

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

(November 2, 2021)

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 both US excess deaths data from 2017–2020 and S&P 500 data from the 2007–2008 financial crisis.

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., 2019; Kaul et al., 2021; Follain et al., 2021; Liu et al., 2021; Londschien et al., 2021; Rinaldo et al., 2021). 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; Chen et al., 2021; Yu et al., 2021b). 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 (generalised) 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 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 $n \in \mathbb{N}$, let²

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

Since $\sum_{i=n-h+1}^n (X_i - b/2)$ can be viewed as the likelihood ratio statistic for testing the null

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

²In the case of a tie, we choose the smallest h achieving the minimum.

of $\mathcal{N}(0, 1)$ against the 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 maximised. 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.

In terms of the base sequential changepoint detection procedures, we focus on the `ocd` algorithm (short for `online changepoint detection`) introduced by Chen et al. (2021), as well as 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. Our confidence intervals, which we correspondingly denote `ocd_CI` and `ocd_CI'`, inherit this same online property, thereby making them applicable even in very high-dimensional settings and where changes may be rare, so that 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, for a suitable choice of input parameters, the confidence interval has at least nominal coverage. Moreover, Theorem 2 ensures that, with high probability, its length is of the same order as the average detection delay for the base `ocd'` procedure, up to a 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 avoids noise coordinates.

Section 4 is devoted to a study of the numerical performance of our methodological proposals. Our simulations confirm that the `ocd_CI` methodology 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. Moreover, in Sections 4.4.1 and 4.4.2, we illustrate the practical utility of our methods by applying them to both excess death data during the Covid-19 pandemic in the US and S&P 500 data during the 2007–2008 financial crisis.

Proofs are given in Section 5, with auxiliary results deferred to Section 6. An R implementation of our methodology is available at github.com/yudongchen88/ocd_CI.

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 $\text{largmax}_{x \in S} f(x) := \max \arg \max_{x \in S} f(x)$ for the smallest and largest maximisers 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 $A^{\cdot,j} := (A^{1,j}, \dots, A^{d_1,j})^\top \in \mathbb{R}^{d_1}$ and $A^{-j,\cdot} := (A^{1,j}, \dots, A^{j-1,j}, A^{j+1,j}, \dots, A^{d_1,j})^\top \in \mathbb{R}^{d_1-1}$. 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 time³ N by which we believe a change has occurred. 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 algorithm of [Chen et al. \(2021\)](#), which forms part of Algorithm 1, relies on a lower bound $\beta > 0$ for the ℓ_2 -norm of the vector of mean change and sets of signed scales \mathcal{B} and \mathcal{B}_0 defined in terms of β . From our perspective, the key aspects of this multiscale algorithm are that, in addition to returning a stopping time N as output, it produces a matrix of residual tail lengths $(t_{N,b}^j)_{j \in [p], b \in \mathcal{B} \cup \mathcal{B}_0}$ with $t_{N,b}^j := \text{sargmax}_{0 \leq h \leq N} \sum_{i=N-h+1}^N (X_i^j - b/2)$ (similarly to (1)), an ‘anchor’ coordinate $\hat{j} \in [p]$, a signed anchor scale $\hat{b} \in \mathcal{B}$ and a tail partial sum vector $A_{N,\hat{b}}^{\cdot,\hat{j}} \in \mathbb{R}^p$ with j th component $A_{N,\hat{b}}^{j,\hat{j}} := \sum_{i=N-t_{N,\hat{b}}^{\hat{j}}+1}^N X_i^j$. The intuition is that the anchor coordinate and signed 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, $A_{N,\hat{b}}^{\cdot,\hat{j}}$ aggregates the last $t_{N,\hat{b}}^{\hat{j}}$ observations in each coordinate, thereby 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 centred close to $\theta^j (t_{N,\hat{b}}^{\hat{j}})^{1/2}$ for $j \in [p] \setminus \{\hat{j}\}$,

³Here 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 .

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 be normally distributed around this centering value, with unit variance. 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 , and therefore plays the role of the lower bound b from the univariate problem discussed in the introduction. Finally, then, our confidence interval can be 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 .

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 are $O(p^2 \log(ep))$ regardless of the observation history, so it satisfies the condition to be an online algorithm, as discussed in the introduction.

2.1 A slight variant of the `ocd_CI` algorithm

While our experience is that `ocd_CI` performs very well empirically, in our theoretical analysis it turns out to be easier to study a slight variant of this algorithm, denoted `ocd_CI'`. There are two main differences between the algorithms. First, in `ocd_CI`, the base changepoint detection procedure is `ocd`, while in `ocd_CI'`, we use the `ocd'` procedure of Chen et al. (2021) instead. This latter algorithm is designed to avoid difficulties caused by adversarial pre-change observations that may lead to lengthy response delays for the `ocd` procedure. In particular, for each $j \in [p]$ and $b \in \mathcal{B}$, instead of using the final $t_{n,b}^j$ observations at time n to construct test statistics based on $A_{n,b}^j$, the `ocd'` procedure aggregates over a reduced number $\tau_{n,b}^j$ of observations to obtain test statistics based on $\Lambda_{n,b}^j$, where $\tau_{n,b}^j$ is constructed in an online manner to lie in the interval $[t_{n,b}^j/2, 3t_{n,b}^j/4]$ for $t_{n,b}^j \geq 2$. Even though the reduced tail lengths may lead to a slight deterioration in empirical performance, provided no change has been declared by time z , they guarantee that from a later time of the form $z + O(b^{-2})$, the last $\tau_{n,b}^j$ observations consist entirely of post-change data.

Second, in `ocd_CI'`, we allow the practitioner to observe 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

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

Input: $X_1, X_2, \dots \in \mathbb{R}^p$ observed sequentially, $\beta > 0$, $a \geq 0$, $T^{\text{diag}}, T^{\text{off}} > 0$, $d_1, d_2 > 0$

Set: $b_{\min} = \frac{\beta}{\sqrt{2^{\lceil \log_2(2p) \rceil} \log_2(2p)}}$, $\mathcal{B}_0 = \{\pm b_{\min}\}$, $\mathcal{B} = \{\pm 2^{\ell/2} b_{\min} : \ell = 1, \dots, \lceil \log_2(2p) \rceil\}$,

$n = 0$, $A_b = \mathbf{0} \in \mathbb{R}^{p \times p}$ 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 $A_b^{:,j} \leftarrow 0$

$E_b^{:,j} \leftarrow A_b^{:,j} / (t_b^j \vee 1)^{1/2}$

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

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

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

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

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

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

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

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

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

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 `ocd.CI'` 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 the `ocd.CI'` algorithm is given in Algorithm 2; its computational complexity per new observation, and storage requirements, remain $O(p^2 \log(ep))$.

