

High-dimensional changepoint estimation with heterogeneous missingness

Bertille Follain^{1,2} | Tengyao Wang^{3,4} | Richard J. Samworth¹

¹Statistical Laboratory, University of Cambridge, Cambridge, Cambridgeshire, UK

²Ecole Normale Supérieure, PSL Research University, INRIA, Paris, France

³Department of Statistics, London School of Economics and Political Science, London, London, UK

⁴Department of Statistical Science, University College London, London, UK

Correspondence

Tengyao Wang, Department of Statistics, London School of Economics and Political Science, London, London, UK.
Email: t.wang59@lse.ac.uk

Abstract

We propose a new method for changepoint estimation in partially observed, high-dimensional time series that undergo a simultaneous change in mean in a sparse subset of coordinates. Our first methodological contribution is to introduce a ‘MissCUSUM’ transformation (a generalisation of the popular cumulative sum statistics), that captures the interaction between the signal strength and the level of missingness in each coordinate. In order to borrow strength across the coordinates, we propose to project these MissCUSUM statistics along a direction found as the solution to a penalised optimisation problem tailored to the specific sparsity structure. The changepoint can then be estimated as the location of the peak of the absolute value of the projected univariate series. In a model that allows different missingness probabilities in different component series, we identify that the key interaction between the missingness and the signal is a weighted sum of squares of the signal change in each coordinate, with weights given by the observation probabilities. More specifically, we prove that the angle between the estimated and oracle projection directions, as well as the changepoint location error, are controlled with high probability by the sum of two terms, both involving this weighted sum of squares, and representing the error incurred due to noise and the error due to missingness respectively. A lower bound confirms that our changepoint estimator,

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* published by John Wiley & Sons Ltd on behalf of Royal Statistical Society.

which we call `MissInspect`, is optimal up to a logarithmic factor. The striking effectiveness of the `MissInspect` methodology is further demonstrated both on simulated data, and on an oceanographic data set covering the Neogene period.

KEYWORDS

change point estimation, high-dimensional data, missing data, segmentation, sparsity

1 | INTRODUCTION

The Big Data era offers the exciting prospect of being able to transform our understanding of many scientific phenomena, but at the same time many traditional statistical techniques may perform poorly, or may no longer be computable at all, when applied to contemporary data challenges. A core assumption that underpins much of statistical theory, as well as the way in which we think about statistical modelling, is that our data are realisations of independent and identically distributed random variables. However, practical experience reveals that this is typically unrealistic for modern data sets, and developing methods and theory to handle departures from this important but limited setting represents a key theme for the field.

In contexts where data are collected over time, one of the simplest generalisations of an independent and identically distributed data stream is given by *change point* models. Here, we postulate that our data may be segmented into shorter, homogeneous series. Of course, the structural break, or change point, between these series is often of interest in applications, such as distributed denial of service monitoring of network traffic (Peng et al., 2004), disease progression tracking via the alignment of electronic medical records (Huopaniemi et al., 2014) and the analysis of ‘shocks’ in stock price data (Chen & Gupta, 1997).

Another issue that turns out to be critical in working with Big Data in practice is that of missing data. One reason for this is that when each observation is high-dimensional, it is frequently the case that most or even every observation has missingness in some coordinates; thus a complete-case analysis, which simply discards such observations, is unviable (Zhu et al., 2019).

The aim of this paper is to study the core, high-dimensional change point problem of a sparse change in mean, but where our data are corrupted by missingness. In fact, in cases where our data arise as discrete observations of several continuous processes, the observation times in different coordinates may not be the same, and such a setting also fits within our framework. A key feature of both our methodology and theory is that we wish to be able to handle *heterogeneous* missingness, that is, where the levels of missingness may differ across coordinates. Specifically, our primary theoretical goal is to understand the way in which the missingness *interacts* with the signal strengths in the different series to determine the difficulty of the problem.

In Section 2, we consider a setting where the practitioner has access to a partially observed $p \times n$ data matrix, where p is the number of series (coordinates) being monitored, and n is the number of time points. We seek to identify a time at which the p -dimensional mean vector changes, in at least one coordinate. One of the key ideas that underpins our methodological contribution is to define a new version of the popular cumulative sum (CUSUM) transformation (Page, 1955)

that is able to handle the missingness appropriately. This operation, which we refer to as the MissCUSUM transformation, returns a $p \times (n - 1)$ matrix, and the intuition is that in coordinates that undergo a change in mean, the transformed series should peak in absolute value near the change-point. One of the main advantages of our proposal is that it avoids the need to impute missing data¹.

Since the changepoint location is shared across the signal coordinates, it is natural to seek to borrow strength across the different data streams to estimate the changepoint. To this end, our next goal is to estimate a projection direction, in order to convert the MissCUSUM transformation into a univariate CUSUM series. Such a projection direction should ideally maximise the signal-to-noise ratio of the projected series. When the data are fully observed, the oracle projection direction turns out to be the leading left singular vector of the (rank one) CUSUM transformation of the mean matrix. This facilitates estimation approaches based on entrywise ℓ_1 -penalised M -estimation, as in the `inspect` algorithm of Wang and Samworth (2018). A crucial difference when we have to handle missing data, however, is that the MissCUSUM transformation of the mean matrix is no longer of rank one, which means that the entrywise ℓ_1 -penalty no longer adequately captures the sparsity structure of the vector of mean change. Instead, we introduce a new optimisation problem that penalises the ℓ_1 -norm of the leading left singular vector of a rank one approximation of the MissCUSUM transformation. This methodological proposal, which we call `MissInspect`, leads to considerably improved performance. Implementation code for our method is available in the GitHub repository <https://github.com/wangtengyao/MissInspect> and in the R package `InspectChangepoint` (Wang et al., 2022).

A further benefit of the `MissInspect` methodology is that it is amenable to theoretical analysis. In particular, we study a missing completely at random (MCAR) model with *row homogeneous* missingness; in other words, the observation probability remains constant in each row, but may vary arbitrarily across rows. In Proposition 1 in Section 3, we provide a high-probability bound on the angle between the estimated and oracle projection directions. Theorems 1 and 2 then establish high-probability bounds on the accuracy of the estimated changepoint location whenever the estimated and oracle projection directions are sufficiently well aligned, for a sample splitting variant of our algorithm. Theorem 1 provides a very general guarantee, while Theorem 2 establishes a faster rate whenever the observation probability in each row satisfies a lower bound. This faster rate comprises two terms, representing the error incurred due to noise in the observations, and the error due to missingness respectively. The key quantity in both of these terms turns out to be a weighted Euclidean norm of the vector of mean change, where the weights are given by the observation probabilities in each row. This weighted average therefore captures the interaction between the signal strength and the missingness probabilities, and suggests that our analysis handles effectively the heterogeneity of the missingness across rows (a more naive analysis would see the worst-case observation probability appearing in the bounds). This intuition is confirmed by our minimax lower bound (Theorem 3), which indicates that the `MissInspect` algorithm attains the minimax rate of convergence in all problem parameters, up to a logarithmic factor.

