

Robust Narrowest Significance Pursuit: Inference for Multiple Change-Points in the Median

Piotr Fryzlewicz

To cite this article: Piotr Fryzlewicz (15 Mar 2024): Robust Narrowest Significance Pursuit: Inference for Multiple Change-Points in the Median, Journal of Business & Economic Statistics, DOI: [10.1080/07350015.2024.2316103](https://doi.org/10.1080/07350015.2024.2316103)

To link to this article: <https://doi.org/10.1080/07350015.2024.2316103>



© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 15 Mar 2024.



Submit your article to this journal [↗](#)



Article views: 298



View related articles [↗](#)



View Crossmark data [↗](#)

Robust Narrowest Significance Pursuit: Inference for Multiple Change-Points in the Median

Piotr Fryzlewicz

Department of Statistics, London School of Economics, London, UK

ABSTRACT

We propose Robust Narrowest Significance Pursuit (RNSP), a methodology for detecting localized regions in data sequences which each must contain a change-point in the median, at a prescribed global significance level. RNSP works by fitting the postulated constant model over many regions of the data using a new sign-multiresolution sup-norm-type loss, and greedily identifying the shortest intervals on which the constancy is significantly violated. By working with the signs of the data around fitted model candidates, RNSP fulfils its coverage promises under minimal assumptions, requiring only sign-symmetry and serial independence of the signs of the true residuals. In particular, it permits their heterogeneity and arbitrarily heavy tails. The intervals of significance returned by RNSP have a finite-sample character, are unconditional in nature and do not rely on any assumptions on the true signal. Code implementing RNSP is available at <https://github.com/pfryz/nsp>.

ARTICLE HISTORY

Received October 2022
Accepted February 2024

KEYWORDS

Confidence intervals;
Narrowest-over-threshold;
Post-inference selection;
Post-selection inference;
Structural breaks

1. Introduction

The problem of uncertainty quantification for possibly multiple parameter changes in time- or space-ordered data is motivated by the practical question of whether any suspected changes reflect real structural changes in the underlying stochastic model, or are due to random fluctuations. Approaches to this problem include confidence sets associated with simultaneous multiscale change-point estimation (Frick, Munk, and Sieling 2014; Pein, Sieling, and Munk 2017; Vanegas, Behr, and Munk 2022), post-selection inference (Hyun, G'Sell, and Tibshirani 2018; Duy et al. 2020; Hyun et al. 2021; Jewell, Fearnhead, and Witten 2022), inference without selection and post-inference selection via Narrowest Significance Pursuit (Fryzlewicz 2024), asymptotic confidence intervals conditional on the estimated change-point locations (Bai and Perron 1998; Eichinger and Kirch 2018; Cho and Kirch 2022), False Discovery Rate (Hao, Niu, and Zhang 2013; Li and Munk 2016; Cheng, He, and Schwartzman 2020), and Bayesian inference (Lavielle and Lebarbier 2001; Fearnhead 2006). These approaches go beyond mere change-point detection and offer statistical significance statements regarding the existence and locations of change-points in the statistical model underlying the data.

In this article, we are concerned with the following problem: given a sequence of noisy observations, automatically determine localized regions of the data which each must contain a change-point in the median, at a prescribed global significance level α . The methodology we introduce, referred to as Robust Narrowest Significance Pursuit (RNSP), achieves this for the piecewise-constant median model, capturing changes in the level of the median. By its algorithmic construction, RNSP offers exact

finite-sample coverage guarantees with practically no distributional assumptions on the data, other than serial independence of the signs of the true residuals and their sign-symmetry, a weak requirement which is immaterial for continuous distributions (as all, even non-symmetric, continuous distributions are sign-symmetric). In contrast to the existing literature, RNSP requires no knowledge of the distribution on the part of the user, and permits arbitrarily heavy tails, heterogeneity over time and/or lack of symmetry. In addition, RNSP is able to handle distributions that are continuous, continuous with mass points, or discrete, where these properties may also vary over the signal. The execution of RNSP does not rely on having an estimate of the number or locations of change-points. Critical values needed by RNSP do not depend on the noise distribution, and can be accurately approximated analytically. RNSP explicitly targets the shortest possible intervals of significance. It is worth noting, however, that our large-sample consistency result for RNSP, shown in Section 4, relies on stronger assumptions than its finite-sample coverage properties, discussed in Sections 2 and 3. We now situate RNSP in the context of the existing literature.

Heterogeneous Simultaneous Multiscale Change Point Estimator, abbreviated as H-SMUCE (Pein, Sieling, and Munk 2017), an extension of SMUCE (Frick, Munk, and Sieling 2014), is a change-point detector in the heterogeneous Gaussian piecewise-constant model $Y_i = f_i + \sigma_i \varepsilon_i$, where f_i is a piecewise-constant signal, ε_i are iid $N(0, 1)$ variables, and σ_i can only change when f_i does. In Section A.1 of their work, the authors provide an algorithmic construction of confidence intervals for the locations of the change-points in f_i , which involves screening the data for short intervals over which a constant signal fit is

unsuitable and they must therefore contain change-points. Crucially, this algorithmic construction relies on the knowledge of scale-dependent critical values (for measuring the unsuitability of a locally constant fit), which are not available analytically but only by simulation, and therefore the method does not extend automatically to unknown noise distributions (as the analyst needs to know what distribution to sample from). In [Section 5](#), we show that H-SMUCE suffers from inflated Type I error rates in the sense that the thus-constructed confidence intervals, in the examples of Gaussian models shown, do not all contain at least one true change-point each in more than $100(1 - \alpha)\%$ of the cases, contrary to what this algorithmic construction sets out to do. H-SMUCE is also prone to failing if the model is mis-specified, for example if the distribution of the data has a mass point (which is unsurprising in view of its assumption of Gaussianity).

Multiscale Quantile Segmentation (Vanegas, Behr, and Munk 2022, MQS), another extension of SMUCE, is a procedure for detecting possibly multiple changes in a given piecewise-constant quantile of the input data sequence, which includes the median as a special case. MQS estimates the number of change-points in the quantile function as the minimum among those candidate fits for which the empirical residuals pass a certain multiscale test at level α , where the empirical residuals are defined as binary exceedance sequences of the data over the level defined by each candidate fit. Working with such binary exceedance sequences means that MQS makes no distributional assumptions on the data other than their serial independence. Like SMUCE and H-SMUCE, MQS then defines a confidence set around the estimated signals as a set of all feasible (in the same sense) signal fits at level α . This then enables the conceptual and algorithmic construction of asymptotic simultaneous confidence intervals for the change-point locations, which are guaranteed to (each) contain a change-point with probability $1 - \alpha + o(1)$. Chen, Shah, and Samworth (2014) provide a critique of SMUCE from the inferential point of view, which also applies to H-SMUCE and MQS. In [Section 5](#), we illustrate the advantages and disadvantages of MQS, as implemented in the R package `mqs`. In contrast to MQS, the coverage guarantees offered by RNSP are wholly based on exact inequalities and therefore hold for any sample size, which is reflected in its performance shown in [Section 5](#).

The (non-robust) Narrowest Significance Pursuit (NSP; Fryzlewicz, 2024) is able to handle heterogeneous data in a variety of models, including the piecewise-constant signal plus independent noise model. However, NSP requires that the noise, if heterogeneous, is within the domain of attraction of the normal distribution and is symmetric, neither of which is assumed in RNSP. The self-normalized statistic used in NSP includes a term resembling an estimate of the local variance of the noise which, however, is only unbiased under the null hypothesis of no change-point being present locally. The fact that the same term over-estimates the variance under the alternative hypothesis, reduces the power of the detection statistic, which leads to typically long intervals of significance. We illustrate this issue in [Section 5](#) and show that RNSP offers a significant improvement.

Bai and Perron (1998, 2003), working with least-squares estimation of multiple change-points in regression models under possible heterogeneity of the errors, describe a procedure for

computing confidence intervals conditional on detection, with asymptotic validity guarantees. Crucially, it does not take into account the uncertainty associated with detection, which can be considerable especially for the more difficult problems (e.g., see the ‘‘US ex-post real interest rate’’ case study in Bai and Perron (2003), where there is genuine uncertainty between models with 2 and 3 change-points; we revisit this example in [Section 6.1](#)). By contrast, RNSP produces intervals of significant change in the median that are not conditional on detection and have a finite-sample nature.

The article is organized as follows. [Section 2](#) motivates RNSP and sets out its general algorithmic framework. [Section 3](#) describes how RNSP measures the local deviation from model constancy and gives finite-sample theoretical performance guarantees for RNSP. [Section 4](#) quantifies the large-sample behavior of RNSP. [Section 5](#) contains numerical examples and comparisons. [Section 6](#) includes examples showing the practical usefulness of RNSP. [Section 7](#) concludes with a brief discussion. Software implementing RNSP is available at <https://github.com/pfryz/nsp>.

2. Motivation and Review of RNSP Algorithmic Setting

2.1. RNSP: Context and Modus Operandi

RNSP discovers regions in the data in which the *median* departs from constancy, at a certain global significance level. This is in contrast to NSP (Fryzlewicz 2024), which targets the *mean*. RNSP does not make moment assumptions about the data, and therefore the median is a natural measure of data centrality. We now review the components of the algorithmic framework that are shared between NSP and RNSP, with the generic measurement of local deviation from the constant model as one of its building blocks. In [Section 3](#), we introduce the particular way in which local deviation from the constant model is measured in RNSP, which is appropriate for the median and hence fundamentally different from NSP.

RNSP operates in the signal plus noise model

$$Y_t = f_t + Z_t, \quad t = 1, \dots, T, \quad (1)$$

in which the signal $\{f_t\}_{t=1}^T$ and the variables $\{Z_t\}_{t=1}^T$ satisfy the assumptions below. Define $\text{sign}(x) = \mathbb{I}(x > 0) - \mathbb{I}(x < 0)$, where $\mathbb{I}(\cdot)$ is the indicator function.

Assumption 2.1. In (1), f_t is a piecewise-constant vector with an unknown number N and locations $0 = \eta_0 < \eta_1 < \dots < \eta_N < \eta_{N+1} = T$ of change-points. (The location η_j is a change-point if $f_{\eta_j} \neq f_{\eta_j+1}$.)

Assumption 2.2. (a) $\forall t, M(Z_t) = 0$, where M is the median operator. (If the median is nonunique, we require $0 \in M(Z_t)$.)

(b) The variables $\{Z_t\}_{t=1}^T$ are sign-symmetric, that is $P(Z_t > 0) = P(Z_t < 0)$, $\forall t$.

(c) The variables $\{\text{sign}(Z_t)\}_{t=1}^T$ are mutually independent.

While [Assumption 2.1](#) holds throughout the article, [Assumption 2.2](#) is not formally needed until [Section 3.2](#). [Section 4](#) provides additional results under extra assumptions.