3 Theoretical analysis

Throughout this section, we will assume that the sequential observations X_1, X_2, \dots are independent, and that there exist $z \in \mathbb{N}_0$ and $\theta = (\theta^1, \dots, \theta^p)^\top \neq 0$ for which $X_1, \dots, X_z \sim \mathcal{N}_p(0, I_p)$ and $X_{z+1}, X_{z+2}, \dots \sim \mathcal{N}_p(\theta, I_p)$. We let $\vartheta := \|\theta\|_2$, and write $\mathbb{P}_{z,\theta}$ 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 $s \in \{2^0, 2^1, \dots, 2^{\lfloor \log_2(p) \rfloor}\}$ 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.

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.

Theorem 1. *Let $p \geq 2$. Fix $\alpha \in (0, 1)$ and $\gamma \geq 1$ and assume that $z \leq 2\alpha\gamma$. Then there exist universal constants $C_1, C_2 > 0$, such that with inputs $(X_t)_{t \in \mathbb{N}}$, $0 < \beta \leq \vartheta$, $a = C_1 \sqrt{\log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$, $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, $\ell \geq 0$, $d_1 = C_2 a$ and $d_2 = 4d_1^2$ in Algorithm 2, the output confidence interval \mathcal{C} satisfies*

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

As mentioned in Section 2.1, our coverage guarantee in Theorem 1 holds even with $\ell = 0$, i.e. with no additional sampling. The condition $z \leq 2\alpha\gamma$ ensures that the probability of a false alarm is at most $\alpha/2$, so that $\mathbb{P}_{z,\theta}(N \leq z) \leq \alpha/2$.

We now provide a guarantee on the length of the `ocd.CI'` confidence interval.

Theorem 2. *Assume that θ has an effective sparsity of $s := s(\theta) \geq 2$. Fix $\alpha \in (0, 1)$ and $\gamma \geq 1$, and assume that $z \leq 2\alpha\gamma$. Then there exist universal constants $C_1, C_2, C_3, C_4 > 0$ such that, with inputs $(X_t)_{t \in \mathbb{N}}$, $0 < \beta \leq \vartheta$, $a = C_1 \sqrt{\log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$,*

Algorithm 2: Pseudo-code for the ocd_CI' algorithm, a slight variant of ocd_CI

Input: $X_1, X_2 \dots \in \mathbb{R}^p$ observed sequentially, $\beta > 0$, $a \geq 0$, $T^{\text{diag}}, T^{\text{off}} > 0$, $d_1, d_2 > 0$ and $\ell \in \mathbb{N}_0$

Set: $b_{\min} = \frac{\beta}{\sqrt{2^{\lfloor \log_2(2p) \rfloor} \log_2(2p)}}$, $\mathcal{B}_0 = \{\pm b_{\min}\}$, $\mathcal{B} = \{\pm 2^{\ell/2} b_{\min} : \ell = 1, \dots, \lfloor \log_2(2p) \rfloor\}$,
 $n = 0$, $A_b = \Lambda_b = \tilde{\Lambda}_b = \mathbf{0} \in \mathbb{R}^{p \times p}$ 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 $A_b^{:,j} \leftarrow A_b^{:,j} + X_n$

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 $\Lambda_b^{:,j} \leftarrow \Lambda_b^{:,j} \delta + \tilde{\Lambda}_b^{:,j} (1 - \delta) + X_n$

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

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

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

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

$E_b^{:,j} \leftarrow \Lambda_b^{:,j} / (\tau_b^j \vee 1)^{1/2}$

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

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

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

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

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

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

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

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

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

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

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

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

$T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, $d_1 = C_2 a$, $d_2 = 4d_1^2$ and $\ell \geq C_3 \left(\frac{a^2 s \log_2(2p)}{\beta^2} + 1\right)$ in Algorithm 2, the length L of the output confidence interval \mathcal{C} satisfies

$$\mathbb{P}_{z,\theta} \left\{ L > C_4 \left(\frac{s \log_2(2p) \log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}{\beta^2} + 1 \right) \right\} \leq \alpha. \quad (2)$$

The main conclusion of Theorem 2 is that, with high probability, the length of the confidence interval is of the same order, up to a logarithmic factor, as the average detection delay guarantee for the `ocd'` procedure (Chen et al., 2020, Theorem 4). Note that the choices of inputs in Theorem 2 are identical to those in Theorem 1, except that we now ask for some additional observations after the changepoint declaration, the number of which is of the same order of magnitude as the length of the interval.

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 2, 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 2. 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. *Assume the conditions of Theorem 1.*

(a) *There exist universal constants $C_1, C_2 > 0$, such that with inputs $(X_t)_{t \in \mathbb{N}}$, $0 < \beta \leq \vartheta$, $a = C_1 \sqrt{\log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$, $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, $\ell \geq 0$ and $d_1 = C_2 a$ in Algorithm 2, we have*

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

(b) *Assume further that θ has effective sparsity $s := s(\theta) \geq 2$. There exist universal constants $C_1, C_2, C_3 > 0$ such that, with inputs $(X_t)_{t \in \mathbb{N}}$, $0 < \beta \leq \vartheta$, $a = C_1 \sqrt{\log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$, $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, $d_1 = C_2 a$ and $\ell \geq C_3 \left(\frac{a^2 s \log_2(2p)}{\beta^2} + 1\right)$ in Algorithm 2, we have*

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

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} .

4 Numerical studies

In this section, we study the empirical performance of the `ocd.CI` algorithm. Recall that in `ocd.CI`, the off-diagonal statistics Q_b^j are computed using tail partial sums of length t_b^j and that we do not have any extra sampling beyond the time of declaration that a change has occurred.

4.1 Tuning parameters

Chen et al. (2021) 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. (2021), we take $a = \sqrt{2 \log p}$ in our simulations.

This leaves us with the choice of tuning parameters d_1 and d_2 . As suggested by Theorems 1 and 2, we take $d_2 = 4d_1^2$. Finally, again as suggested by our theory, we take d_1 to be of the form $d_1 = c\sqrt{\log(p/\alpha)}$, and then tune the parameter $c > 0$ through Monte Carlo simulation, as we now describe. We considered the parameter settings $p \in \{100, 500\}$, $s \in \{2, \lfloor \sqrt{p} \rfloor, p\}$, $\vartheta \in \{2, 1, 1/2\}$, $\alpha = 0.05$, $\beta \in \{2\vartheta, \vartheta, \vartheta/2\}$, $\gamma = 30000$ 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.




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

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 significantly 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. Theorem 2, as well as Chen et al. (2021, Theorem 4), indicate 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 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.

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 $p = 100$, $s \in \{5, 50\}$, $\vartheta \in \{1, 2\}$, and for three different signal shapes: in the uniform, inverse square root and harmonic cases, we took $\theta \propto (\mathbb{1}_{\{j \in [s]\}})_{j \in [p]}$, $\theta \propto (j^{-1/2} \mathbb{1}_{\{j \in [s]\}})_{j \in [p]}$ and $\theta \propto (j^{-1} \mathbb{1}_{\{j \in [s]\}})_{j \in [p]}$ respectively. As inputs to the algorithm, we set $a = \sqrt{2 \log p}$, $\alpha = 0.05$, $d_1 = \sqrt{2 \log(p/\alpha)}$, $\beta = \vartheta$, and, motivated by Theorem 3, took an additional $\ell = \lceil a^2 s \beta^{-2} \log_2(2p) \rceil$ post-declaration observations in constructing the support estimates. The results reported in Table 2 provide empirical confirmation of the support recovery properties claimed in Theorem 3.

