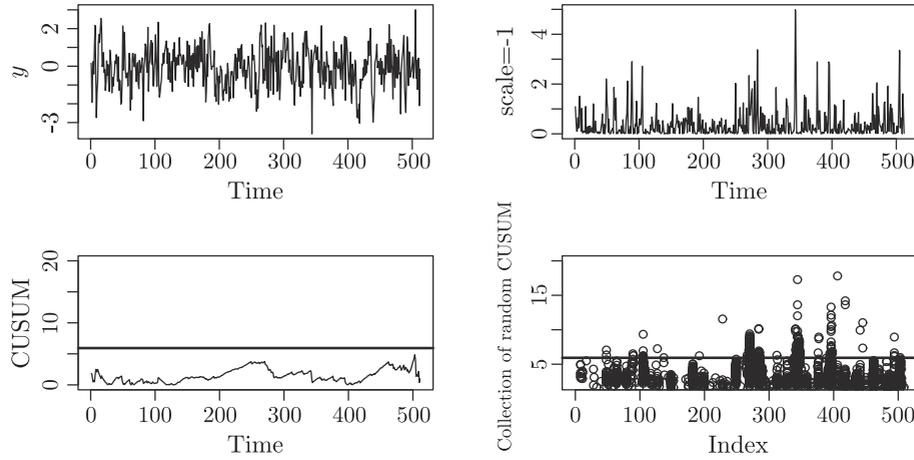


Figure 1. A simulated series (top-left) of an AR(1) model $y_t = \phi_t y_{t-1} + \varepsilon_t$ with $\phi_t = (0.5, 0.0)$ and change-points at $\{50, 100, \dots, 450\}$. The Wavelet Periodogram at scale -1 (top-right). The CUSUM statistic of scale -1 (bottom-left) as in the BS method; the black horizontal line is threshold $C \log(T)$ where T is the sample size (selection of C is discussed in Section 4.4). The rescaled \mathbb{Y}_{s_m, e_m}^b for $m \in \mathcal{M}_{s, e}$ and $b \in s_m, \dots, e_m - 1$ (bottom-right) as in the WBS method; the black horizontal line is the same threshold.

where c_\star is a constant, $c_\star \in [2/3, 1)$. The conditions (2.3) and (2.5) do not appear in Fryzlewicz (2014), but they necessary in the multiplicative model (2.1).

- The obtained CUSUM values are rescaled and tested against a threshold ω_T . This ensures that with probability tending to one with the sample size, only the significant change-points survive. The choice of the threshold ω_T is discussed in Section 4. If the obtained CUSUM statistic is significant the search is continued to the left and to the right of b_0 ; otherwise the algorithm stops. This step differs from the original WBS method of Fryzlewicz (2014) in that the CUSUM statistics are rescaled using q_{s_m, e_m} so that ω_T does not depend on $\sigma_{t, T}^2$.

Sampling distributions other than the uniform are possible and lead to practically the same theoretical results. The uniform distribution plays a special role here as it provides a natural and fair subsampling of the set of *all* possible sub-intervals of $[s, e]$.

3. Locally Stationary Wavelets and the Multiplicative Model

In this section we consider the LSW modelling paradigm of Nason, Von Sachs and Kroisandt (2000). The LSW process enables a time-scale decomposition of

a process and thus permits a rigorous estimation of the evolutionary wavelet spectrum and the local autocovariance; it can be seen as an alternative to the Fourier based approach for modelling time series.

Following Fryzlewicz and Nason (2006), 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

$$X_{t,T} = \sum_{i=-\infty}^{-1} \sum_{k=-\infty}^{\infty} W_i\left(\frac{k}{T}\right) \psi_{i,t-k} \xi_{i,k}. \quad (3.1)$$

Here $i \in -1, -2, \dots$ and $k \in \mathbb{Z}$ are, respectively, scale and location parameters, $(\psi_{i,0}, \dots, \psi_{i,\mathcal{L}-1})$ are discrete, real-valued, compactly supported, non-decimated wavelet vectors with support length $\mathcal{L} = O(2^{-i})$, and the $\xi_{i,k}$ are zero-mean, orthonormal, identically distributed random variables. In this set-up we replace the Lipschitz-continuity constraint on $W_i(z)$ by piecewise constant constraint that allows us to model a process whose second-order structure evolves in a piecewise constant manner over time with a finite but unknown number of change-points. Let L_i be the total magnitude of change-points in $W_i^2(z)$, then the functions $W_i(z)$ satisfy

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

The simplest type of a wavelet system that can be used in (3.1) is that of the Haar wavelets. Specifically,

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

for $i = -1, -2, \dots, k \in \mathbb{Z}$ where $\mathbb{I}_A(k)$ is 1 if $k \in A$ and 0 otherwise. Small absolute values of the scale parameter i denote ‘‘fine’’ scales, while large ones denote ‘‘coarser’’ scales. In fine scales the wavelet vectors are most oscillatory and localised, in coarser scales one has longer, less oscillatory wavelet vectors. Throughout, we only use Haar wavelets, using any other compactly supported wavelets would be less straightforward due to the unavailability of a closed formula for their coefficient values.

We take that the $\xi_{i,k}$ are distributed as $N(0, 1)$ even though extensions to other cases are possible but would entail consideration of quadratic forms of correlated non-Gaussian variables.

Of main interest in the LSW set-up is the Evolutionary Wavelet Spectrum (EWS) $S_i(z) = W_i^2(z)$, $i = -1, -2, \dots$, defined on the rescaled-time interval $z \in [0, 1]$. The estimation of the EWS is done through the wavelet periodogram (Nason, Von Sachs and Kroisandt (2000)).

Definition. Let $X_{t,T}$ be an LSW process constructed using the wavelet system ψ . The triangular stochastic array

$$I_{t,T}^{(i)} = \left| \sum_s X_{s,T} \psi_{i,s-t} \right|^2 \quad (3.2)$$

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

The wavelet periodogram is a convenient statistic: wavelet periodograms are fast to compute; for Gaussian processes X_t , they arise as χ_1^2 -type sequences that are easier to segment than, for example, empirical autocovariance sequences of the type $\{X_t X_{t+\tau}\}_t$; wavelets “decorrelate” a wide range of time series dependence structures; and the expectations of wavelet periodograms encode, in a one-to-one way, the entire autocovariance structure of a time series, so it suffices to estimate change-points in those expectations to obtain segmentation of the autocovariance structure of X_t , our ultimate goal.