Section 4 explores the empirical performance of our `MissInspect` methodology. We study the ability of the algorithm to estimate both the oracle projection direction and the changepoint

¹In fact, our initial approach to this problem was to consider iterating between imputating missing entries using row means on either side of a potential changepoint, and then updating the current changepoint location estimate using the imputed data matrix. This turns out to perform poorly, because the imputation step tends to reinforce bias in the changepoint estimate, leading to the iterations becoming stuck (potentially far from the true location) very quickly; see Section 4.3.

location, and compare with three alternative approaches. The first two use different imputation strategies in combination with the `inspect` algorithm, while the third is an adaptation of the method of Londschien et al. (2021), which was originally proposed in the context of change-point problems in dynamic graphical models. We find that the `MissInspect` algorithm has very strong numerical performance, and this provides further evidence of its practical utility in missing data settings. In this section, we also present an application of the `MissInspect` methodology to detect changes in oceanographic currents from carbon isotope measurements extracted from cores drilled into the ocean floor. Section 5 discusses various methodological and theoretical extensions of our proposal to more complicated problems, such as the estimation of multiple changepoints, or more general data generating and missingness mechanisms. Proofs of our main results are given in Section 6, with auxiliary results deferred to the Appendix in the supplementary material.

The study of changepoint problems dates at least back to Page (1955), and has since found applications in many different areas, including genetics (Olshen et al., 2004; Zhang et al., 2010), disease outbreak watch (Sparks et al., 2010), aerospace engineering (Henry et al., 2010) and functional magnetic resonance imaging studies (Aston & Kirch, 2013), in addition to those already mentioned. Entry points to the literature include Csörgő and Horváth (1997) and Horváth and Rice (2014). In high-dimensional changepoint settings, where we may have a sparsity assumption on the coordinates of change, prior work includes Bai (2010), Zhang et al. (2010), Horváth and Hušková (2012), Cho and Fryzlewicz (2014), Chan and Walther (2015), Jirak (2015), Cho (2016), Soh and Chandrasekaran (2017), Wang and Samworth (2018), Enikeeva and Harchaoui (2019), Padilla et al. (2022) and Liu et al. (2021). Our focus in this work is on the *offline* version of the changepoint estimation problem, where the practitioner sees the whole data set prior to determining a changepoint location. The corresponding online version, where data are observed sequentially and the challenge is to declare a change as soon as possible after it has occurred, has also received attention in recent years; see, for example, Mei (2010), Xie and Siegmund (2013), Chan (2017) and Chen et al. (2022). In addition to Londschien et al. (2021) discussed above, the only works of which we are aware on changepoint estimation with missing data are those of Xie et al. (2013), Cao et al. (2019) and Enikeeva and Klopp (2021). Xie et al. (2013) study a situation where partially observed sequential data lie close to a time-varying, low-dimensional submanifold embedded within an ambient space; Cao et al. (2019) consider a sketching approach to online changepoint detection with missing data, while Enikeeva and Klopp (2021) study the detection of structural breaks in dynamic networks with missing links.

We conclude this section by introducing some notation that is used throughout the paper. Given $n \in \mathbb{N}$, we let $[n] := \{1, \dots, n\}$. For a vector $u = (u_1, \dots, u_M)^\top \in \mathbb{R}^M$, a matrix $A = (A_{ij}) \in \mathbb{R}^{M \times N}$ and for $r \in [1, \infty)$, we write $\|u\|_r := \left(\sum_{i=1}^M |u_i|^r\right)^{1/r}$ and $\|A\|_r := \left(\sum_{i \in [M]} \sum_{j \in [N]} |A_{ij}|^r\right)^{1/r}$ for their entrywise ℓ_r -norms, as well as $\|u\|_\infty := \max_{i \in [M]} |u_i|$ and $\|A\|_\infty := \max_{i \in [M], j \in [N]} |A_{ij}|$. Writing $\sigma_1(A), \dots, \sigma_s(A)$ for the non-zero singular values of A , where $s := \text{rank}(A)$, we let $\|A\|_{\text{op}} := \max_{i \in [s]} \sigma_i(A)$, $\|A\|_* := \sum_{i=1}^s \sigma_i(A)$ and $\|A\|_F := \|A\|_2 = \left\{\sum_{i=1}^s \sigma_i(A)^2\right\}^{1/2}$ denote its operator, nuclear and Frobenius norms respectively. We also write $\|u\|_0 := \sum_{i=1}^M \mathbb{1}_{\{u_i \neq 0\}}$. Given $\mathbf{q} = (q_1, \dots, q_M)^\top \in [0, 1]^M$, we write $\sqrt{\mathbf{q}} := (\sqrt{q_1}, \dots, \sqrt{q_M})^\top$ and let $\|u\|_{r, \mathbf{q}} := \left(\sum_{i=1}^M |u_i|^r q_i\right)^{1/r}$. We denote by $\text{diag}(u)$ the $M \times M$ diagonal matrix with u as its diagonal. For $S \subseteq [M]$ and $T \subseteq [N]$, we write $u_S := (u_i : i \in S)^\top \in \mathbb{R}^{|S|}$ and write $A_{S, T} \in \mathbb{R}^{|S| \times |T|}$ for the sub-matrix of A obtained by extracting the rows and columns with indices in S and T respectively. For two

matrices $A, B \in \mathbb{R}^{M \times N}$, we denote their trace inner product as $\langle A, B \rangle := \text{tr}(A^\top B)$. We also denote their Hadamard product as $A \circ B \in \mathbb{R}^{M \times N}$. For non-zero vectors $u, v \in \mathbb{R}^M$, we write

$$\angle(u, v) := \cos^{-1} \left(\frac{|\langle u, v \rangle|}{\|u\|_2 \|v\|_2} \right)$$

for the acute angle bounded between them. We let $\mathbb{B}^M := \{\mathbb{R}^M : \|x\|_2 \leq 1\}$ and $\mathbb{S}^{M-1} := \{x \in \mathbb{R}^M : \|x\|_2 = 1\}$ denote the unit Euclidean ball and sphere in \mathbb{R}^M respectively, and define $\mathbb{S}^{M-1}(k) := \{x \in \mathbb{S}^{M-1} : \|x\|_0 \leq k\}$. Given positive sequences $(a_n), (b_n)$, we write $a_n \lesssim b_n$ to mean that there exists a universal constant $C > 0$ such that $a_n \leq C b_n$ for all n .

2 | MISSINSPECT METHODOLOGY

Throughout this work, we will assume that the practitioner has access to a partially observed $p \times n$ data matrix. We will denote the full data matrix as $X = (X_{j,t}) \in \mathbb{R}^{p \times n}$, and let $\Omega = (\omega_{j,t}) \in \{0, 1\}^{p \times n}$ denote the *revelation matrix*, so that $\omega_{j,t} = 1$ if $X_{j,t}$ is observed, and is equal to zero otherwise. Formally, then, we can regard the observed data as $(X \circ \Omega, \Omega)$; note here, that since the practitioner has access to the matrix Ω , they are able to distinguish between an observed zero and a zero caused by missingness in $X \circ \Omega$. For our theoretical analysis, the \circ notation is a convenient way of avoiding the need to introduce an ‘NA’ category for missing values.