Finally in this section, we consider the extent to which the additional 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 $\ell = 0$ (i.e., no additional sampling) 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 Theorem 3(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

p	s	ϑ	β	Detection Delay	Coverage (%)	CI Length
100	2	2	4	9.8 _(0.1)	96.2 _(0.4)	20.1 _(0.7)
100	2	2	2	12.6 _(0.1)	97.0 _(0.4)	33.7 _(0.7)
100	2	2	1	14.1 _(0.1)	97.9 _(0.3)	80.8 _(1.0)
100	2	1	2	34.2 _(0.3)	95.8 _(0.4)	66.1 _(1.0)
100	2	1	1	44.2 _(0.3)	97.5 _(0.4)	122.0 _(1.4)
100	2	1	0.5	52.0 _(0.4)	97.4 _(0.4)	309.1 _(2.0)
100	10	2	4	14.7 _(0.1)	96.0 _(0.4)	32.5 _(0.8)
100	10	2	2	15.7 _(0.1)	97.4 _(0.4)	38.4 _(0.8)
100	10	2	1	15.9 _(0.1)	97.0 _(0.4)	80.2 _(1.1)
100	10	1	2	52.6 _(0.5)	96.2 _(0.4)	114.0 _(1.5)
100	10	1	1	56.9 _(0.4)	97.1 _(0.4)	142.5 _(1.8)
100	10	1	0.5	60.2 _(0.4)	98.2 _(0.3)	301.1 _(1.6)
100	100	2	4	27.2 _(0.2)	96.1 _(0.4)	77.6 _(0.9)
100	100	2	2	27.7 _(0.2)	96.0 _(0.4)	81.8 _(1.0)
100	100	2	1	28.2 _(0.2)	97.5 _(0.3)	99.4 _(1.3)
100	100	1	2	100.7 _(0.8)	94.7 _(0.5)	292.8 _(3.5)
100	100	1	1	100.5 _(0.9)	96.3 _(0.4)	296.0 _(3.4)
100	100	1	0.5	103.2 _(0.8)	97.3 _(0.4)	365.9 _(2.8)
500	2	2	4	11.3 _(0.1)	97.2 _(0.4)	23.1 _(0.7)
500	2	2	2	15.8 _(0.1)	97.7 _(0.3)	45.2 _(0.9)
500	2	2	1	17.7 _(0.1)	97.5 _(0.4)	117.3 _(1.0)
500	2	1	2	41.5 _(0.3)	97.3 _(0.4)	81.8 _(1.2)
500	2	1	1	55.0 _(0.4)	96.8 _(0.4)	168.9 _(1.5)
500	2	1	0.5	64.6 _(0.5)	98.1 _(0.3)	445.0 _(1.7)
500	22	2	4	23.6 _(0.2)	96.3 _(0.4)	55.4 _(1.0)
500	22	2	2	25.0 _(0.2)	97.0 _(0.4)	60.3 _(0.8)
500	22	2	1	25.5 _(0.2)	98.1 _(0.3)	119.7 _(0.8)
500	22	1	2	88.1 _(0.7)	97.0 _(0.4)	203.5 _(2.1)
500	22	1	1	91.9 _(0.6)	97.8 _(0.3)	229.7 _(2.2)
500	22	1	0.5	94.9 _(0.6)	98.3 _(0.3)	462.8 _(1.4)
500	500	2	4	79.8 _(0.6)	95.0 _(0.5)	238.9 _(2.7)
500	500	2	2	80.3 _(0.6)	95.8 _(0.4)	245.7 _(2.6)
500	500	2	1	80.9 _(0.6)	97.5 _(0.4)	250.2 _(2.5)
500	500	1	2	290.5 _(2.3)	94.5 _(0.5)	819.7 _(7.9)
500	500	1	1	291.4 _(2.3)	95.2 _(0.5)	831.1 _(7.5)
500	500	1	0.5	297.3 _(2.3)	98.1 _(0.3)	875.0 _(6.7)

Table 1: Estimated coverage, average length of the `ocd.CI` confidence interval and average detection delay over 2000 repetitions, with standard errors in brackets. Other parameters: $\gamma = 30000$, $z = 1000$, $\alpha = 0.05$, $a = \sqrt{2 \log p}$, $c = 0.5$, $d_1 = c\sqrt{\log(p/\alpha)}$, $d_2 = 4d_1^2$.

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)

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

to be selected with higher probability than noise coordinates.

4.4 Real data examples

4.4.1 US Covid-19 data

In this section, we apply `ocd.CI` to a dataset of weekly deaths in the United States between January 2017 and June 2020⁴. The data up to 29 June 2019 are treated as our training data. For each of the 50 states, as well as Washington, D.C. ($p = 51$), we pre-process the data as follows. To remove the seasonal effect, 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 stabilise 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 standardise the transformed weekly excess deaths using the mean and standard deviation of the transformed data over the training period. The standardised, transformed data are plotted in Figure 3 for 12 states. When applying `ocd.CI` to these data, we take $a = \sqrt{2 \log p}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$, $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, $d_1 = 0.5\sqrt{\log(p/\alpha)}$ 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

⁴Available at: https://www.cdc.gov/nchs/nvss/vsrr/covid19/excess_deaths.htm.




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 $\hat{\mathcal{S}} \cup \{j\}$ 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: $p = 100$, $\beta = \vartheta = 2$, $\ell = 0$, $a = \sqrt{2 \log p}$.

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 standardising 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 2017⁵.