$\{Z_t\}_{t=1}^T$ do not have to be identically distributed, and can have arbitrary mass atoms, or none, as long as their distributions satisfy [Assumption 2.2](#). The distribution(s) of $\{Z_t\}_{t=1}^T$ can be unknown to the analyst, and we do not impose moment assumptions. Any zero-median continuous distribution (even one with an asymmetric density function) is sign-symmetric. The requirement of the independence of $\{\text{sign}(Z_t)\}_{t=1}^T$ is weaker than that of the independence of $\{Z_t\}_{t=1}^T$ itself: for example if Z_t is a (G)ARCH process driven by symmetric, independent innovations, then $\text{sign}(Z_t)$ is serially independent, while Z_t is not. While the results of [Section 3](#) require the independence of $\text{sign}(Z_t)$, those in [Section 4](#) require the independence of Z_t ; see [Section 4](#) for details.

RNSP achieves the high level of generality in terms of the permitted distributions of Z_t thanks to its use of the sign transformation. The use of 0 in $\text{sign}(x)$ is critical for RNSP's objective to provide exact finite-sample coverage guarantees also for discrete distributions and continuous distributions with mass points, an aspect we discuss in [Section 3.1.1](#). The sign function is a critical building block of various procedures for nonparametric change-point testing and estimation; key references include [Bhattacharyya and Johnson \(1968\)](#), [Carlstein \(1988\)](#), and [Dümbgen \(1991\)](#).

2.2. (R)NSP: Generic Algorithm

The generic algorithmic framework underlying both RNSP and the non-robust NSP in [Fryzlewicz \(2024\)](#) is the same and is based on recursively searching for the shortest sub-samples in the data with globally significant deviations from the baseline model. In this section, we introduce this shared generic framework. In the following sections, we show how RNSP diverges from NSP through its use of a robust measure of deviation from the baseline model, suitable for the broad class of distributions specified in [Assumption 2.2](#).

We start with a pseudocode definition of the RNSP algorithm, in the form of a recursively defined function RNSP. In its arguments, $[s, e]$ is the current interval under consideration and at the start of the procedure, we have $[s, e] = [1, T]$; Y (of length T) is as in the model formula (1); M is the minimum guaranteed number of sub-intervals of $[s, e]$ drawn (unless the number of all sub-intervals of $[s, e]$ is less than M , in which case drawing M sub-intervals would mean repetition); λ_α is the threshold corresponding to the global significance level α (typical values for α would be 0.05 or 0.1) and τ_L (respectively τ_R) is a functional parameter used to specify the maximum extent of overlap of the left (respectively right) sub-interval of $[s, e]$ searched next after having identified a region of significance within $[s, e]$, if any. The no-overlap case would correspond to $\tau_L = \tau_R \equiv 0$. In each recursive call on a generic interval $[s, e]$, RNSP adds to the set \mathcal{S} any globally significant local regions (intervals) of the data identified within $[s, e]$ on which Y is deemed to depart significantly (at global level α) from the baseline constant model. We provide more details underneath the pseudocode below. In the remainder of the article, the subscript $_{[s,e]}$ relates to a constant indexed by the interval $[s, e]$ whose value will be clear from the context.

```

1: function RNSP( $s, e, Y, M, \lambda_\alpha, \tau_L, \tau_R$ )
2:   if  $e - s < 1$  then
3:     STOP
4:   end if
5:   if  $M \geq \frac{1}{2}(e - s + 1)(e - s)$  then
6:      $M := \frac{1}{2}(e - s + 1)(e - s)$ 
7:     draw all intervals  $[s_m, e_m] \subseteq [s, s + 1, \dots, e]$ ,  $m = 1, \dots, M$ , s.t.  $e_m - s_m \geq 1$ 
8:   else
9:     draw a representative (see description below) sample of intervals  $[s_m, e_m] \subseteq [s, s + 1, \dots, e]$ ,  $m = 1, \dots, M$ , s.t.  $e_m - s_m \geq 1$ 
10:  end if
11:  for  $m \leftarrow 1, \dots, M$  do
12:     $D_{[s_m, e_m]} := \text{DEVIATIONFROMCONSTANTMODEL}(s_m, e_m, Y)$ 
13:  end for
14:   $\mathcal{M}_0 := \arg \min_m \{e_m - s_m : m = 1, \dots, M; D_{[s_m, e_m]} > \lambda_\alpha\}$ 
15:  if  $|\mathcal{M}_0| = 0$  then
16:    STOP
17:  end if
18:   $m_0 := \text{ANYOF}(\arg \max_m \{D_{[s_m, e_m]} : m \in \mathcal{M}_0\})$ 
19:   $[\tilde{s}, \tilde{e}] := \text{SHORTESTSIGNIFICANTSUBINTERVAL}(s_{m_0}, e_{m_0}, Y, M, \lambda_\alpha)$ 
20:  add  $[\tilde{s}, \tilde{e}]$  to the set  $\mathcal{S}$  of significant intervals
21:  RNSP( $s, \tilde{s} + \tau_L(\tilde{s}, \tilde{e}, Y), Y, M, \lambda_\alpha, \tau_L, \tau_R$ )
22:  RNSP( $\tilde{e} - \tau_R(\tilde{s}, \tilde{e}, Y), e, Y, M, \lambda_\alpha, \tau_L, \tau_R$ )
23: end function

```

The RNSP algorithm is launched by the pair of calls: $\mathcal{S} := \emptyset$; $\text{RNSP}(1, T, Y, M, \lambda_\alpha, \tau_L, \tau_R)$. On completion, the output of RNSP is in the set \mathcal{S} ; when the context requires it, we write the output as $\mathcal{S}\{\text{RNSP}(1, T, Y, M, \lambda_\alpha, \tau_L, \tau_R)\}$. We now comment on the RNSP function line by line. In lines 2–4, execution is terminated for intervals that are too short. In lines 5–10, a check is performed to see if M is at least as large as the number of all sub-intervals of $[s, e]$. If so, then M is adjusted accordingly, and all sub-intervals are stored in $\{[s_m, e_m]\}_{m=1}^M$. Otherwise, a sample of M sub-intervals $[s_m, e_m] \subseteq [s, e]$ is drawn in which s_m and e_m are all possible pairs from an (approximately) equispaced grid on $[s, e]$ which permits at least M such sub-intervals. The ability to adjust the M parameter offers the users a choice between a faster but less thorough procedure (for lower values of M) and a slower but more accurate one (for higher values of M). The reason for not necessarily using the maximum possible number of intervals is that this may increase the computation time beyond what is acceptable to the user. All examples in the article use $M = 1000$.

In lines 11–13, each sub-interval $[s_m, e_m]$ is checked to see to what extent the response on this sub-interval (denoted by $Y_{s_m:e_m}$) deviates from the baseline constant model. This core step of the RNSP algorithm will be described in more detail in [Section 3](#).

In line 14, the measures of deviation obtained in line 12 are tested against threshold λ_α , chosen to guarantee the global significance level α . How to choose λ_α is independent of the distribution of Z_t if it is continuous, and there is also a simple distribution-independent choice of λ_α for discrete distributions and continuous distributions with probability masses; see [Section 3.2](#). The shortest sub-interval(s) $[s_m, e_m]$ for which the test

rejects the baseline model at global level α are collected in set \mathcal{M}_0 . In lines 15–17, if \mathcal{M}_0 is empty, then the procedure decides that it has not found regions of significant deviations from the constant model on $[s, e]$, and stops on this interval as a consequence. Otherwise, in line 18, the procedure continues by choosing the sub-interval, from among the shortest significant ones, on which the deviation from the baseline constant model has been the largest. The chosen interval is denoted by $[s_{m_0}, e_{m_0}]$.

In line 19, $[s_{m_0}, e_{m_0}]$ is searched for the shortest sub-interval on which the hypothesis of the baseline model is rejected locally at a global level α . Such a sub-interval certainly exists, as $[s_{m_0}, e_{m_0}]$ itself has this property. The structure of this search again follows the workflow of the RNSP procedure; it proceeds by executing lines 2–18 of RNSP, but with s_{m_0}, e_{m_0} in place of s, e . The chosen interval is denoted by $[\tilde{s}, \tilde{e}]$. This second-stage search is important to RNSP’s pursuit to produce short intervals: indeed, if the sample of intervals $[s_m, e_m]$ contained insufficiently short intervals (perhaps because an insufficiently large M was chosen), then, without the second-stage search in line 19, the intervals of significance returned by RNSP might be overly long. The second-stage search in line 19 can be seen as a guard against a small M , or in other words against an insufficiently fine original grid of interval endpoints. In line 20, the selected interval $[\tilde{s}, \tilde{e}]$ is added to the output set \mathcal{S} .

Because of the second-stage search in line 19, pre-drawing the intervals $[s_m, e_m]$ prior to launching the RNSP procedure (rather than drawing them recursively on each current interval as it done in the algorithm) is not an option for RNSP. Indeed, in the presence of the second-stage search, the selected interval of significance $[\tilde{s}, \tilde{e}]$ may be misaligned with the initial grid of intervals drawn, in which case the grid of intervals to the left and to the right of $[\tilde{s}, \tilde{e}]$ must be redrawn to avoid leaving unexamined gaps in the data.

In lines 21–22, RNSP is executed recursively to the left and to the right of the detected interval $[\tilde{s}, \tilde{e}]$. However, we optionally allow for some overlap with $[\tilde{s}, \tilde{e}]$. The overlap, if present, is a function of $[\tilde{s}, \tilde{e}]$ and, if it involves detection of the location of a change-point within $[\tilde{s}, \tilde{e}]$, then it is also a function of Y . An example of the relevance of this is given in Section 6.1.

3. Robust NSP: Measuring Deviation from the Constant Model

3.1. Deviation Measure: Motivation, Definition, and Properties

The main structure of the `DEVIATIONFROMCONSTANTMODEL`(s_m, e_m, Y) operation is as follows: (a) Fit the best, in the sense described precisely later, constant model to $Y_{s_m:e_m}$. (b) Examine the signs of the empirical residuals from this fit. If their distribution is deemed to contain a change-point (which indicates that the constant model fit is unsatisfactory on $[s_m, e_m]$ and therefore the model contains a change-point on that interval), the value returned by `DEVIATIONFROMCONSTANTMODEL`(s_m, e_m, Y) should be large; otherwise small.

A key ingredient of our measure of deviation is a multiresolution sup-norm introduced below, used on the signs of the input rather than in the original data domain. Its basic building block

is a scaled partial sum statistic, defined for an arbitrary input sequence $\{x_t\}_{t=1}^T$ by $U_{s,e}(x) = (e-s+1)^{-1/2} \sum_{t=s}^e x_t$. We define the multiresolution sup-norm (Nemirovski 1986; Li 2016) of an input vector x (of length T) with respect to the interval set \mathcal{I} as $\|x\|_{\mathcal{I}} = \max_{[s,e] \in \mathcal{I}} |U_{s,e}(x)|$. The set \mathcal{I} used in RNSP contains intervals at a range of scales and locations. A canonical example of a suitable interval set \mathcal{I} is the set \mathcal{I}^a of all subintervals of $[1, T]$. We will use \mathcal{I}^a in defining the largest acceptable global probability of spurious detection. However, for computational reasons, `DEVIATIONFROMCONSTANTMODEL` will use a smaller interval set (we give the details later). This will not affect the exactness of our coverage guarantees, because, naturally, if $\mathcal{J} \subseteq \mathcal{I}$, then $\|x\|_{\mathcal{J}} \leq \|x\|_{\mathcal{I}}$. We also define the restriction of \mathcal{I} to an arbitrary interval $[s, e]$ as $\mathcal{I}_{[s,e]} = \{[u, v] \subseteq [s, e] : [u, v] \in \mathcal{I}\}$. Note the trivial inequality

$$\|x_{s:e}\|_{\mathcal{I}_{[s,e]}^a} \leq \|x\|_{\mathcal{I}^a} \quad (2)$$

for any $[s, e] \subseteq [1, T]$. When the above multiresolution sup-norm is applied to the signs of the input, as is done in this work, rather than the original input, we refer to it as the sign-multiresolution sup-norm. When applied to the empirical residuals from a candidate constant fit on an interval, it can be viewed as a simple robust multiscale measure of data fidelity of the given candidate fit on all time scales up to the length of the interval.