In our theory, we will regard X as a realisation of a random matrix, whose mean matrix we denote by $\mu = (\mu_1, \dots, \mu_n) \in \mathbb{R}^{p \times n}$. The changepoint structure of μ is encoded via the assumption that there exist $z \in [n - 1]$ and $\mu^{(1)}, \mu^{(2)} \in \mathbb{R}^p$ with $\theta := \mu^{(2)} - \mu^{(1)} \neq 0$ such that

$$\mu_1 = \dots = \mu_z = \mu^{(1)} \text{ and } \mu_{z+1} = \dots = \mu_n = \mu^{(2)}. \tag{1}$$

In Section 3, we will assume that the change in mean is sparse, in the sense that $\|\theta\|_0 \leq k$ for some k that is typically much smaller than p . However, we remark that our methodology is adaptive to this unknown sparsity level.

Our goal is to estimate the changepoint location z . To this end, we first introduce a new version of the CUSUM transformation that is appropriate in our missing data setting. Writing $L_{j,t} := \sum_{r=1}^t \omega_{j,r}$, $R_{j,t} := \sum_{r=n-t+1}^n \omega_{j,r}$ and $N_j := L_{j,n} = R_{j,n}$ for $j \in [p]$ and $t \in [n]$, we define the MissCUSUM transformation $\mathcal{T}_{p,n}^{\text{Miss}} : \mathbb{R}^{p \times n} \times \{0, 1\}^{p \times n} \rightarrow \mathbb{R}^{p \times (n-1)}$ by

$$[\mathcal{T}_{p,n}^{\text{Miss}}(M, \Omega)]_{j,t} := \sqrt{\frac{L_{j,t} R_{j,n-t}}{N_j}} \left(\frac{1}{R_{j,n-t}} \sum_{r=t+1}^n (M \circ \Omega)_{j,r} - \frac{1}{L_{j,t}} \sum_{r=1}^t (M \circ \Omega)_{j,r} \right)$$

when $L_{j,t} > 0$ and $R_{j,n-t} > 0$, and define $[\mathcal{T}_{p,n}^{\text{Miss}}(M, \Omega)]_{j,t} := 0$ otherwise. Since the subscripts p and n of $\mathcal{T}_{p,n}^{\text{Miss}}$ can be inferred from the dimensions of its arguments, we will frequently abbreviate this transformation as $\mathcal{T}^{\text{Miss}}$. We note that this transformation only depends on M through $M \circ \Omega$. In practice, we will always apply this transformation to pairs of the form $(M \circ \Omega, \Omega)$; in other words, an entry of the first argument is zero whenever the corresponding entry of the second argument is zero. When the data matrix is fully observed (i.e. Ω is an all-one matrix), the MissCUSUM transformation reduces to the standard CUSUM transformation $\mathcal{T}(M) := \mathcal{T}^{\text{Miss}}(M, \Omega)$.

A key feature of the MissCUSUM transformation is that it captures the interaction between the signal strength and the number of observations in each coordinate. To illustrate this, we focus on a single (j th) coordinate, and the noiseless setting where $X = \mu$. In this case, the peak level of the absolute MissCUSUM transformation is $|\theta_j| \sqrt{\frac{L_{j,z} R_{j,n-z}}{N_j}}$. Since

$$\sqrt{\frac{\min(L_{j,z}, R_{j,n-z})}{2}} \leq \sqrt{\frac{L_{j,z} R_{j,n-z}}{N_j}} \leq \sqrt{\min(L_{j,z}, R_{j,n-z})},$$

we see that the peak level in the j th coordinate is controlled by the absolute mean change $|\theta_j|$, together with the effective sample size $\min(L_{j,z}, R_{j,n-z})$.

Another interesting property of the MissCUSUM transformation is that the multivariate setting allows us to borrow strength across the different coordinates to compensate for some of the missingness. To see this, note that the MissCUSUM transformation is piecewise constant in each coordinate. In particular, even in the noiseless setting, the absolute MissCUSUM series will typically not have a unique maximiser in each coordinate, but combining the information across coordinates allows us to pin down the changepoint location to an interval of length

$$\min \left\{ t > z : \sum_{j=1}^p \omega_{j,t} \neq 0 \right\} - \max \left\{ t \leq z : \sum_{j=1}^p \omega_{j,t} \neq 0 \right\}.$$

This will often be a shorter interval than one would obtain from any of the individual component series.

The next step of the MissInspect algorithm is to use the MissCUSUM transformation to find a good projection direction $\hat{v} \in \mathbb{S}^{p-1}$. The idea is that even though it is not possible to project the data along \hat{v} , due to the missingness, we can nevertheless compute the univariate series $((v^\top T_\Omega)_t)_{t \in [n-1]}$, where $T_\Omega := \mathcal{T}^{\text{Miss}}(X \circ \Omega, \Omega)$. Writing $A_\Omega := \mathcal{T}^{\text{Miss}}(\mu \circ \Omega, \Omega)$ and $A := \mathcal{T}(\mu)$, if each column of X has identity covariance matrix, then for a generic projection direction $v \in \mathbb{S}^{p-1}$, we find that $\mathbb{E}\{(v^\top T_\Omega)_t \mid \Omega\} = (v^\top A_\Omega)_t$, and $\text{Var}\{(v^\top T_\Omega)_t \mid \Omega\} = 1$. In Proposition 2 below, we will show that in a heterogeneous missingness model in which each entry in the j th coordinate is observed with probability q_j , and $\mathbf{q} := (q_1, \dots, q_p)^\top$, the matrix A_Ω can be well approximated by the rank one matrix $(\text{diag} \sqrt{\mathbf{q}}) A = (\theta \circ \sqrt{\mathbf{q}}) \gamma^\top$, where

$$\gamma := \frac{1}{\sqrt{n}} \left(\sqrt{\frac{1}{n-1}}(n-z), \dots, \sqrt{\frac{z-1}{n-z+1}}(n-z), \sqrt{z(n-z)}, \sqrt{\frac{n-z-1}{z+1}}z, \dots, \sqrt{\frac{1}{n-1}}z \right)^\top \in \mathbb{R}^{n-1}$$

attains its peak in absolute value at the true changepoint location z . Substituting this rank one approximation into the expression for $\mathbb{E}\{(v^\top T_\Omega)_t \mid \Omega\} = (v^\top A_\Omega)_t$, suggests that an oracle projection direction is a unit vector in the direction of $\theta \circ \sqrt{\mathbf{q}}$, which is the leading left singular vector of $(\text{diag} \sqrt{\mathbf{q}}) A$.