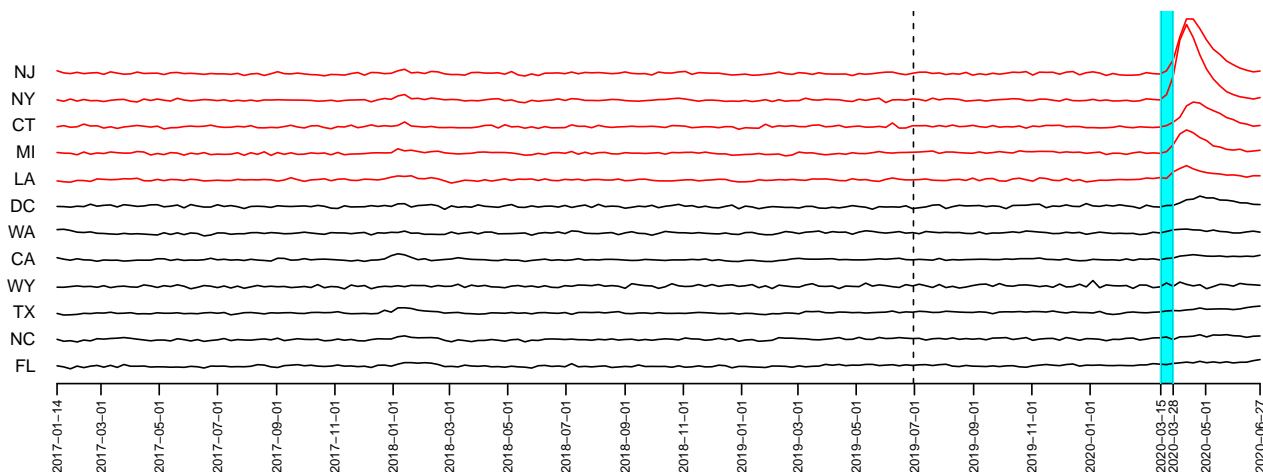


Figure 3: Standardised, transformed weekly excess death data from 12 states (including Washington, D.C.). The monitoring period starts from 30 June 2019 (dashed line). The data from the states in the support estimate are shown in red. The confidence interval [8 March 2020, 28 March 2020] is shown in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

4.4.2 S&P 500 data

We now use `ocd_CI` to study market movements leading up to the financial crisis of 2007–2008. We selected the $p = 254$ stocks that were both in the S&P 500 listing and were traded throughout the period from 1 January 2006 to 31 December 2007. The historical price data were downloaded from finance.yahoo.com using the `quantmod` R package (Ryan et al., 2020); a similar dataset was studied by Cai and Wang (2021). For each stock, we compute the daily logarithmic returns from the adjusted closing prices. We use the data from 2006 as the training data and standardise the entire data according to the mean and variance over the training period. Since these logarithmic returns have a heavy-tailed distribution, we clip the standardised data at ± 3 .

When applying the `ocd_CI` procedure to this dataset, we used the same input parameters as in Section 4.4.1. So as to be able to use `ocd_CI` repeatedly to identify multiple changes, we also set a cool-down period of 10 trading days (i.e. the monitoring resets and restarts 10 trading days after a change is declared). This allows the market to recover from any loss (or gain) from the previous change so that the same market movement is not identified as more

⁵See <https://www.cdc.gov/flu/about/season/flu-season-2017-2018.htm>

than one changepoint. The first four changes were declared on 27 February 2007, 24 May 2007, 24 July 2007 and 8 August 2007, with corresponding confidence intervals shown in Figure 4. This figure also depicts the relative sector impact of each change by showing the percentage of stocks in each sector (according to the Global Industry Classification Standard⁶) that belongs to the estimated support of a changepoint. In particular, the first and last identified changepoints are primarily associated with changes in Real Estate stocks; these correspond to an HSBC announcement indicating loan losses on subprime mortgages in February 2007, and American Home Mortgage Investment Corporation filing for bankruptcy in August 2007 respectively (Hausman and Johnston, 2014).




Figure 4: Heatmap of sector impact of the four changepoints in the S&P 500 data identified by `ocd_CI`, measured as a proportion of the stocks in a sector that appear in the support estimate of the changepoint. The confidence intervals for each of the changepoints are given on the left.

5 Proofs of main results

Proof of Theorem 1. Fix $n > z$, $j \in [p]$, $b \in \mathcal{B}$ and $j' \in [p] \setminus \{j\}$. We assume, without loss of generality, that $\theta^{j'} \geq 0$. The case $\theta^{j'} < 0$ can be analysed similarly. Recall that b_{\min} , defined in Algorithm 2, is the smallest positive scale in $\mathcal{B} \cup \mathcal{B}_0$, and write $b'_{\text{aux}} := \min\{b \in (\mathcal{B} \cup \mathcal{B}_0) \cap (0, \infty) : b \geq \theta^{j'}\}$. Then we have $\Lambda_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'} \mid \tau_{n,b}^j \sim \mathcal{N}(\theta^{j'} \min\{n + \ell -$

⁶See <https://www.msci.com/our-solutions/indexes/gics>

$z, \tau_{n,b}^j + \ell\}, \tau_{n,b}^j + \ell)$. Thus, recalling the definition of $\hat{\mathcal{S}}$ and $\tilde{b}^{j'}$ from Algorithm 2, we have

$$\begin{aligned}
& \mathbb{P}(\{j' \in \hat{\mathcal{S}}\} \cap \{\tilde{b}^{j'} \notin (0, \theta^{j'})\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}) \\
&= \mathbb{E}\left\{\mathbb{P}\left(\{j' \in \hat{\mathcal{S}}\} \cap \{\tilde{b}^{j'} \notin (-b_{\min}, b_{\text{aux}}^{j'})\} \cap \{N = n, \hat{j} = j, \hat{b} = b\} \mid \tau_{n,b}^j\right)\right\} \\
&\leq \mathbb{E}\left\{\mathbb{P}\left(\Lambda_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'} \geq b_{\text{aux}}^{j'}(\tau_{n,b}^j + \ell) + d_1(\tau_{n,b}^j + \ell)^{1/2} \mid \tau_{n,b}^j\right)\right\} \\
&\quad + \mathbb{E}\left\{\mathbb{P}\left(\Lambda_{n,b}^{j',j} + \sum_{i=n+1}^{n+\ell} X_i^{j'} \leq -b_{\min}(\tau_{n,b}^j + \ell) - d_1(\tau_{n,b}^j + \ell)^{1/2} \mid \tau_{n,b}^j\right)\right\} \\
&\leq \mathbb{E}\{\bar{\Phi}((b_{\text{aux}}^{j'} - \theta^{j'}) (\tau_{n,b}^j + \ell)^{1/2} + d_1)\} + \mathbb{E}\{\bar{\Phi}((b_{\min} + \theta^{j'}) (\tau_{n,b}^j + \ell)^{1/2} + d_1)\} \\
&\leq 2\bar{\Phi}(d_1). \tag{3}
\end{aligned}$$

Moreover, by a similar argument to (13) in the proof of Proposition 4, for $b \in (0, \theta^{j'})$, we have

$$\mathbb{P}(n - t_{n,b}^{j'} - d_2/b^2 > z) \leq 2\bar{\Phi}\left(\frac{\sqrt{d_2}}{b}(\theta^{j'} - b/2)\right) \leq 2\bar{\Phi}(\sqrt{d_2}/2). \tag{4}$$

Combining (3) and (4), we have

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

Now, write

$$r_0 := \left(\frac{24T^{\text{off}}}{\vartheta^2} \vee \frac{12(a^2 \vee 8 \log 2)s}{\vartheta^2} \vee \frac{128(T^{\text{diag}} + \log(8/\alpha))s}{\beta^2}\right) \log_2(2p) + 2. \tag{5}$$