We now define the deviation measure $D_{[s_m, e_m]} := \text{DEVIATIONFROMCONSTANTMODEL}(s_m, e_m, Y)$, which satisfies the property that if there is no change-point on the interval $[s_m, e_m]$, then it is guaranteed that

$$D_{[s_m, e_m]} \leq \|\text{sign}(Z_{s_m:e_m})\|_{\mathcal{I}_{[s_m, e_m]}^a}. \quad (3)$$

The discussion below assumes that there is no change-point in $[s_m, e_m]$. For the true constant signal $f_{s_m:e_m}$, denote $f_{[s_m, e_m]} := f_{s_m} = \dots = f_{e_m}$. There are only at most $2(e_m - s_m) + 3$ different possible constants $\tilde{f}_{[s_m, e_m]}$ leading to different sequences $\{\text{sign}(Y_t - \tilde{f}_{[s_m, e_m]})\}_{t=s_m}^{e_m}$. To see this, sort the values of $Y_{s_m:e_m}$ in nondecreasing order to create $Y_{(1)}, Y_{(2)}, \dots, Y_{(e_m-s_m+1)}$. Take candidate constants $\tilde{f}_{[s_m, e_m]}^{(j)}$, $j = 1, \dots, 2(e_m - s_m) + 3$, defined as follows.

$$\begin{aligned} \tilde{f}_{[s_m, e_m]}^{(1)} &< Y_{(1)} \quad (\text{but otherwise arbitrary}) \\ \tilde{f}_{[s_m, e_m]}^{(2)} &= Y_{(1)} \\ \tilde{f}_{[s_m, e_m]}^{(3)} &= \frac{1}{2}(Y_{(1)} + Y_{(2)}) \\ \tilde{f}_{[s_m, e_m]}^{(4)} &= Y_{(2)} \\ \tilde{f}_{[s_m, e_m]}^{(5)} &= \frac{1}{2}(Y_{(2)} + Y_{(3)}) \\ &\vdots \\ \tilde{f}_{[s_m, e_m]}^{(2(e_m-s_m)+2)} &= Y_{(e_m-s_m+1)} \\ \tilde{f}_{[s_m, e_m]}^{(2(e_m-s_m)+3)} &> Y_{(e_m-s_m+1)} \quad (\text{but otherwise arbitrary}). \end{aligned} \quad (4)$$

We have the following simple result; the proof is trivial and we omit it.

Proposition 3.1. Under Assumption 2.1, assume no change-point in $[s_m, e_m]$ and denote $f_{[s_m, e_m]} := f_{s_m} = \dots = f_{e_m}$. Let

the constants $\tilde{f}_{[s_m, e_m]}^{(j)}$, $j = 1, \dots, 2(e_m - s_m) + 3$ be defined as in (4). There exists a $j_0 \in \{1, \dots, 2(e_m - s_m) + 3\}$ such that

$$\{\text{sign}(Y_t - f_{[s_m, e_m]})\}_{t=s_m}^{e_m} = \{\text{sign}(Y_t - \tilde{f}_{[s_m, e_m]}^{(j_0)})\}_{t=s_m}^{e_m}. \quad (5)$$

We now define our measure of deviation $D_{[s_m, e_m]}$, and prove its key property as a corollary to Proposition 3.1.

Definition 3.1. Let the constants $\tilde{f}_{[s_m, e_m]}^{(j)}$, $j = 1, \dots, 2(e_m - s_m) + 3$ be defined as in (4). We define

$$D_{[s_m, e_m]} := \min_{j \in \{1, \dots, 2(e_m - s_m) + 3\}} \|\text{sign}(Y_{s_m:e_m} - \tilde{f}_{[s_m, e_m]}^{(j)})\|_{\mathcal{I}_{[s_m, e_m]}^a}. \quad (6)$$

$D_{[s_m, e_m]}$ tries all possible baseline constant model fits on $[s_m, e_m]$ and chooses the one that minimizes the sign-multiresolution norm of the residuals, to ensure that the finite-sample coverage guarantees hold, as we will see below.

Corollary 3.1. Under Assumption 2.1, for any interval $[s_m, e_m]$ on which there is no change-point, we have

$$D_{[s_m, e_m]} \leq \|\text{sign}(Z_{s_m:e_m})\|_{\mathcal{I}_{[s_m, e_m]}^a}. \quad (7)$$

In other words, the deviation measure defined in (6) satisfies the desired property (3).

Proof. Let the index j_0 be as in the statement of Proposition 3.1. We have

$$\begin{aligned} D_{[s_m, e_m]} &= \min_{j \in \{1, \dots, 2(e_m - s_m) + 3\}} \|\text{sign}(Y_{s_m:e_m} - \tilde{f}_{[s_m, e_m]}^{(j)})\|_{\mathcal{I}_{[s_m, e_m]}^a} \\ &\leq \|\text{sign}(Y_{s_m:e_m} - \tilde{f}_{[s_m, e_m]}^{(j_0)})\|_{\mathcal{I}_{[s_m, e_m]}^a} \\ &= \|\text{sign}(Y_{s_m:e_m} - f_{[s_m, e_m]})\|_{\mathcal{I}_{[s_m, e_m]}^a} \\ &= \|\text{sign}(Z_{s_m:e_m})\|_{\mathcal{I}_{[s_m, e_m]}^a}. \end{aligned} \quad \square$$

This leads to the following guarantee for the RNSP algorithm.

Theorem 3.1. Let Assumption 2.1 hold, and let $\mathcal{S} = \{S_1, \dots, S_R\}$ be the set of intervals returned by the RNSP algorithm. We have $P(\exists i = 1, \dots, R \forall j = 1, \dots, N [\eta_j, \eta_j + 1] \not\subseteq S_i) \leq P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > \lambda_\alpha)$.

Proof. On the set $\|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \lambda_\alpha$, each interval S_i must contain a change-point as if it did not, then by Corollary 3.1 and inequality (2), we would have to have

$$D_{S_i} \leq \|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \lambda_\alpha. \quad (8)$$

However, the fact that S_i was returned by RNSP means, by line 14 of the RNSP algorithm, that $D_{S_i} > \lambda_\alpha$, which contradicts (8). This completes the proof. \square

Theorem 3.1 should be read as follows. Let $\alpha = P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > \lambda_\alpha)$. For a set of intervals returned by RNSP, we are guaranteed, with probability of at least $1 - \alpha$, that there is at least one change-point in each of these intervals. Therefore, $\mathcal{S} = \{S_1, \dots, S_R\}$ can be interpreted as an automatically chosen set of *regions (intervals) of significance* in the data. In the no-change-point case ($N = 0$), the probability of obtaining one of

more intervals of significance ($R \geq 1$) is bounded from above by $P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > \lambda_\alpha)$. **Theorem 3.1** is of a finite-sample character and holds exactly and for any sample size. Moreover, it is independent of the form of the innovations Z . Assumptions on Z will only be needed in controlling the term $P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > \lambda_\alpha)$; we defer this to Section 3.2.

We emphasize that **Theorem 3.1** does not promise to detect all the change-points, or to do so asymptotically as the sample size gets larger: this would be impossible without assumptions on the strength of the change-points (involving spacing between neighboring change-points and the sizes of the jumps). This aspect of RNSP is investigated in Section 4. Instead, **Theorem 3.1** promises that every interval of significance returned by RNSP must contain at least one change-point each, with a certain global probability adjustable by the user. Therefore, one particular implication of **Theorem 3.1** is that we must have

$$P(N \geq R) \geq 1 - P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > \lambda_\alpha). \quad (9)$$

The intervals of significance returned by RNSP have an “unconditional confidence interval” interpretation: they are not conditional on any prior detection event, but indicate regions in the data each of which must unconditionally contain at least one change in the underlying signal f_t , with a global probability of at least $1 - \alpha$. Therefore, as in NSP (Fryzlewicz 2024), RNSP can be viewed as performing “inference without selection” (where “inference” refers to producing the RNSP intervals of significance and “selection” to the estimation of change-point locations, absent from RNSP). This viewpoint also enables “post-inference selection” or “in-inference selection” if the exact change-point locations (if any) are to be estimated within the RNSP intervals of significance after or during the execution of RNSP.

3.1.1. Deviation Measure: Discussion

Achieving computational savings without affecting coverage guarantees. The operation of trying each constant $\tilde{f}_{[s_m, e_m]}^{(j)}$ in (6) is fast, but in order to accelerate it further, we introduce the two computational savings below, which do not increase $D_{[s_m, e_m]}$ and therefore respect the inequality (7) and hence also our coverage guarantees in Theorem 3.1.

Reducing the set $\mathcal{I}_{[s_m, e_m]}^a$. To accelerate the computation of (6), we replace the set $\mathcal{I}_{[s_m, e_m]}^a$ in $D_{[s_m, e_m]}$ with the set $\mathcal{I}_{[s_m, e_m]}^{lr} := \mathcal{I}_{[s_m, e_m]}^l \cup \mathcal{I}_{[s_m, e_m]}^r$ (with l and r standing for left and right, respectively), where $\mathcal{I}_{[s_m, e_m]}^l = \{[s_m, s_m + 1], [s_m, s_m + 2], \dots, [s_m, e_m]\}$ and $\mathcal{I}_{[s_m, e_m]}^r = \{[s_m, e_m], [s_m + 1, e_m], \dots, [e_m - 1, e_m]\}$. This reduces the cardinality of the set of intervals included in $D_{[s_m, e_m]}$ from $O((e_m - s_m)^2)$ to $O(e_m - s_m)$. As $\mathcal{I}_{[s_m, e_m]}^{lr} \subseteq \mathcal{I}_{[s_m, e_m]}^a$ and hence $\|\cdot\|_{\mathcal{I}_{[s_m, e_m]}^{lr}} \leq \|\cdot\|_{\mathcal{I}_{[s_m, e_m]}^a}$, the results of Corollary 3.1 and

Theorem 3.1 remain unchanged for the thus-reduced $D_{[s_m, e_m]}$. On the other hand, $\mathcal{I}_{[s_m, e_m]}^{lr}$ has been defined in this particular way so as not to compromise the detection power in the piecewise-constant signal model. To see this, consider the following illustrative example. Suppose $Y_t = f_t$ (noiseless case) and $f_t = 0$ for $t = 1, \dots, 50$ and $f_t = 1$ for $t = 51, \dots, 100$. On $[s_m, e_m] = [1, 100]$, the baseline constant signal level fitted is $\tilde{f}_{[1, 100]} = 1/2$ and we have $\text{sign}(Y_t - \tilde{f}_{[1, 100]}) = -1$ for $t = 1, \dots, 50$; $\text{sign}(Y_t - \tilde{f}_{[1, 100]}) = 1$ for $t = 51, \dots, 100$.

