

Inference in high-dimensional online changepoint detection

Yudong Chen^{*,†}, Tengyao Wang[†] and Richard J. Samworth^{*}

^{*} Statistical Laboratory, University of Cambridge

[†] Department of Statistics, London School of Economics

(February 17, 2023)

Abstract

We introduce and study two new inferential challenges associated with the sequential detection of change in a high-dimensional mean vector. First, we seek a confidence interval for the changepoint, and second, we estimate the set of indices of coordinates in which the mean changes. We propose an online algorithm that produces an interval with guaranteed nominal coverage, and whose length is, with high probability, of the same order as the average detection delay, up to a logarithmic factor. The corresponding support estimate enjoys control of both false negatives and false positives. Simulations confirm the effectiveness of our methodology, and we also illustrate its applicability on the US excess deaths data from 2017–2020.

1 Introduction

The real-time monitoring of evolving processes has become a characteristic feature of 21st century life. Watches and defibrillators track health data, Covid-19 case numbers are reported on a daily basis and financial decisions are made continuously based on the latest market movements. Given that changes in the dynamics of such processes are frequently of great interest, it is unsurprising that the area of changepoint detection has undergone a renaissance over the last 5–10 years.

One of the features of modern datasets that has driven much of the recent research in changepoint analysis is high dimensionality, where we monitor many processes simultaneously, and seek to borrow strength across the different series to identify changepoints. The nature of series that are tracked in applications, as well as the desire to evade to the greatest extent possible the curse of dimensionality, means that it is commonly assumed that the signal of interest is relatively sparse, in the sense that only a small proportion of the constituent series undergo a change. Furthermore, the large majority of these works have focused on the retrospective (or *offline*) challenges of detecting and estimating changes after seeing all of the available data (e.g. [Chan and Walther, 2015](#); [Cho and Fryzlewicz, 2015](#); [Jirak, 2015](#); [Cho, 2016](#); [Soh and Chandrasekaran, 2017](#); [Wang and Samworth, 2018](#); [Enikeeva and Harchaoui,](#)

2019; Padilla et al., 2021; Kaul et al., 2021; Liu et al., 2021; Londschien et al., 2021; Rinaldo et al., 2021; Follain et al., 2022). Nevertheless, the related problem where one observes data sequentially and seeks to declare changes as soon as possible after they have occurred, is nowadays receiving increasing attention (e.g. Kirch and Stoehr, 2019; Dette and Gösmann, 2020; Gösmann et al., 2020; Yu et al., 2021b; Chen et al., 2022). Although the focus of our review here has been on recent developments, including finite-sample results in multivariate and high-dimensional settings, we also mention that changepoint analysis has a long history (e.g. Page, 1954). Entry points to this classical literature include Csörgő and Horváth (1997) and Horváth and Rice (2014). For univariate data, sequential changepoint detection is also studied under the banner of statistical process control (Duncan, 1952; Tartakovsky et al., 2014). In the field of high-dimensional statistical inference more generally, uncertainty quantification has become a major theme over the last decade, originating with influential work on the debiased Lasso in (generalized) linear models (Javanmard and Montanari, 2014; van de Geer et al., 2014; Zhang and Zhang, 2014), and subsequently developed in other settings (e.g. Janková and van de Geer, 2015; Yu et al., 2021a).

The aim of this paper is to propose methods to address two new inferential challenges associated with the high-dimensional, sequential detection of a sparse change in mean. The first is to provide a confidence interval for the location of the changepoint, while the second is to estimate the signal set of indices of coordinates that undergo the change. Despite the importance of uncertainty quantification and signal support recovery in changepoint applications, neither of these problems has previously been studied in the multivariate sequential changepoint detection literature, to the best of our knowledge. Of course, one option here would be to apply an offline confidence interval construction (e.g., Kaul et al., 2021) after a sequential procedure has declared a change. However, this would be to ignore the essential challenge of the sequential nature of the problem, whereby one wishes to avoid storing all historical data, to enable inference to be carried out in an *online* manner. By this we mean that the computational complexity for processing each new observation, as well as the storage requirements, should depend only on the number of bits needed to represent the new data point observed¹. The online requirement turns out to impose severe restrictions on the class of algorithms available to the practitioner, and lies at the heart of the difficulty of the problem.

To give a brief outline of our construction of a confidence interval with guaranteed $(1 - \alpha)$ -level coverage, consider for simplicity the univariate setting, where $(X_n)_{n \in \mathbb{N}}$ form a sequence of independent random variables with $X_1, \dots, X_z \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $X_{z+1}, X_{z+2}, \dots \stackrel{\text{iid}}{\sim} \mathcal{N}(\theta, 1)$. Without loss of generality, we assume that $\theta > 0$. Suppose that θ is known to be at least $b > 0$ and, for $n \in \mathbb{N}$, define *residual tail lengths*

$$t_{n,b} := \operatorname{argmax}_{0 \leq h \leq n} \sum_{i=n-h+1}^n (X_i - b/2). \quad (1)$$

In the case of a tie, we choose the smallest h achieving the maximum. Since $\sum_{i=n-h+1}^n (X_i - b/2)$ can be viewed as the likelihood ratio statistic for testing the null of $\mathcal{N}(0, 1)$ against the

¹Here, we ignore the errors in rounding real numbers to machine precision; thus, we do not distinguish between continuous random variables and quantized versions where the data have been rounded to machine precision.

alternative of $\mathcal{N}(b, 1)$ using X_{n-h+1}, \dots, X_n , the quantity $t_{n,b}$ is the tail length for which the likelihood ratio statistic is maximized. If N is the stopping time defining a good sequential changepoint detection procedure, then, intuitively, $N - t_{N,b}$ should be close to the true changepoint location z , and almost pivotal. This motivates the construction of a confidence interval of the form $[\max\{N - t_{N,b} - g(\alpha, b), 0\}, N]$, where we control the tail probability of the distribution of $N - t_{N,b}$ to choose $g(\alpha, b)$ so as to ensure the desired coverage. In the multivariate case, considerable care is required to handle the post-selection nature of the inferential problem, as well as to determine an appropriate left endpoint for the confidence interval. For this latter purpose, we only assume a lower bound on the Euclidean norm of the vector of mean change, and employ a delicate multivariate and multiscale aggregation scheme; see Section 2 for details, as well as Section 3.5 for further discussion.

The procedure for the inference tasks discussed above, which we call `ocd_CI` (short for **online changepoint detection Confidence Intervals**), can be run in conjunction with any base sequential changepoint detection procedure. However, we recommend using the `ocd` algorithm introduced by Chen et al. (2022), or its variant `ocd'`, which provides guarantees on both the average and worst-case detection delays, subject to a guarantee on the *patience*, or average false alarm rate under the null hypothesis of no change. Crucially, these are both online algorithms. The corresponding inferential procedures inherit this same online property, thereby making them applicable even in very high-dimensional settings and where changes may be rare, so we may need to see many new data points before declaring a change.

In Section 3 we study the theoretical performance of the `ocd_CI` procedure. In particular, we prove in Theorem 1 that, whenever the base sequential detection procedure satisfies a patience and detection delay condition, the confidence interval has at least nominal coverage for a suitable choice of input parameters. Theorem 2 provides a corresponding guarantee on the length of the interval. In Section 3.3, we show that by using `ocd'` as the base procedure, the aforementioned patience and detection delay condition is indeed satisfied. As a result, the output confidence interval has guaranteed nominal coverage and the length of the interval is of the same order as the average detection delay for the base `ocd'` procedure, up to a poly-logarithmic factor. This is remarkable in view of the intrinsic challenge that the better such a changepoint detection procedure performs, the fewer post-change observations are available for inferential tasks.

A very useful byproduct of our `ocd_CI` methodology is that we obtain a natural estimate of the set of signal coordinates (i.e. those that undergo change). In Theorem 3, we prove that, with high probability, it is able both to recover the effective support of the signal (see Section 3.1 for a formal definition), and to avoid noise coordinates. We then broaden the scope of applicability of our methodology in Section 3.4 by relaxing our distributional assumptions to deal with sub-Gaussian or sub-exponential data. Finally, in Section 3.5, we introduce a modification of our algorithm that permits an arbitrarily loose lower bound $\beta > 0$ on the Euclidean norm of the vector of mean change to be employed, with only a logarithmic increase in the confidence interval length guarantee and the computational cost.

An attraction of our theoretical results is that we are able to handle arbitrary spatial (cross-sectional) dependence between the different coordinates of our data stream. On the other hand, two limitations of our analysis for practical use are that real data may exhibit both heavier than sub-exponential tails and temporal dependence. While a full theoretical

analysis of the `ocd_CI` algorithm in these contexts appears to be challenging, we have made some practical suggestions regarding these issues in Sections 3.4 and 4.4 respectively.

Section 4 is devoted to a study of the numerical performance of our methodological proposals. Our simulations confirm that the `ocd_CI` methodology (with the `ocd` base procedure) attains the desired coverage level across a wide range of parameter settings, that the average confidence interval length is of comparable order to the average detection delay and that our support recovery guarantees are validated empirically. We further demonstrate the way in which naive application of offline methods may lead to poor performance in this problem. Moreover, in Section 4.4, we apply our methods to excess death data from the Covid-19 pandemic in the US. Proofs, auxiliary results, extensions to sub-Gaussian and sub-exponential settings and additional simulation results are provided in the supplementary material (Section 5).

We conclude this introduction with some notation used throughout the paper. We write \mathbb{N}_0 for the set of all non-negative integers. For $d \in \mathbb{N}$, we write $[d] := \{1, \dots, d\}$. Given $a, b \in \mathbb{R}$, we denote $a \vee b := \max(a, b)$ and $a \wedge b := \min(a, b)$. For a set S , we use $\mathbb{1}_S$ and $|S|$ to denote its indicator function and cardinality respectively. For a real-valued function f on a totally ordered set S , we write $\text{sargmax}_{x \in S} f(x) := \min \text{argmax}_{x \in S} f(x)$ and $\text{largmax}_{x \in S} f(x) := \max \text{argmax}_{x \in S} f(x)$ for the smallest and largest maximizers of f in S , and define $\text{sargmin}_{x \in S} f(x)$ and $\text{largmin}_{x \in S} f(x)$ analogously. For a vector $v = (v^1, \dots, v^M)^\top \in \mathbb{R}^M$, we define $\|v\|_0 := \sum_{i=1}^M \mathbb{1}_{\{v^i \neq 0\}}$, $\|v\|_2 := \{\sum_{i=1}^M (v^i)^2\}^{1/2}$ and $\|v\|_\infty := \max_{i \in [M]} |v^i|$. In addition, for $j \in [M]$, we define $\|v^{-j}\|_2 := \{\sum_{i:i \neq j} (v^i)^2\}^{1/2}$. For a matrix $A = (A^{i,j}) \in \mathbb{R}^{d_1 \times d_2}$ and $j \in [d_2]$, we write $A^{:,j} := (A^{1,j}, \dots, A^{d_1,j})^\top \in \mathbb{R}^{d_1}$ and $A^{-j,j} := (A^{1,j}, \dots, A^{j-1,j}, A^{j+1,j}, \dots, A^{d_1,j})^\top \in \mathbb{R}^{d_1-1}$. We use $\Phi(\cdot)$, $\bar{\Phi}(\cdot)$ and $\phi(\cdot)$ to denote the distribution function, survivor function and density function of the standard normal distribution respectively. For two real-valued random variables U and V , we write $U \geq_{\text{st}} V$ or $V \leq_{\text{st}} U$ if $\mathbb{P}(U \leq x) \leq \mathbb{P}(V \leq x)$ for all $x \in \mathbb{R}$. We adopt conventions that an empty sum is 0 and that $\min \emptyset := \infty$, $\max \emptyset := -\infty$.

2 Confidence interval construction and support estimation methodology

In the multivariate sequential changepoint detection problem, we observe p -variate observations X_1, X_2, \dots in turn, and seek to report a stopping time N by which we believe a change has occurred. Here and throughout, a stopping time is understood to be with respect to the natural filtration, so that the event $\{N = n\}$ belongs to the σ -algebra generated by X_1, \dots, X_n . The focus of this work is on changes in the mean of the underlying process, and we denote the time of the changepoint by z . Moreover, since our primary interest is in high-dimensional settings, we will also seek to exploit sparsity in the vector of mean change. Given $\alpha \in (0, 1)$, then, our primary goal is to construct a confidence interval $\mathcal{C} \equiv \mathcal{C}(X_1, \dots, X_N, \alpha)$ with the property that $z \in \mathcal{C}$ with probability at least $1 - \alpha$.

For $i \in \mathbb{N}$ and $j \in [p]$, let X_i^j denote the j th coordinate of X_i . The `ocd_CI` algorithm relies on a lower bound $\beta > 0$ for the ℓ_2 -norm of the vector of mean change, sets of signed

scales \mathcal{B} and \mathcal{B}_0 defined in terms of β and a base sequential changepoint detection procedure CP. As CP processes each new data vector, we update the matrix of residual tail lengths $(t_{n,b}^j)_{j \in [p], b \in \mathcal{B} \cup \mathcal{B}_0}$ with $t_{n,b}^j := \text{sargmax}_{0 \leq h \leq n} \sum_{i=n-h+1}^n (X_i^j - b/2)$, as well as corresponding tail partial sum vectors $(A_{n,b}^{j,j'})_{j \in [p], b \in \mathcal{B} \cup \mathcal{B}_0}$, where $A_{n,b}^{j',j} := \sum_{i=n-t_{n,b}^j+1}^n X_i^{j'}$.

After the base procedure CP declares a change at a stopping time N , we identify an ‘‘anchor’’ coordinate $\hat{j} \in [p]$ and a signed anchor scale $\hat{b} \in \mathcal{B}$, where

$$(\hat{j}, \hat{b}) := \underset{(j,b) \in [p] \times \mathcal{B}}{\text{argmax}} \sum_{j' \in [p] \setminus \{j\}} \frac{(A_{N,b}^{j',j})^2}{t_{N,b}^j \vee 1} \mathbb{1}_{\{|A_{N,b}^{j',j}| \geq a \sqrt{t_{N,b}^j \vee 1}\}}.$$

The intuition is that the anchor coordinate and signed anchor scale are chosen so that the final $t_{N,\hat{b}}^{\hat{j}}$ observations provide the best evidence among all of the residual tail lengths against the null hypothesis of no change. Meanwhile, $A_{N,\hat{b}}^{\hat{j}}$ aggregates the last $t_{N,\hat{b}}^{\hat{j}}$ observations in each coordinate, providing a measure of the strength of this evidence against the null.

The main idea of our confidence interval construction is to seek to identify coordinates with large post-change signal. To this end, observe when $t_{N,\hat{b}}^{\hat{j}}$ is not too much larger than $N - z$, the quantity $E_{N,\hat{b}}^{j,\hat{j}} := A_{N,\hat{b}}^{j,\hat{j}} / (t_{N,\hat{b}}^{\hat{j}} \vee 1)^{1/2}$ should be centered close to $\theta^j (t_{N,\hat{b}}^{\hat{j}})^{1/2}$ for $j \in [p] \setminus \{\hat{j}\}$, with variance close to 1. Indeed, if \hat{j} , \hat{b} , N and $t_{N,\hat{b}}^{\hat{j}}$ were fixed, and if $0 < t_{N,\hat{b}}^{\hat{j}} \leq N - z$, then the former quantity would have unit variance around this centering value. The random nature of these quantities, however, introduces a post-selection inference aspect to the problem. Nevertheless, by choosing an appropriate threshold value $d_1 > 0$, we can ensure that with high probability, when $j \neq \hat{j}$ is a noise coordinate, we have $|E_{N,\hat{b}}^{j,\hat{j}}| < d_1$, and when $j \neq \hat{j}$ is a coordinate with sufficiently large signal, there exists a signed scale $b \in (\mathcal{B} \cup \mathcal{B}_0) \cap [-|\theta^j|, |\theta^j|]$, having the same sign as θ^j , for which $|E_{N,\hat{b}}^{j,\hat{j}}| - |b| (t_{N,\hat{b}}^{\hat{j}})^{1/2} \geq d_1$. In fact, such a signed scale, if it exists, can always be chosen to be from \mathcal{B}_0 . As a convenient byproduct, the set of indices j for which the latter inequality holds, which we denote as $\hat{\mathcal{S}}$, forms a natural estimate of the set of coordinates in which the mean change is large.

For each $j \in \hat{\mathcal{S}}$, there exists a largest scale $b \in (\mathcal{B} \cup \mathcal{B}_0) \cap (0, \infty)$ for which $|E_{N,\hat{b}}^{j,\hat{j}}| - b(t_{N,\hat{b}}^{\hat{j}})^{1/2} \geq d_1$. We denote the signed version of this quantity, where the sign is chosen to agree with that of $E_{N,\hat{b}}^{j,\hat{j}}$, by \tilde{b}^j ; this can be regarded as a shrunken estimate of θ^j , so plays the role of the lower bound b from the univariate problem discussed in the introduction. Finally, then, our confidence interval is constructed as the intersection over indices $j \in \hat{\mathcal{S}}$ of the confidence interval from the univariate problem in coordinate j , with signed scale \tilde{b}^j .

As a device to facilitate our theory, the `ocd_CI` algorithm allows the practitioner the possibility of observing a further ℓ observations after the time of changepoint declaration, before constructing the confidence interval. The additional observations are used to determine the anchor coordinate \hat{j} and scale \hat{b} , as well as the estimated support $\hat{\mathcal{S}}$ and the estimated scale \tilde{b}^j for each $j \in \hat{\mathcal{S}}$. Thus, the extra sampling is used to guard against an unusually early changepoint declaration that leaves very few post-change observations for inference. Nevertheless, we will see in Theorem 1 below that the output confidence interval has guaranteed

nominal coverage even with $\ell = 0$, so that additional observations are only used to control the length of the interval. In fact, even for this latter aspect, the numerical evidence presented in Section 4 indicates that $\ell = 0$ provides confidence intervals of reasonable length in practice. Similarly, Theorem 3 ensures that with high probability, our support estimate $\hat{\mathcal{S}}$ contains no noise coordinates (i.e. has false positive control) even with $\ell = 0$, so that the extra sampling is only used to provide false negative control.

Pseudo-code for this `ocd_CI` confidence interval construction is given in Algorithm 1, where we suppress the n dependence on quantities that are updated at each time step. The computational complexity per new observation, as well as the storage requirements, of this algorithm is equal to the sum of the corresponding quantities for the CP base procedure and $O(p^2 \log(ep))$ regardless of the observation history. Thus the `ocd_CI` method inherits the property of being an online algorithm, as discussed in the introduction, from any online CP base procedure.