For the corresponding problem with fully observed data, Wang and Samworth (2018) proposed a semi-definite relaxation technique to estimating the oracle projection direction. Unfortunately, since A_Ω is not a rank one matrix when some data are missing, this relaxation turns out to be too coarse, and a new approach is required. Motivated by the fact that $\theta \circ \sqrt{\mathbf{q}}$ has the same sparsity pattern as θ , and viewing T_Ω as a perturbation of $(\text{diag} \sqrt{\mathbf{q}}) A$, we propose to estimate the oracle projection direction by solving the following optimisation problem:

$$(\hat{v}, \hat{w}) \in \operatorname{argmax}_{(\tilde{v}, \tilde{w}) \in \mathbb{B}^p \times \mathbb{B}^{n-1}} \{ \langle T_\Omega, \tilde{v} \tilde{w}^\top \rangle - \lambda \|\tilde{v}\|_1 \}, \tag{2}$$

where $\lambda > 0$ is a tuning parameter to be specified later. Here, with a suitable choice of λ , the ℓ_1 penalty on \tilde{v} in Equation (2) exploits the sparsity of the oracle projection direction to allow for consistent estimation of $(\theta \circ \sqrt{\mathbf{q}}) / \|\theta \circ \sqrt{\mathbf{q}}\|_2$, even when the dimension p is large, as will be shown in Proposition 1 in Section 3. A further advantage of Equation (2) over the semi-definite relaxation approach is that it directly exploits the row sparsity pattern of the rank one matrix $(\theta \circ \sqrt{\mathbf{q}}) \gamma^\top$, as opposed to just the overall entrywise sparsity of this matrix. Using the estimated oracle projection direction \hat{v} , we can project the MissCUSUM transformation T_Ω of $(X \circ \Omega, \Omega)$, and estimate the changepoint by the location of the maximum absolute value in the univariate projected series. Pseudocode for the MissInspect algorithm is given in Algorithm 1.

Algorithm 1: Pseudocode of the MissInspect algorithm

- Input:** $X_\Omega = X \circ \Omega \in \mathbb{R}^{p \times n}$, $\Omega \in \{0, 1\}^{p \times n}$, $\lambda > 0$
 1 $T_\Omega \leftarrow \mathcal{T}^{\text{Miss}}(X_\Omega, \Omega)$;
 2 Find $(\hat{v}, \hat{w}) \in \operatorname{argmax}_{\tilde{v} \in \mathbb{B}^p, \tilde{w} \in \mathbb{B}^{n-1}} \{ \langle T_\Omega, \tilde{v} \tilde{w}^\top \rangle - \lambda \|\tilde{v}\|_1 \}$;
 3 $\hat{z} \leftarrow \operatorname{median}(\operatorname{argmax}_{t \in [n-1]} |(\hat{v}^\top T_\Omega)_t|)$;
Output: \hat{z}
-

The optimisation problem in Step 2 of Algorithm 1 is bi-concave in (\tilde{v}, \tilde{w}) ; that is, the objective is concave in \tilde{v} for every fixed \tilde{w} and concave in \tilde{w} for every fixed \tilde{v} . Hence, we can alternate between optimising over \tilde{v} and \tilde{w} in (2). By inspecting the Karush–Kuhn–Tucker conditions as in Lemma 1 in the supplementary material, we see that when $\lambda < \|T_\Omega\|_{2 \rightarrow \infty}$, both steps of each iteration have closed form expressions; this motivates the iterative procedure to optimise (2) given in Algorithm 2. In that algorithm, we define the soft-thresholding function $\operatorname{soft} : \mathbb{R}^p \times [0, \infty) \rightarrow \mathbb{R}^p$ such that for $v = (v_1, \dots, v_p)^\top \in \mathbb{R}^p$, we have $(\operatorname{soft}(v, \lambda))_j = \operatorname{sgn}(v_j) \max\{|v_j| - \lambda, 0\}$ for $j \in [p]$. We remark that T_Ω is known to the practitioner, so we can always choose $\lambda < \|T_\Omega\|_{2 \rightarrow \infty}$. As usual for such iterative algorithms for bi-concave optimisation, the objective increases at each iteration; empirically, we have not observed any convergence issues. In Step 3 of Algorithm 1, we take the median of the maximisers of the series $\{ |(\hat{v}^\top T_\Omega)_t| : t \in [n - 1] \}$ for definiteness, although our theory would hold for any other (measurable) choice of element in this maximising set.

Algorithm 2: Pseudocode for an iterative procedure optimising (2)

- Input:** $T_\Omega \in \mathbb{R}^{p \times (n-1)}$, $\lambda \in (0, \|T_\Omega\|_{2 \rightarrow \infty})$
 1 $\tilde{v} \leftarrow$ leading left singular vector of T_Ω ;
 2 **repeat**
 3 $\tilde{w} \leftarrow \frac{T_\Omega^\top \tilde{v}}{\|T_\Omega^\top \tilde{v}\|_2}$;
 4 $\tilde{v} \leftarrow \frac{\operatorname{soft}(T_\Omega \tilde{w}, \lambda)}{\|\operatorname{soft}(T_\Omega \tilde{w}, \lambda)\|_2}$;
 5 **until convergence**;
Output: $(\hat{v}, \hat{w}) = (\tilde{v}, \tilde{w})$
-

We conclude this section by illustrating the MissInspect algorithm in action in Figure 1. Here, with $n = 250$ and $p = 100$, we generated n independent p -variate Gaussian observations with mean structure (1) and identity covariance matrix. We took $z = 100$, and $\theta = (\vartheta \mathbf{1}_k / k^{1/2}, \mathbf{0}_{p-k})^\top$, with $k = 10$ and $\vartheta = 2$. Thus, the first 10 coordinates represent signals, while the remaining 90




FIGURE 1 MissInspect algorithm in action. Top-left: visualisation of the data matrix with $p = 100$ and $n = 250$, where each column represents a p -dimensional observation and missing entries are shown in white. Darker colours indicate larger values. Time runs from left to right, and a change in mean occurs at time 100 in each of the first ten rows. Top-right: visualisation of the MissCUSUM transformation of the data. Bottom-left: the first five rows of the MissCUSUM matrix are plotted in colour, and the black curve shows the projected MissCUSUM series, which is maximised at the estimated changepoint location of 90 (black dashed line). The true changepoint is shown as a grey solid line. Bottom-right: histogram of estimated changepoints over 1000 repetitions from the same data setting; a log-concave estimated density is shown in red.

are noise coordinates. All entries of our data matrix were observed independently (and independently of the data), with probability 0.2. The top panels display visualisations of the data and the MissCUSUM transformation respectively. In the bottom-left panel, the coloured lines are the first five components of the MissCUSUM transformation; we see that these traces are piecewise constant, with jumps at observed data points. Even though each of these five is obtained from a signal coordinate, the locations of the peaks of these individual series would not yield very reliable changepoint estimates, both because the noise introduces considerable variability (e.g. the peak of the purple series runs from time 217 to 228), and because the missingness can lead to fairly long stretches where these series are constant. Nevertheless, once all 100 series are aggregated appropriately by our MissInspect algorithm, the resulting black trace does have a sharper peak close to the true changepoint. The bottom-right plot shows two non-parametric density estimates of the estimated changepoint locations from the MissInspect procedure over 1000 repetitions from this data generating mechanism; the first is a histogram, which requires the choice of a binwidth, while the second is the log-concave maximum likelihood estimator (Cule et al., 2010; Dümbgen & Rufibach, 2009), which is fully automatic. Both indicate a sharp peak for the density close to the true changepoint; in the latter case, the mode is exactly at 100.