By a union bound and Proposition 6, we have

$$\begin{aligned}
& \mathbb{P}\left(N - \min_{j \in \hat{\mathcal{S}}} \left\{t_{N,\tilde{b}^j}^j + \frac{d_2}{(\tilde{b}^j)^2}\right\} > z\right) \\
&\leq \mathbb{P}(N > z + r_0) \\
&\quad + \sum_{n=z+1}^{z+\lfloor r_0 \rfloor} \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{j'=1}^p \mathbb{P}\left(\{j' \in \hat{\mathcal{S}}\} \cap \left\{n - t_{n,\tilde{b}^{j'}}^{j'} - \frac{d_2}{(\tilde{b}^{j'})^2} > z\right\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}\right) \\
&\leq \exp\left\{-\frac{\beta^2(r_0 - 1)}{48 \log_2(2p)}\right\} + p \exp\left\{-\frac{\vartheta^2(r_0 - 1)}{128 \log_2(2p)}\right\} + 4p^2 \log_2^2(4p) r_0 \{\bar{\Phi}(d_1) + \bar{\Phi}(\sqrt{d_2}/2)\}.
\end{aligned}$$

Therefore, for sufficiently large $C_1 > 0$ and $C_2 > 0$, the choice of d_1 and d_2 in the statement of the theorem ensures that

$$\mathbb{P}\left(N - \min_{j \in \hat{\mathcal{S}}} \left\{t_{N,\tilde{b}^j}^j + \frac{d_2}{(\tilde{b}^j)^2}\right\} > z\right) \leq \alpha/2.$$

Combining this with the fact that $\mathbb{P}(N \leq z) \leq z/(4\gamma) \leq \alpha/2$, which follows from Lemma 7 when $C_1 \geq \sqrt{8}$, we deduce the result. \square

Proof of Theorem 2. Denote $\ell_0 := C_3 \left(\frac{a^2 s \log_2(2p)}{\beta^2} + 1 \right)$. Since the output of Algorithm 2 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 $b^1 \geq b^2 \geq \beta/\sqrt{s \log_2(2p)} \geq \sqrt{2}b_{\min}$. For $C_5 > 0$, let

$$r := \frac{C_5 a^2 s \log_2(2p)}{\beta^2} + 2, \quad u := \frac{\ell_0 \beta^2}{80s \log_2(2p)} \quad \text{and} \quad \delta := \frac{a}{2\sqrt{r + \ell}}.$$

Now define the following events:

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

Finally, we denote event

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

with

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

Henceforth, we will assume without loss of generality that $C_2 \geq 1/2$. Then, on the event $\bigcap_{k=0}^4 \Omega_k$, we have

$$\begin{aligned} L &= \min_{j \in \hat{\mathcal{S}}} \left\{ t_{N,\tilde{b}^j}^j + \frac{d_2}{(\tilde{b}^j)^2} \right\} \wedge N \leq N - z + \frac{2(u + d_2)}{(b^2)^2} \leq r + \frac{\ell_0}{40} + \frac{2sd_2 \log_2(2p)}{\beta^2} \\ &\leq C_1^2 \left(C_5 + \frac{C_3}{40} + 8C_2^2 \right) \left(\frac{s \log_2(2p) \log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}{\beta^2} + 1 \right). \end{aligned}$$

Let $C_4 := C_1^2(C_5 + \frac{C_3}{40} + 8C_2^2)$. Then

$$\mathbb{P} \left(L > C_4 \left(\frac{s \log_2(2p) \log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}{\beta^2} + 1 \right) \right) \leq \mathbb{P} \left(\bigcup_{k=0}^4 \Omega_k^c \right). \quad (6)$$

By choosing $C_1 \geq \sqrt{8}$ and choosing C_5 to be a sufficiently large universal constant, we have $r \geq \left(\frac{24T^{\text{off}}}{\vartheta^2} \vee \frac{12a^2s}{\vartheta^2} \vee \frac{24T^{\text{diag}s}}{\beta^2}\right) \log_2(2p) + 2$, so we may apply Proposition 6 and Lemma 7 to deduce that

$$\begin{aligned} \mathbb{P}(\Omega_0^c) &= \mathbb{P}(N > z + r) + \mathbb{P}(N \leq z) \leq 2p \exp\left\{-\frac{\beta^2(r-1)}{128 \log_2(2p)}\right\} + \frac{z}{4\gamma} \\ &\leq 2p \exp\left\{-\frac{C_5 C_1^2 s \log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}{128}\right\} + \frac{z}{4\gamma}. \end{aligned} \quad (7)$$

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 5 (taking $\mu = -b/2$) and a union bound, we have

$$\mathbb{P}(\Omega_0 \cap \Omega_1^c) \leq 2p \log_2(4p) e^{-u/8} = 2p \log_2(4p) \exp\left\{-\frac{\ell_0 \beta^2}{640s \log_2(2p)}\right\}. \quad (8)$$

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

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

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

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

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

$$\mathbb{P}(\Omega_0 \cap \Omega_2^c) \leq 2rp^2 \log_2(4p) e^{-a^2/8} \leq 2rp^2 \log_2(4p) \{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2/8}. \quad (9)$$

By Chen et al. (2021, Lemma 9 in the supplement), we have $t_{n,b}^j/2 \leq \tau_{n,b}^j \leq t_{n,b}^j$ for all $n \in \mathbb{N}_0$, $b \in \mathcal{B} \cup \mathcal{B}_0$ and $j \in [p]$. Moreover, when $C_3 \geq 80(C_5 \vee 2)$, we have $\ell_0 \geq 80r$. Define $b_* := \beta/\sqrt{s \log_2(2p)} \in \mathcal{B}$, so 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

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

where the final inequality follows from Lemma 8(a), applied with $U = \sum_{i=n-\tau_{n,b_*}^j}^z X_i^{j'}$, $V = \sum_{i=n-\tau_{n,b}^j}^{n-\tau_{n,b_*}^j} X_i^{j'}$, $Y = \sum_{i=z+1}^{n+\ell} X_i^{j'}$, $\alpha = \theta^{j'}$, $\phi_1 = z - n + \tau_{n,b_*}^j$, $\phi_2 = z - n + \tau_{n,b}^j$, $\phi_3 = n - z + \ell$ and $\kappa = \ell/80$, as well as Lemma 8(b), with $U = \sum_{i=n-\tau_{n,b}^j}^z X_i^{j'}$, $Y = \sum_{i=n-\tau_{n,b_*}^j}^{n+\ell} X_i^{j'}$, $Z = \sum_{i=z+1}^{n-\tau_{n,b}^j} X_i^{j'}$, $\alpha = \theta^{j'}$, $\phi_1 = z - n + \tau_{n,b}^j$, $\phi_3 = \ell + \tau_{n,b_*}^j$, $\phi_4 = n - z - \tau_{n,b_*}^j$ and $\kappa = \ell/80$. Observe that $\tilde{Q}_{n,\hat{b}}^j \geq \tilde{Q}_{n,b_*}^j$. Thus, by a union bound, we have

$$\begin{aligned}
& \mathbb{P}(\Omega_0 \cap \Omega_1 \cap \Omega_2 \cap \Omega_3^c) \\
& \leq \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{n=z+1}^{z+r} \mathbb{P}\left(\{N = n\} \cap \Omega_1 \cap \Omega_2 \cap \{\tilde{Q}_{n,b}^j \geq \tilde{Q}_{n,b_*}^j\} \cap \{\tau_{n,b}^j > n - z + \ell/20\}\right) \\
& \leq 2rp^2 \log_2(2p) \exp\left(-\frac{\ell a^2}{3840(r + \ell)}\right). \tag{10}
\end{aligned}$$