Algorithm 1: Pseudo-code for the confidence interval construction algorithm `ocd_CI`

Input: $X_1, X_2, \dots \in \mathbb{R}^p$ observed sequentially, $\beta > 0$, $a \geq 0$, an online changepoint detection procedure CP, $d_1, d_2 > 0$ and $\ell \in \mathbb{N}_0$

Set: $b_{\min} = \frac{\beta}{\sqrt{2^{\lfloor \log_2(2p) \rfloor} \log_2(2p)}}$, $\mathcal{B}_0 = \{\pm b_{\min}\}$,

$\mathcal{B} = \{\pm 2^{m/2} b_{\min} : m = 1, \dots, \lfloor \log_2(2p) \rfloor\}$, $n = 0$, $A_b = \mathbf{0} \in \mathbb{R}^{p \times p}$ and $t_b = 0 \in \mathbb{R}^p$ for all $b \in \mathcal{B} \cup \mathcal{B}_0$

repeat

$n \leftarrow n + 1$

 observe new data vector X_n and update CP with X_n

for $(j, b) \in [p] \times (\mathcal{B} \cup \mathcal{B}_0)$ **do**

$t_b^j \leftarrow t_b^j + 1$

$A_b^{j,j} \leftarrow A_b^{j,j} + X_n$

if $b A_b^{j,j} - b^2 t_b^j / 2 \leq 0$ **then**

$t_b^j \leftarrow 0$ and $A_b^{j,j} \leftarrow 0$

until CP declares a change;

Observe ℓ new data vectors $X_{n+1}, \dots, X_{n+\ell}$

Set $E_b^{j',j} \leftarrow \frac{A_b^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'}}{\sqrt{(t_b^j + \ell) \vee 1}}$ for $j', j \in [p]$, $b \in \mathcal{B} \cup \mathcal{B}_0$

Compute $Q_b^j \leftarrow \sum_{j' \in [p] \setminus \{j\}} (E_b^{j',j})^2 \mathbb{1}_{\{|E_b^{j',j}| \geq a\}}$ for $j \in [p]$, $b \in \mathcal{B}$

$(\hat{j}, \hat{b}) \leftarrow \operatorname{argmax}_{(j,b) \in [p] \times \mathcal{B}} Q_b^j$

$\hat{\mathcal{S}} \leftarrow \left\{ j \in [p] \setminus \{\hat{j}\} : |E_b^{j,\hat{j}}| - b_{\min} (t_b^j + \ell)^{1/2} \geq d_1 \right\}$

for $j \in \hat{\mathcal{S}}$ **do**

$\tilde{b}^j \leftarrow \operatorname{sgn}(E_b^{j,\hat{j}}) \max \left\{ b \in (\mathcal{B} \cup \mathcal{B}_0) \cap (0, \infty) : |E_b^{j,\hat{j}}| - b (t_b^j + \ell)^{1/2} \geq d_1 \right\}$

Output: Confidence interval $\mathcal{C} = \left[\max \left\{ n - \min_{j \in \hat{\mathcal{S}}} \left\{ t_{\tilde{b}^j}^j + \frac{d_2}{(\tilde{b}^j)^2} \right\}, 0 \right\}, n \right]$

A natural choice for the base online changepoint detection procedure CP is the `ocd` algo-

rithm, or its variant `ocd'`, introduced by [Chen et al. \(2022\)](#). Both are online algorithms, with computational complexity per new observation and storage requirements of $O(p^2 \log(ep))$. The `ocd'` base procedure is considered for the theoretical analysis in [Section 3](#) due to its known patience and detection delay guarantees, while we prefer `ocd` for numerical studies and practical use. For the reader's convenience, the `ocd` and `ocd'` algorithms are provided as [Algorithms 2](#) and [3](#) respectively in [Section 5.3](#).

3 Theoretical analysis

Throughout this section, we will assume that the sequential observations X_1, X_2, \dots are independent, and that for some unknown covariance matrix $\Sigma \in \mathbb{R}^{p \times p}$ whose diagonal entries are all equal to 1, there exist $z \in \mathbb{N}_0$ and $\theta = (\theta^1, \dots, \theta^p)^\top \neq 0$ for which $X_1, \dots, X_z \sim \mathcal{N}_p(0, \Sigma)$ and $X_{z+1}, X_{z+2}, \dots \sim \mathcal{N}_p(\theta, \Sigma)$. We let $\vartheta := \|\theta\|_2$, and write $\mathbb{P}_{z, \theta, \Sigma}$ for probabilities computed under this model, though in places we omit the subscripts for brevity. Define the *effective sparsity* of θ , denoted $s(\theta)$, to be the smallest $s \in \{2^0, 2^1, \dots, 2^{\lceil \log_2(p) \rceil}\}$ such that the corresponding *effective support* $\mathcal{S}(\theta) := \{j \in [p] : |\theta^j| \geq \|\theta\|_2 / \sqrt{s \log_2(2p)}\}$ has cardinality at least $s(\theta)$. Thus, the sum of squares of coordinates in the effective support of θ has the same order of magnitude as $\|\theta\|_2^2$, up to logarithmic factors. Moreover, if at most s components of θ are non-zero, then $s(\theta) \leq s$, and the equality is attained when, for example, all non-zero coordinates have the same magnitude.

For $r > 0$ and an online changepoint detection procedure CP characterized by an extended stopping time N , we define

$$g(r; N) := \sup_{z \in \mathbb{N}_0} \mathbb{P}_{z, \theta, \Sigma}(N > z + r). \quad (2)$$

3.1 Coverage Probability and Length of the Confidence Interval

The following theorem shows that the confidence interval constructed in the `ocd_CI` algorithm has the desired coverage level whenever the base online changepoint detection procedure satisfies a patience and detection delay condition.

Theorem 1. *Let $p \geq 2$ and fix $\alpha \in (0, 1)$. Suppose that $\vartheta \geq \beta > 0$. Let CP be an online changepoint procedure characterized by an extended stopping time N satisfying*

$$\mathbb{P}_{z, \theta, \Sigma}(N \leq z) + g(r; N) + 4rp^2 \log_2^2(4p) e^{-r\beta^2/(8s \log_2(2p))} \leq \frac{3}{4}\alpha \quad (3)$$

for some $r \geq 1$. Then, with inputs $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $a \geq 0$, CP, $d_1 = \sqrt{\frac{5r\beta^2}{9s \log_2(2p)}}$, $d_2 = 4d_1^2$ and $\ell \geq 0$, the output confidence interval \mathcal{C} from [Algorithm 1](#) satisfies

$$\mathbb{P}_{z, \theta, \Sigma}(z \in \mathcal{C}) \geq 1 - \alpha.$$

As mentioned in [Section 2](#), our coverage guarantee in [Theorem 1](#) holds even with $\ell = 0$, i.e. with no additional sampling. Condition [\(3\)](#) places a joint assumption on the base

changepoint procedure CP and the parameter r , the latter of which appears in the inputs d_1 and d_2 of Algorithm 1. The first term on the left-hand side of (3) is the false alarm rate of the stopping time N associated with CP. The second term can be interpreted as an upper bound on the probability of the detection delay of N being larger than r , and in addition we also need r to be at least of order s/β^2 up to logarithmic factors for the third term to be sufficiently small. See Section 3.3 for further discussion, where in particular we provide a choice of r for which (3) holds with the `ocd'` base procedure.

We now provide a guarantee on the length of the `ocd_CI` confidence interval.

Theorem 2. *Fix $\alpha \in (0, 1)$. Assume that θ has an effective sparsity of $s := s(\theta) \geq 2$ and that $\vartheta \geq \beta > 0$. Let CP be an online changepoint detection procedure characterized by an extended stopping time N that satisfies (3) for some $r \geq 1$. Then there exists a universal constant $C > 0$ such that, with inputs $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $a = C\sqrt{\log(rp/\alpha)}$, CP, $d_1 = \sqrt{\frac{5r\beta^2}{9s \log_2(2p)}}$, $d_2 = 4d_1^2$, $\ell \geq 80r$, the length L of the output confidence interval \mathcal{C} satisfies*

$$\mathbb{P}_{z, \theta, \Sigma}(L > 8r) \leq \alpha.$$

As mentioned following Theorem 1, we can take r to be the maximum of an appropriate quantile of the detection delay distribution of CP and a quantity that is of order s/β^2 up to logarithmic factors. The main conclusion of Theorem 2 is that, with high probability, the length of the confidence interval is then of this same order r . Whenever the quantile of the detection delay distribution achieves the maximum above — which is the case, up to logarithmic factors, for the `ocd'` base procedure (see Proposition 5) — we can conclude that with high probability, the length of the `ocd_CI` confidence interval is of the same order as this detection delay quantile (which is the best one could hope for). Note that the choices of inputs in Theorem 2 are identical to those in Theorem 1, except that we now ask for order r additional observations after the changepoint declaration.

3.2 Support Recovery

Recall the definition of $\mathcal{S}(\theta)$ from the beginning of this section, and denote $\mathcal{S}_\beta(\theta) := \{j \in [p] : |\theta^j| \geq b_{\min}\}$, where b_{\min} , defined in Algorithm 1, is the smallest positive scale in $\mathcal{B} \cup \mathcal{B}_0$. We will suppress the dependence on θ of both these quantities in this subsection. Theorem 3 below provides a support recovery guarantee for $\hat{\mathcal{S}}$, defined in Algorithm 1. Since neither $\hat{\mathcal{S}}$ nor the anchor coordinate \hat{j} defined in the algorithm depend on d_2 , we omit its specification; the choices of other tuning parameters mimic those in Theorems 1 and 2.

Theorem 3. *Let $p \geq 2$ and fix $\alpha \in (0, 1)$. Suppose that $\vartheta \geq \beta > 0$. Let CP be an online changepoint detection procedure characterized by an extended stopping time N that satisfies (3) for some $r \geq 1$.*

(a) *Then, with inputs $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $a \geq 0$, CP, $d_1 = \sqrt{\frac{5r\beta^2}{9s \log_2(2p)}}$, $\ell \geq 0$, we have*

$$\mathbb{P}_{z, \theta, \Sigma}(\hat{\mathcal{S}} \subseteq \mathcal{S}_\beta) \geq 1 - \alpha.$$

(b) Now assume that θ has an effective sparsity of $s := s(\theta) \geq 2$. Then there exists a universal constant $C > 0$ such that, with inputs $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $a = C\sqrt{\log(rp/\alpha)}$, CP , $d_1 = \sqrt{\frac{5r\beta^2}{9s \log_2(2p)}}$, $\ell \geq 80r$, we have

$$\mathbb{P}_{z, \theta, \Sigma}(\hat{\mathcal{S}} \cup \{\hat{j}\} \supseteq \mathcal{S}) \geq 1 - \alpha.$$

Note that $\mathcal{S} \subseteq \mathcal{S}_\beta \subseteq \{j \in [p] : \theta^j \neq 0\}$. Thus, part (a) of the theorem reveals that with high probability, our support estimate $\hat{\mathcal{S}}$ does not contain any noise coordinates. Part (b) offers a complementary guarantee on the inclusion of all “big” signal coordinates, provided we augment our support estimate with the anchor coordinate \hat{j} . See also the further discussion of this result following Proposition 4 below and in Section 3.3.

We now turn our attention to the optimality of our support recovery algorithm, by establishing a complementary minimax lower bound on the performance of any support estimator. In fact, we can already establish this optimality by restricting the cross-sectional covariance matrix to be the identity matrix. Thus, given $\theta \in \mathbb{R}^p$ and $z \in \mathbb{N}_0$, we write $\mathbb{P}_{z, \theta}$ for a probability measure under which $(X_n)_{n \in \mathbb{N}}$ are independent with $X_n \sim \mathcal{N}_p(\theta \mathbb{1}_{\{n > z\}}, I_p)$. For $r > 0$ and $m \in [p] \cup \{0\}$, write

$$\Theta_{r, m} := \{\theta \in \mathbb{R}^p : |\{j \in [p] : |\theta^j| \leq 1/(8\sqrt{r})\}| \geq m\}.$$

Define \mathcal{T} to be the set of stopping times with respect to the natural filtration $(\mathcal{F}_n)_{n \in \mathbb{N}_0}$, and set

$$\mathcal{T}_{r, m} := \left\{ N \in \mathcal{T} : \sup_{z \in \mathbb{N} \cup \{0\}, \theta \in \Theta_{r, m}} \mathbb{P}_{z, \theta}(N > z + r) \leq \frac{1}{4} \right\}.$$

Write $2^{[p]}$ for the power set of $[p]$, equipped with the symmetric difference metric $d : (A, B) \mapsto |(A \setminus B) \cup (B \setminus A)|$. For any stopping time N , denote

$$\mathcal{J}_N := \{\psi : (\mathbb{R}^p)^\infty \rightarrow 2^{[p]} : \psi \text{ is } \mathcal{F}_N\text{-measurable}\},$$

where we recall that ψ is said to be \mathcal{F}_N -measurable if for any $A \in 2^{[p]}$ and $n \in \mathbb{N}_0$, we have that $\psi^{-1}(A) \cap \{N = n\}$ is \mathcal{F}_n -measurable.

Proposition 4. *For $r > 0$ and $m \geq 15$, we have*

$$\inf_{N \in \mathcal{T}_{r, m}} \inf_{\psi \in \mathcal{J}_N} \sup_{z \in \mathbb{N}_0, \theta \in \Theta_{r, m}} \mathbb{E}_{z, \theta} d(\psi(X_1, X_2, \dots), \text{supp}(\theta)) \geq \frac{m}{32}.$$

This proposition considers any support estimation algorithm obtained from a stopping time in $\mathcal{T}_{r, m}$, and we note that such a competing procedure is even allowed to store all data up to this stopping time, in contrast to our online algorithm. This result can be interpreted as an optimality guarantee for the support recovery property of the `ocd.CI` algorithm presented in Theorem 3(b), provided that the base procedure N belongs to the class $\mathcal{T}_{r, m}$, and that N and r satisfy (3). See Section 3.3 below for further discussion.

3.3 Using ocd' as the Base Procedure

In this subsection, we provide a value of r that suffices for condition (3) to hold when we take our base procedure to be ocd' . For the convenience of the reader, this algorithm is presented as Algorithm 3 in Section 5.3, where we also provide interpretation to the input parameters \tilde{a} , T^{diag} and T^{off} .

Proposition 5. *Fix $\alpha \in (0, 1)$ and $\gamma > 0$. Assume that θ has an effective sparsity of $s := s(\theta) \geq 2$, that $\vartheta \geq \beta > 0$ and that $z \leq 2\alpha\gamma$. Then with inputs $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $\tilde{a} = \sqrt{2 \log\{16p^2\gamma \log_2(2p)\}}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$ and $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$ in the ocd' procedure, there exists a universal constant $C' > 0$ such that for all*

$$r \geq \frac{C' s \log_2(2p) \log\{p\gamma\alpha^{-1}(\beta^{-2} \vee 1)\}}{\beta^2} + 2 =: r_1, \quad (4)$$

the output N satisfies (3).

By combining Proposition 5 with Theorems 1, 2 and 3 respectively, we immediately arrive at the following corollary.

Corollary 6. *Fix $\alpha \in (0, 1)$, $\gamma > 0$. Assume that θ has an effective sparsity of $s := s(\theta) \geq 2$, that $\vartheta \geq \beta > 0$ and that $z \leq 2\alpha\gamma$. Let $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $\tilde{a} = \sqrt{2 \log\{16p^2\gamma \log_2(2p)\}}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$ and $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$ be the inputs of the ocd' procedure. Then the following statements hold:*

(a) *With extra inputs $a \geq 0$, $\text{CP} = \text{ocd}'$, $d_1 = \sqrt{\frac{5r_1\beta^2}{9s \log_2(2p)}}$, $d_2 = 4d_1^2$ and $\ell \geq 0$ for Algorithm 1, the output confidence interval \mathcal{C} and the support estimate $\hat{\mathcal{S}}$ satisfy $\mathbb{P}_{z, \theta, \Sigma}(z \in \mathcal{C}) \geq 1 - \alpha$ and $\mathbb{P}_{z, \theta, \Sigma}(\hat{\mathcal{S}} \subseteq \mathcal{S}_\beta) \geq 1 - \alpha$.*

(b) *There exists a universal constant $C > 0$ such that, with extra inputs $a = C\sqrt{\log(r_1 p / \alpha)}$, $\text{CP} = \text{ocd}'$, $d_1 = \sqrt{\frac{5r_1\beta^2}{9s \log_2(2p)}}$, $d_2 = 4d_1^2$, $\ell \geq 80r_1$ for Algorithm 1, the length L of the output confidence interval \mathcal{C} and the support estimate satisfy $\mathbb{P}_{z, \theta, \Sigma}(L > 8r_1) \leq \alpha$ and $\mathbb{P}_{z, \theta, \Sigma}(\hat{\mathcal{S}} \cup \{\hat{j}\} \supseteq \mathcal{S}) \geq 1 - \alpha$.*

Corollary 6 reveals that, when ocd' is used as the base procedure, the ocd_CI methodology provides guaranteed confidence interval coverage. Moreover, up to poly-logarithmic factors, with an additional $O(1 \vee (s/\beta^2))$ post-change observations, the ocd_CI interval length is of the same order as the average detection delay. In terms of signal recovery, Corollary 6(b) shows that with high probability, ocd_CI with inputs as given in that result selects all signal coordinates whose magnitude exceeds $\vartheta/s^{1/2}$, up to logarithmic factors. Focusing on the case $\beta = \vartheta$ and where s/ϑ^2 is bounded away from zero for simplicity of discussion (though see also Section 3.5 for discussion of the effect of the choice of β), Proposition 10 also reveals that the ocd' base procedure belongs to $\mathcal{T}_{r, m}$ with r of order s/ϑ^2 , up to logarithmic factors, and $m = |\{j : |\theta^j| \leq 1/(8\sqrt{r})\}|$. On the other hand, Proposition 4 shows that any such support estimation algorithm makes on average a non-vanishing fraction of errors in distinguishing between noise coordinates and signals that are below the level $\vartheta/s^{1/2}$, again up to logarithmic factors. In other words, with high probability, the ocd_CI algorithm with base procedure ocd' selects all signals that are strong enough (up to logarithmic factors) to be reliably detected, while at the same time including no noise coordinates (see Corollary 6(a)).

3.4 Relaxing the Gaussianity assumption

It is natural to ask to what extent the theory of Sections 3.1, 3.2 and 3.3 can be generalised beyond the Gaussian setting. The purpose of this subsection, then, is to describe how our earlier results can be modified to handle both sub-Gaussian and sub-exponential data. Recall that we say a random variable Z with $\mathbb{E}Z = 0$ is *sub-Gaussian* with variance parameter $\sigma^2 > 0$ if $\mathbb{E}e^{\lambda Z} \leq e^{\sigma^2 \lambda^2 / 2}$ for all $\lambda \in \mathbb{R}$, and is *sub-exponential* with variance parameter $\sigma^2 > 0$ and rate parameter $A > 0$ if $\mathbb{E}e^{\lambda Z} \leq e^{\sigma^2 \lambda^2 / 2}$ for all $|\lambda| \leq A$.

We first consider the sub-Gaussian setting where $X_1, \dots, X_z, X_{z+1} - \theta, X_{z+2} - \theta, \dots$ are independent, each having sub-Gaussian components with variance parameter 1. Note that this data generating mechanism no longer requires all pre-change observations to be identically distributed, and likewise the post-change observations need not all have the same distribution. We assume that the base changepoint procedure, characterized by an extended stopping time N , satisfies a slightly strengthened version of (3), namely that

$$\mathbb{P}_{z, \theta, \Sigma}(N \leq z) + g(r; N) + 100rp^2 \log_2^3(4p)(p\beta^{-2} \vee 1)e^{-r\beta^2 / (8s \log_2(2p))} \leq \frac{3}{4}\alpha. \quad (5)$$

for some $r \geq 1$. Under (5), Theorems 1, 2 and 3 hold with the same choices of input parameters. Moreover, the `ocd'` base procedure satisfies the conclusion of Proposition 5, i.e. there exists a universal constant $C' > 0$ such that (5) holds for $r \geq r_1 = r_1(C')$ in (4), provided that we use the modified input $\tilde{a} = \sqrt{2 \log\{32p^2\gamma \log_2(2p)\}}$.