3 | THEORETICAL GUARANTEES

We will focus our theoretical analysis on the single changepoint setting, in order to try to articulate more clearly the way that the coordinate-wise signal-to-noise ratio and missingness mechanism interact to determine both the performance of the `MissInspect` algorithm and the fundamental difficulty of the problem. Moreover, we assume that the revelation matrix $\Omega = (\omega_{j,t}) \in \{0, 1\}^{n \times p}$ has a *row-homogeneous* distribution, in the sense that there exists a vector $\mathbf{q} = (q_1, \dots, q_p)^\top \in (0, 1]^p$ such that $\omega_{j,t} \sim \text{Bern}(q_j)$, independently for all $j \in [p]$ and $t \in [n]$. We will refer to \mathbf{q} as the *observation rate vector*. Such a row-homogeneous assumption may be appropriate, for instance, in applications where each component series is measured by a separate device with its own observation rate. As for the data, we will assume that the columns $(X_t)_{t \in [n]}$ of the data matrix $X = (X_{j,t})_{j \in [p], t \in [n]} \in \mathbb{R}^{n \times p}$ satisfy

$$X_t \sim \mathcal{N}_p(\mu_t, \sigma^2 I_p), \text{ independently for } t \in [n], \tag{3}$$

where μ_1, \dots, μ_n satisfy Equation (1).

For $n \in \mathbb{N}, z \in [n - 1], \theta = (\theta_1, \dots, \theta_p)^\top \in \mathbb{R}^p, \sigma > 0$ and $\mathbf{q} = (q_1, \dots, q_p)^\top \in (0, 1]^p$ we write $P_{n,p,z,\theta,\sigma,\mathbf{q}}$ for the joint distribution of (X, Ω) , where X and Ω are independent, where X satisfies (3) with the vector of mean change $\theta := \mu^{(2)} - \mu^{(1)} \in \mathbb{R}^p \setminus \{0\}$ satisfying $\|\theta\|_0 \leq k$, and where Ω has a row-homogeneous distribution with observation rate vector $\mathbf{q} = (q_1, \dots, q_p)^\top \in (0, 1]^p$. We write $\tau := n^{-1} \min(z, n - z)$. Recall our notation $\|\theta\|_{2,\mathbf{q}}^2 := \sum_{j=1}^p \theta_j^2 q_j$, a quantity that captures a key interaction between the signal strength and observation rate. Our first result below shows that the projection direction \hat{v} obtained from Step 2 of Algorithm 1 is closely aligned with $\theta \circ \sqrt{\mathbf{q}}$, which, as argued in Section 2, can be regarded as an oracle projection direction.

Proposition 1 *Let $(X, \Omega) \sim P_{n,p,z,\theta,\sigma,\mathbf{q}}$ and let (\hat{v}, \hat{w}) be obtained from Step 2 in Algorithm 1, applied with inputs $X_\Omega = X \circ \Omega, \Omega$ and $\lambda \geq 2\sigma \sqrt{n \log(pn)}$. Then*

$$\mathbb{P} \left\{ \sin \angle(\hat{v}, \theta \circ \sqrt{\mathbf{q}}) \leq \frac{32\lambda \sqrt{k}}{n\tau \|\theta\|_{2,\mathbf{q}}} + \frac{112\|\theta\|_2}{\tau \|\theta\|_{2,\mathbf{q}}} \sqrt{\frac{6 \log(kn)}{n}} \right\} \geq 1 - \frac{6}{kn}.$$

Considering the case $\lambda = 2\sigma \sqrt{n \log(pn)}$ for simplicity, Proposition 1 reveals that, with high probability, the sine of the acute angle between \hat{v} and $\theta \circ \sqrt{\mathbf{q}}$ is controlled by the sum of two terms: the first of these represents the estimation error caused by the noise in the data we observe, and we see that $\|\theta\|_{2,\mathbf{q}}/\sigma$ can be thought of as an effective signal-to-noise ratio. On the other hand, the second term reflects the error due to our incomplete observations (and would be present even in the noiseless case with $\sigma = 0$); here $\|\theta\|_{2,\mathbf{q}}^2/\|\theta\|_2^2$ may be regarded as a *signal-weighted observation probability*. Observe that in this discussion we assume that σ is known; extensions to the case of unknown σ , as well as more general cross-sectional covariance structures, are discussed in Section 5.

From a theoretical point of view, the fact that \hat{v} is estimated using the entire available data set X_Ω makes it difficult to analyse the post-projection noise structure. For this reason, in the analysis below, we work with a sample-splitting variant of Algorithm 1, as given in Algorithm 3. Here, the projection direction \hat{v} is estimated using only the observed data at odd-numbered time points, and the `MissCUSUM` transformation of the observed data at even-numbered time points is then projected along \hat{v} to obtain the final estimate of the changepoint location. We emphasise, however, that while this sample splitting facilitates an informative theoretical analysis, in practice

we recommend using the full data set to estimate \hat{v} as in Algorithm 1. Indeed, it is this version of MissInspect that we used throughout our numerical studies in Section 4.

Algorithm 3: Pseudo-code for the sample-splitting variant of the MissInspect algorithm.

Input: $X_\Omega = X \circ \Omega \in \mathbb{R}^{p \times n}$, $\Omega \in \{0, 1\}^{p \times n}$, $\lambda > 0$

- 1 $n_1 \leftarrow \lfloor n/2 \rfloor$;
- 2 Let $\Omega^{(1)} \in \{0, 1\}^{p \times n_1}$ and $\Omega^{(2)} \in \{0, 1\}^{p \times n_1}$ denote the matrices formed from the first n_1 odd and the n_1 even numbered columns of Ω respectively;
- 3 Let $X_\Omega^{(1)} \in \mathbb{R}^{p \times n_1}$ and $X_\Omega^{(2)} \in \mathbb{R}^{p \times n_1}$ denote the matrices formed from the first n_1 odd and the n_1 even numbered columns of X_Ω respectively;
- 4 $T_\Omega^{(1)} \leftarrow \mathcal{T}^{\text{Miss}}(X_\Omega^{(1)}, \Omega^{(1)}) \in \mathbb{R}^{p \times (n_1-1)}$;
- 5 $T_\Omega^{(2)} \leftarrow \mathcal{T}^{\text{Miss}}(X_\Omega^{(2)}, \Omega^{(2)}) \in \mathbb{R}^{p \times (n_1-1)}$;
- 6 $(\hat{v}, \hat{w}) \leftarrow \operatorname{argmax}_{\hat{v} \in \mathbb{B}^p, \hat{w} \in \mathbb{B}^{n-1}} \{ \langle T_\Omega^{(1)}, \hat{v} \hat{w}^\top \rangle - \lambda \|\hat{v}\|_1 \}$;
- 7 $\hat{z} \leftarrow 2 \operatorname{median}(\operatorname{argmax}_{t \in [n_1-1]} |(\hat{v}^\top T_\Omega^{(2)})_t|)$;

Output: \hat{z}

Theorem 1 is our first main result on the performance of Algorithm 3 in contexts where our data are generated from the single changepoint, row-homogeneous model $P_{n,p,z,\theta,\sigma,\mathbf{q}}$.