We recall two definitions from Nason, Von Sachs and Kroisandt (2000): the autocorrelation wavelets $\Psi_i(\tau) = \sum_k \psi_{i,k} \psi_{i,k-\tau}$ and the autocorrelation wavelet inner product matrix $A_{i,k} = \sum_\tau \Psi_i(\tau) \Psi_k(\tau)$. Fryzlewicz and Nason (2006) show that $\mathbb{E}I_{t,T}^{(i)}$ is “close” (in the sense that the integrated squared bias converges 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 N change-points, whose set is denoted by \mathcal{N} . Every change-point in the autocovariance structure of the time series results in a change-point in at least one of the $\beta_i(z)$; therefore, detecting a change-point in the wavelet periodogram implies a change-point in the autocovariance structure of the process.

In addition, each wavelet periodogram ordinate is a squared wavelet coefficient of a standard Gaussian time series, and satisfies

$$I_{t,T}^{(i)} = \mathbb{E}I_{t,T}^{(i)} Z_{t,T}^2, \quad (3.3)$$

where $\{Z_{t,T}\}_{t=0}^{T-1}$ are autocorrelated standard normal variables (equivalently, the distribution of the squared wavelet coefficient $I_{t,T}^{(i)}$ is that of a scaled χ_1^2 variable). Then, the quantities $I_{t,T}^{(i)}$ and $\mathbb{E}I_{t,T}^{(i)}$ can be seen as special cases of $Y_{t,T}^2$ and $\sigma_{t,T}^2$, respectively, of the multiplicative model (2.1). To enable the application of the model (3.3) in this context, we need a condition.

(A0): $\sigma_{t,T}^2$ is deterministic and “close” to a piecewise constant function $\sigma^2(t/T)$ (apart from intervals around the discontinuities in $\sigma^2(t/T)$ which have length at

most $K2^{-i}$) in the sense that $T^{-1} \sum_{t=0}^{T-1} |\sigma_{t,T}^2 - \sigma^2(t/T)|^2 = o(\log^{-1} T)$, where the rate of convergence comes from the integrated squared bias between $\beta_i(t/T)$ and $\mathbb{E}I_{t,T}^{(i)}$ (see Fryzlewicz and Nason (2006)).

4. The Algorithm

In this section we present the WBS algorithm within the framework of the LSW model. Consider the CUSUM-type statistic

$$\mathbb{Y}_{s_m, e_m}^{b(i)} = \sqrt{\frac{e_m - b}{n(b - s_m + 1)}} \sum_{t=s_m}^b I_{t,T}^{(i)} - \sqrt{\frac{b - s_m + 1}{n(e_m - b)}} \sum_{t=b+1}^{e_m} I_{t,T}^{(i)}, \quad (4.1)$$

where the subscript $(\cdot)_m$ denotes an element chosen randomly from the set $\{0, \dots, T-1\}$ as in (2.3), $n = e_m - s_m + 1$, and $I_{t,T}^{(i)}$ are the wavelet periodogram ordinates at scale i that form the multiplicative model $I_{t,T}^{(i)} = \mathbb{E}I_{t,T}^{(i)} Z_{t,T}^2$ discussed in Section 3. The likely location of a change-point b_0 is then given by (2.4).

The following stages summarise the recursive procedure.

Stage I: Start with $s = 1$ and $e = T$.

Stage II: Examine whether $h_{m_0} = |\mathbb{Y}_{s_{m_0}, e_{m_0}}^{b_0}| / q_{s_{m_0}, e_{m_0}} > \omega_T = C \log(T)$, where $q_{s_{m_0}, e_{m_0}} = \sum_{t=s_{m_0}}^{e_{m_0}} I_{t,T}^{(i)} / n_{m_0}$, $n_{m_0} = e_{m_0} - s_{m_0} + 1$, and m_0, b_0 is as in (2.4); C is a parameter that remains constant and only varies between scales. Define $h'_{m_0} = h_{m_0} \mathbb{I}(h_{m_0} > \omega_T)$ where $\mathbb{I}(\cdot)$ is 1 if the inequality is satisfied and 0 otherwise.

Stage III: If $h'_{m_0} > 0$, then add b_0 to the set of estimated change-points; otherwise if $h'_{m_0} = 0$, stop the algorithm.

Stage IV: Repeat stages II-III to each of the two segments $(s, e) = (1, b_0)$ and $(s, e) = (b_0 + 1, T)$ if their length is more than Δ_T .

The choice of parameters C and Δ_T is described in Section 4.4. In addition to the random intervals $[s_m, e_m]$ we also include into $\mathcal{M}_{s,e}$ the index (labelled 0) corresponding to the interval $[s, e]$. This does *not* mean that the WBS procedure “includes” the classical BS, as at the first stage the WBS and BS are not guaranteed to locate the same change-point (even if WBS also examines the full interval $[s, e]$), so the two procedures can “go their separate ways” after examining the first full interval. The reason for manually including the full interval $[s, e]$ is that if there is at most one change-point in $[s, e]$, considering the entire interval $[s, e]$ is an optimal thing to do.

We expect that finer scales will be more useful in detecting the number and locations of the change-points in $\mathbb{E}I_{t,T}^{(i)}$, for as we move to coarser scales the

autocorrelation within $I_{t,T}^{(i)}$ becomes stronger and the intervals on which a wavelet periodogram sequence is not piecewise constant become longer. Hence, we select the scale $i < -I^*$, where $I^* = \lfloor \alpha \log \log T \rfloor$ and $\alpha \in (0, 3\lambda]$ for $\lambda > 0$, such that the consistency of our method is retained.