Generalising these ideas further, now consider the model where $X_1, \dots, X_z, X_{z+1} - \theta, X_{z+2} - \theta, \dots$ are independent, each having sub-exponential components with variance parameter 1 and rate parameter $A > 0$. In this setting, provided the base procedure satisfies (5) for some $r \geq 1$ and $\vartheta \leq \sqrt{2A^2 \log_2(2p)}$, Theorems 1, 2 and 3 hold when we redefine $a := C \max\{\sqrt{\log(rp/\alpha)}, \frac{1}{A} \log(rp/\alpha)\}$ and $d_1 := \max\left\{\sqrt{\frac{5r\beta^2}{9s \log_2(2p)}}, \frac{5r\beta^2}{9As \log_2(2p)}\right\}$. Furthermore, with the modified input $\tilde{a} = \sqrt{2 \log\{32p^2\gamma \log_2(2p)\}} \vee \frac{2}{A} \log\{32p^2\gamma \log_2(2p)\}$, the `ocd'` base procedure satisfies the conclusion of Proposition 5 for

$$r \geq \frac{C' s \log_2(2p) \log^2\{p\gamma\alpha^{-1}(\beta^{-2} \vee 1)\} \max\{1, A^{-2} \log(p\gamma)\}}{\beta^2} + 2,$$

where $C' > 0$ is a universal constant.

The claims made in the previous two paragraphs are justified in Section 5.4. These results confirm the flexible scope of the `ocd_CI` methodology beyond the original Gaussian setting, at least as far as sub-exponential tails are concerned. Where data may exhibit heavier tails than this, clipping (truncation) and quantile transformations may represent viable ways to proceed, though further research is needed to confirm the theoretical validity of such approaches.

3.5 Confidence interval construction with unknown signal size

In some settings, an experienced practitioner may have a reasonable idea of the magnitude ϑ of the ℓ_2 -norm of the vector of mean change that would be of interest to them, and this

would facilitate a choice of a lower bound β for ϑ in Algorithm 1. However, it is also worth considering the effect of this choice, and the extent to which its impact can be mitigated.

We first remark that the coverage probability guarantee for the `ocd.CI` interval in Corollary 6 remains valid for any (arbitrarily loose) lower bound β on ϑ . The only issue is in terms of power: if β is chosen to be too small, then both the average detection delay and the high-probability bound on the length of the confidence interval may be inflated. In the remainder of this subsection, then, we describe a simple modification to Algorithm 1 that permits a loose lower bound β to be employed that retains coverage validity with only a logarithmic effect on the high-probability bound on the length of the confidence interval. The only other price we pay is that the computational cost increases as β decreases, as we describe below.

Our change to Algorithm 1 is as follows: we replace the definition of b_{\min} by setting

$$b_{\min} = \frac{\beta}{\sqrt{2^{\lceil \log_2(2p) \rceil} \log_2(2p)}} \wedge \frac{1}{\sqrt{2}},$$

set $M = \lceil 2 \log_2(1/b_{\min}) \rceil$ and define $\mathcal{B} = \{\pm 2^{m/2} b_{\min} : m \in [M]\}$ in both the `ocd'` base procedure and in Algorithm 1. The rest of algorithm remains as previously stated. Thus, if we choose a conservative (very small) β , then the effect of the modification is to increase the number of scales on which we search for a change, so that the largest element of \mathcal{B} is of order 1. In order to state our theoretical results for this modified algorithm, first define $b_{\text{opt}} := \max\left\{b \in \mathcal{B} \cap (0, \infty) : b \leq \frac{\vartheta}{\sqrt{s \log_2(2p)}}\right\}$, which satisfies $b_{\text{opt}} \geq \frac{\vartheta}{\sqrt{2s \log_2(2p)}} \wedge 1$. Under the same assumptions as in Proposition 5, and modifying the inputs to $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $\tilde{a} = \sqrt{2 \log(16p^2 \gamma M)}$, $T^{\text{diag}} = \log(16p\gamma(M+1))$ and $T^{\text{off}} = 8 \log(16p\gamma M)$, it can be shown using very similar arguments to those in the proof of Proposition 5 that there exists a universal constant $C' > 0$ such that with $r \geq C' b_{\text{opt}}^{-2} \log(p\gamma\alpha^{-1}(\beta^{-2} \vee 1)) =: r_1$, the output N satisfies

$$\mathbb{P}_{z, \theta, \Sigma}(N \leq z) + g(r; N) + 4rp^2(M+1)^2 e^{-rb_{\text{opt}}^2/8} \leq \frac{3}{4}\alpha.$$

With this in place, we can derive Corollary 7 below, which is the analog of Corollary 6 for our modified algorithm.

Corollary 7. *Fix $\alpha \in (0, 1)$, $\gamma > 0$. Assume that θ has an effective sparsity of $s := s(\theta) \geq 2$, that $\vartheta \geq \beta > 0$ and that $z \leq 2\alpha\gamma$. Let $(X_t)_{t \in \mathbb{N}}$, $\beta > 0$, $\tilde{a} = \sqrt{2 \log(16p^2 \gamma M)}$, $T^{\text{diag}} = \log(16p\gamma(M+1))$ and $T^{\text{off}} = 8 \log(16p\gamma M)$ be the inputs of the modified `ocd'` procedure. Let $d_1 = \sqrt{5C' \log(p\gamma\alpha^{-1}(\beta^{-2} \vee 1))}/9$ and $d_2 = 4d_1^2$. Then the following statements hold:*

(a) *With extra inputs $\text{CP} = \text{ocd}'$, $a \geq 0$, and $\ell \geq 0$ for the modified Algorithm 1, the output confidence interval \mathcal{C} and the support estimate $\hat{\mathcal{S}}$ satisfy $\mathbb{P}_{z, \theta, \Sigma}(z \in \mathcal{C}) \geq 1 - \alpha$ and $\mathbb{P}_{z, \theta, \Sigma}(\hat{\mathcal{S}} \subseteq \mathcal{S}_\beta) \geq 1 - \alpha$.*

(b) *There exists a universal constant $C > 0$ such that, with extra inputs $\text{CP} = \text{modified ocd}'$, $a = C \sqrt{\log(p\gamma\alpha^{-1}(\beta^{-2} \vee 1))}$, and $\ell \geq 80r_1$ for the modified Algorithm 1, the length L of the output confidence interval \mathcal{C} and the support estimate satisfy*

$$\mathbb{P}_{z, \theta, \Sigma}\left(L > 8C' \max\left\{\frac{2s \log_2(2p)}{\vartheta^2}, 1\right\} \log(p\gamma\alpha^{-1}(\beta^{-2} \vee 1))\right) \leq \alpha \quad (6)$$

and $\mathbb{P}_{z,\theta,\Sigma}(\hat{\mathcal{S}} \cup \{\hat{j}\} \supseteq \mathcal{S}) \geq 1 - \alpha$.

The main difference between Corollaries 7 and 6 concerns the high-probability guarantees on the length of the confidence interval. Ignoring logarithmic factors, with high probability the length of the confidence interval in the modified algorithm is at most of order $(s/\vartheta^2) \vee 1$, whereas for the original algorithm it was of order $(s/\beta^2) \vee 1$. Thus, the modified algorithm has the significant advantage of enabling a conservative choice of β with only a logarithmic effect on the length guarantee relative to an oracle procedure with knowledge of $\|\theta\|_2$. The computational complexity per new observation and the storage requirements of this modified algorithm are $O\left(p^2(\log(ep) + \log(1/\beta))\right)$, so the order of magnitude is increased relative to the original algorithm only in an asymptotic regime where β is small by comparison with $1/p^K$ for every $K > 0$. Moreover, the modified algorithm still does not require storage of historical data and the computational time per new observation after observing n observations does not increase with n . Nevertheless, since the computational complexity now depends on β , the modified algorithm does not strictly satisfy our definition of an online algorithm given in introduction.

4 Numerical studies

In this section, we study the empirical performance of the `ocd_CI` algorithm. Throughout this section, by default, the `ocd_CI` algorithm is used in conjunction with the recommended base online changepoint detection procedure $\text{CP} = \text{ocd}$.

4.1 Tuning Parameters

Chen et al. (2022) found that the theoretical choices of thresholds T^{diag} and T^{off} for the `ocd` procedure were a little conservative, and therefore recommended determining these thresholds via Monte Carlo simulation; we replicate the method for choosing these thresholds described in their Section 4.1. Likewise, as in Chen et al. (2022), we take $a = \tilde{a} = \sqrt{2 \log p}$ in our simulations.

For d_1 and d_2 , as suggested by our theory, we take $d_2 = 4d_1^2$, and take d_1 to be of the form $d_1 = c\sqrt{\log(p/\alpha)}$. Here, we tune the parameter $c > 0$ through Monte Carlo simulation, as we now describe. We considered the parameter settings $p \in \{100, 500\}$, $s \in \{2, \lfloor \sqrt{p} \rfloor, p\}$, $\vartheta \in \{2, 1, 1/2\}$, $\Sigma = I_p$, $\alpha = 0.05$, $\beta \in \{2\vartheta, \vartheta, \vartheta/2\}$, $\gamma = 30000$ and $z = 500$. Then, with θ generated as ϑU , where U is uniformly distributed on the union of all s -sparse unit spheres in \mathbb{R}^p (independent of our data), we studied the coverage probabilities, estimated over 2000 repetitions as c varies, of the `ocd_CI` confidence interval for data generated according to the Gaussian model defined at the beginning of Section 3. Figure 1 displays a subset of the results (the omitted curves were qualitatively similar). On this basis, we recommend $c = 0.5$ as a safe choice across a wide range of data generating mechanisms, and we used this value of c throughout our confidence interval simulations.

The previous two paragraphs, in combination with Algorithms 1 and 2, provide the practical implementation of the `ocd_CI` algorithm that we use in our numerical studies and that

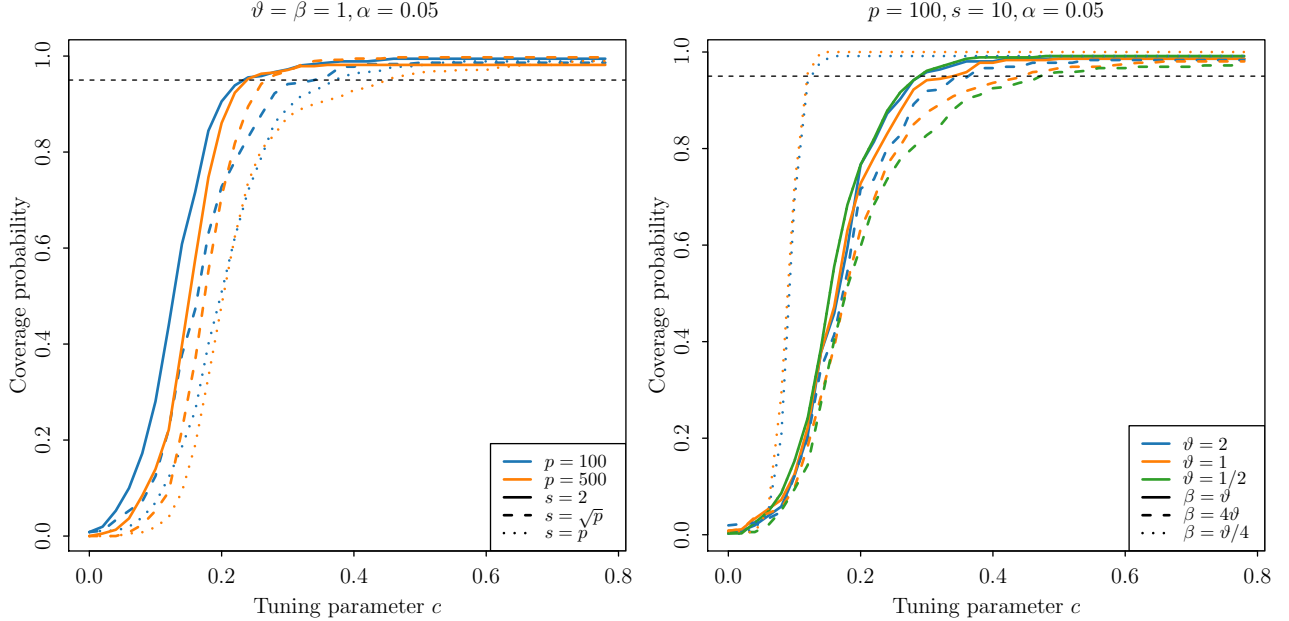


Figure 1: Coverage probabilities of the `ocd_CI` confidence interval as the parameter c , involved in the choice of tuning parameter d_1 , varies.

we recommend for practitioners. The only quantity that remains for the practitioner to input (other than the data) is β , which represents a lower bound on the Euclidean norm of the vector of mean change. Fortunately, this description makes β easily interpretable by practitioners. In cases where an informed default choice is not available, practitioners can make a conservative (very small) choice and use an increased grid of scales to with only a small inflation in the confidence interval length guarantee and computational cost; see Section 3.5.

4.2 Coverage Probability and Interval Length

In Table 1, we present the detection delay of the `ocd` procedure, as well as the coverage probabilities and average confidence interval lengths of the `ocd_CI` procedure, all estimated over 2000 repetitions, with the same set of parameter choices and data generating mechanism as in Section 4.1. From this table, we see that the coverage probabilities are at least at the nominal level (up to Monte Carlo error) across all settings considered. Underspecification of β means that the grid of scales that can be chosen for indices in $\hat{\mathcal{S}}$ is shifted downwards, and therefore increases the probability that \tilde{b}^j will underestimate θ^j for $j \in \hat{\mathcal{S}}$. In turn, this leads to a slight conservativeness for the coverage probability (and corresponding increased average confidence interval length). On the other hand, overspecification of β yields a shorter interval on average, though these were nevertheless able to retain the nominal coverage in all cases considered.

Another interesting feature of Table 1 is to compare the average confidence interval lengths with the corresponding average detection delays. Corollary 6(b), as well as Chen et al. (2022, Theorem 4), indicates that both of these quantities are of order $(s/\beta^2) \vee 1$, up to polylogarithmic factors in p and γ , but of course whenever the confidence interval includes

Table 1: Estimated coverage and average length of the `ocd_CI` confidence interval and average detection delay over 2000 repetitions, with standard errors in brackets. Other parameters: $\gamma = 30000$, $z = 1000$, $\Sigma = I_p$, $\alpha = 0.05$, $a = \tilde{a} = \sqrt{2\log p}$, $c = 0.5$, $d_1 = c\sqrt{\log(p/\alpha)}$, $d_2 = 4d_1^2$. For comparison, we also present the corresponding estimated coverage probabilities and average lengths of the procedure based on [Kaul et al. \(2021\)](#), as described in Section 4.2.

p	s	ϑ	β	ocd_CI			Kaul_et_al	
				Delay	Coverage (%)	CI Length	Coverage (%)	CI Length
100	2	2	4	9.8 _(0.1)	96.2 _(0.4)	20.1 _(0.7)	83.2 _(0.8)	732.7 _(9.6)
100	2	2	2	12.6 _(0.1)	97.0 _(0.4)	33.7 _(0.7)	82.5 _(0.8)	474.9 _(11.0)
100	2	2	1	14.1 _(0.1)	97.9 _(0.3)	80.8 _(1.0)	83.5 _(0.8)	341.4 _(10.4)
100	2	1	2	34.2 _(0.3)	95.8 _(0.4)	66.1 _(1.0)	76.6 _(0.9)	399.3 _(10.5)
100	2	1	1	44.2 _(0.3)	97.5 _(0.4)	122.0 _(1.4)	80.5 _(0.9)	123.8 _(6.5)
100	2	1	0.5	52.0 _(0.4)	97.4 _(0.4)	309.1 _(2.0)	81.5 _(0.9)	90.8 _(5.4)
100	10	2	4	14.7 _(0.1)	96.0 _(0.4)	32.5 _(0.8)	80.2 _(0.9)	636.4 _(10.5)
100	10	2	2	15.7 _(0.1)	97.4 _(0.4)	38.4 _(0.8)	77.4 _(0.9)	537.9 _(10.9)
100	10	2	1	15.9 _(0.1)	97.0 _(0.4)	80.2 _(1.1)	80.8 _(0.9)	542.4 _(10.9)
100	10	1	2	52.6 _(0.5)	96.2 _(0.4)	114.0 _(1.5)	75.8 ₍₁₎	342.4 _(10.1)
100	10	1	1	56.9 _(0.4)	97.1 _(0.4)	142.5 _(1.8)	73.9 ₍₁₎	262.6 _(9.1)
100	10	1	0.5	60.2 _(0.4)	98.2 _(0.3)	301.1 _(1.6)	75.9 ₍₁₎	248.3 _(8.9)
100	100	2	4	27.2 _(0.2)	96.1 _(0.4)	77.6 _(0.9)	68.2 _(1.0)	533.9 _(10.7)
100	100	2	2	27.7 _(0.2)	96.0 _(0.4)	81.8 _(1.0)	71.3 _(1.0)	537.7 _(10.8)
100	100	2	1	28.2 _(0.2)	97.5 _(0.3)	99.4 _(1.3)	71.8 _(1.0)	556.0 _(10.7)
100	100	1	2	100.7 _(0.8)	94.7 _(0.5)	292.8 _(3.5)	87.7 _(0.7)	850.5 _(9.5)
100	100	1	1	100.5 _(0.9)	96.3 _(0.4)	296.0 _(3.4)	88.0 _(0.7)	863.7 _(9.3)
100	100	1	0.5	103.2 _(0.8)	97.3 _(0.4)	365.9 _(2.8)	89.3 _(0.7)	876.8 _(9.1)
500	2	2	4	11.3 _(0.1)	97.2 _(0.4)	23.1 _(0.7)	92.0 _(0.6)	958.7 _(4.3)
500	2	2	2	15.8 _(0.1)	97.7 _(0.3)	45.2 _(0.9)	83.3 _(0.8)	806.4 _(8.7)
500	2	2	1	17.7 _(0.1)	97.5 _(0.4)	117.3 _(1.0)	79.9 _(0.9)	624.9 _(10.7)
500	2	1	2	41.5 _(0.3)	97.3 _(0.4)	81.8 _(1.2)	80.0 _(0.9)	774.9 _(9.4)
500	2	1	1	55.0 _(0.4)	96.8 _(0.4)	168.9 _(1.5)	73.0 ₍₁₎	275.9 _(9.4)
500	2	1	0.5	64.6 _(0.5)	98.1 _(0.3)	445.0 _(1.7)	75.6 ₍₁₎	186.1 _(8.0)
500	22	2	4	23.6 _(0.2)	96.3 _(0.4)	55.4 _(1.0)	87.0 _(0.8)	884.9 _(7.3)
500	22	2	2	25.0 _(0.2)	97.0 _(0.4)	60.3 _(0.8)	85.5 _(0.8)	864.2 _(7.8)
500	22	2	1	25.5 _(0.2)	98.1 _(0.3)	119.7 _(0.8)	83.6 _(0.8)	823.0 _(8.6)
500	22	1	2	88.1 _(0.7)	97.0 _(0.4)	203.5 _(2.1)	77.2 _(0.9)	645.0 _(11.0)
500	22	1	1	91.9 _(0.6)	97.8 _(0.3)	229.7 _(2.2)	76.2 ₍₁₎	562.8 _(11.1)
500	22	1	0.5	94.9 _(0.6)	98.3 _(0.3)	462.8 _(1.4)	75.5 ₍₁₎	538.3 _(11.2)
500	500	2	4	79.8 _(0.6)	95.0 _(0.5)	238.9 _(2.7)	88.5 _(0.7)	913.0 _(8.0)
500	500	2	2	80.3 _(0.6)	95.8 _(0.4)	245.7 _(2.6)	90.3 _(0.7)	928.8 _(7.7)
500	500	2	1	80.9 _(0.6)	97.5 _(0.4)	250.2 _(2.5)	90.6 _(0.7)	928.3 _(7.7)
500	500	1	2	290.5 _(2.3)	94.5 _(0.5)	819.7 _(7.9)	95.2 _(0.5)	1189.4 _(7.3)
500	500	1	1	291.4 _(2.3)	95.2 _(0.5)	831.1 _(7.5)	94.3 _(0.5)	1204.9 _(7.0)
500	500	1	0.5	297.3 _(2.3)	98.1 _(0.3)	875.0 _(6.7)	94.6 _(0.5)	1207.4 _(6.8)

the changepoint, its length must be at least as long as the detection delay. Nevertheless, in most settings, it is only 2 to 3 times longer on average, and in all cases considered was less than 7 times longer on average. Moreover, we can also observe that the confidence interval length increases with s and decreases with β , as anticipated by our theory.