Theorem 1 *Suppose $(X, \Omega) \sim P_{n,p,z,\theta,\sigma,\mathbf{q}}$. Assume for simplicity that n and z are even. Let \hat{z} be the output of Algorithm 3 with inputs $X \circ \Omega$, Ω and $\lambda = 2\sigma\sqrt{n \log(pn)}$. There exist universal constants $C, C' > 0$ such that whenever*

$$\frac{C'}{\tau} \sqrt{\frac{\log(pn)}{n}} \left(\frac{\sigma\sqrt{k}}{\|\theta\|_{2,\mathbf{q}}} + \frac{\|\theta\|_2}{\|\theta\|_{2,\mathbf{q}}} \right) \leq \frac{1}{2}, \quad (4)$$

we have

$$\mathbb{P} \left\{ \frac{|\hat{z} - z|}{n\tau} > C \sqrt{\frac{\log(kn)}{n\tau}} \left(\frac{\sigma}{\|\theta\|_{2,\mathbf{q}}} + \frac{\|\theta\|_2}{\|\theta\|_{2,\mathbf{q}}} \right) \right\} \leq \frac{22}{n}.$$

Condition (4) ensures that the projection direction \hat{v} obtained in Step 6 of Algorithm 3 has non-trivial correlation with the oracle projection direction $\theta \circ \sqrt{\mathbf{q}}$; cf. Proposition 1. An attractive feature of Theorem 1 is the way that the interaction between the signal strength and the observation rate vector is captured through $\|\theta\|_{2,\mathbf{q}}$. As mentioned in the introduction, this weighted average provides much greater understanding of the influence of missingness on the performance of the MissInspect algorithm than more naive bounds that depend on the worst-case missingness probability across all rows. For instance, we see that a high degree of missingness in noise or weak signal coordinates may not have too much of a detrimental effect on performance compared with complete observation of these coordinates. See also Theorem 3 below for confirmation of the way in which $\|\theta\|_{2,\mathbf{q}}$ also controls the fundamental difficulty of the problem (not just for our procedure).

A further attraction of Theorem 1 is the absence of any condition on the number of observations in each row. On the other hand, it turns out that if the expected number of observations in each row is at least k/τ^2 (up to logarithmic factors), then we can obtain a substantially improved bound on the rate of estimation of Algorithm 3.

Theorem 2 Suppose $(X, \Omega) \sim P_{n,p,z,\theta,\sigma,\mathbf{q}}$. Assume for simplicity that n and z are even. Let \hat{z} be the output of Algorithm 3 with inputs $X \circ \Omega, \Omega$ and $\lambda = 2\sigma\sqrt{n \log(pn)}$. Define

$$\rho := \frac{1}{\tau} \sqrt{\frac{\log(pn)}{n}} \left(\frac{\sigma\sqrt{k}}{\|\theta\|_{2,\mathbf{q}}} + \frac{\|\theta\|_2}{\|\theta\|_{2,\mathbf{q}}} \right).$$

Then there exist universal constants $c, C_1, C_2 > 0$ such that if $\rho \leq c$ and $n\tau^2 \min_{j \in [p]} q_j \geq C_1 k \log(pn)$, then

$$\mathbb{P} \left\{ \frac{|\hat{z} - z|}{n\tau} > \frac{C_2 \log(pn)}{n\tau} \left(\frac{\sigma^2}{\|\theta\|_{2,\mathbf{q}}^2} + \frac{\|\theta\|_\infty^2}{\|\theta\|_{2,\mathbf{q}}^2} \right) \right\} \leq \frac{23}{n}.$$

The rate obtained in Theorem 2 is essentially the square of that obtained in Theorem 1. In fact, an additional improvement is the reduction of $\|\theta\|_2$ in the second term to $\|\theta\|_\infty$. Again, we see the decomposition of the estimation error into terms reflecting the noise in the observed data and the incompleteness of the observations respectively.

As a complement to Theorem 2, we now present a minimax lower bound, which studies the fundamental limits of the expected estimation error that are achievable by any algorithm. We write $\tilde{\mathcal{Z}}$ for the set of estimators of z , that is, the set of Borel measurable functions $\hat{z} : \mathbb{R}^{n \times p} \times \{0, 1\}^{n \times p} \rightarrow [n - 1]$.

Theorem 3 Let $M \geq 1$ satisfy $\|\theta\|_\infty \leq M \min_{j \in [p]: \theta_j \neq 0} |\theta_j|$. If $\max\{\sigma^2, \|\theta\|_\infty^2 / (2M^2)\} \geq \|\theta\|_{2,\mathbf{q}}^2$, then there exists $c > 0$, depending only on M , such that for $n \geq 3$,

$$\inf_{\tilde{\mathcal{Z}}} \max_{z \in [n-1]} \mathbb{E}_P \frac{|\hat{z}(X \circ \Omega, \Omega) - z|}{n} \geq \frac{c}{n} \min \left\{ \frac{\sigma^2}{\|\theta\|_{2,\mathbf{q}}^2} + \frac{\|\theta\|_\infty^2}{\|\theta\|_{2,\mathbf{q}}^2}, n \right\}.$$

Theorem 3 reveals that the `MissInspect` algorithm as given in Algorithm 3 attains the minimax optimal estimation error rate up to logarithmic factors in all of the parameters of the problem, at least in settings where the signals are of comparable magnitude. Note that the `MissInspect` algorithm also matches (deterministically) the second term in the minimum in Theorem 3, because it trivially satisfies $|\hat{z} - z| \leq n - 2$. The form of the lower bound in Theorem 3 confirms that $\|\theta\|_{2,\mathbf{q}}$ is the correct functional of the mean change vector θ and observation \mathbf{q} for capturing the difficulty of the changepoint estimation problem in our missing data setting.

4 | NUMERICAL STUDIES

4.1 | Choice of tuning parameter

The tuning parameter choice of $\lambda = 2\sigma\sqrt{n \log(pn)}$ is convenient in our theoretical analysis. However, this choice often turns out to be slightly too conservative in practice, so to explore this, we considered the output of Algorithm 2 for a range of λ values, under several different settings of n, p, k, θ and \mathbf{q} . Figure 2 displays the mean angle between the estimated projection direction \hat{v} from Equation (2) and the oracle projection direction $\theta \circ \sqrt{\mathbf{q}} / \|\theta \circ \sqrt{\mathbf{q}}\|_2$ as a function of λ in two such