In stage II, we rescale the statistic h_{m_0} before we test it against the threshold. This division plays the role of stabilising the variance, which is exact in the multiplicative model in which the observations are independent, over intervals where the variance is constant. In all other cases, the variance stabilisation cannot be guaranteed to be exact, but in the case where the process under consideration is stationary over the given interval, the cancellation of the variance parameter still takes place and therefore the distribution of the rescaled CUSUM is a function of the *autocorrelation* of the process, rather than its entire *autocovariance*. This reduces the difficulty in choosing the threshold parameter ω_T and one can still hope to obtain “universal” thresholds that work well over a wide range of dependence structures. Exact variance stabilisation in the non-independent case would require estimating what is referred to as the “long-run variance” parameter, the variance of the sample mean of a time series, which is a difficult problem in time series analysis. If we were to pursue it, the estimation error would likely not make it worthwhile - we choose this rescaling method as a compromise between doing nothing and having to estimate the long-run variance. In Section 4 of the supplementary material, we illustrate the essence of this issue, and provide a simple but informative example showing that the variance stabilisation, performed as described above, is desirable, despite this “non-exactness” problem.

Horváth, Horváth and Hušková (2008) propose a similar type of CUSUM statistic that does not require an estimate of the variance of a stochastic process by using the ratio of the maximum of two local means. The authors apply the method to detect a single change-point in the mean of a stochastic process under independent, correlated, or heteroscedastic error settings.

4.1. Technical assumptions and consistency

In this section we present the consistency theorem for the WBS algorithm for the total number N and locations of the change-points $0 < \eta_1 < \dots < \eta_N < T - 1$ with $\eta_0 = 0$ and $\eta_{N+1} = T$. We need some assumptions.

(A1): $0 < \sigma^2(t/T) < \sigma^* < \infty$ where $\sigma^* \leq \max_{t,T} \sigma^2(t/T)$. The number of change-points N in (2.1) is unknown and allowed to increase with T , only the minimum distance between the change-points can restrict the maximum number of N .

(A2): $\{Z_{t,T}\}_{t=0}^{T-1}$ is a sequence of standard Gaussian variables and the autocorrelation function $\rho(\tau) = \sup_{t,T} |\text{cor}(Z_{t,T}, Z_{t+\tau,T})|$ is absolutely summable.

(A3): The distance between any two adjacent change-points satisfies $\min_{r=1,\dots,N+1} |\eta_r - \eta_{r-1}| \geq \delta_T$, where $\delta_T \geq C \log^2 T$ for a large enough C .

(A4): The magnitudes of the change-points satisfy $\inf_{1 \leq r \leq N} |\sigma((\eta_r + 1)/T) - \sigma(\eta_r/T)| \geq \sigma_*$ where $\sigma_* > 0$.

(A5): $\Delta_T \asymp \delta_T$ where Δ_T as defined in (2.3).

Theorem 1. *Let $Y_{t,T}^2$ follow model (2.1), and suppose that Assumptions (A0)–(A5) hold. With the number of change-points in $\sigma^2(t/T)$ as N and the locations of those change-points as η_1, \dots, η_N , let \hat{N} and $\hat{\eta}_1, \dots, \hat{\eta}_N$ be the number and locations of the change-points (in ascending order) estimated by the Wild Binary Segmentation algorithm. There exist constants C_1 and C_2 such that if $C_1 \log T \leq \omega_T \leq C_2 \sqrt{\delta_T}$, then $P(\mathcal{Z}_T) \rightarrow 1$, where*

$$\mathcal{Z}_T = \{\hat{N} = N; \max_{r=1,\dots,N} |\hat{\eta}_r - \eta_r| \leq C \log^2 T\}$$

for a certain $C > 0$. The guaranteed speed of convergence of $P(\mathcal{Z}_T)$ to 1 is no faster than $T\delta_T^{-1}(1 - \delta_T^2(1 - \bar{c})^2 T^{-2}/9)^M$ where M is the number of random draws and $\bar{c} = 3 - 2/c_*$.

The rate of convergence for the estimated change-points obtained for the BS method by Cho and Fryzlewicz (2015) is $\mathcal{O}(\sqrt{T} \log^{(2+\vartheta)} T)$ and $\mathcal{O}(\log^{(2+\vartheta)} T)$ for $\vartheta > 0$ when δ_T is $T^{3/4}$ and T , respectively. In the WBS setting, the rate is square logarithmic when δ_T is of order $\log^2 T$, which represents an improvement. In addition, the lower threshold is always of order $\log T$ regardless of the minimum space between the change-points.

We now discuss the issue of the minimum number M of random draws needed to ensure that the bound on the speed of convergence of $P(\mathcal{Z}_T)$ to 1 in Theorem 1 is suitably small. Suppose that we wish to ensure that

$$T\delta_T^{-1} \left(1 - \delta_T^2(1 - \bar{c})^2 \frac{T^{-2}}{9}\right)^M \leq T^{-1}.$$

As $\log(1 - y) \approx -y$ around $y = 0$, this is (practically) equivalent to

$$M \geq \frac{9T^2}{\delta_T^2(1 - \bar{c})^2} \log(T^2 \delta_T^{-1}).$$

In the “easiest” case, $\delta_T \sim T$, this results in a logarithmic number of draws, which leads to particularly low computational complexity. The required M progressively increases as δ_T decreases; our recommendations for the choice of M are discussed in Section 4.4.

4.2. Simultaneous across-scale post-processing

Theorem 1 covers the multiplicative model (2.1). We now consider change-point detection in the full model (3.1). To accomplish this we propose two methods.

Method 1: The search for further change-points in each interval (s_m, e_m) proceeds to the next scale $i - 1$ only if no change-points are detected at scale i on that interval. This ensures that the finest scales are preferred (since change-points detected at the finest scales are likely to be more accurate) and one only moves to coarser if necessary. Cho and Fryzlewicz (2012) use a similar technique to combine across scales change-points, but involving an extra parameter. Their method will be used as a benchmark for our first type of across-scale post-processing.