For comparison, we also present the corresponding coverage probabilities and average lengths of confidence intervals obtained using an offline procedure as described in the introduction. More precisely, after the `ocd` algorithm has declared a change, we treat the data up to the stopping time as an offline dataset, and apply the `inspect` algorithm (Wang and Samworth, 2018), followed by the one-step refinement of Kaul et al. (2021), to construct an estimate, \hat{z}^{KFJS} , of the changepoint location. As recommended by Kaul et al. (2021), we obtain an estimator $\hat{\vartheta}^{\text{KFJS}}$ of ϑ using the ℓ_2 -norm of the soft-thresholded difference in empirical mean vectors before and after \hat{z}^{KFJS} , with the soft-thresholding parameter chosen via the Bayesian Information Criterion. The final confidence interval is then of the form $[\hat{z}^{\text{KFJS}} - q^{\alpha/2}/(\hat{\vartheta}^{\text{KFJS}})^2, \hat{z}^{\text{KFJS}} + q^{\alpha/2}/(\hat{\vartheta}^{\text{KFJS}})^2]$, where $q^{\alpha/2}$ is the $1 - \alpha/2$ quantile of the distribution of the (almost surely unique) maximizer of a two-sided Brownian motion with a triangular drift as given by Kaul et al. (2021, Theorem 3.1). In particular, we have $q^{0.025} = 11.03$. The last two columns of Table 1 reveal that both the coverage probabilities and confidence interval lengths from this procedure are disappointing and not competitive with those of the `ocd_CI` algorithm. There are two main reasons for this: first, the nature of the online problem means that the changepoint is often located near the right-hand end of the dataset up to the stopping time; on the other hand, the theoretical guarantees of Kaul et al. (2021) are obtained under an asymptotic setting where the fraction of data either side of the change is bounded away from zero. Thus, the estimated changepoint from the one-step refinement is often quite poor. Moreover, the estimated magnitude of change, $\hat{\vartheta}^{\text{KFJS}}$, is often a significant underestimate of ϑ due to the soft-thresholding operation, and this can lead to substantially inflated confidence interval lengths. We emphasize that the Kaul et al. (2021) procedure was not designed for use in this online setting, but we nevertheless present these results to illustrate the fact that the naive application of offline methods in sequential problems may fail badly.

While Table 1 covers the most basic setting for our methodology, our theory in Section 3 applies equally well to data with spatial dependence across different coordinates. To assess whether this theory carries over to empirical performance, Table 3 in the supplement (Section 5.5) presents corresponding coverage probabilities and lengths for the `ocd_CI` procedure with the cross-sectional covariance matrix $\Sigma = (\Sigma_{jk})_{j,k \in [p]}$ taken to be Toeplitz with parameter $\rho \in \{0.5, 0.75\}$; in other words, $\Sigma_{jk} = \rho^{|j-k|}$. The results are again encouraging: the coverage remains perfectly satisfactory in all settings considered, and moreover, the lengths of the confidence intervals are very similar to those in Table 1.

4.3 Support Recovery

We now turn our attention to the empirical support recovery properties of the quantity $\hat{\mathcal{S}}$ (in combination with the anchor coordinate \hat{j}) computed in the `ocd_CI` algorithm. In Table 2, we present the probabilities, estimated over 500 repetitions, that $\hat{\mathcal{S}} \subseteq \mathcal{S}_\beta$ and that $\hat{\mathcal{S}} \cup \{\hat{j}\} \supseteq \mathcal{S}$ for $p = 100$, $s \in \{5, 50\}$, $\vartheta \in \{1, 2\}$, $\Sigma = I_p$, and for three different signal shapes: in the uniform, inverse square root and harmonic cases, we took $\theta \propto (\mathbb{1}_{\{j \in [s]\}})_{j \in [p]}$, $\theta \propto$

Table 2: Estimated support recovery probabilities (with standard errors in brackets). Other parameters: $p = 100$, $\Sigma = I_p$, $\alpha = 0.05$, $a = \tilde{a} = \sqrt{2 \log p}$, $d_1 = \sqrt{2 \log(p/\alpha)}$, $\beta = \vartheta$, and with an additional $\ell = \lceil 2s \log_2(2p) \log(p) \beta^{-2} \rceil$ post-declaration observations.

s	ϑ	Signal Shape	$\hat{\mathcal{S}} \subseteq \mathcal{S}_\beta$ (%)	$\hat{\mathcal{S}} \cup \{\hat{j}\} \supseteq \mathcal{S}$ (%)
5	2	uniform	99.8 _(0.2)	97.6 _(0.7)
5	1	uniform	100.0 _(0.0)	97.6 _(0.7)
50	2	uniform	100.0 _(0.0)	95.6 _(0.9)
50	1	uniform	100.0 _(0.0)	97.8 _(0.7)
5	2	inv sqrt	99.6 _(0.3)	96.6 _(0.8)
5	1	inv sqrt	100.0 _(0.0)	98.8 _(0.5)
50	2	inv sqrt	100.0 _(0.0)	99.8 _(0.2)
50	1	inv sqrt	100.0 _(0.0)	100.0 _(0.0)
5	2	harmonic	100.0 _(0.0)	97.6 _(0.7)
5	1	harmonic	99.6 _(0.3)	97.8 _(0.7)
50	2	harmonic	100.0 _(0.0)	99.4 _(0.3)
50	1	harmonic	100.0 _(0.0)	100.0 _(0.0)

$(j^{-1/2} \mathbb{1}_{\{j \in [s]\}})_{j \in [p]}$ and $\theta \propto (j^{-1} \mathbb{1}_{\{j \in [s]\}})_{j \in [p]}$ respectively. As inputs to the algorithm, we set $a = \tilde{a} = \sqrt{2 \log p}$, $\alpha = 0.05$, $d_1 = \sqrt{2 \log(p/\alpha)}$, $\beta = \vartheta$, and, motivated by Corollary 6, took an additional $\ell = \lceil 2s \beta^{-2} \log_2(2p) \log p \rceil$ post-declaration observations in constructing the support estimates. The results reported in Table 2 provide empirical confirmation of the support recovery properties claimed in Corollary 6.

Finally in this section, we consider the extent to which the additional ℓ observations are necessary in practice to provide satisfactory support recovery. In the left panel of Figure 2, we plot Receiver Operating Characteristic (ROC) curves to study the estimated support recovery probabilities with $\ell = 0$ as a function of the input parameter d_1 , which can be thought of as controlling the trade-off between $\mathbb{P}(\hat{\mathcal{S}} \cup \{\hat{j}\} \supseteq \mathcal{S})$ and $\mathbb{P}(\hat{\mathcal{S}} \subseteq \mathcal{S}_\beta)$. The fact that the triangles in this plot are all to the left of the dotted vertical line confirms the theoretical guarantee provided in Corollary 6(a), which holds with $d_1 = \sqrt{2 \log(p/\alpha)}$, and even with $\ell = 0$); the less conservative choice $d_1 = \sqrt{2 \log p}$, which roughly corresponds to an average of one noise coordinate included in $\hat{\mathcal{S}}$, allows us to capture a larger proportion of the signal. From this panel, we also see that additional sampling is needed to ensure that, with high probability, we recover all of the true signals. This is unsurprising: for instance, with a uniform signal shape and $s = 50$, it is very unlikely that all 50 signal coordinates will have accumulated such similar levels of evidence to appear in $\hat{\mathcal{S}} \cup \{\hat{j}\}$ by the time of declaration. The right panel confirms that, with an inverse square root signal shape, the probability that we capture each signal increases with the signal magnitude, and that even small signals tend to be selected with higher probability than noise coordinates.

4.4 US Covid-19 Data Example

In this section, we apply `ocd_CI` to a dataset of weekly deaths in the United States between January 2017 and June 2020 (available at: <https://www.cdc.gov/nchs/nvss/vsrr/>




Figure 2: Support recovery properties of `ocd_CI`. In the left panel, we plot ROC curves for three different signal shapes and for sparsity levels $s \in \{5, 50\}$. The triangles and circles correspond to points on the curves with $d_1 = \sqrt{2 \log(p/\alpha)}$ (with $\alpha = 0.05$), and $d_1 = \sqrt{2 \log p}$ respectively. The dotted vertical line corresponds to $\mathbb{P}(\hat{\mathcal{S}} \subseteq \mathcal{S}_\beta) = 1 - \alpha$. In the right panel, we plot the proportion of 500 repetitions for which each coordinate belongs to $\hat{\mathcal{S}} \cup \{\hat{j}\}$ with $d_1 = \sqrt{2 \log p}$; here, the $s = 20$ signals have an inverse square root shape, and are plotted in red; noise coordinates are plotted in black. Other parameters for both panels: $p = 100$, $\Sigma = I_p$, $\beta = \vartheta = 2$, $\ell = 0$, $a = \tilde{a} = \sqrt{2 \log p}$.

covid19/excess_deaths.htm). The data up to 29 June 2019 are treated as our training data. An obvious discrepancy between underlying dynamics of these weekly deaths and the conditions assumed in our theory in Section 3 is temporal dependence, particularly induced by seasonal and weather effects. Although we can never hope to remove this dependence entirely, we seek to mitigate its impact by pre-processing the data as follows: for each of the 50 states, as well as Washington, D.C. ($p = 51$), we first estimate the “seasonal death curve”, i.e. the mean death numbers for each day of the year, for each state. The seasonal death curve is estimated by first splitting the weekly death numbers evenly across the seven relevant days, and then estimating the average number of deaths on each day of the year from these derived daily death numbers using a Gaussian kernel with a bandwidth of 20 days. As the death numbers follow an approximate Poisson distribution, we apply a square-root transformation to stabilize the variance; more precisely, the transformed weekly excess deaths are computed as the difference of the square roots of the weekly deaths and the predicted weekly deaths from the seasonal death curve. Finally, we standardize the transformed weekly excess deaths using the mean and standard deviation of the transformed data over the training period. The standardized, transformed data are plotted in Figure 3 for 12 states.




Figure 3: Standardized, transformed weekly excess death data from 12 states (including Washington, D.C.). The monitoring period starts from 30 June 2019 (dashed line). The data from the states in the support estimate are shown in red. The confidence interval [8 March 2020, 28 March 2020] is shown in the light blue shaded region.

When applying `ocd_CI` to these data, we take $a = \tilde{a} = \sqrt{2 \log p}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$, $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, $d_1 = 0.5\sqrt{\log(p/\alpha)}$ and $d_2 = 4d_1^2$, with $\alpha = 0.05$, $\beta = 50$ and $\gamma = 1000$. On the monitoring data (from 30 June 2019), the `ocd_CI` algorithm declares a change on the week ending 28 March 2020, and provides a confidence interval from the week ending 21 March 2020 to the week ending 28 March 2020. This coincides with the beginning of the first wave of Covid-19 deaths in the United States. The algorithm also identifies New York, New Jersey, Connecticut, Michigan and Louisiana as the estimated support of the change. Interestingly, if we run the `ocd_CI` procedure from the beginning of the training data period (while still standardizing as before, due to the lack of available data prior to 2017),

it identifies a subtler change on the week ending 6 January 2018, with a confidence interval of [17 December 2017, 6 January 2018]. This corresponds to a bad influenza season at the end of 2017 (see, <https://www.cdc.gov/flu/about/season/flu-season-2017-2018.htm>). Despite the natural interpretation of these findings, we recognize that the model in Section 3 under which we proved our theoretical results cannot capture the full complexity of the temporal dependence in this dataset even after our pre-processing transformations. A complete theoretical analysis of the performance of `ocd_CI` in time-dependent settings is challenging and beyond the scope of the current work; in practical applications, we advise careful modeling of this dependence to facilitate the construction of appropriate residuals for which the main effects of this dependence have been removed.

5 Supplementary material

In this section, we provide proofs of our main results (Section 5.1), auxiliary results and their proofs (Section 5.2), pseudocode for the `ocd` and `ocd'` base procedures (Section 5.3), extensions of our results to sub-Gaussian and sub-exponential settings (Section 5.4) and additional simulation results (Section 5.5). Throughout this section, we use \mathbb{P} instead of $\mathbb{P}_{z,\theta,\Sigma}$ when it is clear from the context.

5.1 Proofs of main results

Proof of Theorem 1. Fix $r \geq 1$ that satisfies the assumption (3) in the theorem, $n > z, j \in [p]$, $b \in \mathcal{B}$ and $j' \in [p] \setminus \{j\}$. We assume, without loss of generality, that $\theta^{j'} \geq 0$. The case $\theta^{j'} < 0$ can be analyzed similarly. Recall that b_{\min} , defined in Algorithm 1, is the smallest positive scale in $\mathcal{B} \cup \mathcal{B}_0$, and write $b_{\text{aux}}^{j'} := \min\{b \in (\mathcal{B} \cup \mathcal{B}_0) \cap (0, \infty) : b \geq \theta^{j'}\}$. Then we have $A_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'} \mid t_{n,b}^j \sim \mathcal{N}(\theta^{j'} \min\{n + \ell - z, t_{n,b}^j + \ell\}, t_{n,b}^j + \ell)$. Thus, recalling the definition of $\hat{\mathcal{S}}$ and $\tilde{b}^{j'}$ from Algorithm 1, we have

$$\begin{aligned}
& \mathbb{P}(\{j' \in \hat{\mathcal{S}}\} \cap \{\tilde{b}^{j'} \notin (0, \theta^{j'})\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}) \\
&= \mathbb{E}\left\{\mathbb{P}\left(\{j' \in \hat{\mathcal{S}}\} \cap \{\tilde{b}^{j'} \notin (-b_{\min}, b_{\text{aux}}^{j'})\} \cap \{N = n, \hat{j} = j, \hat{b} = b\} \mid t_{n,b}^j\right)\right\} \\
&\leq \mathbb{E}\left\{\mathbb{P}\left(A_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'} \geq b_{\text{aux}}^{j'}(t_{n,b}^j + \ell) + d_1(t_{n,b}^j + \ell)^{1/2} \mid t_{n,b}^j\right)\right\} \\
&\quad + \mathbb{E}\left\{\mathbb{P}\left(A_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'} \leq -b_{\min}(t_{n,b}^j + \ell) - d_1(t_{n,b}^j + \ell)^{1/2} \mid t_{n,b}^j\right)\right\} \\
&\leq \mathbb{E}\{\bar{\Phi}((b_{\text{aux}}^{j'} - \theta^{j'})(t_{n,b}^j + \ell)^{1/2} + d_1)\} + \mathbb{E}\{\bar{\Phi}((b_{\min} + \theta^{j'})(t_{n,b}^j + \ell)^{1/2} + d_1)\} \\
&\leq 2\bar{\Phi}(d_1). \tag{7}
\end{aligned}$$

Moreover, by a similar argument to (28) in the proof of Proposition 8, for $b \in (0, \theta^{j'})$, we have

$$\mathbb{P}(n - t_{n,b}^{j'} - d_2/b^2 > z) \leq 2\bar{\Phi}\left(\frac{\sqrt{d_2}}{b}(\theta^{j'} - b/2)\right) \leq 2\bar{\Phi}(\sqrt{d_2}/2). \tag{8}$$

Combining (7) and (8), we have

$$\begin{aligned}
& \mathbb{P}\left(\{j' \in \hat{\mathcal{S}}\} \cap \{n - t_{n, \tilde{b}^{j'}}^{j'} - d_2/(\tilde{b}^{j'})^2 > z\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}\right) \\
& \leq \mathbb{P}\left(\{j' \in \hat{\mathcal{S}}\} \cap \{\tilde{b}^{j'} \notin (0, \theta^{j'})\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}\right) + \sum_{b \in (\mathcal{B} \cup \mathcal{B}_0) \cap (0, \theta^{j'})} 2\bar{\Phi}(\sqrt{d_2}/2) \\
& \leq 2\bar{\Phi}(d_1) + 2\log_2(4p)\bar{\Phi}(\sqrt{d_2}/2) \leq 2\log_2(4p)e^{-5r\beta^2/(18s\log_2(2p))},
\end{aligned}$$

where the last inequality follows from the choice of d_1 and d_2 in the statement of the theorem and the standard Gaussian tail bound used at the end of the proof of Lemma 9. By a union bound and (3), we have

$$\begin{aligned}
\mathbb{P}(z \notin \mathcal{C}) & \leq \mathbb{P}(N \leq z) + \mathbb{P}\left(N - \min_{j \in \hat{\mathcal{S}}} \left\{t_{N, \tilde{b}^j}^j + \frac{d_2}{(\tilde{b}^j)^2}\right\} > z\right) \\
& \leq \mathbb{P}(N \leq z) + \mathbb{P}(N > z + r) \\
& \quad + \sum_{n=z+1}^{z+[r]} \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{j'=1}^p \mathbb{P}\left(\{j' \in \hat{\mathcal{S}}\} \cap \left\{n - t_{n, \tilde{b}^{j'}}^{j'} - \frac{d_2}{(\tilde{b}^{j'})^2} > z\right\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}\right) \\
& \leq \mathbb{P}(N \leq z) + g(r; N) + 4rp^2 \log_2^2(4p)e^{-5r\beta^2/(18s\log_2(2p))} \leq \alpha,
\end{aligned}$$

as required. \square

Proof of Theorem 2. Fix $r \geq 1$ that satisfies the assumption (3). Denote $\ell_0 := 80r$. Then $\ell \geq \ell_0$. Since the output of Algorithm 1 remains unchanged if we replace $(X_t^j : t \in \mathbb{N})$ by $(-X_t^j : t \in \mathbb{N})$ for any fixed j , we may assume without loss of generality that $\theta^1 \geq \theta^2 \geq \vartheta/\sqrt{s\log_2(2p)}$. For $j \in \{1, 2\}$, we denote $b^j := \max\{b \in \mathcal{B} \cup \mathcal{B}_0 : b \leq \theta^j\}$. Since $\vartheta \geq \beta$ and $s \leq 2^{\lceil \log_2(p) \rceil}$, we have $b^1 \geq b^2 \geq \beta/\sqrt{s\log_2(2p)} =: b_*$. Denote

$$u := \frac{\ell_0\beta^2}{80s\log_2(2p)} = rb_*^2 \quad \text{and} \quad \delta := \frac{a}{2\sqrt{r+\ell}}.$$