In this setting, the two multiresolution sup-norms: $\|\text{sign}(Y_t - \tilde{f}_{[1,100]})\|_{\mathcal{I}_{[1,100]}^r}$ and $\|\text{sign}(Y_t - \tilde{f}_{[1,100]})\|_{\mathcal{I}_{[1,100]}^a}$ are identical, equal to $\sqrt{50}$ and achieve this value for the intervals $[1, 50]$ and $[51, 100]$, members of both $\mathcal{I}_{[1,100]}^a$ and $\mathcal{I}_{[1,100]}^r$. This simple example illustrates the wider phenomenon that if there is a single change-point in f_t on a generic interval $[s_m, e_m]$ under consideration, then in the noiseless case the multiresolution norm over the set $\mathcal{I}_{[s_m, e_m]}^a$ is maximized at one of the “left” or “right” intervals in $\mathcal{I}_{[s_m, e_m]}^r$, and we are happy to sacrifice potential negligible differences in the noisy case in exchange for the substantial computational savings.

Limiting interval lengths. In practice, the analyst may not be interested in excessively long RNSP intervals of significance and may therefore wish to ignore intervals for which $e_m - s_m > L$ for a user-specified maximum length L .

Validity for non-continuously distributed data. The following three aspects of RNSP ensure the validity of its coverage guarantees in the presence of mass points in Z_t or if the distribution of Z_t is discrete: (a) the fact that the sign function defined in Section 2.1 returns zero if its argument is zero, (b) the fact that the test levels $\tilde{f}_{[s_m, e_m]}^{(l)}$ (definition (4)) are placed at data points, and (c) the fact that these test levels are placed in between the sorted data points. Indeed, in the absence of (a), (b), or (c), it is easy to construct simple discrete distributions of Z_t for which $D_{[s_m, e_m]}$ would be spuriously large in the absence of change-points on $[s_m, e_m]$.

3.2. Evaluation and Bounds for $\|\text{sign}(Z)\|_{\mathcal{I}^a}$

To make Theorem 3.1 operational, we need to obtain an understanding of the distribution of $\|\text{sign}(Z)\|_{\mathcal{I}^a}$ so we are able to choose λ_α such that $P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > \lambda_\alpha) = \alpha$ (or approximately so) for a desired global significance level α .

Initially we consider Z_t such that $P(\text{sign}(Z_t) = 1) = P(\text{sign}(Z_t) = -1) = 1/2$ (the general case $P(\text{sign}(Z_t) = 0) \geq 0$ is covered in the next paragraph). One simple way of determining the distribution of $\|\text{sign}(Z)\|_{\mathcal{I}^a}$ for any finite T is by simulation; this would only need to be done once for every T and the quantiles stored for fast access. Another approach is asymptotic and proceeds as follows. From Theorem 1.1 in Kabluchko and Wang (2014) (which applies to sequences of serially independent symmetric Bernoulli variables as explained in Section 1.5.1 of that work; our Assumption 2.2 means that result is applicable in our context), we have

$$\lim_{T \rightarrow \infty} P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > a_T + \tau/a_T) = 1 - \exp(-2\Lambda \exp(-\tau)), \quad (10)$$

where $a_T = \{2 \log(T \log^{-1/2} T)\}^{1/2}$ and Λ is a constant. As the theoretical calculation of Λ in Kabluchko and Wang (2014) contains an error, we use simulation over a range of values of T and τ to determine a suitable value of Λ as 0.274. The practical choice of the significance threshold λ_α then proceed as follows: (a) fix α to the desired level, for example 0.05 or 0.1; (b) obtain the value of τ by equating $1 - \exp(-2\Lambda \exp(-\tau)) = \alpha$; (c) obtain $\lambda_\alpha = a_T + \tau/a_T$. While this approach is asymptotic in nature (note the limit as $T \rightarrow \infty$ in (10)), we observe that the finite-sample agreement of $P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > a_T + \tau/a_T)$ with its limit in (10) is excellent even for small sample sizes. If the user

chooses to pursue this route of obtaining λ_α , this will be the only asymptotic component of RNSP.

Suppose now that $P(\text{sign}(Z_t) = 0) = \rho_t \geq 0$; note that the sign-symmetry Assumption 2.2(b) implies $P(\text{sign}(Z_t) = 1) = P(\text{sign}(Z_t) = -1) = (1 - \rho_t)/2$. Construct the variable $\tilde{Z}_t = Z_t \mid Z_t \neq 0$. As $P(\text{sign}(\tilde{Z}_t) = 1) = P(\text{sign}(\tilde{Z}_t) = -1) = 1/2$, the limiting statement (10) applies to $\text{sign}(\tilde{Z}_t)$. However, we have the double inequality

$$\|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \|\text{sign}(\tilde{Z})\|_{\mathcal{I}^a} \leq \|\text{sign}(\tilde{Z})\|_{\mathcal{I}^a}, \quad (11)$$

with $I = [1, 2, \dots, T_1]$, where $T_1 = |\{t \in [1, \dots, T] : Z_t \neq 0\}|$. The first inequality in (11) holds because every constituent partial sum of $\|\text{sign}(Z)\|_{\mathcal{I}^a}$ has a corresponding larger or equal in magnitude partial sum in $\|\text{sign}(\tilde{Z})\|_{\mathcal{I}^a}$ constructed by removing the zeros from its numerator and decreasing (or not increasing) its denominator as it contains fewer (or the same number of) terms. As an illustrative example, suppose the sequence of $\text{sign}(Z_t)$ starts $-1, 0, 1, 1$. The absolute partial sum $|-1 + 0 + 1 + 1|/\sqrt{4}$, a constituent of $\|\text{sign}(Z)\|_{\mathcal{I}^a}$, is majorized by the absolute partial sum $|-1 + 1 + 1|/\sqrt{3}$, a constituent of $\|\text{sign}(\tilde{Z})\|_{\mathcal{I}^a}$, where the latter sum has been constructed by removing the 0 from $-1, 0, 1, 1$ and adjusting for the number of terms (now 3 instead of 4). The second inequality in (11) holds simply because $T_1 \leq T$. The implication of (11) is that $\|\text{sign}(Z)\|_{\mathcal{I}^a}$ for $\rho_t \geq 0$ is majorized by $\|\text{sign}(Z)\|_{\mathcal{I}^a}$ for $\rho_t = 0$, the case handled by (10). Therefore, the threshold λ_α obtained as a consequence of (10) can also meaningfully be applied in the general case $\rho_t \geq 0$.

4. Detection Consistency and Lengths of RNSP Intervals

This section shows the large-sample consistency of RNSP in detecting change-points, and the rates at which the lengths of the RNSP intervals contract (relative to T), as T increases. To simplify our technical arguments, we consider a version of the RNSP algorithm that considers all subintervals of $[1, T]$. Our focus on iid Z_t 's in this section is mainly due to our ability to rely on the Dvoretzky-Kiefer-Wolfowitz inequality in the proof of Corollary 4.1; however, note that the finite-sample result of Theorem 4.1 holds regardless of the dependence structure of Z_t . We focus on continuously-distributed Z_t 's as this results in notationally much less involved arguments regarding the minimum signal strength required. We first introduce some essential notation, and then state our assumption and the result. For each change-point η_j , define

$$\Delta_j = \min\{P\{Z_t \in (-|f_{\eta_j} - f_{\eta_{j+1}}|/2, 0)\}, P\{Z_t \in (0, |f_{\eta_j} - f_{\eta_{j+1}}|/2)\}, \quad (12)$$

$$\bar{d}_j = \bar{d}_j(\lambda, \lambda_\alpha) = \left\lceil \left(\frac{2\lambda + \lambda_\alpha}{2\Delta_j} \right)^2 + 1 \right\rceil. \quad (13)$$

In addition, for any process V_t , define $\epsilon_t^V(w) = \mathbb{I}(V_t - w > 0) - P(V_t - w > 0)$.

Assumption 4.1. (a) The variables $\{Z_t\}_{t=1}^T$ are mutually independent.
(b) The variables $\{Z_t\}_{t=1}^T$ are identically distributed.

- (c) The distribution of Z_1 is continuous.
 (d) With the notation $\eta_0 = 0$ and $\eta_{N+1} = T$, we have $\eta_{j+1} - \eta_j \geq 2\bar{d}_{j+1} + 2\bar{d}_j - 2$ for $j = 1, \dots, N-1$, and $\eta_1 - \eta_0 \geq 2\bar{d}_1 - 1$ as well as $\eta_{N+1} - \eta_N \geq 2\bar{d}_N - 1$.

Our first result below is of a finite-sample nature.

Theorem 4.1. Let Assumptions 2.1, 2.2(a) and 4.1(b),(c),(d) hold. On the set defined by the intersection of the events $\|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \lambda_\alpha$ and $\max_{s,e} \sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Z(w) \right| \leq \lambda$, a version of the RNSP algorithm that considers all intervals, executed with no overlaps and with threshold λ_α , returns exactly N intervals of significance $[s_1, e_1] < \dots < [s_N, e_N]$ such that $\eta_j \in [s_j, e_j - 1]$ and $e_j - s_j + 1 \leq 2\bar{d}_j$ for $j = 1, \dots, N$.

Theorem 4.1 leads to the following corollary giving a large-sample consistency result.

Corollary 4.1. Let the assumptions of Theorem 4.1 and Assumption 4.1(a) hold. Let $\lambda_\alpha = (1 + \delta)\{2 \log T\}^{1/2}$ and $\lambda = (1 + \delta) \log^{1/2} T$, for any $\delta > 0$. Let \mathcal{S} denote the set of intervals of significance $[s_1, e_1] < \dots < [s_R, e_R]$ returned by RNSP algorithm that considers all intervals, executed with no overlaps and with threshold λ_α . We have

$$P \left\{ R = N \quad \wedge \quad \forall j = 1, \dots, N \right. \\ \left. \eta_j \in [s_j, e_j - 1] \quad \wedge \quad e_j - s_j + 1 \leq 2\bar{d}_j \right\} \rightarrow 1$$

as $T \rightarrow \infty$.

This and Corollary 4.2 below are the only large-sample results of the article; the others are of a finite-sample character. Setting $\lambda_\alpha = (1 + \delta)\{2 \log T\}^{1/2}$ and $\lambda = (1 + \delta) \log^{1/2} T$, for any $\delta > 0$, causes the probabilities of the events $\|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \lambda_\alpha$ and $\max_{s,e} \sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Z(w) \right| \leq \lambda$ in Theorem 4.1 (respectively) to converge to one. Note that here, λ_α does not depend on α and is of a higher order of magnitude than the (α -dependent) λ_α of Section 3.2. Paraphrasing, this is to say that we need the global significance level α to tend to zero with T in order to obtain large-sample consistency; a result in line with analogous results in Pein, Sieling, and Munk (2017) and Vanegas, Behr, and Munk (2022).