When $C_3 \geq 144C_2^2 \vee 80C_5 \vee 160$, we have $\ell_0 \geq 80r$ and $d_1 \leq \frac{\sqrt{\ell_0\beta}}{12\sqrt{s \log_2(2p)}} \leq b^1 \sqrt{\ell}/12$. Thus, on $\Omega_0 \cap \Omega_3$, we have

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

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

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

Here, in the final bound, we have used the facts that $\Xi_{n,b}^{1,j} \mid \tau_{n,b}^j \sim \mathcal{N}(\theta^1 \min\{(n + \ell - z)(\tau_{n,b}^j +$

$\ell)^{-1/2}, (\tau_{n,b}^j + \ell)^{1/2}\}, 1)$ and that

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

when $\tau_{n,b}^j \leq \ell/16$, as well as the standard Gaussian tail bound used at the end of the proof of Lemma 5. By a similar argument, we also have for any $z < n \leq z + r$ and $b \in \mathcal{B}$ that

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

Thus, by a union bound, we have

$$\begin{aligned} \mathbb{P}(\Omega_0 \cap \Omega_3 \cap \Omega_4^c) &= \mathbb{P}(\Omega_0 \cap \Omega_3 \cap \Omega_{4,1}^c \cap \Omega_{4,2}^c) \\ &\leq \sum_{j=2}^p \sum_{b \in \mathcal{B}} \sum_{n=z+1}^{z+[r]} \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b\} \cap \Omega_{4,1}^c) \\ &\quad + \sum_{b \in \mathcal{B}} \sum_{n=z+1}^{z+[r]} \mathbb{P}(\Omega_3 \cap \{N = n, \hat{j} = 1, \hat{b} = b\} \cap \Omega_{4,2}^c) \\ &\leq rp \log_2(2p) e^{-d_1^2/2}. \end{aligned} \tag{11}$$

Hence by substituting (7), (8), (9), (10) and (11) into (6), we conclude that, by increasing the universal constant $C_1 > 0$ if necessary,

$$\begin{aligned} & \mathbb{P}\left(L > C_4 \left(\frac{s \log_2(2p) \log\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}}{\beta^2} + 1\right)\right) \\ & \leq 2p\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2 C_5 s/128} + \frac{z}{4\gamma} + 2p \log_2(4p) \{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2 C_3/640} \\ & \quad + 2rp^2 \log_2(4p) \{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2/8} + 2rp^2 \log_2(2p) \{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2/3888} \\ & \quad + rp \log_2(2p) \{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2 C_2^2/2} \\ & \leq \alpha, \end{aligned}$$

as required. \square

Proof of Theorem 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 (3), we have, for $n > z, j \in [p], b \in \mathcal{B}$ and $j' \in \mathcal{S}_\beta^c$, that

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

Hence, recalling the definition of r_0 from (5), by Lemma 7 applied with $C_1 \geq \sqrt{8}$, a union bound and Proposition 6, we have

$$\begin{aligned} \mathbb{P}_{z,\theta}(\hat{\mathcal{S}} \not\subseteq \mathcal{S}_\beta) &\leq \mathbb{P}_{z,\theta}(N \leq z) + \mathbb{P}_{z,\theta}(N > z + r_0) \\ &\quad + \sum_{n=z+1}^{z+[r_0]} \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{j' \in \mathcal{S}_\beta^c} \mathbb{P}_{z,\theta}(\{j' \in \hat{\mathcal{S}}\} \cap \{N = n, \hat{j} = j, \hat{b} = b\}) \\ &\leq \frac{\alpha}{2} + \exp\left\{-\frac{\beta^2(r_0 - 1)}{48 \log_2(2p)}\right\} + p \exp\left\{-\frac{\vartheta^2(r_0 - 1)}{128 \log_2(2p)}\right\} + 4p^2 \log_2(2p)r_0\bar{\Phi}(d_1) \leq \alpha, \end{aligned}$$

where the final bound follows as in the proof of Theorem 1.

(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 (10) that when $C_3 \geq 144C_2^2 \vee 80C_5 \vee 160$, we have $\tau_{N,\hat{b}}^{\hat{j}} \leq \ell/16$ 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 $\Xi_{n,b}^{j',j}$ from Algorithm 2. Then, for any $z < n \leq z + r$, $j \in [p]$, $j' \in \mathcal{S} \setminus \{j\}$ and $b \in \mathcal{B}$, we have

$$\begin{aligned} &\mathbb{P}_{z,\theta}(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b, j' \notin \hat{\mathcal{S}}\} \mid X_1^j, X_2^j, \dots) \\ &= \mathbb{P}_{z,\theta}\left(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b\} \cap \{|\Xi_{n,b}^{j',j}| < b_{\min} \sqrt{\tau_{n,b}^j + \ell} + d_1\} \mid X_1^j, X_2^j, \dots\right) \\ &\leq \mathbb{P}_{z,\theta}\left(\{\tau_{n,b}^j \leq \ell/16\} \cap \left\{|\Xi_{n,b}^{j',j}| - b_{\min} \sqrt{\tau_{n,b}^j + \ell} < d_1\right\} \mid X_1^j, X_2^j, \dots\right) \leq \frac{1}{2} e^{-d_1^2/2}, \quad (12) \end{aligned}$$

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

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

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

$$\begin{aligned} &\mathbb{P}_{z,\theta}(\hat{\mathcal{S}} \cup \{\hat{j}\} \not\subseteq \mathcal{S}) \\ &\leq \mathbb{P}_{z,\theta}(\Omega_0^c) + \mathbb{P}_{z,\theta}(\Omega_0 \cap \Omega_1^c) + \mathbb{P}_{z,\theta}(\Omega_0 \cap \Omega_2^c) + \mathbb{P}_{z,\theta}(\Omega_0 \cap \Omega_1 \cap \Omega_2 \cap \Omega_3^c) \\ &\quad + \sum_{n=z+1}^{z+[r]} \sum_{j=1}^p \sum_{b \in \mathcal{B}} \sum_{j' \in \mathcal{S} \setminus \{j\}} \mathbb{P}_{z,\theta}(\Omega_3 \cap \{N = n, \hat{j} = j, \hat{b} = b, j' \notin \hat{\mathcal{S}}\}) \\ &\leq 2p\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2 C_5 s/128} + \frac{z}{4\gamma} + 2p \log_2(4p)\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2 C_3/640} \\ &\quad + 2rp^2 \log_2(4p)\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2/8} + 2rp^2 \log_2(2p)\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2/3888} \\ &\quad + rp^2 \log_2(2p)\{p\gamma(\beta^{-2} \vee 1)\alpha^{-1}\}^{-C_1^2 C_2^2/2} \leq \alpha, \end{aligned}$$

where the penultimate inequality follows from (7), (8), (9), (10) and (12), and the last inequality follows by choosing the universal constant $C_1 > 0$ to be sufficiently large. \square