Now define the following events:

$$\begin{aligned}
\Omega_0 & := \{z < N \leq z + r\} \\
\Omega_1 & := \{t_{N,b}^j \leq N - z + ub^{-2} \text{ for all } j \in [p] \text{ and } b \in \mathcal{B} \cup \mathcal{B}_0\}, \\
\Omega_2 & := \left\{ \left| A_{N,b}^{j',j} + \sum_{i=N+1}^{N+\ell} X_i^{j'} \right| < a\sqrt{t_{N,b}^j + \ell} \text{ for all } b \in \mathcal{B} \cup \mathcal{B}_0, j \in [p] \right. \\
& \quad \left. \text{and all } j' \in [p] \setminus \{j\} \text{ with } |\theta^{j'}| \leq \delta \right\}, \\
\Omega_3 & := \{t_{N,\hat{b}}^{\hat{j}} \leq N - z + \ell/20\}.
\end{aligned}$$

Finally, we denote event

$$\Omega_4 := \Omega_{4,1} \cup \Omega_{4,2},$$

with

$$\begin{aligned}\Omega_{4,1} &:= \{\hat{j} \neq 1, 1 \in \hat{\mathcal{S}}, \tilde{b}^1 \geq b^1/\sqrt{2}\} \\ \Omega_{4,2} &:= \{\hat{j} = 1, 2 \in \hat{\mathcal{S}}, \tilde{b}^2 \geq b^2/\sqrt{2}\}.\end{aligned}$$

We note that for $j \in \{1, 2\}$, on $\Omega_{4,j}$, we have $\tilde{b}^j \in \mathcal{B} \cup \mathcal{B}_0$. Then, on the event $\bigcap_{k=0}^4 \Omega_k$, we have

$$L = \min_{j \in \hat{\mathcal{S}}} \left\{ t_{N, \tilde{b}^j}^j + \frac{d_2}{(\tilde{b}^j)^2} \right\} \wedge N \leq N - z + \frac{2(u + d_2)}{(b^2)^2} \leq 3r + \frac{2d_2}{b_*^2} \leq 8r.$$

Thus, it suffices to control the probability of $\cup_{k=0}^4 \Omega_k^c$. First, we note that

$$\mathbb{P}(\Omega_0^c) \leq g(r; N) + \mathbb{P}(N \leq z). \quad (9)$$

On Ω_0 , we have for any $j \in [p]$ and $b \in \mathcal{B} \cup \mathcal{B}_0$ that

$$\begin{aligned}t_{N,b}^j &= \operatorname{sargmax}_{0 \leq h \leq N} \sum_{i=N-h+1}^N b(X_i^j - b/2) \leq \operatorname{sargmax}_{N-z \leq h \leq N} \sum_{i=N-h+1}^N b(X_i^j - b/2) \\ &= N - z + \operatorname{sargmax}_{0 \leq h \leq z} \sum_{i=z-h+1}^z b(X_i^j - b/2).\end{aligned}$$

Thus, by Lemma 9 (taking $\mu = -b/2$) and a union bound, we have

$$\mathbb{P}(\Omega_0 \cap \Omega_1^c) \leq 2p \log_2(4p) e^{-u/8} = 2p \log_2(4p) e^{-rb_*^2/8}. \quad (10)$$

Now observe that, for all $z < n \leq z + r, b \in \mathcal{B} \cup \mathcal{B}_0, j \in [p]$ and $j' \in [p] \setminus \{j\}$, we have

$$A_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'} \Big| t_{n,b}^j \sim \mathcal{N}\left(\theta^{j'} \{\ell + \min(t_{n,b}^j, n - z)\}, t_{n,b}^j + \ell\right). \quad (11)$$

Hence, when $|\theta^{j'}| \leq \delta$, we have that

$$\mathbb{P}\left(\left|A_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'}\right| \geq a\sqrt{t_{n,b}^j + \ell} \Big| t_{n,b}^j\right) \leq \mathbb{P}(|Y_1| \geq a) \leq 2\mathbb{P}(Y_1 \geq a) \leq e^{-a^2/8}, \quad (12)$$

where $Y_1 \sim \mathcal{N}(\delta\sqrt{n - z + \ell}, 1)$, and where the last inequality follows from the relation $a = 2\delta\sqrt{r + \ell}$. Thus, by a union bound, we have

$$\mathbb{P}(\Omega_0 \cap \Omega_2^c) \leq 2rp^2 \log_2(4p) e^{-a^2/8}. \quad (13)$$

Recall that $u = \ell_0 b_*^2 / 80$. We therefore have for any $z < n \leq z + r$, $j \in [p]$ and $b \in \mathcal{B}$ that

$$\begin{aligned}
& \mathbb{P}\left(\{N = n\} \cap \Omega_1 \cap \Omega_2 \cap \{Q_{n,b}^j \geq Q_{n,b_*}^j\} \cap \{t_{n,b}^j > n - z + \ell/20\} \mid X_1^j, X_2^j, \dots\right) \\
& \leq \mathbb{P}\left(\bigcup_{\substack{j' \in [p] \setminus \{j\}: \\ |\theta^{j'}| > \delta}} \{|E_{n,b}^{j',j}| \geq |E_{n,b_*}^{j',j}|\} \cap \{N = n\} \cap \Omega_1 \cap \Omega_2 \cap \{t_{n,b}^j > n - z + \ell/20\} \mid X_1^j, X_2^j, \dots\right) \\
& \leq \sum_{\substack{j' \in [p] \setminus \{j\}: \\ |\theta^{j'}| > \delta}} \mathbb{P}\left(\{|E_{n,b}^{j',j}| \geq |E_{n,b_*}^{j',j}|\} \cap \{n - z < t_{n,b_*}^j \leq n - z + \ell/80\} \right. \\
& \quad \left. \cap \{t_{n,b}^j > n - z + \ell/20\} \mid X_1^j, X_2^j, \dots\right) \\
& \quad + \sum_{\substack{j' \in [p] \setminus \{j\}: \\ |\theta^{j'}| > \delta}} \mathbb{P}\left(\{|E_{n,b}^{j',j}| \geq |E_{n,b_*}^{j',j}|\} \cap \{t_{n,b_*}^j \leq n - z\} \cap \{t_{n,b}^j > n - z + \ell/20\} \mid X_1^j, X_2^j, \dots\right) \\
& \leq p \exp\left(-\frac{\ell \delta^2}{960}\right) = p \exp\left(-\frac{\ell a^2}{3840(r + \ell)}\right), \tag{14}
\end{aligned}$$

where the final inequality follows from Lemma 11(a), applied with $U = \sum_{i=n-t_{n,b_*}^j+1}^z X_i^{j'}$, $V = \sum_{i=n-t_{n,b}^j+1}^{n-t_{n,b_*}^j} X_i^{j'}$, $Y = \sum_{i=z+1}^{n+\ell} X_i^{j'}$, $\alpha = \theta^{j'}$, $\phi_1 = z - n + t_{n,b_*}^j$, $\phi_2 = z - n + t_{n,b}^j$, $\phi_3 = n - z + \ell$ and $\kappa = \ell/80$, as well as Lemma 11(b), with $U = \sum_{i=n-t_{n,b}^j+1}^z X_i^{j'}$, $Y = \sum_{i=n-t_{n,b_*}^j+1}^{n+\ell} X_i^{j'}$, $Z = \sum_{i=z+1}^{n-t_{n,b_*}^j} X_i^{j'}$, $\alpha = \theta^{j'}$, $\phi_1 = z - n + t_{n,b}^j$, $\phi_3 = \ell + t_{n,b_*}^j$, $\phi_4 = n - z - t_{n,b_*}^j$ and $\kappa = \ell/80$. Observe that $Q_{n,\hat{b}}^j \geq Q_{n,b_*}^j$. Thus, by a union bound, we have

$$\begin{aligned}
& \mathbb{P}(\Omega_0 \cap \Omega_1 \cap \Omega_2 \cap \Omega_3^c) \\
& \leq \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{n=z+1}^{z+[r]} \mathbb{P}\left(\{N = n\} \cap \Omega_1 \cap \Omega_2 \cap \{Q_{n,b}^j \geq Q_{n,b_*}^j\} \cap \{t_{n,b}^j > n - z + \ell/20\}\right) \\
& \leq 2rp^2 \log_2(2p) \exp\left(-\frac{\ell a^2}{3840(r + \ell)}\right) \leq 2rp^2 \log_2(2p) e^{-a^2/3888}, \tag{15}
\end{aligned}$$

where the last inequality follows from $r = \ell_0/80 \leq \ell/80$. Recall that $d_1^2 = 5rb_*^2/9 \leq (b^1)^2 \ell/144$. Thus, on $\Omega_0 \cap \Omega_3$, we have

$$t_{N,\hat{b}}^j \leq N - z + \ell/20 \leq r + \ell/20 \leq \ell/16.$$

Hence, for any $z < n \leq z + r$, $j \in [p] \setminus \{1\}$ and $b \in \mathcal{B}$, we have

$$\begin{aligned}
& \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b\} \cap \Omega_{4,1}^c \mid X_1^j, X_2^j, \dots) \\
& \leq \mathbb{P}\left(\{t_{n,b}^j \leq \ell/16\} \cap \left\{E_{n,b}^{1,j} - b^1 \sqrt{(t_{n,b}^j + \ell)/2} < d_1\right\} \mid X_1^j, X_2^j, \dots\right) \leq \frac{1}{2} e^{-d_1^2/2}. \tag{16}
\end{aligned}$$

Here, in the final bound, we have used the facts that $E_{n,b}^{1,j} | t_{n,b}^j \sim \mathcal{N}(\theta^1 \min\{(n + \ell - z)(t_{n,b}^j + \ell)^{-1/2}, (t_{n,b}^j + \ell)^{1/2}\}, 1)$ and that

$$\begin{aligned} & \theta^1 \min\{(n + \ell - z)(t_{n,b}^j + \ell)^{-1/2}, (t_{n,b}^j + \ell)^{1/2}\} - b^1 \sqrt{(t_{n,b}^j + \ell)/2} \\ & \geq \frac{4\theta^1 \sqrt{\ell}}{\sqrt{17}} - \frac{b^1 \sqrt{17\ell}}{4\sqrt{2}} \geq \frac{b^1 \sqrt{\ell}}{6} \geq 2d_1, \end{aligned}$$

when $t_{n,b}^j \leq \ell/16$, as well as the standard Gaussian tail bound used at the end of the proof of Lemma 9. By a similar argument, we also have for any $z < n \leq z + r$ and $b \in \mathcal{B}$ that

$$\begin{aligned} & \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = 1, \hat{b} = b\} \cap \Omega_{4,2}^c | X_1^1, X_2^1, \dots) \\ & \leq \mathbb{P}\left(\{t_{n,b}^1 \leq \ell/16\} \cap \left\{E_{n,b}^{2,1} - b^2 \sqrt{(t_{n,b}^1 + \ell)/2} < d_1\right\} | X_1^1, X_2^1, \dots\right) \leq \frac{1}{2} e^{-d_1^2/2}. \quad (17) \end{aligned}$$

Thus, by a union bound, we have

$$\begin{aligned} \mathbb{P}(\Omega_0 \cap \Omega_3 \cap \Omega_4^c) &= \mathbb{P}(\Omega_0 \cap \Omega_3 \cap \Omega_{4,1}^c \cap \Omega_{4,2}^c) \\ &\leq \sum_{j=2}^p \sum_{b \in \mathcal{B}} \sum_{n=z+1}^{z+[r]} \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b\} \cap \Omega_{4,1}^c) \\ &\quad + \sum_{b \in \mathcal{B}} \sum_{n=z+1}^{z+[r]} \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = 1, \hat{b} = b\} \cap \Omega_{4,2}^c) \\ &\leq rp \log_2(2p) e^{-5rb_*^2/18}. \quad (18) \end{aligned}$$

Hence combining (9), (10), (13), (15) and (18), we conclude that

$$\mathbb{P}(L > 8r) \leq g(r; N) + \mathbb{P}(N \leq z) + 4rp^2 \log_2(4p) e^{-a^2/3888} + 3rp \log_2(4p) e^{-rb_*^2/8} \leq \alpha,$$

where the last inequality follows from the choice of a with a sufficiently large universal constant C and (3). \square

Proof of Theorem 3. Fix $r \geq 1$ that satisfies the assumption (3).

(a) For $j' \in \mathcal{S}_\beta^c$, we have $|\theta^{j'}| < b_{\min}$, so the event $\{|\tilde{b}^{j'}| \leq |\theta^{j'}|\}$ is empty. Thus by (7), we have, for $n > z$, $j \in [p]$, $b \in \mathcal{B}$ and $j' \in \mathcal{S}_\beta^c$, that

$$\mathbb{P}(\{j' \in \hat{\mathcal{S}}\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}) \leq 2\bar{\Phi}(d_1).$$

Hence, by a union bound, we have

$$\begin{aligned} \mathbb{P}(\hat{\mathcal{S}} \not\subseteq \mathcal{S}_\beta) &\leq \mathbb{P}(N \leq z) + \mathbb{P}(N > z + r) \\ &\quad + \sum_{n=z+1}^{z+[r]} \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{j' \in \mathcal{S}_\beta^c} \mathbb{P}(\{j' \in \hat{\mathcal{S}}\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}) \\ &\leq \mathbb{P}(N \leq z) + g(r; N) + 4rp^2 \log_2(2p) \bar{\Phi}(d_1) \leq \alpha, \end{aligned}$$

as required.

(b) We use the events $\Omega_0, \Omega_1, \Omega_2, \Omega_3$ defined in the proof of Theorem 2. Recall from the argument immediately below (15) that we have $t_{N, \hat{b}}^{\hat{j}} \leq \ell/16$ and $d_1 \leq \min_{j' \in \mathcal{S}} |\theta^{j'}| \sqrt{\ell}/12$ on $\Omega_0 \cap \Omega_3$. Recall also the definition of $E_{n,b}^{j',j}$ from Algorithm 1. Then, for any $z < n \leq z+r$, $j \in [p]$, $j' \in \mathcal{S} \setminus \{j\}$ and $b \in \mathcal{B}$, we have

$$\begin{aligned} & \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b, j' \notin \hat{\mathcal{S}}\} \mid X_1^j, X_2^j, \dots) \\ &= \mathbb{P}\left(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b\} \cap \{|E_{n,b}^{j',j}| < b_{\min} \sqrt{t_{n,b}^j + \ell} + d_1\} \mid X_1^j, X_2^j, \dots\right) \\ &\leq \mathbb{P}\left(\{t_{n,b}^j \leq \ell/16\} \cap \left\{|E_{n,b}^{j',j}| - b_{\min} \sqrt{t_{n,b}^j + \ell} < d_1\right\} \mid X_1^j, X_2^j, \dots\right) \leq \frac{1}{2} e^{-d_1^2/2}, \end{aligned} \quad (19)$$

where, in the final bound, we have used the facts that $E_{n,b}^{j',j} \mid t_{n,b}^j \sim \mathcal{N}(\theta^{j'} \min\{(n+\ell-z)(t_{n,b}^j + \ell)^{-1/2}, (t_{n,b}^j + \ell)^{1/2}\}, 1)$ and that

$$\begin{aligned} & |\theta^{j'}| \min\{(n+\ell-z)(t_{n,b}^j + \ell)^{-1/2}, (t_{n,b}^j + \ell)^{1/2}\} - b_{\min} \sqrt{t_{n,b}^j + \ell} \\ &\geq \frac{4|\theta^{j'}| \sqrt{\ell}}{\sqrt{17}} - \frac{b_{\min} \sqrt{17\ell}}{4\sqrt{2}} \geq \frac{|\theta^{j'}| \sqrt{\ell}}{6} \geq 2d_1, \end{aligned}$$

when $t_{n,b}^j \leq \ell/16$. Hence

$$\begin{aligned} & \mathbb{P}(\hat{\mathcal{S}} \cup \{\hat{j}\} \not\subseteq \mathcal{S}) \\ &\leq \mathbb{P}(\Omega_0^c) + \mathbb{P}(\Omega_0 \cap \Omega_1^c) + \mathbb{P}(\Omega_0 \cap \Omega_2^c) + \mathbb{P}(\Omega_0 \cap \Omega_1 \cap \Omega_2 \cap \Omega_3^c) \\ &\quad + \sum_{n=z+1}^{z+[r]} \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{j' \in \mathcal{S} \setminus \{j\}} \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b, j' \notin \hat{\mathcal{S}}\}) \\ &\leq g(r; N) + \mathbb{P}(N \leq z) + 4rp^2 \log_2(4p) e^{-a^2/3888} + 3rp^2 \log_2(4p) e^{-r\beta^2/(8s \log_2(2p))} \leq \alpha, \end{aligned}$$

where the penultimate inequality follows from (9), (10), (13), (15) and (19), and the last inequality follows from the choice of a with a sufficiently large universal constant C and (3). \square

Proof of Proposition 4. Fix $N \in \mathcal{T}_{r,m}$ and $\psi \in \mathcal{J}_N$. We denote by $\mathbb{P}_{z,\theta}^{(n_0)}$ the restriction of $\mathbb{P}_{z,\theta}$ to the filtration $\mathcal{F}_{n_0} := \sigma(X_1, \dots, X_{n_0})$. Denote

$$\tilde{\Theta} := \{\theta \in \mathbb{R}^p : \theta^j \in \{0, 1/(8\sqrt{r})\}, |\text{supp}(\theta)| = m\},$$

and let $\tilde{\Theta}_{\text{pa}} \subseteq \tilde{\Theta}$ be an $(m/4)$ -packing set with respect to the symmetric difference metric defined above, i.e. for any $\theta, \tilde{\theta} \in \tilde{\Theta}_{\text{pa}}$, we have $d(\text{supp}(\theta), \text{supp}(\tilde{\theta})) > m/4$. We also have $\text{KL}(\mathbb{P}_{z,\theta}^{(z+r)}, \mathbb{P}_{z,\tilde{\theta}}^{(z+r)}) = r \|\theta - \tilde{\theta}\|_2^2/2 \leq m/64$.