Method 2: The method simultaneously joins the estimated change-points across all the scales such that all the information from every scale is combined. Motivated by Cho and Fryzlewicz (2015), who propose an alternative aggregation method to these of Groen, Kapetanios and Price (2013) in order to detect change-points in the second order structure of a high-dimensional time series, we take

$$\mathbb{Y}_t^{thr} = \sum_{i=-I^*}^{-1} \mathcal{Y}_t^{(i)} \mathbb{I}(\mathcal{Y}_t^{(i)} > \omega_T^{(i)}) \text{ for } i = -1, \dots, -I^*, \quad (4.2)$$

where $\mathcal{Y}_t^{(i)} = |\mathbb{Y}_{s_m, e_m}^{b(i)}|/q_{s_m, e_m}^{(i)}$. This statistic differs from that of Cho and Fryzlewicz (2015) in that it applies across the scales $i = -1, -2, \dots, -I^*$ of a univariate time series.

The algorithm is identical to the algorithm in Section 4 except for replacing (4.1) with (4.2). In addition, if the obtained $\mathbb{Y}_t^{thr} > 0$ there is no need to test further for the significance of b_0 .

Theorem 2. *Let X_t follow model (3.1), and suppose that Assumptions (A0)–(A5) for $\sigma^2(t/T)$ hold for each $\beta_i(z)$. With the number of change-points in $\beta_i(z)$ as N and the locations of those change-points as $\theta_1, \dots, \theta_N$, let \hat{N} and $\hat{\theta}_1, \dots, \hat{\theta}_N$ be the number and locations of the change-points (in ascending order) estimated by the across-scale post-processing Method 1 or 2. There exist constants C_3 and C_4 such that if $C_3 \log T \leq \omega_T \leq C_4 \delta_T$, then $P(\mathcal{U}_T) \rightarrow 1$, where*

$$\mathcal{U}_T = \left\{ \hat{N} = N; \max_{r=1, \dots, N} |\hat{\theta}_r - \theta_r| \leq C' \log^2 T \right\}$$

for a certain $C' > 0$, where the guaranteed speed of convergence is the same as that in Theorem 1.

Methods 1 and 2 achieve the same rate of convergence for the estimated change-points; their relative performance is empirically examined in Section 5.

4.3. Post-processing

In order to control the number of change-points estimated from the WBS algorithm, and to reduce the risk of over-segmentation, we propose a post-processing method similar to Cho and Fryzlewicz (2012) and Inclan and Tiao

(1994). More specifically, we compare every change-point against the adjacent ones using the CUSUM statistic, making sure that (2.5) is satisfied. Thus, for a set $\hat{\mathcal{N}} = \{\hat{\theta}_0, \dots, \hat{\theta}_{N+1}\}$ where $\hat{\theta}_0 = 0$ and $\hat{\theta}_{N+1} = T$, we test whether $\hat{\theta}_r$ satisfies

$$\mathbb{Y}_t^{thr} = \sum_{i=-I^*}^{-1} \mathcal{Y}_t^{(i)} \mathbb{I}(\mathcal{Y}_t^{(i)} > \omega_T^{(i)}) > 0 \text{ for } i = -1, \dots, -I^*,$$

where $\mathcal{Y}_t^{(i)} = |\mathbb{Y}_{\hat{\theta}_{r-1}, \hat{\theta}_{r+1}}^{\hat{\theta}_r(i)}| / q_{\hat{\theta}_{r-1}, \hat{\theta}_{r+1}}^{(i)}$ and

$$\max \left(\frac{\hat{\theta}_{r+1} - \hat{\theta}_r}{\hat{\theta}_{r+1} - \hat{\theta}_{r-1} + 1}, \frac{\hat{\theta}_r - \hat{\theta}_{r-1} + 1}{\hat{\theta}_{r+1} - \hat{\theta}_{r-1} + 1} \right) \leq c_*. \quad (4.3)$$

If $\mathbb{Y}_t^{thr} = 0$ then change-point $\hat{\theta}_r$ is temporarily eliminated from set $\hat{\mathcal{N}}$. In the next run, when considering change-point $\hat{\theta}_{r+1}$, the adjacent change-points are $\hat{\theta}_{r-1}$ and $\hat{\theta}_{r+2}$. When the post-processing finishes its cycle all temporarily eliminated change-points are reconsidered using as adjacent change-points those that have survived the first cycle. It is necessary for $\hat{\theta}_r$ to satisfy (4.3) with its adjacent estimated change-points $\hat{\theta}_{r-1}$ and $\hat{\theta}_{r+1}$, otherwise it is never eliminated. The algorithm is terminated when the set of change-points does not change.

4.4. Choice of threshold and parameters

In this section we present choices of the parameters involved in the algorithms. From Theorems 1 and 2 we have that the threshold ω_T includes the constant $C^{(i)}$ that varies between the scales. The values of $C^{(i)}$ are the same for all the methods presented, either BS/WBS or the Methods 1 and 2 in Section 4.2. Therefore, we can use the thresholds by Cho and Fryzlewicz (2012) who conduct experiments to establish the value of the threshold parameter under the null hypothesis of no change-points such that when the obtained statistic exceeds the threshold the null hypothesis is rejected. However, we go to the experiments described below.

We generate a vector $\mathbf{X} \sim N(0, \Sigma)$ where the covariance matrix $\Sigma = (\sigma_{\kappa, \kappa'})_{\kappa, \kappa'=1}^T$ and $\sigma_{\kappa, \kappa'} = \rho^{|\kappa - \kappa'|}$. Then we find v that maximises (4.1). The ratio

$$C_T^{(i)} = \mathbb{Y}_v^{(i)} (\log T)^{-1} \left(\sum_{t=1}^T I_{t,T}^{(i)} \right)^{-1} T$$

gives us an insight into the magnitude of parameter $C^{(i)}$. We repeated this experiment for different values of ρ and, for every scale i , we selected $C^{(i)}$ as the 95% quantile. The same values were used for the post-processing method explained in Section 4.3. Our results indicated that $C^{(i)}$ tends to increase as we move

to coarser scales due to the increasing dependence in the wavelet periodogram sequences. Since our method applies to non-dyadic structures it is reasonable to propose a general rule that will apply in most cases. To accomplish this we repeated the simulation study above for $T = 50, 100, \dots, 6,000$. Then, for each scale i we fitted the regression

$$C^{(i)} = c_0^{(i)} + c_1^{(i)}T + c_2^{(i)}\frac{1}{T} + c_3^{(i)}T^2 + \varepsilon.$$

The adjusted R^2 was above 90% for all the scales. Having estimated the values for $\hat{c}_0^{(i)}, \hat{c}_1^{(i)}, \hat{c}_2^{(i)}, \hat{c}_3^{(i)}$ we were able to use fitted values for any sample size T . For samples larger than $T = 6,000$, we used the same $C^{(i)}$ values as for $T = 6,000$.