6 Auxiliary results

Proposition 4. *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 5 in (13).*

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. (2021, 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

$$\mathbb{P}(N - t_{N,b} - y \geq z) \leq \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq z+y} R_{n,b} = 0\right) \leq \mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} U_{n,b} \leq 0\right) \leq 2\bar{\Phi}\left(\sqrt{y}(\theta - b/2)\right), \quad (13)$$

where the last inequality follows from Lemma 5. 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 5. *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

$$\mathbb{P}\left(\inf_{n \in \mathbb{N}_0: n \geq y} U_n \leq 0 \mid U_y = u\right) \leq \mathbb{P}\left\{\inf_{t \in [y, \infty)} (B_t + t\mu) \leq 0 \mid B_y = u\right\} \leq e^{-2u\mu}, \quad (14)$$

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 (14) 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 6 and Lemma 7 below, we assume the Gaussian data generating mechanism given at the beginning of Section 3.

Proposition 6. *Assume that θ has an effective sparsity of $s := s(\theta) \geq 2$. Then, the right endpoint N of the interval output from the `ocd_CI'` algorithm, with inputs $(X_t)_{t \in \mathbb{N}}$, $0 < \beta \leq \vartheta$, $a > 0$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$ and $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, satisfies*

$$\mathbb{P}_{z, \theta}(N > z + r) \leq \exp\left\{-\frac{\beta^2(r-1)}{48 \log_2(2p)}\right\} + p \exp\left\{-\frac{\vartheta^2(r-1)}{128 \log_2(2p)}\right\},$$

for all $r \geq \left\{\frac{24T^{\text{off}} \log_2(2p)}{\vartheta^2} \vee \frac{12(a^2 \vee 8 \log 2)s \log_2(2p)}{\vartheta^2} \vee \frac{24T^{\text{diag}} s \log_2(2p)}{\beta^2}\right\} + 2$.

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

$$\mathcal{J} := \left\{j \in [p] : \theta^j/b_* \geq 1 \text{ and } |\theta^j| \leq \vartheta/\sqrt{2}\right\} \quad (15)$$

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. (2021, Proposition 8), we have on the event $\{N > z\}$ that

$$q(X_1, \dots, X_z, \theta) := \max\{t_{z, b_*}^j : j \in \mathcal{J}\} \leq \frac{8T^{\text{diag}} s \log_2(2p)}{\beta^2}. \quad (16)$$

We now fix

$$r \geq \left\{\frac{24T^{\text{off}} \log_2(2p)}{\vartheta^2} \vee \frac{12(a^2 \vee 8 \log 2)s \log_2(2p)}{\vartheta^2} \vee \frac{24T^{\text{diag}} s \log_2(2p)}{\beta^2}\right\} + 2 =: r_0. \quad (17)$$

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. \(2021, Lemma 2\)](#) to $t_{z+[r],b_*}^j$, we have for $j \in \mathcal{J}$ that

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

Recall that $X_{z+1}, X_{z+2}, \dots \stackrel{\text{iid}}{\sim} \mathcal{N}_p(\theta, I_p)$. Hence, by applying [Chen et al. \(2021, Lemma 6\(b\)\)](#) with $a = 0, b = |b_*|/2$ and $c = [r]/3$, we have

$$\begin{aligned} \mathbb{P}_{z,\theta} \left(\bigcap_{j \in \mathcal{J}} (\Omega_r^j)^c \right) &= \prod_{j \in \mathcal{J}} \mathbb{P}_{z,\theta} \left(t_{z+[r],b_*}^j \leq \frac{2[r]}{3} \right) \\ &\leq \prod_{j \in \mathcal{J}} \mathbb{P}_{z,\theta} \left(\operatorname{largmax}_{0 \leq h \leq [r]} \sum_{i=z+1}^{z+h} -b_*(X_i^j - b_*/2) \geq \frac{[r]}{3} \right) \\ &\leq \prod_{j \in \mathcal{J}} \mathbb{P}_{z,\theta} \left(\sup_{h \geq [r]/3} \sum_{i=z+1}^{z+h} -\operatorname{sgn}(b_*)(X_i^j - b_*/2) \geq 0 \right) \\ &\leq \exp(-|\mathcal{J}|b_*^2[r]/24) \leq \exp(-sb_*^2[r]/48). \end{aligned} \quad (18)$$

We now work on the event Ω_r^j , for some fixed $j \in \mathcal{J}$. We note that (17) guarantees that $r \geq 2$, and thus $t_{z+[r],b_*}^j \geq \lceil 2[r]/3 \rceil \geq 2$. Then, by (16) and (17), we have $r_0 > 3t_{z,b_*}^j$, and hence by [Chen et al. \(2021, Lemma 9\)](#),

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

We conclude that

$$2/3 \leq [r]/3 < \tau_{z+[r],b_*}^j \leq [r]. \quad (19)$$

Recall that $\Lambda_{z+[r],b_*}^{j,j} \in \mathbb{R}^p$ records the tail CUSUM statistics with tail length $\tau_{z+[r],b_*}^j$. We observe by (19) that only post-change observations are included in $\Lambda_{z+[r],b_*}^{j,j}$. Hence we have that

$$\Lambda_{z+[r],b_*}^{k,j} \mid \tau_{z+[r],b_*}^j \stackrel{\text{iid}}{\sim} \mathcal{N}(\theta^k \tau_{z+[r],b_*}^j, \tau_{z+[r],b_*}^j) \quad (20)$$

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

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

has cardinality at least $s - 1$. Hence, by (19), for all $k \in \mathcal{L}^j$,

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

We then observe, from (17), that

$$\tilde{a}_r > 2(a \vee \sqrt{8 \log 2}). \quad (21)$$

Hence, from (20), we have for all $k \in \mathcal{L}^j$ that

$$\mathbb{P}_{z,\theta} \left(\Omega_r^j \cap \left\{ |\Lambda_{z+[r],b_*}^{k,j}| < \frac{1}{2} \tilde{a}_r \sqrt{\tau_{z+[r],b_*}^j} \right\} \middle| \tau_{z+[r],b_*}^j \right) \leq \frac{1}{2} e^{-\tilde{a}_r^2/8} =: q_r.$$

We denote

$$U^j := \left\{ k \in \mathcal{L}^j : |\Lambda_{z+[r],b_*}^{k,j}| < \frac{1}{2} \tilde{a}_r \sqrt{\tau_{z+[r],b_*}^j} \right\}.$$