Enumerate $\tilde{\Theta}_{\text{pa}} = \{\theta_{(1)}, \theta_{(2)}, \dots, \theta_{(|\tilde{\Theta}_{\text{pa}}|)}\}$. Let $\phi^* := \text{sargmin}_{k \in [|\tilde{\Theta}_{\text{pa}}|]} d(\psi, \text{supp}(\theta_{(k)}))$. Note that ϕ^* is also \mathcal{F}_N -measurable. Then for any $z \in \mathbb{N}_0$,

$$\begin{aligned}
\sup_{\theta \in \Theta_{r,m}} \mathbb{E}_{z,\theta} d(\psi, \text{supp}(\theta)) &\geq \frac{m}{8|\tilde{\Theta}_{\text{pa}}|} \sum_{k=1}^{|\tilde{\Theta}_{\text{pa}}|} \mathbb{P}_{z,\theta_{(k)}} \left(d(\psi, \text{supp}(\theta_{(k)})) > \frac{m}{8} \right) \\
&\geq \frac{m}{8|\tilde{\Theta}_{\text{pa}}|} \sum_{k=1}^{|\tilde{\Theta}_{\text{pa}}|} \mathbb{P}_{z,\theta_{(k)}} (\phi^* \neq k) \\
&= \frac{m}{8} \left\{ 1 - \frac{1}{|\tilde{\Theta}_{\text{pa}}|} \sum_{k=1}^{|\tilde{\Theta}_{\text{pa}}|} \mathbb{P}_{z,\theta_{(k)}} (\phi^* = k) \right\} \\
&\geq \frac{m}{8} \left\{ \frac{3}{4} - \frac{1}{|\tilde{\Theta}_{\text{pa}}|} \sum_{k=1}^{|\tilde{\Theta}_{\text{pa}}|} \mathbb{P}_{z,\theta_{(k)}} (\phi^* = k, N \leq z+r) \right\} \\
&= \frac{m}{8} \left\{ \frac{3}{4} - \frac{1}{|\tilde{\Theta}_{\text{pa}}|} \sum_{k=1}^{|\tilde{\Theta}_{\text{pa}}|} \mathbb{P}_{z,\theta_{(k)}}^{(z+r)} (\phi^* = k, N \leq z+r) \right\}. \tag{20}
\end{aligned}$$

Now set

$$\tilde{\phi}^* := \begin{cases} \phi^* & \text{if } N \leq z+r \\ 1 & \text{if } N > z+r. \end{cases}$$

Then $\tilde{\phi}^*$ is \mathcal{F}_{z+r} -measurable and by Fano's inequality (Yu, 1997, Lemma 3), we have

$$\begin{aligned}
\frac{1}{|\tilde{\Theta}_{\text{pa}}|} \sum_{k=1}^{|\tilde{\Theta}_{\text{pa}}|} \mathbb{P}_{z,\theta_{(k)}}^{(z+r)} (\phi^* = k, N \leq z+r) &\leq \frac{1}{|\tilde{\Theta}_{\text{pa}}|} \sum_{k=1}^{|\tilde{\Theta}_{\text{pa}}|} \mathbb{P}_{z,\theta_{(k)}}^{(z+r)} (\tilde{\phi}^* = k) \\
&\leq \frac{\log 2 + |\tilde{\Theta}_{\text{pa}}|^{-2} \sum_{j,k=1}^{|\tilde{\Theta}_{\text{pa}}|} \text{KL}(\mathbb{P}_{z,\theta_{(j)}}^{(z+r)}, P_{z,\theta_{(k)}}^{(z+r)})}{\log |\tilde{\Theta}_{\text{pa}}|} \\
&\leq \frac{\log 2 + m/64}{\log |\tilde{\Theta}_{\text{pa}}|}. \tag{21}
\end{aligned}$$

By Massart (2007, Lemma 4.7), there exists an $(m/4)$ -packing set with

$$\log |\tilde{\Theta}_{\text{pa}}| \geq m/8. \tag{22}$$

Combining (20), (21) and (22), we conclude that

$$\begin{aligned}
\sup_{z \in \mathbb{N}_0, \theta \in \Theta_{r,m}} \mathbb{E}_{z,\theta} d(\psi, \text{supp}(\theta)) &\geq \frac{m}{8} \left(\frac{3}{4} - \frac{\log 2 + m/64}{m/8} \right) \\
&\geq \frac{m}{8} \left(\frac{3}{4} - \frac{8 \log 2}{m} - \frac{1}{8} \right) \geq \frac{m}{32},
\end{aligned}$$

where we have used the assumption that $m \geq 15$ in the final inequality. \square

Proof of Proposition 5. Following the proof of [Chen et al. \(2022, Theorem 1\(b\)\)](#) up to, but not including, (16), we have for every $j \in [p]$, $b \in \mathcal{B} \cup \mathcal{B}_0$ that

$$\mathbb{P}\left(\max_{1 \leq n \leq z} (bA_{n,b}^{j,j} - b^2 t_{n,b}^j / 2) \geq T^{\text{diag}}\right) \leq 1 - (1 - e^{-T^{\text{diag}}})^z \leq ze^{-T^{\text{diag}}}. \quad (23)$$

It follows by a union bound that

$$\begin{aligned} \mathbb{P}\left(\max_{1 \leq n \leq z} S_n^{\text{diag}} \geq T^{\text{diag}}\right) &= \mathbb{P}\left(\max_{1 \leq n \leq z} \max_{j \in [p]} \max_{b \in \mathcal{B} \cup \mathcal{B}_0} (bA_{n,b}^{j,j} - b^2 t_{n,b}^j / 2) \geq T^{\text{diag}}\right) \\ &\leq zp|\mathcal{B} \cup \mathcal{B}_0|e^{-T^{\text{diag}}} \leq \frac{\alpha}{4}. \end{aligned} \quad (24)$$

Next, for every $j \in [p]$, $j' \in [p] \setminus \{j\}$, $b \in \mathcal{B}$ and $n \in [z]$, we have $\Lambda_{n,b}^{j',j} \mid \tau \sim \mathcal{N}(0, \tau_{n,b}^j)$, so $(\Xi_{n,b}^{j',j})^2 \mid \tau_{n,b}^j \leq_{\text{st}} \chi_1^2$. Thus, by another union bound, we have

$$\begin{aligned} \mathbb{P}\left(\max_{n \in [z]} S_n^{\text{off}} \geq T^{\text{off}}\right) &= \mathbb{P}\left(\max_{n \in [z]} \max_{j \in [p]} \max_{b \in \mathcal{B}} \sum_{j' \in [p] \setminus \{j\}} (\Xi_{n,b}^{j',j})^2 \mathbb{1}_{\{|\Xi_{n,b}^{j',j}| \geq \tilde{a}\}} \geq T^{\text{off}}\right) \\ &\leq \mathbb{P}\left(\max_{n \in [z]} \max_{j \in [p]} \max_{b \in \mathcal{B}} \max_{j' \in [p] \setminus \{j\}} |\Xi_{n,b}^{j',j}| \geq \tilde{a}\right) \\ &\leq zp^2|\mathcal{B}|e^{-\tilde{a}^2/2} \leq \frac{\alpha}{4}. \end{aligned} \quad (25)$$

From (24) and (25), we deduce that

$$\mathbb{P}(N \leq z) \leq \alpha/2. \quad (26)$$

On the other hand, for a sufficiently large universal constant $C' > 0$, we have

$$\begin{aligned} r_1 &= \frac{C' s \log_2(2p) \log\{p\gamma\alpha^{-1}(\beta^{-2} \vee 1)\}}{\beta^2} + 2 \\ &\geq \left\{ \frac{24T^{\text{off}} \log_2(2p)}{\vartheta^2} \vee \frac{12\tilde{a}^2 s \log_2(2p)}{\vartheta^2} \vee \frac{24T^{\text{diag}} s \log_2(2p)}{\beta^2} \right\} + 2. \end{aligned}$$

Thus, by [Proposition 10](#) and by increasing the value of C if necessary, we have for every $r \geq r_1$ that

$$\begin{aligned} g(r; N) + 4rp^2 \log_2^2(4p) e^{-r\beta^2/(8s \log_2(2p))} &\leq 5rp^2 \log_2^2(4p) \exp\left\{-\frac{r\beta^2}{48s \log_2(2p)}\right\} \\ &\leq \frac{240sp^2 \log_2^3(4p)}{\beta^2} \frac{r\beta^2}{48s \log_2(2p)} \exp\left\{-\frac{r\beta^2}{48s \log_2(2p)}\right\} \\ &\leq \frac{240sp^2 \log_2^3(4p)}{\beta^2} \exp\left\{-\frac{r\beta^2}{96s \log_2(2p)}\right\} \leq \frac{\alpha}{4}, \end{aligned} \quad (27)$$

where the penultimate inequality follows from the fact that $xe^{-x} \leq e^{-x/2}$ for $x \geq 0$. The desired result follows by combining (26) and (27). \square

5.2 Auxiliary results

Proposition 8. *Let X_1, X_2, \dots be independent random variables with $X_1, \dots, X_z \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $X_{z+1}, X_{z+2}, \dots \stackrel{\text{iid}}{\sim} \mathcal{N}(\theta, 1)$. Assume that $0 < b \leq \theta$ and let $t_{n,b}$ be defined as in (1) for $n \in \mathbb{N}$. Then for any $\alpha \in (0, 1)$, and any stopping time N satisfying $\mathbb{P}(N < z) \leq \alpha/2$, we have that the confidence interval*

$$\mathcal{C}_0 := \left[N - t_{N,b} - \frac{4\{\Phi^{-1}(1 - \alpha/4)\}^2}{b^2}, N \right]$$

satisfies $\mathbb{P}(z \in \mathcal{C}_0) \geq 1 - \alpha$.

Remark. *We could also replace $4\{\Phi^{-1}(1 - \alpha/4)\}^2/b^2$ by $8 \log(2/\alpha)/b^2$ in the confidence interval construction, if we apply the final bound from Lemma 9 in (28) of the proof below.*

Proof. For $n \in \mathbb{N}$, define $R_{n,b} := \max\{R_{n-1,b} + b(X_n - b/2), 0\}$, with $R_{0,b} = 0$. By Chen et al. (2022, Lemma 2 in the supplement), we have $t_{N,b} = \min\{i : 0 \leq i \leq N, R_{N-i,b} = 0\} = \text{sargmax}_{0 \leq h \leq N} \sum_{i=N-h+1}^N b(X_i - b/2)$. Let $U_{n,b} := \sum_{i=z+1}^{z+n} (X_i - b/2)$ for $n \in \mathbb{N}$, with $U_{0,b} := 0$. Then $R_{n+z,b} \geq bU_{n,b}$ for all $n \in \mathbb{N}$. Hence, for $y \in [0, \infty)$, we have

$$\mathbb{P}(N - t_{N,b} - y \geq z) \leq \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq z+y} R_{n,b} = 0\right) \leq \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} U_{n,b} \leq 0\right) \leq 2\bar{\Phi}\left(\sqrt{y}(\theta - b/2)\right), \quad (28)$$

where the last inequality follows from Lemma 9. Thus, if we choose $y = 4\{\Phi^{-1}(1 - \alpha/4)\}^2/b^2$, then we are guaranteed that $\mathbb{P}(N - t_{N,b} - y > z) \leq \alpha/2$. Combining this with the assumption that $\mathbb{P}(N < z) \leq \alpha/2$, the desired result follows. \square

Lemma 9. *Let $Y_1, Y_2, \dots \stackrel{\text{iid}}{\sim} \mathcal{N}(\mu, 1)$. Define $U_n := \sum_{i=1}^n Y_i$ for $n \in \mathbb{N}_0$, and let $\xi := \text{sargmin}_{n \in \mathbb{N}_0} \mu U_n$. Then, for $y \in [0, \infty)$, we have*

$$\mathbb{P}(\xi \geq y) \leq \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} \mu U_n \leq 0\right) \leq 2\bar{\Phi}(\sqrt{y}|\mu|) \leq e^{-y\mu^2/2}.$$

Proof. The first inequality holds since $\mu U_\xi \leq \mu U_0 = 0$. For the second and third inequalities, we may assume without loss of generality that $\mu > 0$, since the result is clear when $\mu = 0$, and if $\mu < 0$ then the result will follow from the corresponding result with $\mu > 0$ by setting $Y'_i := -Y_i$ for $i \in \mathbb{N}$. Note that $(U_n - n\mu)_{n \in \mathbb{N}_0}$ is a standard Gaussian random walk starting at 0. Let $(B_t)_{t \in [0, \infty)}$ denote a standard Brownian motion starting at 0. Then, we have for any $y \in \mathbb{N}_0$ and $u > 0$ that

$$\mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} U_n \leq 0 \mid U_y = u\right) \leq \mathbb{P}\left\{\inf_{t \in [y, \infty)} (B_t + t\mu) \leq 0 \mid B_y = u\right\} \leq e^{-2u\mu}, \quad (29)$$

where the final inequality follows from Siegmund (1986, Proposition 2.4 and Equation (2.5)).

Thus, for $y \in [0, \infty)$, we have

$$\begin{aligned} \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} U_n \leq 0\right) &= \mathbb{P}(U_{\lceil y \rceil} \leq 0) + \mathbb{E}\left\{\mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq \lceil y \rceil} U_n \leq 0 \mid U_{\lceil y \rceil}\right) \mathbb{1}_{\{U_{\lceil y \rceil} > 0\}}\right\} \\ &\leq \bar{\Phi}\left(\sqrt{\lceil y \rceil} \mu\right) + \int_0^\infty \frac{1}{\sqrt{2\pi \lceil y \rceil}} e^{-\frac{(u - \lceil y \rceil \mu)^2}{2 \lceil y \rceil}} e^{-2u\mu} du \\ &= 2\bar{\Phi}\left(\sqrt{\lceil y \rceil} \mu\right) \leq 2\bar{\Phi}(\sqrt{y} \mu) \leq e^{-y\mu^2/2}, \end{aligned}$$

where the first inequality follows from (29) and the fact that $U_{\lceil y \rceil} \sim \mathcal{N}(\lceil y \rceil \mu, \lceil y \rceil)$ and the last inequality follows from the standard normal distribution tail bound $\bar{\Phi}(x) \leq e^{-x^2/2}/2$ for $x \geq 0$. \square

In Proposition 10, we assume the Gaussian data generating mechanism given at the beginning of Section 3, and show that for the ocd' base procedure, the quantity $g(r; N)$ from (2) has essentially the same form as the final term in (3).

Proposition 10. *Assume that θ has an effective sparsity of $s := s(\theta) \geq 2$. Then, the output N from ocd', with inputs $(X_t)_{t \in \mathbb{N}}$, $0 < \beta \leq \vartheta$, $\tilde{a} > 0$, $T^{\text{diag}} > 0$ and $T^{\text{off}} > 0$, satisfies*

$$\mathbb{P}_{z, \theta, \Sigma}(N > z + r) \leq p \exp\left\{-\frac{\beta^2(r-1)}{24s \log_2(2p)}\right\},$$

for all $r \geq \left\{\frac{24T^{\text{off}} \log_2(2p)}{\vartheta^2} \vee \frac{12\tilde{a}^2 s \log_2(2p)}{\vartheta^2} \vee \frac{24T^{\text{diag}} s \log_2(2p)}{\beta^2}\right\} + 2$.

Proof. For $\theta \in \mathbb{R}^p$ with effective sparsity $s(\theta)$, there is at most one coordinate in θ of magnitude larger than $\vartheta/\sqrt{2}$, so there exists $b_* \in \{\beta/\sqrt{s(\theta) \log_2(2p)}, -\beta/\sqrt{s(\theta) \log_2(2p)}\} \subseteq \mathcal{B}$ such that

$$\mathcal{J} := \left\{j \in [p] : \theta^j/b_* \geq 1 \text{ and } |\theta^j| \leq \vartheta/\sqrt{2}\right\}$$

has cardinality at least $s(\theta)/2$. Note that the condition $\theta^j/b_* \geq 1$ above ensures that $\{\theta^j : j \in \mathcal{J}\}$ all have the same sign as b_* . By Chen et al. (2022, Proposition 8), we have on the event $\{N > z\}$ that

$$q(X_1, \dots, X_z, \theta) := \max\{t_{z, b_*}^j : j \in \mathcal{J}\} \leq \frac{8T^{\text{diag}} s \log_2(2p)}{\beta^2}. \quad (30)$$

We now fix

$$r \geq \left\{\frac{24T^{\text{off}} \log_2(2p)}{\vartheta^2} \vee \frac{12\tilde{a}^2 s \log_2(2p)}{\vartheta^2} \vee \frac{24T^{\text{diag}} s \log_2(2p)}{\beta^2}\right\} + 2 =: r_0. \quad (31)$$

For $j \in \mathcal{J}$, define the event

$$\Omega_r^j := \{t_{z + \lceil r \rceil, b_*}^j > 2 \lceil r \rceil / 3\}.$$

By applying [Chen et al. \(2022, Lemma 2\)](#) to $t_{z+\lfloor r \rfloor, b_*}^j$, we have for $j \in \mathcal{J}$ that

$$\begin{aligned} t_{z+\lfloor r \rfloor, b_*}^j &= \operatorname{sargmax}_{0 \leq h \leq z+\lfloor r \rfloor} \sum_{i=z+\lfloor r \rfloor-h+1}^{z+\lfloor r \rfloor} b_*(X_i^j - b_*/2) \geq \operatorname{sargmax}_{0 \leq h \leq \lfloor r \rfloor} \sum_{i=z+\lfloor r \rfloor-h+1}^{z+\lfloor r \rfloor} b_*(X_i^j - b_*/2) \\ &= \operatorname{sargmax}_{0 \leq h \leq \lfloor r \rfloor} \sum_{i=z+1}^{z+\lfloor r \rfloor-h} -b_*(X_i^j - b_*/2) = \lfloor r \rfloor - \operatorname{largmax}_{0 \leq h \leq \lfloor r \rfloor} \sum_{i=z+1}^{z+h} -b_*(X_i^j - b_*/2). \end{aligned}$$

Recall that $X_{z+1}^j, X_{z+2}^j, \dots \stackrel{\text{iid}}{\sim} \mathcal{N}(\theta^j, 1)$. Hence, by applying [Lemma 9](#) with $\mu = |b_*|/2$ and $y = \lfloor r \rfloor/3$, we have for each $j \in \mathcal{J}$ that

$$\begin{aligned} \mathbb{P}\{(\Omega_r^j)^c\} &= \mathbb{P}\left(t_{z+\lfloor r \rfloor, b_*}^j \leq \frac{2\lfloor r \rfloor}{3}\right) \\ &\leq \mathbb{P}\left(\operatorname{largmax}_{0 \leq h \leq \lfloor r \rfloor} \sum_{i=z+1}^{z+h} -b_*(X_i^j - b_*/2) \geq \frac{\lfloor r \rfloor}{3}\right) \\ &\leq \mathbb{P}\left(\sup_{h \geq \lfloor r \rfloor/3} \sum_{i=z+1}^{z+h} -\operatorname{sgn}(b_*)(X_i^j - b_*/2) \geq 0\right) \leq \exp(-b_*^2 \lfloor r \rfloor/24). \end{aligned} \quad (32)$$

We now work on the event Ω_r^j , for some fixed $j \in \mathcal{J}$. We note that [\(31\)](#) guarantees that $r \geq 2$, and thus $t_{z+\lfloor r \rfloor, b_*}^j \geq \lceil 2\lfloor r \rfloor/3 \rceil \geq 2$. Then, by [\(30\)](#) and [\(31\)](#), we have $r_0 > 3t_{z, b_*}^j$, and hence by [Chen et al. \(2022, Lemma 9\)](#),

$$\frac{\lfloor r \rfloor}{3} < \frac{t_{z+\lfloor r \rfloor, b_*}^j}{2} \leq \tau_{z+\lfloor r \rfloor, b_*}^j \leq \frac{3t_{z+\lfloor r \rfloor, b_*}^j}{4} \leq \frac{3(t_{z, b_*}^j + r)}{4} < r.$$