We briefly comment on what the result of Corollary 4.1 means for the minimum signal strength Assumption 4.1(iii), and for the localization rates of the RNSP algorithm in detecting the change-points. If the distribution of Z_t does not vary with T (this section already assumes that it does not vary with t), and if the jump sizes $|f_{\eta_j} - f_{\eta_{j-1}}|$ are bounded from below by a positive constant independent of j and T , then Δ_j (formula (12)) is also bounded from below by a positive constant independent of j and T . By formula (13), the assumptions on λ, λ_α in Corollary 4.1 then imply $\bar{d}_j = \Theta(\log T)$ (where Θ should be read “of the exact order”). Assumption 4.1(iii) then requires that the spacings between the change-points be at least of order $\log T$. Corollary 4.1 states that the length of each RNSP interval of significance, $e_j - s_j + 1$, is, with global probability approaching 1, at most of order $\log T$. These minimum-spacing assumptions and the implied lengths of the localization intervals are near-optimal and the same as those in the non-robust literature, see

for example Theorem 1 in Baranowski, Chen, and Fryzlewicz (2019) and Corollary 4 in Fryzlewicz (2024), as well as the associated discussions. However, the results of this section also permit $\Delta_j \rightarrow 0$ with T , which will, naturally, affect the above minimum-spacing requirements and localization rates as stipulated by formulas (12) and (13) and Assumption 4.1(iii).

The consistency of RNSP in the sense of Corollary 4.1 implies the consistency of any pointwise estimators $\hat{\eta}_j$ contained within the RNSP intervals of significance $[s_j, e_j]$, with the localization rate of $\hat{\eta}_j$ bounded from above by the length of the interval $[s_j, e_j]$. In particular, the near-optimality of the lengths of the RNSP intervals $[s_j, e_j]$ (as discussed in the preceding paragraph) automatically implies the near-optimality of the localization rate of $\hat{\eta}_j$. Indeed, since by Corollary 4.1, $[s_j, e_j - 1]$ is guaranteed to contain η_j (on an event of high probability), for any estimator $\hat{\eta}_j \in [s_j, e_j - 1]$, we must have $|\hat{\eta}_j - \eta_j| \leq |e_j - s_j - 1|$ (on the same event). In other words, the rate with which $[s_j, e_j]$ contract (relative to T) is inherited by any estimator $\hat{\eta}_j \in [s_j, e_j - 1]$; this applies even to naive estimators constructed for example as the middle points of their corresponding RNSP intervals $[s_j, e_j]$, that is $\hat{\eta}_j = \lfloor (s_j + e_j)/2 \rfloor$. More refined estimators, for example one based on the CUSUM maximization of the signs of the data around their median (Sen and Srivastava 1975) within each RNSP interval, can also be used and will also automatically inherit the consistency and rate. Sections 5 and 6 illustrate both of these estimators. Trivially, in light of Corollary 4.1, the set of estimated change-points $\{\hat{\eta}_j\}_{j=1}^R$ is consistent in the Hausdorff measure for $\{\eta_j\}_{j=1}^N$, the set of true change-points in f_t .

Yet another implication of Theorem 4.1 appears below.

Corollary 4.2. Let the assumptions of Corollary 4.1 hold. Let λ_α be such that $P(\|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \lambda_\alpha) \geq 1 - \alpha$ and let $\lambda = (1 + \delta) \log^{1/2} T$, for any $\delta > 0$. Let \mathcal{S} denote the set of intervals of significance $[s_1, e_1] < \dots < [s_R, e_R]$ returned by RNSP algorithm that considers all intervals, executed with no overlaps and with threshold λ_α . We then have

$$\liminf_{T \rightarrow \infty} P \left\{ R = N \quad \wedge \quad \forall j = 1, \dots, N \right. \\ \left. \eta_j \in [s_j, e_j - 1] \quad \wedge \quad e_j - s_j + 1 \leq 2\bar{d}_j \right\} \geq 1 - \alpha.$$

The computation of the deviation measure $D_{[s,e]}$ for all subintervals $[s, e]$ of $[1, T]$, as the results of this section require, is a $O(T^3)$ operation. The cubic complexity should not surprise in the context of a robust method that considers all intervals, as there are $O(T^2)$ intervals to consider and $O(T)$ binary exceedance levels within each interval. In practice, we use three devices to reduce the computational complexity of RNSP: (a) using a fixed (i.e., unchanging with T) value of M , (b) restricting the length of intervals under consideration to $\leq L$, and (c) reducing the set $\mathcal{I}_{[s,e]}^a$. Items (b) and (c) are described in more detail in Section 3.1.1.

5. Numerical Illustrations

In this section, we demonstrate numerically that the guarantee offered by Theorem 3.1 holds for RNSP in practice over a variety of homogeneous and heterogeneous models for which the variables Z_t satisfy Assumption 2.2. We also investigate the circumstances under which similar guarantees are not

Table 1. Null models for the comparative simulation study in Section 5.

Model name	Sample path execution in R
Plain Gauss	<code>rnorm(100)</code>
Plain Gauss Long	<code>rnorm(1000)</code>
Plain Poisson	<code>as.numeric(rpois(200, 1))</code>
Heterogeneous Gauss	<code>c(rep(1, 100), rep(8, 50), rep(1, 100)) * rnorm(250)</code>
Symmetric Bernoulli	<code>as.numeric(rbinom(200, 1, .5))</code>
Plain Cauchy	<code>rcauchy(100, 0)</code>
Mix 1	<code>sample(3, size=300, replace=TRUE, prob=c(.35, .3, .35)) -> xx</code> <code>xx[xx != 2] <- rnorm(sum(xx != 2))</code>
Mix 2	<code>rpois(200, 5)+rnorm(200)/30</code>

Table 2. Percentage of times, out of 200 simulated sample paths of each null model, that the respective method indicates no intervals of significance at level $\alpha = 0.1$ (nominal coverage = 90%).

Model	RNSP	H-SMUCE	MQS-R	MQS-K	SN-NSP
Plain Gauss	100.0	99.0	99.5	99.5	100.0
Plain Gauss Long	100.0	96.5	98.5	98.5	100.0
Plain Poisson	98.5	0.0	99.5	99.5	1.0
Heterogeneous Gauss	99.5	98.0	99.5	99.5	91.0
Symmetric Bernoulli	95.5	0.0	99.5	99.5	35.0
Plain Cauchy	99.5	100.0	98.0	98.0	99.5
Mix 1	100.0	0.0	99.0	99.0	99.5
Mix 2	98.5	88.0	96.5	96.5	100.0

offered by H-SMUCE (Pein, Sieling, and Munk 2017), MQS (Vanegas, Behr, and Munk 2022) or the self-normalized version of NSP (SN-NSP), suitable for heterogeneous data (Fryzlewicz 2024). In this section, we use the acronyms RNSP and SN-NSP to denote the versions of these respective procedures with no interval overlaps, that is $\tau_L = \tau_R = 0$. Later in this section, we introduce notation for versions with overlaps. Both RNSP and SN-NSP use $M = 1000$ intervals, the default setting. For H-SMUCE, the function call we use is `stepR::stepFit(x, alpha=0.1, family="hsmuce", confband=TRUE)`. The `type` parameter in MQS specifies the loss function for their final estimate with multiscale constraints; we denote by MQS-R the result of `mqs::mqse(x, alpha=0.1, conf=TRUE, type="runs")` and by MQS-K the result of `mqs::mqse(x, alpha=0.1, conf=TRUE, type="koenker")`. We use the following package versions: `mqs v1.0`, `stepR v2.1-8`, `nsp v1.0.0`.

We begin with null models, by which we mean models (1) for which f_t is constant throughout, that is $N = 0$. For null models, Theorem 3.1 promises that RNSP at level α returns no intervals of significance with probability at least $1 - \alpha$. In this section, we use $\alpha = 0.1$. There are analogous parameters in H-SMUCE, MQS and SN-NSP, and they are also set to 0.1. However, while in both RNSP and SN-NSP, the parameter α is responsible for finite-sample coverage guarantees (for both null and non-null models), in MQS and H-SMUCE it is responsible for similar but only asymptotic coverage guarantees, as $T \rightarrow \infty$. For H-SMUCE, this latter point is clarified (jointly) in Theorem 5 of Pein, Sieling, and Munk (2017) and in Section A.1 of the online supplement to that work, and for MQS – in Theorem 2.3 of Vanegas, Behr, and Munk (2022) and in Section S.4 of the online supplement to that work.

The null models used are listed in Table 1, and Table 2 shows the associated results. RNSP, MQS-R and MQS-K keep the nominal size well across all the models considered, returning no

intervals of significance at least 95% of the time in all situations. H-SMUCE behaves correctly for the three Gaussian models, but fails for the discrete distributions and model Mix 1, which contains mass points. It is unexpectedly successful in the Plain Cauchy model, but this is perhaps because it has very limited detection power in the Cauchy model with change-points (more on this model below). It also underperforms slightly for model Mix 2, which is continuous (and within the domain of attraction of the Gaussian distribution) but multimodal. SN-NSP fails for the discrete distributions, which is a consequence of the (asymptotically guaranteed) closeness of the self-normalized deviation measure to the appropriate functional of the Wiener process not kicking in in these instances (due to the relatively small sample sizes).

We now discuss performance for signals with change-points ($N > 0$). Table 3 defines our models; the model labeled MQS.easy is taken from <https://github.com/ljvanegas/mqs/blob/master/mqs.ipynb> and MQS.hard and MQS.vhard are its lower signal-to-noise versions. Theorem 3.1 promises that any intervals of significance returned by RNSP at levels α are such that, with probability at least $1 - \alpha$, they each contain at least one true change-point. In addition to RNSP, H-SMUCE, MQS-R, MQS-K, and SN-NSP, we also test versions of RNSP and SN-NSP with the following overlap functions:

$$\begin{aligned}\tau_L(\tilde{s}, \tilde{e}) &= \lfloor (\tilde{s} + \tilde{e})/2 \rfloor - \tilde{s}, \\ \tau_R(\tilde{s}, \tilde{e}) &= \lfloor (\tilde{s} + \tilde{e})/2 \rfloor + 1 - \tilde{e}.\end{aligned}\quad (14)$$

This setting means that upon detecting a generic interval of significance $[\tilde{s}, \tilde{e}]$ within $[s, e]$, the RNSP and SN-NSP algorithms continue on the left interval $[s, \lfloor (\tilde{s} + \tilde{e})/2 \rfloor]$ and the right interval $[\lfloor (\tilde{s} + \tilde{e})/2 \rfloor + 1, e]$ (recall that the no-overlap case results uses the left interval $[s, \tilde{s}]$ and the right interval $[\tilde{e}, e]$). We denote the versions of the two procedures with the overlaps as above by RNSP-O and SN-NSP-O, respectively. As before, we set $\alpha = 0.1$ for all methods tested.

For each model and method tested, we evaluate the following metrics, which collectively promote the detection of genuine, and the non-detection of spurious, intervals of significance: [coverage] the empirical coverage (i.e., whether at least $(1 - \alpha)100\%$ of the simulated sample paths are such that any intervals of significance returned contain at least one true change-point each); [prop. gen. int.] if any intervals are returned, the proportion of those that are genuine (i.e., the proportion of those intervals returned that contain at least one true change-point); [no. gen. int.] the number of genuine intervals (i.e., the number of those intervals returned that contain at least one true change-point); and [av. gen. int. len.] the average length

Table 3. Non-null models for the comparative simulation study in Section 5.