Based on empirical evidence (see the supplementary material) we selected the scale I^* by setting $\lambda = 0.7$. In stage III of the algorithm, the procedure was terminated when either the CUSUM statistic did not exceed a certain threshold or the length of the respective segment was Δ_T . This also defined the minimum length of a favourable draw from (2.3). We chose Δ_T to be of the same order as δ_T since this is the lowest permissible order of magnitude according to (A5). Practically, we found that the choice $\Delta_T = \lfloor \log^2 T/3 \rfloor$ worked well. In addition, a simulation study found in the supplementary material provides empirical arguments for the choice $c_* = 0.75$. The main idea of this parameter is to ensure that long enough stretches of data are included in the computation of our CUSUM statistics, else the computed CUSUM statistics will be too variable to be reliable. This is particularly important in the autocorrelated multiplicative setting where there tends to be a large amount of noise, so the use of such a parameter is needed to suppress the variance of the CUSUM statistics. Finally, our recommendation for the parameter M is 3,500 when T does not exceed 10,000. These values are used in the remainder of the paper.

5. Simulation Study

We present a set of simulation studies to assess the performance of our methods. In all the simulations we took sample size to be 1,024 over 100 iterations. In the supplementary material smaller and larger sample sizes are also considered. For comparison we also report the performance of the method of Cho and Fryzlewicz (2012) – henceforth CF – using the default values specified in their paper. BS1 and BS2 refer to Method 1 and Method 2 of aggregation (as described in Section 4.2) using the BS technique. WBS1 and WBS2 refer to the Method 1 and Method 2 of aggregation (as in Section 4.2) using the Wild Binary Segmentation technique. To briefly illustrate computation times, our code, executed on a standard PC, runs in approximately 25 seconds for a time series of length 10,000 with 10 change-points.

Table 1. Stationary processes results. For all the models the sample size was 1,024 and there were no change-points. Figures show the number of occasions the methods detected change-points with the universal thresholds $C^{(i)}$ obtained as described in Section 4.4. Figures in brackets are the number of occasions the methods detected change-points with the thresholds $C^{(i)}$ obtained as described in Section 5.1.

Model	BS1	WBS1	BS2	WBS2	CF
S1: iid standard normal	1 [0]	3 [2]	0 [0]	1 [0]	4
S2: AR(1) with parameter 0.9	3 [1]	5 [1]	1 [1]	5 [1]	9
S3: AR(1) with parameter -0.9	58 [0]	93 [0]	46 [0]	48 [5]	79
S4: MA(1) with parameter 0.8	2 [3]	7 [4]	3 [3]	1 [0]	7
S5: MA(1) with parameter -0.8	2 [0]	4 [2]	4 [0]	0 [0]	7
S6: ARMA(1,0,2) with AR= $\{-0.4\}$ and MA= $\{-0.8, 0.4\}$	8 [0]	27 [0]	8 [0]	8 [0]	25
S7: AR(2) with parameters 1.39 and -0.96	88 [3]	99 [4]	88 [3]	88 [5]	96

5.1. Models with no change-points

We simulated stationary time series with innovations $\varepsilon_t \sim N(0, 1)$ and we report the number of occasions (out of 100) the methods incorrectly rejected the null hypothesis of no change-points. The models S1-S7 (Table 1) are taken from Nason (2013).

The results of Table 1 indicate our methods' better performance than that of Cho and Fryzlewicz (2012), apart from models S3 and S7 where all methods incorrectly reject the null hypothesis on many occasions. A visual inspection of an AR(1) process with $\phi = -0.9$ would confirm that this type of process exhibits a "clustering behaviour" which mimics changing variance. Hence, the process is interpreted as non-stationary by the wavelet periodogram resulting in erroneous outcomes. A similar argument is valid for S7 model. To correct for that limitation, $C^{(i)}$ should be chosen with care. Higher values ensure that the null hypothesis is not rejected frequently. This is achieved by not using universal thresholds (as shown in Section 4.4) but calculating them for every instance. Specifically, given a time series y_t we fit an AR(p) model. Then we generated 100 instances of the same length and with the same AR(p) coefficients. Similarly with Section 4.4 we selected $C^{(i)}$ as the 95% quantile. This procedure was more computationally intensive but improved the method significantly; see the figures in brackets (Table 1). An alternative approach in obtaining thresholds, by taking time-averages of spectrum values for each $i = -1, -2, \dots, -I^*$ and then simulating stationary models, described in the supplementary material, does well but not as well as our suggestion above.

5.2. Non-stationary models

We now examine the performance of our method for a set of non-stationary models by using and extending the examples from Cho and Fryzlewicz (2012). Since the WBS method has improved rates of convergence new simulation results are presented that assess how close the change-points are to their estimates. For this reason we report the total number of change-points identified within $\lfloor 5\% \cdot T \rfloor$ from the real ones. Results for $\lfloor 2.5\% \cdot T \rfloor$ distances are reported in the supplementary material.

The accuracy of a method should be also judged in parallel with the total number of change-points identified. We propose a test that tries to accomplish this. Assuming that we define the maximum distance from a real change-point η as d_{\max} , an estimated change-point $\hat{\eta}$ is correctly identified if $|\eta - \hat{\eta}| \leq d_{\max}$ (here within 5% of the sample size). If two (or more) estimated change-points are within this distance then only one change-point, the closest to the real change-point, is classified as correct. The rest are deemed to be false, except if any of these are close to another change-point. An estimator performs well when the hit ratio

$$HR = \frac{\#\text{correct change-points identified}}{\max(N, \hat{N})}$$

is close to 1. By using the term $\max(N, \hat{N})$ we aim to penalise cases where, for example, the estimator correctly identifies a certain number of change-points all within the distance d_{\max} , but $\hat{N} < N$. It also penalises the estimator when $\hat{N} > N$ and all \hat{N} estimated change-points are within distance d_{\max} of the true ones.