We conclude that

$$2/3 \leq \lfloor r \rfloor/3 < \tau_{z+\lfloor r \rfloor, b_*}^j \leq \lfloor r \rfloor. \quad (33)$$

Recall that $\Lambda_{z+\lfloor r \rfloor, b_*}^{j, j}$ records the tail CUSUM statistics with tail length $\tau_{z+\lfloor r \rfloor, b_*}^j$. We observe by [\(33\)](#) that only post-change observations are included in $\Lambda_{z+\lfloor r \rfloor, b_*}^{j, j}$. Hence we have that

$$\Lambda_{z+\lfloor r \rfloor, b_*}^{j', j} \mid \tau_{z+\lfloor r \rfloor, b_*}^j \stackrel{\text{ind}}{\sim} \mathcal{N}(\theta^{j'} \tau_{z+\lfloor r \rfloor, b_*}^j, \tau_{z+\lfloor r \rfloor, b_*}^j) \quad (34)$$

for $j' \in [p] \setminus \{j\}$. By the definition of the effective sparsity of θ , the set

$$\mathcal{L}^j := \left\{ j' \in [p] \setminus [j] : |\theta^{j'}| \geq \frac{\vartheta}{\sqrt{s \log_2(2p)}} \right\}$$

has cardinality at least $s - 1$. Hence, by [\(33\)](#), for all $j' \in \mathcal{L}^j$,

$$|\theta^{j'}| \sqrt{\tau_{z+\lfloor r \rfloor, b_*}^j} > \sqrt{\frac{\vartheta^2 \lfloor r \rfloor}{3s \log_2(2p)}} =: \tilde{a}_r.$$

We then observe, from (31), that

$$\tilde{a}_r > 2\tilde{a}. \quad (35)$$

Hence, from (34), we have for all $j' \in \mathcal{L}^j$ that

$$\mathbb{P}\left(\Omega_r^j \cap \left\{ |\Lambda_{z+\lfloor r \rfloor, b_*}^{j',j}| < \frac{1}{2}\tilde{a}_r \sqrt{\tau_{z+\lfloor r \rfloor, b_*}^j} \right\} \mid \tau_{z+\lfloor r \rfloor, b_*}^j \right) \leq \frac{1}{2}e^{-\tilde{a}_r^2/8}. \quad (36)$$

We denote

$$U^j := \bigcap_{j' \in \mathcal{L}^j} \left\{ |\Lambda_{z+\lfloor r \rfloor, b_*}^{j',j}| \geq \frac{1}{2}\tilde{a}_r \sqrt{\tau_{z+\lfloor r \rfloor, b_*}^j} \right\}.$$

Thus, by a union bound, we have

$$\mathbb{P}(\Omega_r^j \cap (U^j)^c) \leq \frac{p}{2}e^{-\tilde{a}_r^2/8}. \quad (37)$$

Moreover, on the event $\Omega_r^j \cap U^j$, we have

$$\begin{aligned} \sum_{j' \in [p]: j' \neq j} \frac{(\Lambda_{z+\lfloor r \rfloor, b_*}^{j',j})^2}{\tau_{z+\lfloor r \rfloor, b_*}^j} \mathbb{1}_{\{|\Lambda_{z+\lfloor r \rfloor, b_*}^{j',j}| \geq \tilde{a} \sqrt{\tau_{z+\lfloor r \rfloor, b_*}^j}\}} &\geq \sum_{j' \in \mathcal{L}^j} \frac{(\Lambda_{z+\lfloor r \rfloor, b_*}^{j',j})^2}{\tau_{z+\lfloor r \rfloor, b_*}^j} \\ &\geq \frac{\tilde{a}_r^2}{4} |\mathcal{L}^j| \geq \frac{\vartheta^2 \lfloor r \rfloor}{24 \log_2(2p)} \geq T^{\text{off}}, \end{aligned} \quad (38)$$

where the penultimate inequality uses the fact that $|\mathcal{L}^j| \geq s-1$ and the last inequality follows from (31). We now denote

$$\tilde{E}_r^j := \left\{ \sum_{j' \in [p]: j' \neq j} \frac{(\Lambda_{z+\lfloor r \rfloor, b_*}^{j',j})^2}{\tau_{z+\lfloor r \rfloor, b_*}^j} \mathbb{1}_{\{|\Lambda_{z+\lfloor r \rfloor, b_*}^{j',j}| \geq \tilde{a} \sqrt{\tau_{z+\lfloor r \rfloor, b_*}^j}\}} < T^{\text{off}} \right\}.$$

By (38), we have $\Omega_r^j \cap \tilde{E}_r^j \subseteq \Omega_r^j \cap (U^j)^c$. Thus, by (32) and (37) we have that

$$\begin{aligned} \mathbb{P}(N > z+r) &\leq \mathbb{P}(N > z + \lfloor r \rfloor) \leq \mathbb{P}\left(\bigcap_{j \in \mathcal{J}} \tilde{E}_r^j\right) \leq \min_{j \in \mathcal{J}} \mathbb{P}(\tilde{E}_r^j) \\ &\leq \min_{j \in \mathcal{J}} \left\{ \mathbb{P}((U^j)^c \cap \Omega_r^j) + \mathbb{P}((\Omega_r^j)^c) \right\} \\ &\leq \frac{p}{2} \exp\left\{-\frac{\vartheta^2(r-1)}{24s \log_2(2p)}\right\} + \exp\left\{-\frac{\beta^2(r-1)}{24s \log_2(2p)}\right\} \\ &\leq p \exp\left\{-\frac{\beta^2(r-1)}{24s \log_2(2p)}\right\}. \end{aligned}$$

as desired. \square

Lemma 11. *Let $U \sim \mathcal{N}(0, \phi_1)$, $V \sim \mathcal{N}(0, \phi_2 - \phi_1)$, $Y \sim \mathcal{N}(\alpha\phi_3, \phi_3)$ and $Z \sim \mathcal{N}(\alpha\phi_4, \phi_4)$ be independent random variables.*

(a) Assume that $\min\{\phi_2, \phi_3\}/4 \geq \kappa \geq \phi_1 \geq 0$ for some $\kappa > 0$. Then

$$\mathbb{P}\left(\frac{|U+V+Y|}{\sqrt{\phi_2+\phi_3}} \geq \frac{|U+Y|}{\sqrt{\phi_1+\phi_3}}\right) \leq \exp\left(-\frac{\kappa\alpha^2}{6}\right).$$

(b) Assume that $\min\{\phi_1, \phi_3\}/4 \geq \kappa \geq \phi_4 \geq 0$ for some $\kappa > 0$. Then

$$\mathbb{P}\left(\frac{|U+Y+Z|}{\sqrt{\phi_1+\phi_3+\phi_4}} \geq \frac{|Y|}{\sqrt{\phi_3}}\right) \leq \exp\left(-\frac{\kappa\alpha^2}{12}\right).$$

Proof. The case $\alpha = 0$ is trivial in both cases, so without loss of generality, we may assume $\alpha > 0$ throughout the rest of the proof.

(a) Let

$$W_1 := (\sqrt{\phi_2+\phi_3} - \sqrt{\phi_1+\phi_3})(U+Y) - \sqrt{\phi_1+\phi_3}V,$$

so that

$$W_1 \sim \mathcal{N}\left(\alpha\phi_3(\sqrt{\phi_2+\phi_3} - \sqrt{\phi_1+\phi_3}), \{(\sqrt{\phi_2+\phi_3} - \sqrt{\phi_1+\phi_3})^2 + \phi_2 - \phi_1\}(\phi_1 + \phi_3)\right).$$

Hence, by the standard Gaussian tail bound used at the end of the proof of Lemma 9, we have

$$\mathbb{P}(W_1 \leq 0) \leq \frac{1}{2}e^{-\alpha^2/(2w_1)}, \quad (39)$$

where $w_1 := \frac{\phi_1+\phi_3}{\phi_3^2} \left(1 + \frac{\phi_2-\phi_1}{(\sqrt{\phi_2+\phi_3}-\sqrt{\phi_1+\phi_3})^2}\right)$. Then

$$\begin{aligned} w_1 &= \frac{\phi_1+\phi_3}{\phi_3^2} \left(1 + \frac{(\sqrt{\phi_2+\phi_3} + \sqrt{\phi_1+\phi_3})^2}{\phi_2-\phi_1}\right) \\ &\leq \frac{5}{16\kappa} \left(1 + \frac{(\sqrt{8\kappa} + \sqrt{5\kappa})^2}{3\kappa}\right) \leq \frac{3}{\kappa}, \end{aligned} \quad (40)$$

where the first inequality holds because w_1 is increasing in ϕ_1 and decreasing in both ϕ_2 and ϕ_3 . Hence, using the fact that $-(U+V+Y) \leq_{\text{st}} U+V+Y$, as well as (39) and (40), we have

$$\begin{aligned} \mathbb{P}\left(\frac{|U+V+Y|}{\sqrt{\phi_2+\phi_3}} \geq \frac{|U+Y|}{\sqrt{\phi_1+\phi_3}}\right) &\leq \mathbb{P}\left(\left\{\frac{U+Y}{\sqrt{\phi_1+\phi_3}} \leq \frac{U+V+Y}{\sqrt{\phi_2+\phi_3}}\right\} \cap \{U+V+Y \geq 0\}\right) \\ &\quad + \mathbb{P}\left(\left\{\frac{U+Y}{\sqrt{\phi_1+\phi_3}} \leq -\frac{U+V+Y}{\sqrt{\phi_2+\phi_3}}\right\} \cap \{U+V+Y < 0\}\right) \\ &\leq 2\mathbb{P}\left(\frac{U+Y}{\sqrt{\phi_1+\phi_3}} \leq \frac{U+V+Y}{\sqrt{\phi_2+\phi_3}}\right) \\ &= 2\mathbb{P}(W_1 \leq 0) \leq \exp\left(-\frac{\kappa\alpha^2}{6}\right), \end{aligned}$$

as required.

(b) Let

$$W_2 := (\sqrt{\phi_1 + \phi_3 + \phi_4} - \sqrt{\phi_3})Y - \sqrt{\phi_3}(U + Z),$$

so that

$$W_2 \sim \mathcal{N}\left(\alpha\phi_3\sqrt{\phi_1 + \phi_3 + \phi_4} - \alpha(\phi_3 + \phi_4)\sqrt{\phi_3}, \left\{(\sqrt{\phi_1 + \phi_3 + \phi_4} - \sqrt{\phi_3})^2 + \phi_1 + \phi_4\right\}\phi_3\right).$$

Note that the assumption guarantees that $\mathbb{E}(W_2) > 0$. Hence, by the standard Gaussian tail bound used at the end of the proof of Lemma 9, we have

$$\mathbb{P}(W_2 \leq 0) \leq \frac{1}{2}e^{-\alpha^2/(2w_2)}, \quad (41)$$

where

$$w_2 := \frac{(\sqrt{\phi_1 + \phi_3 + \phi_4} - \sqrt{\phi_3})^2 + \phi_1 + \phi_4}{(\sqrt{\phi_3(\phi_1 + \phi_3 + \phi_4)} - \phi_3 - \phi_4)^2}.$$

Calculating the partial derivatives of w_2 with respect to ϕ_1, ϕ_3 and ϕ_4 and simplifying the expressions, we have

$$\begin{aligned} \frac{\partial w_2}{\partial \phi_1} &= \frac{(\phi_3 + \phi_4)\sqrt{\phi_3} - (\phi_3 + 2\phi_4)\sqrt{\phi_1 + \phi_3 + \phi_4}}{\sqrt{\phi_1 + \phi_3 + \phi_4}(\sqrt{\phi_3(\phi_1 + \phi_3 + \phi_4)} - \phi_3 - \phi_4)^3} \leq 0, \\ \frac{\partial w_2}{\partial \phi_3} &= \frac{-(\sqrt{\phi_1 + \phi_3 + \phi_4} - \sqrt{\phi_3})^2 [3\phi_1 + \phi_4 + (\sqrt{\phi_1 + \phi_3 + \phi_4} - \sqrt{\phi_3})^2]}{2\sqrt{\phi_3(\phi_1 + \phi_3 + \phi_4)}(\sqrt{\phi_3(\phi_1 + \phi_3 + \phi_4)} - \phi_3 - \phi_4)^3} \leq 0, \\ \frac{\partial w_2}{\partial \phi_4} &= \frac{2\phi_1(2\sqrt{\phi_1 + \phi_3 + \phi_4} - \sqrt{\phi_3})^2 + 3(\phi_3 + \phi_4)(\sqrt{\phi_1 + \phi_3 + \phi_4} - \sqrt{\phi_3})^2}{2(\phi_1 + \phi_3 + \phi_4)(\sqrt{\phi_3(\phi_1 + \phi_3 + \phi_4)} - \phi_3 - \phi_4)^3} \\ &\quad + \frac{(\phi_1 + \phi_4)(\phi_3 + \phi_4)}{2(\phi_1 + \phi_3 + \phi_4)(\sqrt{\phi_3(\phi_1 + \phi_3 + \phi_4)} - \phi_3 - \phi_4)^3} \geq 0. \end{aligned}$$

Thus w_2 is increasing in ϕ_4 and decreasing in both ϕ_1 and ϕ_3 and hence

$$w_2 \leq \frac{6}{\kappa}. \quad (42)$$

Hence, using the fact that $-(U + Y + Z) \leq_{\text{st}} U + Y + Z$, as well as (41) and (42), we have

$$\begin{aligned} \mathbb{P}\left(\frac{|U + Y + Z|}{\sqrt{\phi_1 + \phi_3 + \phi_4}} \geq \frac{|Y|}{\sqrt{\phi_3}}\right) &\leq \mathbb{P}\left(\left\{\frac{Y}{\sqrt{\phi_3}} \leq \frac{U + Y + Z}{\sqrt{\phi_1 + \phi_3 + \phi_4}}\right\} \cap \{U + Y + Z \geq 0\}\right) \\ &\quad + \mathbb{P}\left(\left\{\frac{Y}{\sqrt{\phi_3}} \leq -\frac{U + Y + Z}{\sqrt{\phi_1 + \phi_3 + \phi_4}}\right\} \cap \{U + Y + Z < 0\}\right) \\ &\leq 2\mathbb{P}\left(\frac{Y}{\sqrt{\phi_3}} \leq \frac{U + Y + Z}{\sqrt{\phi_1 + \phi_3 + \phi_4}}\right) \\ &= 2\mathbb{P}(W_2 \leq 0) \leq \exp\left(-\frac{\kappa\alpha^2}{12}\right), \end{aligned}$$

as required. \square

5.3 The ocd and ocd' base procedures

Algorithms 2 and 3, which are taken from [Chen et al. \(2022\)](#), provide two options for a base online changepoint detection procedure. The ocd algorithm is recommended for practical use, while its variant, the ocd' algorithm, satisfies the key condition (3) that underpins our theoretical results in Section 3.

In order to help make the paper self-contained, and to aid interpretability, we remark that the input parameter \tilde{a} represents a thresholding level employed in the definition of the quantity Q_b^j in Algorithm 2 and \tilde{Q}_b^j in Algorithm 3 that is designed to ensure that we only aggregate over signal coordinates. The input parameters T^{diag} and T^{off} represent critical values for the diagonal and off-diagonal statistics S^{diag} and S^{off} respectively, both defined in these algorithms. In other words, we declare a change as soon as either $S^{\text{diag}} \geq T^{\text{diag}}$ or $S^{\text{off}} \geq T^{\text{off}}$.

Algorithm 2: Pseudo-code of the ocd algorithm

Input: $X_1, X_2, \dots \in \mathbb{R}^p$ observed sequentially, $\beta > 0$, $\tilde{a} \geq 0$, $T^{\text{diag}} > 0$ and $T^{\text{off}} > 0$

Set: $b_{\min} = \frac{\beta}{\sqrt{2^{\lfloor \log_2(2p) \rfloor} \log_2(2p)}}$, $\mathcal{B}_0 = \{\pm b_{\min}\}$,

$\mathcal{B} = \{\pm 2^{m/2} b_{\min} : m = 1, \dots, \lfloor \log_2(2p) \rfloor\}$, $n = 0$, $A_b = \mathbf{0} \in \mathbb{R}^{p \times p}$ and $t_b = 0 \in \mathbb{R}^p$
for all $b \in \mathcal{B} \cup \mathcal{B}_0$

repeat

$n \leftarrow n + 1$

 observe new data vector X_n

for $(j, b) \in [p] \times (\mathcal{B} \cup \mathcal{B}_0)$ **do**

$t_b^j \leftarrow t_b^j + 1$

$A_b^{:,j} \leftarrow A_b^{:,j} + X_n$

if $bA_b^{j,j} - b^2 t_b^j / 2 \leq 0$ **then**

$t_b^j \leftarrow 0$ and $A_b^{:,j} \leftarrow 0$

 compute $Q_b^j \leftarrow \sum_{j' \in [p]: j' \neq j} \frac{(A_b^{j',j})^2}{t_b^j \vee 1} \mathbb{1}_{\{|A_b^{j',j}| \geq \tilde{a} \sqrt{t_b^j}\}}$

$S^{\text{diag}} \leftarrow \max_{(j,b) \in [p] \times (\mathcal{B} \cup \mathcal{B}_0)} (bA_b^{j,j} - b^2 t_b^j / 2)$

$S^{\text{off}} \leftarrow \max_{(j,b) \in [p] \times \mathcal{B}} Q_b^j$

until $S^{\text{diag}} \geq T^{\text{diag}}$ or $S^{\text{off}} \geq T^{\text{off}}$;

Output: $N = n$

5.4 Results under sub-Gaussian and sub-exponential assumptions

This section provides justification for the claimed theoretical results in Section 3.4. We will rely on the following three propositions, the first of which is standard (e.g. [Wainwright, 2019](#), Proposition 2.5 and (2.18)).

Proposition 12 (Hoeffding-type and Bernstein-type tail bound). (a) Let $a_1, \dots, a_n \in \mathbb{R}$ and let X_1, \dots, X_n be independent sub-Gaussian random variables with variance parameter 1.

Algorithm 3: Pseudo-code for the ocd' algorithm, a slight variant of ocd

Input: $X_1, X_2, \dots \in \mathbb{R}^p$ observed sequentially, $\beta > 0$, $\tilde{a} \geq 0$, $T^{\text{diag}} > 0$ and $T^{\text{off}} > 0$

Set: $b_{\min} = \frac{\beta}{\sqrt{2^{\lfloor \log_2(2p) \rfloor} \log_2(2p)}}$, $\mathcal{B}_0 = \{\pm b_{\min}\}$,

$\mathcal{B} = \{\pm 2^{m/2} b_{\min} : m = 1, \dots, \lfloor \log_2(2p) \rfloor\}$, $n = 0$, $A_b = \Lambda_b = \tilde{\Lambda}_b = \mathbf{0} \in \mathbb{R}^{p \times p}$ and $t_b = \tau_b = \tilde{\tau}_b = 0 \in \mathbb{R}^p$ for all $b \in \mathcal{B} \cup \mathcal{B}_0$

repeat

$n \leftarrow n + 1$

observe new data vector X_n

for $(j, b) \in [p] \times (\mathcal{B} \cup \mathcal{B}_0)$ **do**

$t_b^j \leftarrow t_b^j + 1$ and $A_b^{\cdot j} \leftarrow A_b^{\cdot j} + X_n$

set $\delta = 0$ if t_b^j is a power of 2 and $\delta = 1$ otherwise.