Then, by the Chernoff–Hoeffding binomial tail bound (Hoeffding, 1963, Equation (2.1)), we have

$$\begin{aligned} \mathbb{P}_{z,\theta} \left(\Omega_r^j \cap \{U^j \geq |\mathcal{L}^j|/2\} \middle| \tau_{z+[r],b_*}^j \right) &\leq \exp \left\{ -\frac{|\mathcal{L}^j|}{2} \log \left(\frac{1}{4q_r(1-q_r)} \right) \right\} \\ &\leq \exp \left\{ -\frac{|\mathcal{L}^j|}{2} \left(\frac{\tilde{a}_r^2}{8} - \log 2 \right) \right\} \leq \exp \left(-\frac{3|\mathcal{L}^j|\tilde{a}_r^2}{64} \right) \leq \exp \left\{ -\frac{\vartheta^2[r]}{128 \log_2(2p)} \right\}, \end{aligned} \quad (22)$$

where the penultimate inequality follows from (21). Now, on the event $\Omega_r^j \cap \{U^j < |\mathcal{L}^j|/2\}$, we have

$$\begin{aligned} \sum_{j' \in [p]: j' \neq j} \frac{(\Lambda_{z+[r],b_*}^{j',j})^2}{\tau_{z+[r],b_*}^j \vee 1} \mathbb{1}_{\{|\Lambda_{z+[r],b_*}^{j',j}| \geq a \sqrt{\tau_{z+[r],b_*}^j}\}} &\geq \sum_{j' \in [p]: j' \neq j} \frac{(\Lambda_{z+[r],b_*}^{j',j})^2}{\tau_{z+[r],b_*}^j \vee 1} \mathbb{1}_{\{|\Lambda_{z+[r],b_*}^{j',j}| \geq \frac{1}{2} \tilde{a}_r \sqrt{\tau_{z+[r],b_*}^j}\}} \\ &\geq \frac{\tilde{a}_r^2}{4} \left\{ |\mathcal{L}^j| - \left(\left\lceil \frac{|\mathcal{L}^j|}{2} \right\rceil - 1 \right) \right\} = \frac{\tilde{a}_r^2}{4} \left\lceil \frac{|\mathcal{L}^j| + 1}{2} \right\rceil \geq \frac{\vartheta^2[r]}{24 \log_2(2p)} \geq T^{\text{off}}, \end{aligned} \quad (23)$$

where the penultimate inequality uses the fact that $|\mathcal{L}^j| \geq s-1$ and the last inequality follows from (17). We now denote

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

Combining (18), (22) and (23), we deduce that

$$\begin{aligned} \mathbb{P}_{z,\theta}(N > z+r) &\leq \mathbb{P}_{z,\theta}(N > z+[r]) \leq \mathbb{P}_{z,\theta} \left(\bigcap_{j \in \mathcal{J}} (\Omega_r^j)^c \right) + \sum_{j \in \mathcal{J}} \mathbb{P}_{z,\theta}(\tilde{E}_r^j \cap \Omega_r^j) \\ &\leq \mathbb{P}_{z,\theta} \left(\bigcap_{j \in \mathcal{J}} (\Omega_r^j)^c \right) + \sum_{j \in \mathcal{J}} \mathbb{P}_{z,\theta}(\Omega_r^j \cap \{U^j \geq |\mathcal{L}^j|/2\}) \\ &\leq \exp \left\{ -\frac{sb_*^2(r-1)}{48} \right\} + p \exp \left\{ -\frac{\vartheta^2(r-1)}{128 \log_2(2p)} \right\}, \end{aligned}$$

as desired. \square

Lemma 7. Let $p \geq 2$. Then the output N from `ocd'`, with inputs $(X_t)_{t \in \mathbb{N}}$, $0 < \beta \leq \vartheta$, $a \geq \sqrt{8 \log(p-1)}$, $T^{\text{diag}} = \log\{16p\gamma \log_2(4p)\}$ and $T^{\text{off}} = 8 \log\{16p\gamma \log_2(2p)\}$, satisfies

$$\mathbb{P}_{z,\theta}(N \leq z) \leq \frac{z}{4\gamma}.$$

Proof. This follows from [Chen et al. \(2021, \(16\)\)](#) in the proof of Theorem 1 and (31) in the proof of Theorem 3). \square

Lemma 8. 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 5](#), we have

$$\mathbb{P}(W_1 \leq 0) \leq \frac{1}{2}e^{-\alpha^2/(2w_1)}, \quad (24)$$

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

$$\begin{aligned} w_1 &= \frac{\phi_1 + \phi_3}{\phi_3^2} \left(1 + \frac{(\sqrt{\phi_2 + \phi_3} + \sqrt{\phi_1 + \phi_3})^2}{\phi_2 - \phi_1}\right) \\ &\leq \frac{5}{16\kappa} \left(1 + \frac{(\sqrt{8\kappa} + \sqrt{5\kappa})^2}{3\kappa}\right) \leq \frac{3}{\kappa}, \end{aligned} \quad (25)$$

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 (24) and (25), 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

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

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 5, we have

$$\mathbb{P}(W_2 \leq 0) \leq \frac{1}{2}e^{-\alpha^2/(2w_2)}, \quad (26)$$

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

$$w_2 \leq \frac{6}{\kappa}. \quad (27)$$

Hence, using the fact that $-(U + Y + Z) \leq_{st} U + Y + Z$, as well as (26) and (27), 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

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

- Cai, H. and Wang, T. (2021) Estimation of high-dimensional change-points under a group sparsity structure. *arXiv preprint*, arxiv:2107.08724.
- 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. (2020) *ocd: high-dimensional, multiscale online changepoint detection*. R package, available at <https://cran.r-project.org/package=ocd>.
- Chen, Y., Wang, T. and Samworth, R. J. (2021) High-dimensional, multiscale online change-point detection. *J. Roy. Statist. Soc., Ser. B*, to appear.
- Cho, H. (2016) Change-point detection in panel data via double CUSUM statistic. *Elect. 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. (2021) High-dimensional changepoint estimation with heterogeneous missingness. *arXiv preprint*, arxiv:2108.01525.
- Gösmann, J., Stoehr, C., Heiny, J. and Dette, H. (2020) Sequential change point detection in high dimensional time series. *arXiv preprint*, arxiv:2006.00636.
- Hausman, A. and Johnston, W. J. (2014) Timeline of a financial crisis: Introduction to the special issue. *J. Bus. Res.*, **67**, 2667–2670.
- Hoeffding, W. (1963) Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, **58**, 13–30.
- 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 changepoint 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.
- Padilla, O. H. M., Yu, Y., Wang, D. and Rinaldo, A. (2019) Optimal nonparametric multivariate change point detection and localization. *arXiv preprint*, arxiv:1910.13289.
- 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. In *International Conference on Artificial Intelligence and Statistics*, 2089–2097. PMLR.
- Ryan, J. A., Ulrich, J. M., Thielen, W., Teetor, P., Bronder, S. and Ulrich, M. J. M. (2020) *quantmod: quantitative financial modelling framework*. R package, available at <https://cran.r-project.org/package=quantmod>.

- 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. *Appl. Comp. Harm. Anal.*, **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.
- 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, 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.