Tables 2 and 3 summarising the results, and histograms of the estimated change-point locations for every model, can be found in the supplementary material.

Model A: A non-stationary process that includes one AR(1) and two AR(2) processes with two clearly observable change-points

$$y_t = \begin{cases} 0.9y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 1 \leq t \leq 512, \\ 1.68y_{t-1} - 0.81y_{t-2} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 513 \leq t \leq 768, \\ 1.32y_{t-1} - 0.81y_{t-2} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 769 \leq t \leq 1,024. \end{cases}$$

BS2 was the best option, marginally ahead of WBS1 and WBS2. The fact that BS performed well here was not surprising given the fact that the change-points are far apart and prominent. However, it is reassuring to see the WBS methods also performed well.

Model B: A non-stationary process with two less clearly observable change-points

$$y_t = \begin{cases} 0.4y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 1 \leq t \leq 400, \\ -0.6y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 401 \leq t \leq 612, \\ 0.5y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 613 \leq t \leq 1,024. \end{cases}$$

The WBS methods did better, marginally ahead of the BS methods. This is again not unexpected given the fact that the change-points are less prominent than in Model A.

Model C: A non-stationary process with a short segment at the start

$$y_t = \begin{cases} 0.75y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 1 \leq t \leq 50, \\ -0.5y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 51 \leq t \leq 1,024. \end{cases}$$

Here BS2 and CF perform well compared to the BS1, WBS1, and WBS2 methods. It was expected that binary segmentation methods would perform better due to the fact that they start their search assuming a single change-point.

Model D: A non-stationary process similar to model B, but with the two change-points at a short distance from each other. In this model, the two change-points occur are close to each other, (400, 470) instead of (400, 612). The CF method, BS1, and BS2 did not perform well, as the two change-points were detected in less than half of the cases. By contrast, the WBS1 and WBS2 methods achieved high hit ratios (almost double that of the BS methods).

Model E: A highly persistent non-stationary process with time-varying variance

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

The CF and BS1 methods performed well as they detected most of the change-points within 5% distance from the real ones. From our simulations we noticed that in most cases the two change-points were found in the finest scale ($i = -1$). The aggregation Method 2 did not improve the estimation since its purpose is to simultaneously combine the information from different scales. On the other hand, the CF method and Method 1 favoured change-points detected in the finest scales leading to good performance.

Model F: A piecewise constant ARMA(1,1) process

$$y_t = \begin{cases} 0.7y_{t-1} + \varepsilon_t + 0.6\varepsilon_{t-1}, & \text{for } 1 \leq t \leq 125, \\ 0.3y_{t-1} + \varepsilon_t + 0.3\varepsilon_{t-1}, & \text{for } 126 \leq t \leq 532, \\ 0.9y_{t-1} + \varepsilon_t, & \text{for } 533 \leq t \leq 704, \\ 0.1y_{t-1} + \varepsilon_t - 0.5\varepsilon_{t-1}, & \text{for } 704 \leq t \leq 1,024. \end{cases}$$

The first change-point was the least apparent and was left undetected in most cases when applying the CF method. Our methods were capable of capturing this point more frequently, and within 5% from its real position.

Model G: A near-unit-root non-stationary process with time-varying variance

$$y_t = \begin{cases} 0.999y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1) & \text{for } 1 \leq t \leq 200, 401 \leq t \leq 600, \\ & \text{and } 801 \leq t \leq 1,024, \\ 0.999y_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, 1.5^2) & \text{for } 201 \leq t \leq 400 \text{ and } 601 \leq t \leq 800. \end{cases}$$

In this near-unit-root process there are four change-points in its variance. The binary segmentation methods did not perform well as they often missed the middle change-points. Both WBS1 and WBS2 managed to detect most of the change-points achieving a hit ratio almost three times higher than BS2. In almost 70% of the occasions WBS2 detected at least four change-points.

Model H: A non-stationary process similar to model F but with the three change-points at a short distance from each other. In this model the three change-points occur close to each other, $\mathcal{N} = (125, 325, 550)$. The first two failed to be detected by the CF in many instances. By contrast, BS1 and BS2 did well while WBS1 and WBS2 performed slightly better in this case by identifying them more often. This results in a higher hit ratio.

Model I: A non-stationary AR process with many changes within close distances. We simulated instances with five change-points occurring at uniformly distributed positions. We allowed the distances to be as small as 30 and not larger than 100.

In this scenario, CF correctly identified more than four change-points in 15% of the instances, while BS1 and BS2 did in 24% and 23%, respectively. Again, the WBS methods did well in revealing the majority of the change-points and in many cases close to the real ones.

In summary, the WBS methods offer a reliable default choice. In terms of the hit ratio, they performed the best or nearly the best in 7 of the 9 models studied, and did not perform particularly poorly in the other models, especially if the total number of detected change-points is taken into account. BS1, BS2 and CF performed poorly in at least three of the models. In terms of the hit ratio, both BS methods were best or close to it only in two models. Overall, the WBS methods seemed the clear winners here. Our recommendation to the user is to try the WBS2 method first, since overall it appears to be the most reliable one.