Model name	sample path execution in R
Blocks	<pre>blocks <- c(rep(0, 204), rep(14.63795, 62), rep(-3.659487, 41), rep(7.318975, 164), rep(-7.318975, 40), rep(10.97846, 308), rep(-4.391385, 82), rep(3.293539, 430), rep(19.02933, 225), rep(7.684923, 41), rep(15.36985, 61), rep(-3.250278e-15, 390)) blocks + 10 * rnorm(2048)</pre>
Cauchy	<pre>c(rcauchy(100, 1), rcauchy(100, 2), rcauchy(100, 1))</pre>
Bursts	<pre>(c(rep(1, 200), rep(3, 80), rep(1, 200), rep(3, 80), rep(1, 200), rep(4, 40)) * rnorm(800))²</pre>
Poisson	<pre>as.numeric(rpois(350, c(rep(1, 50), rep(4, 50), rep(10, 50), rep(2, 200))))</pre>
MQS.easy	<pre>f <- c(rep(-0.25, 690), rep(0.12, 435), rep(1.07, 85), rep(-0.53, 255), rep(0.60, 75), rep(-0.69, 155), rep(-0.10, 790)) z <- c(rnorm(350, sd = sqrt(1)), rt(1540-350, df = 3)/sqrt(3)*sqrt(0.1), (rchisq(length(mu)-1540, 1)-qchisq(0.5, 1))/sqrt(2)*sqrt(0.05)) f + z</pre>
MQS.hard	<pre>f + 5 * z</pre>
MQS.vhard	<pre>f + 10 * z</pre>

of genuine intervals (i.e., the average length of those intervals returned that contain at least one true change-point). No single measure describes the performance of (any) method accurately, but the four measures in conjunction provide a clear picture. As a cartoon illustration of how naive solutions are able to skew some of these measures but not the others, consider a putative interval estimator that always returns the longest possible interval of $[1, T]$. For a model with ≥ 1 change-points, “coverage” and “prop. gen. int” will return 100 and 1, respectively. However, this naive solution will be penalized by the next two measures, “no. gen. int.” and “av. gen. int. len.” For models with a single change-point, the $[1, T]$ solution will return an interval that in most cases will be unreasonably long, and this will be picked up by the “av. gen. int. len.” measure. In addition, for models with more than one change-point, “no. gen. int.” (equal to one) will be inaccurate.

For completeness, we also show the Mean-Square Errors (MSEs) between the reconstructed signal \hat{f}_t and the truth f_t . Given an increasingly sorted set $\{\hat{\eta}_j\}_{j=1}^R$ of pointwise change-point estimators (for any method), and defining in addition $\hat{\eta}_0 = 0$, $\hat{\eta}_{R+1} = T$, a natural estimate of the signal f_t is a piecewise-constant vector \hat{f}_t such that $\hat{f}_t = \text{med}(Y_{\hat{\eta}_j+1}, \dots, Y_{\hat{\eta}_{j+1}})$ for $t = \hat{\eta}_j + 1, \dots, \hat{\eta}_{j+1}$, $j = 0, \dots, R$, where med is the empirical median operator. HSMUCE comes with its own pointwise change-point estimates. MQS, NSP or RNSP do not automatically provide pointwise location estimates. For these three methods, we use two different approaches to producing pointwise change-point location estimates within each interval of significance: (a) we take the mid-points of the respective intervals of significance and (b) we take the argument-maximum of the absolute value of the CUSUM statistic of the signs of the data around the empirical median within each interval of significance. (Midpoints of RNSP intervals of significance, while seemingly appearing ad hoc as pointwise estimates of change-point location, often behave well empirically, which may be due to the fact that RNSP pursues short intervals, and those tend to be symmetric around the true change-points as this offers the same amount of evidence on either side of the change-points—hence the frequent empirical closeness of RNSP interval midpoints to the truth.)

Table 4 shows the results. H-SMUCE does not perform well in any scenario, not even in the Blocks model, an instance of

the homogeneous Gaussian model, a simple sub-class of the heterogeneous Gaussian model class for which it was specifically designed (where it achieves the coverage of 30, well short of the expected 90). Its coverage of 100 in the Cauchy model is an artifact of the fact that it does not achieve almost any detections over the 100 simulated sample paths (so there are also no spurious detections).

With the exception of MQS.easy, the least challenging model, MQS frequently produces spurious intervals in signals with change-points: the coverage figures for MQS are well short of 90% in all but one models tested. On the upside, in MQS.easy (the only set-up in which MQS achieves the correct coverage), the average length of genuine intervals is around 20% below that of RNSP.

RNSP and RNSP-O significantly outperform SN-NSP and SN-NSP-O in five out of the seven scenarios tested, the only exception being Bursts and MQS.vhard. In Blocks and Cauchy, the RNSP methods achieve more detections and shorter intervals of significance (so better localization). In Poisson, in addition, they achieve much better coverage (the SN-NSP methods are misled by the discrete nature of this relatively low-intensity Poisson dataset, for which their required asymptotics do not kick in, which results in a very large number of spurious detections). However, the SN-NSP methods work better for the Bursts data in the sense that they lead to more detections. The underlying reason is that the signal level in this model is linearly proportional to the standard deviation of the noise, which particularly suits the self-normalized SN-NSP methods. The RNSP methods are the clear winners for the MQS.easy and MQS.hard models. No method performs particularly well for the MQS.vhard model, but the SN-NSP methods achieve more detections than RNSP there, although at the price of the intervals being much longer.

With regards to the MSE (midpoint) and MSE (CUSUM) measures, the clear overall winner is the RNSP-O method with sign-CUSUM localization; however, all four versions of the RNSP approach offer competitive performance.

6. Data Examples

6.1. The U.S. Real Interest Rates

We first reanalyze the time series of U.S. ex-post real interest rate (the three-month treasury bill rate deflated by the

Table 4. Results for each model+method combination, out of 200 simulated sample paths: “coverage” is the percentage of times that the respective model+method combination did not return a spurious interval of significance; “prop. gen. int.” is the average proportion of genuine intervals out of all intervals returned, if any (if none are returned, the corresponding 0/0 ratio is ignored in the average); “no. gen. int.” is the average number of genuine intervals returned; “av. gen. int. len.” is the average length of a genuine interval returned in the respective model+method combination; “MSE (·)” is the MSE of \hat{f}_t constructed as described in the text, with the respective pointwise change-point estimation in brackets.

Model	attribute	RNSP	RNSP-O	H-SMUCE	MQS-R	MQS-K	SN-NSP	SN-NSP-O
Blocks (11 cpts)	coverage	100	100	30	14	14	100	100
	prop. gen. int.	1	1	0.83	0.75	0.75	1	1
	no. gen. int.	6.17	8.19	4.98	4.77	4.77	4.93	6.10
	av. gen. int. len.	90.65	101.07	94.72	88.71	88.71	160.37	182.83
	MSE (midpoint)	15.47	11.87	12.381(*)	15.903	15.903	16.40	15.5
	MSE (CUSUM)	11.92	6.94	–	13.008	13.008	14.01	11.5
Cauchy (2 cpts)	coverage	100	100	100	71	70	100	99.5
	prop. gen. int.	1	1	1	0.61	0.59	1	0.99
	no. gen. int.	0.62	0.78	0.005	0.58	0.56	0.27	0.3
	av. gen. int. len.	119.16	124.32	170	116.59	116.28	159.58	159.79
	MSE (midpoint)	0.210	0.193	0.235(*)	0.210	0.210	0.229	0.227
	MSE (CUSUM)	0.209	0.194	–	0.208	0.208	0.225	0.224
Bursts (5 cpts)	coverage	100	100	43.5	25.5	27	100	100
	prop. gen. int.	1	1	0.79	0.6	0.61	1	1
	no. gen. int.	3.10	4.44	3.19	2.14	2.17	4.18	5.36
	av. gen. int. len.	100.79	107.64	92.85	94.42	94.40	111.61	108.81
	MSE (midpoint)	20.1	16.73	15.935(*)	20.803	20.770	18.22	14.08
	MSE (CUSUM)	17.4	12.37	–	17.722	17.673	13.51	9.97
Poisson (3 cpts)	coverage	100	100	0	76.5	76	0	0
	prop. gen. int.	1	1	0.1	0.87	0.86	0.07	0.07
	no. gen. int.	2.9	3	1.48	2.37	2.35	0.85	1.02
	av. gen. int. len.	36.99	37.48	19.61	32.1	31.97	33.61	32.69
	MSE (midpoint)	0.716	0.636	1.2322(*)	1.165	1.171	2.5122	2.2460
	MSE (CUSUM)	0.339	0.204	–	0.842	0.860	2.1320	1.9296
MQS.easy (6 cpts)	coverage	100	100	53	100	100	99.5	99
	prop. gen. int.	1	1	0.92	1	1	1	1
	no. gen. int.	6	6	5.99	6	6	5.7	6.2
	av. gen. int. len.	45	46	37.33	37	37	125	122.5
	MSE (midpoint)	0.00272	0.00271	0.0011(*)	0.00279	0.00279	0.031	0.0225
	MSE (CUSUM)	0.00076	0.00072	–	0.00084	0.00084	0.012	0.0013
MQS.hard (6 cpts)	coverage	100	100	33.5	38	38	99	99
	prop. gen. int.	1	1	0.66	0.73	0.73	0.99	1
	no. gen. int.	3.5	4.8	2.13	2.94	2.94	1.67	2.7
	av. gen. int. len.	216.6	194.6	282.97	220.89	220.89	464.08	605.3
	MSE (midpoint)	0.081	0.059	0.10(*)	0.080	0.080	0.11	0.094
	MSE (CUSUM)	0.069	0.036	–	0.069	0.069	0.11	0.080
MQS.vhard (6 cpts)	coverage	100	100	47.5	63	63	99	99
	prop. gen. int.	1	1	0.58	0.63	0.63	0.99	0.99
	no. gen. int.	0.69	1	0.79	0.77	0.77	0.97	1.14
	av. gen. int. len.	422.62	485	722.08	501.86	501.86	1149.57	1179.74
	MSE (midpoint)	0.120	0.115	0.128(*)	0.120	0.120	0.129	0.128
	MSE (CUSUM)	0.119	0.113	–	0.121	0.121	0.133	0.133

(*) note HSMUCE uses its own pointwise change-point estimation. Significance level $\alpha = 0.1$ (nominal coverage = 90%). All numbers rounded to two decimal digits except when such rounding takes a positive number to zero.

CPI inflation rate) considered in Garcia and Perron (1996), Bai and Perron (2003), and Fryzlewicz (2024). The dataset is available at <http://qed.econ.queensu.ca/jae/datasets/bai001/>. The time series, shown in the left plot of Figure 1, is quarterly and the range is 1961:1–1986:3, so $t = 1, \dots, T = 103$.