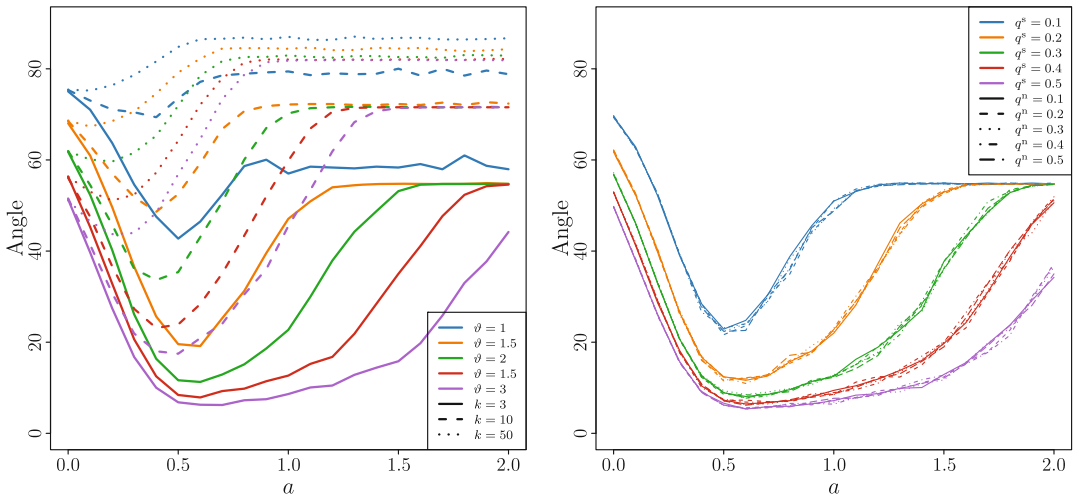


FIGURE 2 Mean angle in degrees (averaged over 200 repetitions) between the oracle projection direction and the estimated projection direction from Algorithm 2 with $\lambda = a\sigma\sqrt{n \log(pn)}$ for $a \in [0, 2]$. Data are generated under (3) with row-homogeneous missingness, independent of the data. Parameters: $n = 1000$, $p = 500$, $z = 400$, $\sigma = 1$. Left panel: all entries observed independently with probability 0.2, $k \in \{3, 10, 50\}$ and $\vartheta \in \{1, 1.5, 2, 2.5, 3\}$. Right panel: $k = 3$, $\vartheta = 2$, signal coordinates are observed with probability $q^s \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ and noise coordinates are observed with probability $q^n \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$.

sets of simulations. In both panels, we set $n = 1000$, $p = 500$, $z = 400$ and took $\lambda = a\sigma\sqrt{n \log(pn)}$ for $a \in [0, 2]$. The vector of mean change is $\theta = \vartheta k^{-1/2}(\mathbf{1}_k^\top, \mathbf{0}_{p-k}^\top)^\top$. Data were observed according to the row-homogeneous missingness model with $q_j = q^s$ if $\theta_j \neq 0$ and $q_j = q^n$ otherwise. In the left panel of the figure, we set $q^s = q^n = 0.2$ and vary $k \in \{3, 10, 50\}$ and $\vartheta \in \{1, 1.5, 2, 2.5, 3\}$, whereas in the right panel, we took $k = 3$, $\vartheta = 2$, $\sigma = 1$ and vary $q^s, q^n \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. We note that the choice $\lambda = 2^{-1}\sigma\sqrt{n \log(pn)}$ performs well in all settings, especially when the signal is relatively sparse. We therefore settled on this choice of λ throughout our numerical studies. It is reassuring to see from the right panel of Figure 2 that the performance of the projection direction estimator has almost no dependence on q^n , as predicted by our Proposition 1 (since $\|\theta\|_{2,q}$ does not depend on q^n).

4.2 | Validation of theoretical results

The aim of this subsection is to provide empirical confirmation of the forms of the bounds obtained in Proposition 1 and Theorem 2. In particular, we would like to verify that the crucial quantity $\|\theta\|_{2,q}$ does indeed capture the appropriate interaction between signal and missingness that determines the performance of the version of the MissInspect algorithm given in Algorithm 1. The two panels of Figure 3 study the angle between the estimated and oracle projection directions, and the estimated changepoint location error respectively. To obtain this figure, we set $n = 1200$, $p = 1000$ and generated data vectors under (3) with every entry observed independently with probability $q \in \{0.1, 0.2, 0.4, 0.8\}$, independent of the data. A single change occurred at $z = 400$ with vector of mean change $\theta = \vartheta k^{-1/2}(\mathbf{1}_k^\top, \mathbf{0}_{p-k}^\top)^\top$ and $k = 3$. We investigated the performance of MissInspect over 200 Monte Carlo repetitions for each of $\vartheta \in \{0.5, 1, 1.5, 2\}$ and $\sigma \in \{0.2, 0.4, 0.8, 1.6\}$. The left panel of Figure 3 shows that the

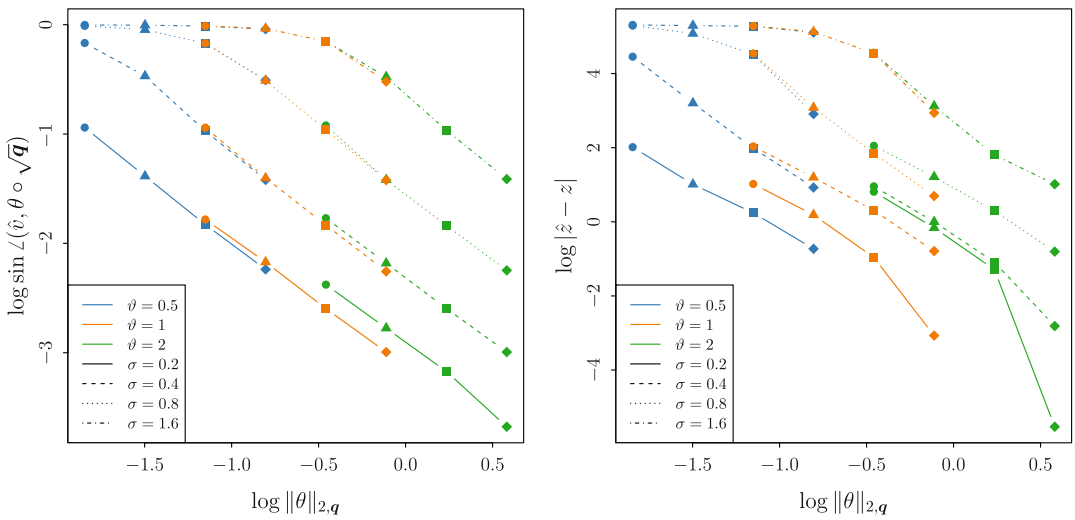


FIGURE 3 Estimation accuracy of `MissInspect` as a function of $\|\theta\|_{2,q}$. Data are generated under (3) with row-homogeneous missingness, independent of the data. Parameters: $n = 1200$, $p = 1000$, $z = 400$, $k = 3$, $\vartheta \in \{0.5, 1, 2\}$, $\sigma \in \{0.2, 0.4, 0.8, 1.6\}$ and $\mathbf{q} = q\mathbf{1}_p$ with $q \in \{0.1, 0.2, 0.4, 0.8\}$. Colours indicate ϑ , line type indicates σ , and the circle, triangle, square and diamond plotting characters correspond to $q = 0.1, 0.2, 0.4, 0.8$ respectively. Left panel: logarithm of the mean sine angle loss (averaged over 200 repetitions) between estimated and oracle projection directions. Right panel: logarithm of mean changepoint location loss (averaged over 200 repetitions).