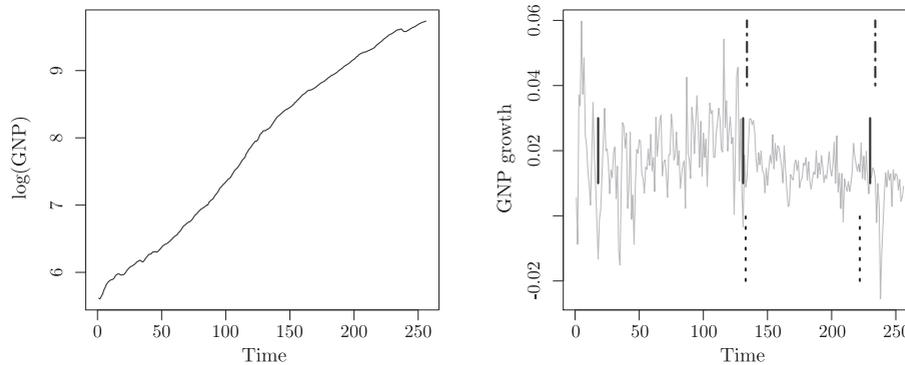


Figure 2. Natural logarithm of the GNP series (left) and its first difference (right). The top, middle and bottom vertical lines are the change-points as estimated by CF, WBS2 and BS2 respectively.

6. Applications

6.1. US Gross National Product series (GNP)

We obtained the GNP time series from the Federal Reserve Bank of St. Louis web page (<http://research.stlouisfed.org/fred2/series/GNP>). The seasonally adjusted and quarterly data is expressed in billions of dollars and spans from 1947:1 until 2013:1 but we only used the last 256 observations. In the left panel of Figure 2 one can see the logarithm of the GNP series. As in Shumway and Stoffer (2011), we only examine the first difference of the logarithm of the GNP (also called the growth rate) since there is an obvious linear trend. In the right panel of the same figure, which illustrates the growth rate, it is visually clear that the series exhibits less variability in its latter portion. We were interested in finding whether our method is capable of spotting this change and/or possibly others.

Applying BS2 and WBS2 (BS1 and WBS1 produced identical results) we found that BS2 detected two change-points $\hat{\eta} = \{133, 222\}$, while the WBS2 detected three at positions $\{18, 131, 230\}$. For comparison, CF detected two possible change-points $\hat{\eta} = \{134, 234\}$. The acf graphs (not shown here) confirm that there may be changes in the autocovariance structure occurring at all of these estimated change-points.

Change-point 18 1953(3), almost exactly coincides with a peak of the GNP growth as decided by the Business Cycle Dating Committee of the National Bureau of Economic Research; their official date is July 1953 (note that cycles do not necessarily overlap with the quarterly publications of the GNP). In addition, change-points 131, 133 and 134 lie within a cycle that peaks in January 1981 and has a trough in November 1982. This cycle corresponds to the start of the Great

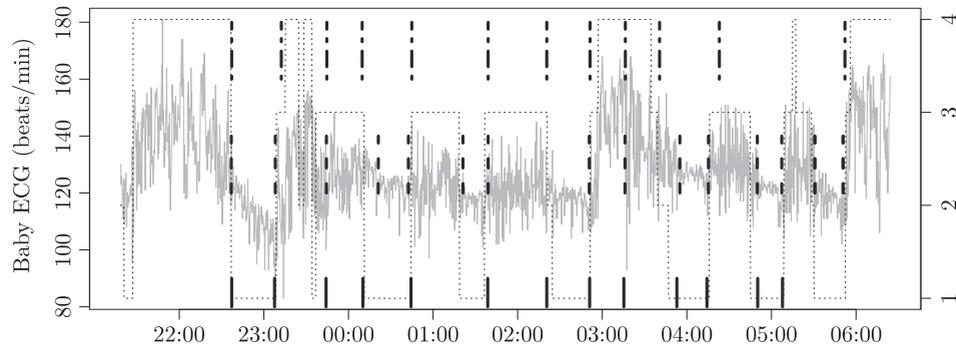


Figure 3. Plot of BabyECG data. The top, middle and bottom vertical lines are the change-points as estimated by CF, WBS2, and BS2, respectively. The horizontal dotted line represents the sleep states: 1 = quiet sleep, 2 = quiet-to-active sleep, 3 = active sleep, 4 =awake.

Moderation (around 1980s), a period that experienced more efficient monetary policy and shocks of small magnitude, see and references therein. Finally, we note that all three methods detected a change-point toward the end of the series - 222, 230, 234 (2004(3), 2006(3) and 2007(3)), respectively. According to e.g., Clark (2009) the Great Moderation had reversed and the decline was offset by negative growth rates due to the recent economic recession.

6.2. Infant Electrocardiogram Data (ECG)

We applied CF, BS2, and WBS2 to the ECG data of an infant, found at the *R* package *wavethresh*. This is a popular example of a non-stationary time series and it has been analysed in e.g. Nason, Von Sachs and Kroisandt (2000). The local segments of possible stationarity indicate the sleep state of the infant, classified on a scale from 1 to 4, see the caption of Figure 3. The same figure plots the time series with the respective estimated change-points (the methods were applied on the first difference so that its mean is approximately zero). All methods identify most of the sleep states and, notably, WBS2 detects an abrupt change of short duration (quiet sleep-awake-quiet sleep) toward the end of the series.

7. Conclusion

The paper has addressed the problem of detecting the change-points in the autocovariance structure of a univariate time series. There are many types of non-stationary time series that require segmentation methods. Using the WBS framework we are able to detect multiple change-points that are small in magnitude and/or close to each other. The simulation study in Section 5 indicates that the WBS mechanism performs well at this task.

Supplementary Materials

The online supplementary material contains additional simulation studies supporting the choice of the default parameters of our procedure, empirical performance evaluation for small and large samples and using other error measures, additional material on the variance stabilization, as well as proofs of our theoretical results.

Acknowledgements

Piotr Fryzlewicz's work was supported by the Engineering and Physical Sciences Research Council grant no. EP/L014246/1.