$\tau_b^j \leftarrow \tau_b^j \delta + \tilde{\tau}_b^j (1 - \delta) + 1$ and $\Lambda_b^{\cdot j} \leftarrow \Lambda_b^{\cdot j} \delta + \tilde{\Lambda}_b^{\cdot j} (1 - \delta) + X_n$

$\tilde{\tau}_b^j \leftarrow (\tilde{\tau}_b^j + 1) \delta$ and $\tilde{\Lambda}_b^{\cdot j} \leftarrow (\tilde{\Lambda}_b^{\cdot j} + X_n) \delta$.

if $bA_b^{j,j} - b^2 t_b^j / 2 \leq 0$ **then**

$t_b^j \leftarrow \tau_b^j \leftarrow \tilde{\tau}_b^j \leftarrow 0$

$A_b^{\cdot j} \leftarrow \Lambda_b^{\cdot j} \leftarrow \tilde{\Lambda}_b^{\cdot j} \leftarrow 0$

$\Xi_b^{\cdot j} \leftarrow \Lambda_b^{\cdot j} / (\tau_b^j \vee 1)^{1/2}$

$\tilde{Q}_b^j \leftarrow \sum_{j' \in [p] \setminus \{j\}} (\Xi_b^{j',j})^2 \mathbb{1}_{\{|\Xi_b^{j',j}| \geq \tilde{a}\}}$

$S^{\text{diag}} \leftarrow \max_{(j,b) \in [p] \times (\mathcal{B} \cup \mathcal{B}_0)} (bA_b^{j,j} - b^2 t_b^j / 2)$

$S^{\text{off}} \leftarrow \max_{(j,b) \in [p] \times \mathcal{B}} \tilde{Q}_b^j$

until $S^{\text{diag}} \geq T^{\text{diag}}$ or $S^{\text{off}} \geq T^{\text{off}}$;

Output: $N = n$

Then

$$\mathbb{P}\left(\sum_{i=1}^n a_i X_i \geq x\right) \leq \exp\left(-\frac{x^2}{2\sum_{i=1}^n a_i^2}\right)$$

for all $x \geq 0$.

(b) Let $a_1, \dots, a_n \in \mathbb{R}$ and let X_1, \dots, X_n be independent sub-exponential random variables with variance parameter 1 and rate parameter $A > 0$. Then

$$\mathbb{P}\left(\sum_{i=1}^n a_i X_i \geq x\right) \leq \exp\left\{-\min\left(\frac{x^2}{2\sum_{i=1}^n a_i^2}, \frac{Ax}{2\max_{i \in [n]} |a_i|}\right)\right\}$$

for all $x \geq 0$.

One special case where we apply this proposition frequently is with $a_1 = \dots = a_n = n^{-1/2}$. The two bounds are then $e^{-x^2/2}$ and $e^{-x(x \wedge A\sqrt{n})/2}$ respectively.

The following proposition can be used in place of Lemma 9 to control excursion probabilities of sub-Gaussian and sub-exponential random walks with drift.

Proposition 13. Let $\mu \in \mathbb{R}$ and let Y_1, Y_2, \dots be independent random variables. Define $U_n := \sum_{i=1}^n Y_i$ for $n \in \mathbb{N}$ with $U_0 := 0$, and let $\xi := \operatorname{sargmin}_{n \in \mathbb{N}_0} \mu U_n$.

(a) Assume that $Y_1 - \mu, Y_2 - \mu, \dots$ are independent sub-Gaussian random variables with variance parameter 1. Then

$$\mathbb{P}(\xi \geq y) \leq \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} \mu U_n \leq 0\right) \leq 3(\mu^{-2} \vee 1)e^{-y\mu^2/2}$$

for $y \in [0, \infty)$.

(b) Assume that $Y_1 - \mu, Y_2 - \mu, \dots$ are independent sub-exponential random variables with variance parameter 1 and rate parameter $A > 0$. Then

$$\mathbb{P}(\xi \geq y) \leq \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} \mu U_n \leq 0\right) \leq 3\left(\frac{1}{\mu^2} \vee \frac{1}{\mu A} \vee 1\right)e^{-y\mu(\mu \wedge A)/2}$$

for $y \in [0, \infty)$.

Proof. Following the same argument at the beginning of the proof of Lemma 9, it suffices to only prove the latter inequality for both results for $\mu > 0$.

(a) By a union bound and Proposition 12(a), we have

$$\begin{aligned} \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} \mu U_n \leq 0\right) &\leq \sum_{n=\lceil y \rceil}^{\infty} \mathbb{P}(U_n \leq 0) = \sum_{n=\lceil y \rceil}^{\infty} \mathbb{P}(U_n - n\mu \leq -n\mu) \leq \sum_{n=\lceil y \rceil}^{\infty} e^{-n\mu^2/2} \\ &\leq \frac{e^{-y\mu^2/2}}{1 - e^{-\mu^2/2}} \leq \frac{e^{-y\mu^2/2}}{\frac{\mu^2}{4 \log 2} \wedge \frac{1}{2}} \leq 3(\mu^{-2} \vee 1)e^{-y\mu^2/2}. \end{aligned}$$

(b) Again, by a union bound and Proposition 12(b), we have

$$\begin{aligned} \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} \mu U_n \leq 0\right) &\leq \sum_{n=\lceil y \rceil}^{\infty} \mathbb{P}(U_n - n\mu \leq -n\mu) \leq \sum_{n=\lceil y \rceil}^{\infty} e^{-n\mu(\mu \wedge A)/2} \\ &\leq 3\left(\frac{1}{\mu^2} \vee \frac{1}{\mu A} \vee 1\right)e^{-y\mu(\mu \wedge A)/2}, \end{aligned}$$

where the last inequality follows from the last three inequalities in the proof for the (a) part, with $\mu(\mu \wedge A)$ taking the place of μ^2 . \square

Our final preparatory result will be used to modify (23) in the proof of Proposition 5, which can be traced back to the proof of Chen et al. (2022, Theorem 1). Let $b \neq 0$, $T^{\text{diag}} > 0$ and let Z_1, Z_2, \dots be independent centered random variables. We define the stopping time

$$N_{\text{os}} := \inf \left\{ n \in \mathbb{N} : b \sum_{t=1}^n (Z_t - b/2) \geq T^{\text{diag}} \right\}, \quad (43)$$

where ‘os’ stands for one-sided. When $Z_1, Z_2, \dots \stackrel{\text{iid}}{\sim} N(0, 1)$, we can use a sequential probability ratio test argument to show that $\mathbb{P}(N_{\text{os}} < \infty) \leq e^{-T^{\text{diag}}}$. The aim of Proposition 14 below is to establish similar bounds under the sub-Gaussian and the sub-exponential distributional assumptions respectively.

Proposition 14. (a) Let $b \neq 0$, $T^{\text{diag}} > 0$ and let Z_1, Z_2, \dots be independent sub-Gaussian random variables each with variance parameter 1. Then N_{os} defined in (43) satisfies

$$\mathbb{P}(N_{\text{os}} < \infty) \leq e^{-T^{\text{diag}}}.$$

(b) Let $T^{\text{diag}} > 0$ and Z_1, Z_2, \dots be independent sub-exponential random variables with variance parameter 1 and rate parameter $A > 0$. Let $b \in [-A, A] \setminus \{0\}$. Then

$$\mathbb{P}(N_{\text{os}} < \infty) \leq e^{-T^{\text{diag}}}.$$

Proof. Without loss of generality, assume $b > 0$ in part (a) and $0 < b < A$ in part (b). The following argument then holds for both cases. Denote $S_n := \sum_{t=1}^n Z_t$ with $S_0 := 0$, $V_n := \exp\{(bS_n - b^2n/2) \wedge T^{\text{diag}}\}$ and $\mathcal{F}_n = \sigma(Z_1, \dots, Z_n)$ with \mathcal{F}_0 defined to be trivial σ -algebra. Then for $n \in \mathbb{N}$, we have

$$\begin{aligned} \mathbb{E}[V_n \mid \mathcal{F}_{n-1}] &= \mathbb{E}[V_n \mathbb{1}_{\{V_{n-1} < e^{T^{\text{diag}}}\}} \mid \mathcal{F}_{n-1}] + \mathbb{E}[V_n \mathbb{1}_{\{V_{n-1} = e^{T^{\text{diag}}}\}} \mid \mathcal{F}_{n-1}] \\ &\leq \mathbb{E}[V_{n-1} e^{bZ_n - b^2/2} \mathbb{1}_{\{V_{n-1} < e^{T^{\text{diag}}}\}} \mid \mathcal{F}_{n-1}] + \mathbb{E}[e^{T^{\text{diag}}} \mathbb{1}_{\{V_{n-1} = e^{T^{\text{diag}}}\}} \mid \mathcal{F}_{n-1}] \\ &\leq V_{n-1} \mathbb{1}_{\{V_{n-1} < e^{T^{\text{diag}}}\}} + e^{T^{\text{diag}}} \mathbb{1}_{\{V_{n-1} = e^{T^{\text{diag}}}\}} = V_{n-1}. \end{aligned}$$

Hence $(V_n)_{n \geq 0}$ is a supermartingale with respect to the filtration $(\mathcal{F}_n)_{n \geq 0}$. Since $|V_n| \leq e^{T^{\text{diag}}}$ for all n , we have by the Optional Stopping Theorem that

$$1 = \mathbb{E}V_0 \geq \mathbb{E}V_{N_{\text{os}}} = \mathbb{E}[V_{N_{\text{os}}} \mathbb{1}_{\{N_{\text{os}} < \infty\}} + V_{N_{\text{os}}} \mathbb{1}_{\{N_{\text{os}} = \infty\}}] \geq e^{T^{\text{diag}}} \mathbb{P}(N_{\text{os}} < \infty),$$

as required. \square

The modifications to the proofs of Theorems 1, 2 and 3, as well as Proposition 5 that are needed in the sub-Gaussian and sub-exponential settings are as follows. In (7), (11), (12), (14), (15), (16), (17), (25), (36), (39) and (41), we apply Proposition 12 in place of the Gaussian tail bounds; in (8), (10) and (32), we apply Proposition 13 in place of Lemma 9; and finally, we apply Proposition 14 to yield the bound corresponding to (23).

5.5 Additional simulation results

Table 3 provides additional simulation results for the `ocd_CI` procedure under spatial dependence. The data generating mechanisms and conclusions from these results are given in Section 4.2.

Table 3: Spatial dependence. Estimated coverage and average length of the `ocd_CI` confidence interval and average detection delay over 2000 repetitions, with standard errors in brackets, under a Toeplitz cross-sectional covariance matrix Σ with entries $\Sigma_{jk} = \rho^{|j-k|}$ for $j, k \in [p]$. Other parameters: $p = 100$, $\beta = \vartheta$, $\gamma = 30000$, $z = 1000$, $\alpha = 0.05$, $a = \tilde{a} = \sqrt{2 \log p}$, $c = 0.5$, $d_1 = c\sqrt{\log(p/\alpha)}$, $d_2 = 4d_1^2$.

ρ	s	ϑ	Detection Delay	Coverage (%)	CI Length
0.5	2	2	13.9 _(0.1)	98.5 _(0.3)	35.5 _(1.0)
0.5	2	1	49.1 _(0.3)	99.0 _(0.2)	125.1 _(1.6)
0.5	2	0.5	172.5 _(1.0)	99.5 _(0.2)	447.0 _(2.8)
0.5	10	2	21.9 _(0.1)	98.7 _(0.3)	42.0 _(0.9)
0.5	10	1	76.1 _(0.5)	98.8 _(0.2)	154.2 _(1.5)
0.5	10	0.5	266.7 _(1.8)	99.0 _(0.2)	566.9 _(3.9)
0.5	100	2	52.1 _(0.3)	98.3 _(0.3)	106.8 _(0.9)
0.5	100	1	187.7 _(1.3)	98.4 _(0.3)	399.5 _(3.3)
0.5	100	0.5	655.3 _(5.0)	98.5 _(0.3)	1366.2 _(10.1)
0.75	2	2	13.9 _(0.1)	96.9 _(0.4)	51.1 _(2.6)
0.75	2	1	47.9 _(0.3)	96.8 _(0.4)	146.0 _(3.3)
0.75	2	0.5	171.5 _(1.1)	97.7 _(0.3)	463.8 _(4.2)
0.75	10	2	21.8 _(0.2)	96.4 _(0.4)	48.6 _(1.7)
0.75	10	1	75.3 _(0.5)	96.7 _(0.4)	165.0 _(2.7)
0.75	10	0.5	266.3 _(1.9)	96.0 _(0.4)	558.9 _(4.5)
0.75	100	2	50.9 _(0.3)	96.8 _(0.4)	106.8 _(1.2)
0.75	100	1	184.8 _(1.4)	95.6 _(0.5)	401.8 _(3.8)
0.75	100	0.5	647.3 _(5.4)	94.6 _(0.5)	1312.3 _(11.2)

Acknowledgements: The research of TW was supported by Engineering and Physical Sciences Research Council (EPSRC) grant EP/T02772X/1 and that of RJS was supported by EPSRC grants EP/P031447/1 and EP/N031938/1, as well as European Research Council Advanced Grant 101019498.

References

- Chan, H. P. and Walther, G. (2015) Optimal Detection of Multi-sample Aligned Sparse Signals. *Ann. Statist.*, **43**, 1865–1895.
- Chen, Y., Wang, T. and Samworth, R. J. (2022) High-dimensional, Multiscale Online Change-point Detection. *J. Roy. Statist. Soc., Ser. B*, **84**, 234–266.
- Cho, H. (2016) Change-point Detection in Panel Data via Double CUSUM Statistic. *Electron. J. Statist.*, **10**, 2000–2038.

- 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.
- Csörgő, M. and Horváth, L. (1997) *Limit Theorems in Change-Point Analysis*. John Wiley and Sons, New York.
- Dette, H. and Gösmann, J. (2020) A Likelihood Ratio Approach to Sequential Change Point Detection for a General Class of Parameters. *J. Amer. Statist. Assoc.*, **115**, 1361–1377.
- Duncan, A. J. (1952) *Quality Control and Industrial Statistics*. Richard D. Irwin Professional Publishing Inc., Chicago.
- Enikeeva, F. and Harchaoui, Z. (2019) High-dimensional Change-point Detection Under Sparse Alternatives. *Ann. Statist.*, **47**, 2051–2079.
- Follain, B., Wang, T. and Samworth, R. J. (2022) High-dimensional Change-point Estimation With Heterogeneous Missingness. *J. Roy. Statist. Soc., Ser. B*, **84**, 1023–1055.
- Gösmann, J., Stoehr, C., Heiny, J. and Dette, H. (2020) Sequential Change Point Detection in High Dimensional Time Series. *arXiv preprint*, arxiv:2006.00636.
- Horváth, L. and Rice, G. (2014) Extensions of Some Classical Methods in Change Point Analysis. *TEST*, **23**, 219–255.
- Janková, J. and van de Geer, S. (2015) Confidence Intervals for High-dimensional Inverse Covariance Estimation. *Electron. J. Statist.*, **9**, 1205–1229.
- Javanmard, A. and Montanari, A. (2014) Confidence Intervals and Hypothesis Testing for High-dimensional Regression. *J. Mach. Learn. Res.*, **15**, 2869–2909.
- Jirak, M. (2015) Uniform Change Point Tests in High Dimension. *Ann. Statist.*, **43**, 2451–2483.
- Kaul, A., Fotopoulos, S. B., Jandhyala, V. K. and Safikhani, A. (2021) Inference on the Change Point Under a High Dimensional Sparse Mean Shift. *Electron. J. Statist.*, **15**, 71–134.
- Kirch, C. and Stoehr, C. (2019) Sequential Change Point Tests Based on U -statistics. *arXiv preprint*, arxiv:1912.08580.
- Liu, H., Gao, C. and Samworth, R. J. (2021) Minimax Rates in Sparse, High-dimensional Change-point Detection. *Ann. Statist.*, **49**, 1081–1112.
- Londschien, M., Kovács, S. and Bühlmann, P. (2021) Change-point Detection for Graphical Models in the Presence of Missing Values. *J. Comput. Graph. Statist.*, **30**, 1–12.
- Massart, P. (2007) *Concentration inequalities and model selection: Ecole d’Eté de Probabilités de Saint-Flour XXXIII-2003*. Springer.

- Padilla, O. H. M., Yu, Y., Wang, D. and Rinaldo, A. (2021) Optimal nonparametric multivariate change point detection and localization. *IEEE Trans. Inform. Theory*, **68**, 1922–1944.
- Page, E. S. (1954) Continuous Inspection Schemes. *Biometrika*, **41**, 100–115.
- Rinaldo, A., Wang, D., Wen, Q., Willett, R. and Yu, Y. (2021) Localizing Changes in High-dimensional Regression Models. *Proc. Mach. Learn. Res.*, 2089–2097.
- Siegmund, D. (1986) Boundary Crossing Probabilities and Statistical Applications. *Ann. Statist.*, **14**, 361–404.
- Soh, Y. S. and Chandrasekaran, V. (2017) High-dimensional Change-point Estimation: Combining Filtering With Convex Optimization. *Applied and Computational Harmonic Analysis*, **43**, 122–147.
- Tartakovsky, A., Nikiforov, I. and Basseville, M. (2014) *Sequential Analysis: Hypothesis testing and Changepoint Detection*. Chapman and Hall, London.
- van de Geer, S., Bühlmann, P., Ritov, Y. and Dezeure, R. (2014) On Asymptotically Optimal Confidence Regions and Tests for High-dimensional Models. *Ann. Statist.*, **42**, 1166–1202.
- Wainwright, M. (2019) *High-dimensional Statistics: A Non-asymptotic Viewpoint*. Cambridge University Press, Cambridge.
- Wang, T. and Samworth, R. J. (2018) High Dimensional Change Point Estimation via Sparse Projection. *J. Roy. Statist. Soc., Ser. B*, **80**, 57–83.
- Yu, B. (1997) Assouad, Fano and Le Cam. In D. Pollard, E. Torgersen and Y. G. L. (eds.), *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*, 423–435. Springer, New York.
- Yu, Y., Bradic, J. and Samworth, R. J. (2021a) Confidence Intervals for High-dimensional Cox Models. *Statist. Sinica*, **31**, 243–267.
- Yu, Y., Padilla, O. H. M., Wang, D. and Rinaldo, A. (2021b) Optimal Network Online Change Point Localisation. *arXiv preprint*, arxiv:2021.05477.
- Zhang, C.-H. and Zhang, S. S. (2014) Confidence Intervals for Low Dimensional Parameters in High Dimensional Linear Models. *J. Roy. Statist. Soc., Ser. B*, **76**, 217–242.