RNSP appears to be an appropriate tool here, as the data displays heterogeneity and possibly some heavy-tailed movements toward the latter part. We run the RNSP algorithm with the default setting of $M = 1000$, with $\alpha = 0.1$ and with overlaps as defined in (14). The procedure returns two intervals of significance: [23, 75] and [65, 91]. These are shown in the left plot of Figure 1, together with their midpoints as well as the maximizers of absolute CUSUMs of signs of the data around the median in each interval of significance. As with any RNSP execution with nonzero overlaps, one question that may be

asked is whether the two intervals may be indicating the same change-point, but this, visually, is unlikely here (the reason for using nonzero overlaps is simply to provide larger samples for RNSP following the detection of the first interval; using zero overlaps means the samples are too short and RNSP with zero overlaps does not pick up the second change-point). Therefore, the solution points to a model with at least two change-points. This is consistent, or at least not inconsistent, with both Garcia and Perron (1996), who also settle on a model with two change-points, and Bai and Perron (2003), who prefer a three-change-point model, not excluded by RNSP here.

The difference between those two earlier analyses and ours is that those two (a) were based on asymptotic arguments (and therefore valid asymptotically, for unspecified large samples) and (b) were conditional in the sense that the confidence regions




Figure 1. Left: The U.S. 3-month ex-post real interest rate time series (black); intervals of significance returned by RNSP (transparent pink); their midpoints (red); argument-maxima of absolute CUSUMs of signs of the data around the median in each interval of significance (blue). See Section 6.1 for a detailed description. Right: The U.S. 1-month real interest rate time series (black); intervals of significance returned by RNSP (transparent pink); their midpoints (red); argument-maxima of absolute CUSUMs of signs of the data around the median in each interval of significance (blue). Superimposed on the bottom of the graph is the probability of recession time series (purple), on the scale of 0 (bottom purple dashed line) to 1 (top purple dashed line). RNSP significance level $\alpha = 0.1$. See Section 6.1 for a detailed description.

for change-point locations in those two works were conditional on the detection event. By contrast, our analysis via RNSP has a finite-sample nature and the intervals of significance have an unconditional character. Importantly, we do not make any distributional assumptions besides independence and sign-symmetry, both of which are likely to be acceptable for this dataset. The analysis via RNSP is unaffected by the likely heterogeneity in the data.

To illustrate RNSP on a more recent dataset of a similar nature, we examine the 1-month U.S. real interest rate, available from <https://fred.stlouisfed.org/series/REAINTRATREARAT1MO>. The time series, shown in the right plot of Figure 1, is monthly and runs from January 1982 to February 2023. We run RNSP with $M = 1000$, $\alpha = 0.1$ and no overlaps. The intervals of significance returned by RNSP appear visually plausible and it is interesting (albeit not unexpected) to observe that the periods of peaking probabilities of recession (data available from: <https://fred.stlouisfed.org/series/RECPROUSM156N>) from year 2000 onwards are wholly contained within RNSP intervals of significance, indicating declines in the 1-month real interest rate. The period of high probability of recession in 1982 does not appear supported in the 1-month real rate data in a way detectable to RNSP. Finally, the period of high probability of recession in 1990 appears to coincide with a falling 1-month rate but RNSP has, in this case, a visually justifiable preference for intervals just before and just after this likely recession period.

6.2. Interest in the Search Term “Data Science”

We analyze the weekly interest in the search term “data science” from Google Trends, in the U.S. state of California. The link

to obtain the data was <https://trends.google.com/trends/explore?date=today%205-yr&geo=US-CA&q=data%20science>. Google Trends describe the data as follows. “Numbers represent search interest relative to the highest point on the chart for the given region and time. A value of 100 is the peak popularity for the term. A value of 50 means that the term is half as popular. A score of 0 means there was not enough data for this term.” Weeks in this data series start on Sundays and the dataset spans the weeks from w/c 28th August 2016 to w/c 15th August 2021 (so almost five years’ worth of data). The observations are discrete (integers from 22 to 100), which would likely pose difficulties for the competing methods as outlined earlier.

We execute the RNSP procedure with the default setting of $M = 1000$ and with $\alpha = 0.1$, with no overlaps, which returns the three intervals of significance shown in Figure 2. The intervals are: w/c 23 April 2017 – w/c 10 June 2018 (interval 1), w/c 24 June 2018 – w/c 17 November 2019 (interval 2), w/c 3 May 2020 – w/c 14 March 2021 (interval 3). While intervals 1 and 2 correspond to likely increases in the median interest, interval 3 clearly corresponds to a likely decrease, and the midpoint of interval 3 (w/c 4 October 2020) visually aligns well with what seems to be a rather sudden drop of interest.

It is difficult to speculate as to possible reasons for this drop of interest. A blog post on the popular website <https://towardsdatascience.com> (<https://bit.ly/3kqwDWO>) reports that “2020 was the first year since 2016, Data Scientist was not the number one job in America, according to Glassdoor’s annual ranking. That title would belong to Front End Engineer, followed by Java Developer, followed by Data Scientist.” However, visually, a similar decline in interest is observed for example in the analogous Google Trends series for the term “Java” (not shown).




Figure 2. Weekly relative interest (top=100) in the search term “data science” in California in weeks from w/c 28 August 2016 to 15 August 2021 (black); intervals of significance returned by RNSP (transparent pink); their midpoints (red); argument-maxima of absolute CUSUMs of signs of the data around the median in each interval of significance (blue). RNSP significance level $\alpha = 0.1$. See Section 6.2 for a detailed description.

7. Discussion

Other quantiles. Note that Corollary 3.1, crucial to the success of RNSP, can be rewritten as

$$D_{[s_m, e_m]} \leq \|\text{sign}\{Y_{s_m:e_m} - Q_{1/2}(Y_{s_m:e_m})\}\|_{\mathcal{I}_{[s_m, e_m]}^a}, \quad (15)$$

where $Q_q(\cdot)$ is the population q -quantile. However, the left-hand side of (15) was not specifically constructed with $q = 1/2$ in

mind and therefore the inequality $D_{[s_m, e_m]} \leq \|\text{sign}\{Y_{s_m:e_m} - Q_q(Y_{s_m:e_m})\}\|_{\mathcal{I}_{[s_m, e_m]}^a}$ is true for any $q \in (0, 1)$. This shows that RNSP can equally be used for significant change detection in any quantile, and not just the median. However, for $q \neq 1/2$, the challenge is to obtain the null distribution of $\|\text{sign}\{Y - Q_q(Y)\}\|_{\mathcal{I}^a}$. Even if this challenge is overcome (e.g., by simulation), RNSP as defined in this work may not be effective for change detection in quantiles “far” from the median, due to the particular way in which $D_{[s_m, e_m]}$ is constructed (involving minimization over all levels). For RNSP to be a successful device for change detection in other quantiles, the definition of $D_{[s_m, e_m]}$ would have to be modified to only minimize over ‘realistic’ candidate levels not far from the population q -quantile under the local null.

Set of feasible signals. It is convenient to define the set of feasible signals f_t^0 at level α with respect to the algorithmic execution $\text{RNSP}(1, T, \cdot, M, \lambda_\alpha, \tau_L, \tau_R)$ by $\mathcal{F}_\alpha = \{f^0 : \mathcal{S}\{\text{RNSP}(1, T, Y - f^0, M, \lambda_\alpha, \tau_L, \tau_R)\} = \emptyset\}$, where λ_α is such that $P(\|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \lambda_\alpha) \geq 1 - \alpha$. That is, \mathcal{F}_α is the set of postulated signals f^0 such that, on fitting f^0 to the data and obtaining the empirical residuals, RNSP cannot distinguish (at level α) the empirical residuals from white noise. For any candidate fit f^0 , it is straightforward to check whether or not it is a member of \mathcal{F}_α by executing $\text{RNSP}(1, T, Y - f^0, M, \lambda_\alpha, \tau_L, \tau_R)$.

Appendix A: Proofs

Proof of Theorem 4.1. Consider initially the case of a single change-point η_1 . RNSP will, among others, consider intervals symmetric about the true change-point, that is $[\eta_1 - d + 1, \eta_1 + d]$, for all appropriate d . Take a constant candidate fit w on the interval $[\eta_1 - d + 1, \eta_1 + d]$ and define $U_t(w) := \text{sign}(Y_t - w) = 2\mathbb{I}(Y_t - w > 0) - 1$ (the latter equality holds due to the continuity of the distribution of Z_t). Assume $w \log f_{\eta_1} > f_{\eta_1+1}$. We have

$$\begin{aligned} D_{[\eta_1-d+1, \eta_1+d]} &\geq \inf_w \frac{1}{\sqrt{d}} \max \left\{ \left| \sum_{t=\eta_1-d+1}^{\eta_1} U_t(w) \right|, \left| \sum_{t=\eta_1+1}^{\eta_1+d} U_t(w) \right| \right\} \\ &= \frac{2}{\sqrt{d}} \inf_w \max \left\{ \left| \sum_{t=\eta_1-d+1}^{\eta_1} \mathbb{I}(Y_t - w > 0) - \frac{1}{2} \right|, \left| \sum_{t=\eta_1+1}^{\eta_1+d} \mathbb{I}(Y_t - w > 0) - \frac{1}{2} \right| \right\} \\ &\geq \frac{2}{\sqrt{d}} \inf_w \max \left\{ \left| \sum_{t=\eta_1-d+1}^{\eta_1} P(Y_t - w > 0) - \frac{1}{2} \right|, \left| \sum_{t=\eta_1+1}^{\eta_1+d} P(Y_t - w > 0) - \frac{1}{2} \right| \right\} \\ &\quad - \frac{2}{\sqrt{d}} \sup_w \max \left\{ \left| \sum_{t=\eta_1-d+1}^{\eta_1} \epsilon_t^Y(w) \right|, \left| \sum_{t=\eta_1+1}^{\eta_1+d} \epsilon_t^Y(w) \right| \right\}. \end{aligned} \quad (16)$$

Now note

$$\begin{aligned} \frac{1}{\sqrt{d}} \sup_w \max \left\{ \left| \sum_{t=\eta_1-d+1}^{\eta_1} \epsilon_t^Y(w) \right|, \left| \sum_{t=\eta_1+1}^{\eta_1+d} \epsilon_t^Y(w) \right| \right\} &\leq \max_{s,e} \sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Y(w) \right| \\ &= \max_{s,e} \sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Z(w) \right| \leq \lambda. \end{aligned}$$

Continuing from (16), this implies

$$\begin{aligned}
 & D_{[\eta_1-d+1, \eta_1+d]} \\
 & \geq \frac{2}{\sqrt{d}} \inf_w \max \left\{ \left| \sum_{t=\eta_1-d+1}^{\eta_1} P(Y_t - w > 0) - \frac{1}{2} \right|, \left| \sum_{t=\eta_1+1}^{\eta_1+d} P(Y_t - w > 0) - \frac{1}{2} \right| \right\} - 2\lambda \\
 & = \frac{2}{\sqrt{d}} \inf_w \max \left\{ \left| \sum_{t=\eta_1-d+1}^{\eta_1} P(Z_t > w - f_{\eta_1}) - \frac{1}{2} \right|, \left| \sum_{t=\eta_1+1}^{\eta_1+d} P(Z_t > w - f_{\eta_1+1}) - \frac{1}{2} \right| \right\} - 2\lambda. \tag{17}
 \end{aligned}$$