References

- Auger, I. and Lawrence, C. (1989). Algorithms for the optimal identification of segment neighborhoods. *Bull. Math. Biol.* **51**, 39-54.
- Bellman, R. and Dreyfus, S. (1966). *Applied Dynamic Programming*, Vol. 2. Princeton University Press, Princeton.
- Berkes, I. Gombay, E. and Horváth, L. (2009). Testing for changes in the covariance structure of linear processes. *J. Statist. Plan. Inference* **139**, 2044-2063.
- Cho, H. and Fryzlewicz, P. (2012). Multiscale and multilevel technique for consistent segmentation of nonstationary time series. *Statist. Sinica* **22**, 207-229.
- Cho, H. and Fryzlewicz, P. (2015). Multiple change-point detection for high-dimensional time series via Sparsified Binary Segmentation. *J. Roy. Statist. Soc. Ser. B* **77**, 475-507.
- Clark, T. (2009). Is the Great Moderation over? An empirical analysis. *Fed. Reserve Bank Kansas City Econ. Rev.* **94**, 5-42.
- Davies, S. and Bland, D. (2010). Interestingness detection in sports audio broadcasts. *IEEE Ninth International Conference on Machine Learning and Applications (ICMLA)*, 643-648.
- Davis, R., Huang, D. and Yao, Y. C. (1995). Testing for a change in the parameter values and order of an autoregressive model. *Ann. Stat.* **23**, 282-304.
- Davis, R., Lee, T. and Rodriguez-Yam, G. (2006). Structural break estimation for nonstationary time series models. *J. Amer. Statist. Assoc.* **101**, 223-239.
- Davis, R., Lee, T. and Rodriguez-Yam, G. (2008). Break detection for a class of nonlinear time series models. *J. Time Ser. Anal.* **29**, 834-867.
- Fryzlewicz, P. (2014). Wild binary segmentation for multiple change-point detection. *Ann. Statist.* **42**, 2243-2281.
- Fryzlewicz, P. and Nason, G. (2006). HaarFisz estimation of evolutionary wavelet spectra. *J. Roy. Statist. Soc. Ser. B* **68**, 311-634.
- Fryzlewicz, P. and Subba Rao, S. (2014). Multiple-change-point detection for auto-regressive conditional heteroscedastic processes. *J. Roy. Statist. Soc. Ser. B* **76**, 903-924.
- Gombay, E. (2008). Change detection in autoregressive time series. *J. Multivariate Anal.* **99**, 451-464.
- Gombay, E. and Serban, D. (2009). Monitoring parameter change in time series models. *J. Multivariate Anal.* **100**, 715-725.

- Groen, J., Kapetanios, G. and Price, S. (2013). Multivariate methods for monitoring structural change. *J. Appl. Econometr.* **28**, 250-274.
- Horváth, L., Horváth, Z. and Hušková, M. (2008). Ratio tests for change point detection. *Inst. Math. Statist.* **1**, 293-304.
- Inclan, 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.
- Jackson, B., Scargle, J., Barnes, D., Arabhi, S., Alt, A., Gioumouisis, P., Gwin, E., Sangtrakulcharoen, P., Tan, L. and Tsai, T. (2005). An algorithm for optimal partitioning of data on an interval. *IEEE Signal Process. Lett.* **12**, 105-108.
- Kay S. (1998). *Fundamentals of Statistical Signal Processing, Volume 2: Detection Theory*. Prentice Hall PTR, New Jersey.
- Killick, R., Fearnhead, P. and Eckley, I. A. (2012). Optimal detection of changepoints with a linear computational cost. *J. Amer. Statist. Assoc.* **107**, 1590-1598.
- Killick, R., Eckley, I. and Jonathan, P. (2013). A wavelet-based approach for detecting changes in second order structure within nonstationary time series cost. *Electron. J. Stat.* **7**, 1167-1183.
- Kim, S., Cho, S. and Lee, S. (2000). On the cusum test for parameter changes in GARCH (1, 1) models. *Commun. Stat. Theory Methods* **29**, 445-462.
- 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 Teyssiere, G. (2006). Detection of multiple change-points in multivariate time series. *LITH Math. J.* **46**, 287-306.
- Lavielle, M. and Teyssiere, G. (2007). Adaptive detection of multiple change-points in asset price volatility. *Long Memory in Economics*, 129-156.
- Lee, S. and Park, S. (2001). The cusum of squares test for scale changes in infinite order moving average processes. *Scand. J. Statist.* **28**, 325-644.
- Lee, S., Na, O. and Na, S. (2003). On the cusum of squares test for variance change in nonstationary and nonparametric time series models. *Ann. Inst. Stat. Math.* **55**, 467-485.
- Mercurio, D. and Spokoiny, V. (2004). Statistical inference for time-inhomogeneous volatility models. *Ann. Inst. Stat. Math.* **32**, 577-602.
- Nason, G. (2013). A test for second-order stationarity and approximate confidence intervals for localized autocovariances for locally stationary time series. *J. Roy. Statist. Soc. Ser. B* **75**, 879-904.
- Nason, G., 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.
- Olshen, A., Venkatraman, E. S., Lucito, R. and Wigler, M. (2004). Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics* **5**, 557-572.
- Ombao, H., Raz, J., von Sachs, R. and Malow, B. (2001). Automatic statistical analysis of bivariate nonstationary time series. *J. Amer. Statist. Assoc.* **96**, 543-560.
- Perron, P. (2006). Dealing with structural breaks. *Palgrave Handbook of Econometrics* **1**, 278-352.
- Schröder, A. L. and Fryzlewicz, P. (2013). Adaptive trend estimation in financial time series via multiscale change-point-induced basis recovery. *Statist. Interface* **6**, 449-461.
- Sen, A. and Srivastava, M. (1975). On tests for detecting change in mean. *Ann. Statist.* **3**, 98-108.

- Shumway, R. and Stoffer, D. (2011). *Time Series Analysis and Its Applications: with R examples*. Springer Science, New York.
- Stărică, C. and Granger, C. (2005). Nonstationarities in stock returns. *Rev. Econom. Statist.* **87**, 503-522.
- Venkatraman, E. (1992). Consistency results in multiple change-point problems. Ph.D. thesis, Department of Statistics. Stanford University.
- Vostrikova, L. (1981). Detecting disorder in multidimensional random processes. *Soviet Math. Dokl.* **24**, 55-59.
- Yao, Y. C. (1988). Estimating the number of change-points via Schwarz'criterion. *Statist. Probab. Lett.* **6**, 181-189.
- Yao, Y. C. and Au, S. T. (1989). Least-squares estimation of a step function. *Sankhyā Ser. A* **51**, 370-381.

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

E-mail: k.k.korkas@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 July 2015; accepted December 2015)