The infimum over w will be achieved if both elements of the maximum are the same (or otherwise it would be possible to alter w slightly to decrease the larger of the two moduli). But this is only possible if $w \in (f_{\eta_1+1}, f_{\eta_1})$ and hence, bearing in mind that Z_t is median-zero, we have

$$\begin{aligned}
 \left| \sum_{t=\eta_1-d+1}^{\eta_1} P(Z_t > w - f_{\eta_1}) - \frac{1}{2} \right| &= dP(Z_t > w - f_{\eta_1}) - d/2 \\
 &= dP\{Z_t \in (w - f_{\eta_1}, 0)\}, \\
 \left| \sum_{t=\eta_1+1}^{\eta_1+d} P(Z_t > w - f_{\eta_1+1}) - \frac{1}{2} \right| &= d/2 - dP(Z_t > w - f_{\eta_1+1}) \\
 &= dP\{Z_t \in (0, w - f_{\eta_1+1})\}.
 \end{aligned}$$

Let w_0 be such that $P\{Z_t \in (w_0 - f_{\eta_1}, 0)\} = P\{Z_t \in (0, w_0 - f_{\eta_1+1})\}$. Since $(w_0 - f_{\eta_1+1}) - (w_0 - f_{\eta_1}) = f_{\eta_1} - f_{\eta_1+1}$, then either $f_{\eta_1} - w_0 \geq (f_{\eta_1} - f_{\eta_1+1})/2$ or $w_0 - f_{\eta_1+1} \geq (f_{\eta_1} - f_{\eta_1+1})/2$. Therefore,

$$\begin{aligned}
 P\{Z_t \in (w_0 - f_{\eta_1}, 0)\} &= P\{Z_t \in (0, w_0 - f_{\eta_1+1})\} \\
 &\geq \min\{P\{Z_t \in (-|f_{\eta_1} - f_{\eta_1+1}|/2, 0)\}, \\
 &\quad P\{Z_t \in (0, |f_{\eta_1} - f_{\eta_1+1}|/2)\}\} = \Delta_1.
 \end{aligned}$$

Continuing from (17), we therefore have

$$D_{[\eta_1-d+1, \eta_1+d]} \geq 2\sqrt{d}\Delta_1 - 2\lambda. \tag{18}$$

But from the definition of the RNSP algorithm (line 14), detection on an interval $[s, e]$ will occur if $D_{[s,e]} > \lambda_\alpha$. Therefore, from (18), detection on $[\eta_1 - d + 1, \eta_1 + d]$ will occur if

$$d > \left(\frac{2\lambda + \lambda_\alpha}{2\Delta_1} \right)^2. \tag{19}$$

As RNSP looks for the shortest intervals of significance, the length of the RNSP interval will not exceed that of $[\eta_1 - d + 1, \eta_1 + d]$, which is $2d$. From (19), its length will therefore be bounded from above by $2 \left\lceil \left(\frac{2\lambda + \lambda_\alpha}{2\Delta_1} \right)^2 + 1 \right\rceil = 2\bar{d}_1$. We now turn our attention to the multiple change-point case. Note that even though the RNSP interval of significance around η_j is guaranteed to be of length at most $2\bar{d}_j$, it will not necessarily be a sub-interval of $[\eta_j - \bar{d}_j + 1, \eta_j + \bar{d}_j]$. Therefore, in order that an interval detection around η_j does not interfere with detections around η_{j-1} or η_{j+1} , the distances $\eta_j - \eta_{j-1}$ and $\eta_{j+1} - \eta_j$ must be suitably long, but this is guaranteed by [Assumption 4.1\(iii\)](#). This completes the proof. \square

Proof of Corollary 4.1. The fact that $P(\|\text{sign}(Z)\|_{\mathcal{I}^a} > \lambda_\alpha) \rightarrow 0$ as $T \rightarrow \infty$ is a simple consequence of Corollary 1 in [Shao \(1995\)](#). We next assess and bound the magnitude of $\sup_w \left| \frac{1}{\sqrt{d}} \sum_{t=s}^{s+d-1} \epsilon_t^Z(w) \right|$.

The Dvoretzky-Kiefer-Wolfowitz inequality (with Massart's optimal constant, see [Massart 1990](#)) implies

$$\begin{aligned}
 & P \left(\sup_w \left| \frac{1}{\sqrt{d}} \sum_{t=s}^{s+d-1} \epsilon_t^Z(w) \right| > \lambda \right) \\
 & = P \left(\sup_w \left| \frac{1}{d} \sum_{t=s}^{s+d-1} \epsilon_t^Z(w) \right| > \lambda d^{-1/2} \right) \leq 2 \exp(-2\lambda^2).
 \end{aligned}$$

This leads to a uniform bound via Bonferroni's correction.

$$\begin{aligned}
 & P \left(\max_{s,e} \sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Z(w) \right| > \lambda \right) \\
 & \leq \sum_{s \leq e} P \left(\sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Z(w) \right| > \lambda \right) \\
 & \leq T(T+1) \exp(-2\lambda^2).
 \end{aligned}$$

For $\lambda = (1 + \delta) \log^{1/2} T$, the above tends to zero if $\delta > 0$. This completes the proof. \square

Proof of Corollary 4.2. As in the proof of [Corollary 4.1](#), setting α at this level means that

$$P \left(\max_{s,e} \sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Z(w) \right| \leq \lambda \right) \rightarrow 1$$

as $T \rightarrow \infty$, and therefore

$$\liminf_{T \rightarrow \infty} P \left(\|\text{sign}(Z)\|_{\mathcal{I}^a} \leq \lambda_\alpha \wedge \max_{s,e} \sup_w \left| \frac{1}{\sqrt{e-s+1}} \sum_{t=s}^e \epsilon_t^Z(w) \right| \leq \lambda \right) \geq 1 - \alpha.$$

[Theorem 4.1](#) implies the result. \square

Supplementary Materials

R code to accompany the paper appears at <https://github.com/pfryz/nsp>.

Disclosure Statement

There are no competing interests to declare.

Funding

The author acknowledges EPSRC grant EP/V053639/1.

References

- Bai, J., and Perron, P. (1998), “Estimating and Testing Linear Models with Multiple Structural Changes,” *Econometrica*, 66, 47–78. [1,2]
- (2003), “Computation and Analysis of Multiple Structural Change Models,” *Journal of Applied Econometrics*, 18, 1–22. [2,10]
- Baranowski, R., Chen, Y., and Fryzlewicz, P. (2019), “Narrowest-Over-Threshold Detection of Multiple Change-Points and Change-Point-Like Features,” *Journal of the Royal Statistical Society, Series B*, 81, 649–672. [7]
- Bhattacharyya, G., and Johnson, R. (1968), “Nonparametric Tests for Shift at an Unknown Time Point,” *The Annals of Mathematical Statistics*, 39, 1731–1743. [3]
- Carlstein, E. (1988), “Nonparametric Change-Point Estimation,” *Annals of Statistics*, 16, 188–197. [3]
- Chen, Y., Shah, R., and Samworth, R. (2014), “Discussion of ‘Multiscale Change Point Inference’ by Frick, Munk and Sieling,” *Journal of the Royal Statistical Society, Series B*, 76, 544–546. [2]
- Cheng, D., He, Z., and Schwartzman, A. (2020), “Multiple Testing of Local Extrema for Detection of Change Points,” *Electronic Journal of Statistics*, 14, 3705–3729. [1]
- Cho, H., and Kirch, C. (2022), “Bootstrap Confidence Intervals for Multiple Change Points based on Moving Sum Procedures,” *Computational Statistics and Data Analysis*, 175, Article no. 107552. [1]
- Dümbgen, L. (1991), “The Asymptotic Behavior of Some Nonparametric Change-Point Estimators,” *Annals of Statistics*, 19, 1471–1495. [3]
- Duy, V. N. L., Toda, H., Sugiyama, R., and Takeuchi, I. (2020), “Computing Valid p -value for Optimal Change-point by Selective Inference Using Dynamizing Programming,” in *Advances in Neural Information Processing Systems*, (Vol. 33), pp. 11356–11367. [1]
- Eichinger, B., and Kirch, C. (2018), “A MOSUM Procedure for the Estimation of Multiple Random Change Points,” *Bernoulli*, 24, 526–564. [1]
- Fearnhead, P. (2006), “Exact and Efficient Bayesian Inference for Multiple Change-point Problems,” *Statistics and Computing*, 16, 203–213. [1]
- Frick, K., Munk, A., and Sieling, H. (2014), “Multiscale Change-Point Inference,” (with Discussion), *Journal of the Royal Statistical Society, Series B*, 76, 495–580. [1]
- Fryzlewicz, P. (2024), “Narrowest Significance Pursuit: Inference for Multiple Change-Points in Linear Models,” *Journal of the American Statistical Association*, to appear. [1,2,3,5,7,8,10]
- Garcia, R., and Perron, P. (1996), “An Analysis of the Real Interest Rate under Regime Shifts,” *Review of Economics and Statistics*, 78, 111–125. [10]
- Hao, N., Niu, Y., and Zhang, H. (2013), “Multiple Change-Point Detection via a Screening and Ranking Algorithm,” *Statistica Sinica*, 23, 1553–1572. [1]
- Hyun, S., G’Sell, M., and Tibshirani, R. (2018), “Exact Post-Selection Inference for the Generalized Lasso Path,” *Electronic Journal of Statistics*, 12, 1053–1097. [1]
- Hyun, S., Lin, K., G’Sell, M., and Tibshirani, R. (2021), “Post-Selection Inference for Change-point Detection Algorithms with Application to Copy Number Variation Data,” *Biometrics*, 77, 1037–1049. [1]
- Jewell, S., Fearnhead, P., and Witten, D. (2022), “Testing for a Change in Mean After Change-point Detection,” *Journal of the Royal Statistical Society, Series B*, 84, 1082–1104. [1]
- Kabluchko, Z., and Wang, Y. (2014), “Limiting Distribution for the Maximal Standardized Increment of a Random Walk,” *Stochastic Processes and their Applications*, 124, 2824–2867. [6]
- Lavielle, M., and Lebarbier, E. (2001), “An Application of MCMC Methods for the Multiple Change-Points Problem,” *Signal Processing*, 81, 39–53. [1]
- Li, H. (2016), “Variational Estimators in Statistical Multiscale Analysis,” PhD thesis, Georg August University of Göttingen. [4]
- Li, H., and Munk, A. (2016), “FDR-Control in Multiscale Change-Point Segmentation,” *Electronic Journal of Statistics*, 10, 918–959. [1]
- Massart, P. (1990), “The Tight Constant in the Dvoretzky-Kiefer-Wolfowitz Inequality,” *The Annals of Probability*, 18, 1269–1283. [13]
- Nemirovski, A. (1986), “Nonparametric Estimation of Smooth Regression Functions,” *Journal of Computer and System Sciences*, 23, 1–11. [4]
- Pein, F., Sieling, H., and Munk, A. (2017), “Heterogeneous Change Point Inference,” *Journal of the Royal Statistical Society, Series B*, 79, 1207–1227. [1,7,8]
- Sen, A., and Srivastava, M. S. (1975), “On Tests for Detecting Change in Mean,” *The Annals of Statistics*, 3, 98–108. [7]
- Shao, Q.-M. (1995), “On a Conjecture of Révész,” *Proceedings of the American Mathematical Society*, 123, 575–582. [13]
- Vanegas, L., Behr, M., and Munk, A. (2022), “Multiscale Quantile Segmentation,” *Journal of the American Statistical Association*, 117, 1384–1397. [1,2,7,8]