logarithm of the mean sine angle loss decreases approximately linearly with $\log \|\theta\|_{2,q}$, with gradient approximately -1 . This is consistent with the conclusion of Proposition 1, which shows that the sine angle loss is controlled with high probability by an upper bound that is inversely proportional to $\|\theta\|_{2,q} = \vartheta q^{1/2}$. Moreover, curves corresponding to $\sigma = 0.2, 0.4, 0.8, 1.6$ are roughly equally spaced on the logarithmic scale, which corresponds to the linear dependence on $\lambda = 2^{-1}\sigma\sqrt{n\log(pn)}$ of the first term in the high-probability bound in Proposition 1. For fixed σ , the blue, orange and green curves are approximately overlapping, especially when the sine angle loss is large. In particular, doubling ϑ and reducing q by a factor of four leaves the sine angle loss virtually unchanged in these settings, which is consistent with the first term in the high probability upper bound in Proposition 1 being the dominant one, with its reciprocal dependence on $\|\theta\|_{2,q} = \vartheta q^{1/2}$. The contribution from the second term in Proposition 1 is still visible in the high signal-to-noise ratio settings, where, for instance when $\sigma = 0.2$, the $\vartheta = 2$ curve (green) lies above the $\vartheta = 1$ curve (orange). This is again consistent with the form of the second term in the bound in Proposition 1, which, in our setting, does not depend on ϑ , but is inversely proportional to $q^{1/2}$.

A similar story emerges in the right panel of Figure 3 for the changepoint location estimator accuracy. Here, for fixed ϑ and σ , most points lie on approximate straight lines with slope -2 , which is in agreement with the $\|\theta\|_{2,q}^{-2}$ dependence in the high probability bound of $|\hat{z} - z|$ in Theorem 2. The σ^2 dependence in the first term of the bound in Theorem 2 is represented by the mostly equi-spaced curves for the four different equi-spaced σ values on the logarithmic scale. The contribution of the second term in the bound in Theorem 2 can be seen from the three curves corresponding to the smallest noise scale $\sigma = 0.2$. Here, the estimation error only improves slightly as ϑ increases, which is in agreement with our theoretical prediction, since, in the setting

of this simulation, the second term in Theorem 2 is proportional to $q^{-1/2}$ and does not depend on ϑ .

4.3 | Comparison with alternative approaches

In this subsection, we compare the performance of `MissInspect` with three alternatives. The first two combine the idea of handling missing data via imputation and the original `inspect` procedure. Specifically, given a data matrix with incomplete observations, we consider applying the `softImpute` procedure of Mazumder et al. (2010), with the maximum matrix rank parameter set to 2 (since $X \circ \Omega$ can be viewed as a perturbation of its mean $(\text{diag} \sqrt{q})\mu$, which has rank 2). We then perform changepoint estimation on the imputed data matrix using the `inspect` procedure of Wang and Samworth (2018), with the suggested regularisation parameter choice therein. This method is denoted as `ImputeInspect`.

An alternative imputation strategy was mentioned briefly in the introduction: after an initial imputation of missing entries with the corresponding row mean, we iterate between applying the `inspect` algorithm to update our estimate of the changepoint location, and then re-imputing the missing entries with the row means to the left and right of the current changepoint location estimate. We call this method `IteratedMeanImputation`.

Our final comparator is an adaptation of the method of Londschiën et al. (2021), who consider dynamic Gaussian graphical models with piecewise constant, but unknown, mean vectors and precision matrices. Although in such a graphical modelling context, it is changes in the precision matrix that are of primary interest, one can nevertheless apply this method with the precision matrices fixed in advance to be the identity matrix. In that case, the method can be regarded as reporting the location of the peak of the absolute value of the ℓ_2 -aggregated `MissCUSUM` transformation of the observed data; equivalently, it seeks the most extreme generalised likelihood ratio statistic for testing the alternative of a change at a particular time point. We refer to this method as `GeneralisedLikelihoodRatio`.

Table 1 compares the performance of the four changepoint estimation algorithms `MissInspect`, `ImputeInspect`, `IteratedMeanImputation` and `GeneralisedLikelihoodRatio` under various settings. Here, we choose $n = 1200$, $p = 2000$, $k \in \{3, \lfloor \sqrt{p} \rfloor, p\}$, $\vartheta \in \{1, 2, 3\}$. Data are generated according to Equation (3) with row-homogeneous missingness, independent of the data. The changepoint occurs at $z = 400$, with vector of mean change having ℓ_2 norm ϑ and proportional to $(1, 2^{-1/2}, \dots, k^{-1/2}, 0, \dots, 0)^\top$. The observation rate vector $\mathbf{q} = (q_1, \dots, q_p)^\top$ is randomly generated, independent of all other sources of randomness, such that $q_j \stackrel{\text{iid}}{\sim} \text{Beta}(10\nu, 10(1-\nu))$ for $j \in [p]$, where $\nu \in \{0.1, 0.5\}$. Since the `MissInspect` and `ImputeInspect` procedures are based on a single projection, we can compare their performance in projection direction estimation; we also compare all methods in terms of their changepoint estimation risk. Note that the oracle projection direction for `ImputeInspect` is parallel to θ (since the imputed matrix has no missing entries), whereas the oracle projection direction of `MissInspect` is parallel to $\theta \circ \sqrt{q}$. We see in Table 1 that `MissInspect` consistently outperforms `ImputeInspect` and `IteratedMeanImputation`, often dramatically, for all observation fractions, sparsity levels and signal strengths considered. The `MissInspect` approach also achieves much smaller average estimation errors than the `GeneralisedLikelihoodRatio` method when there is some sparsity to the signal, with comparable performance when $k = p$. This is unsurprising in view of the fact that

TABLE 1 Location and projection direction estimation errors (averaged over 200 Monte Carlo repetitions) for `MissInspect` (denoted by superscript MI), `ImputeInspect` (denoted by superscript II), `IteratedMeanImputation` (denoted by IMI) and `GeneralisedLikelihoodRatio` (denoted by GLR). Other parameters: $n = 1200, p = 2000, z = 400, q_1, \dots, q_p \stackrel{iid}{\sim} \text{Beta}(10\nu, 10(1 - \nu))$. The best performing method is highlighted in bold in each case

ν	k	ϑ	$\angle(\hat{v}^{\text{MI}}, \theta \circ \sqrt{q})$	$\angle(\hat{v}^{\text{II}}, \theta)$	$ \hat{z}^{\text{MI}} - z $	$ \hat{z}^{\text{II}} - z $	$ \hat{z}^{\text{IMI}} - z $	$ \hat{z}^{\text{GLR}} - z $
0.1	3	1	71.4	86.8	141.7	468.0	184.1	212.4
0.1	3	2	40.6	56.7	36.5	304.8	147.7	139.8
0.1	3	3	26.1	40.1	14.5	257.5	101.0	66.6
0.1	44	1	82.6	88.9	185.9	468.9	187.6	209.4
0.1	44	2	63.5	83.2	66.9	404.5	133.7	118.3
0.1	44	3	49.0	72.8	18.7	308.6	90.8	52.0
0.1	2000	1	86.5	88.2	180.0	485.0	184.1	219.6
0.1	2000	2	76.9	87.6	121.2	457.3	138.9	137.5
0.1	2000	3	67.7	82.9	50.4	376.9	79.2	41.0
0.5	3	1	32.3	81.0	11.9	358.4	150.8	176.0
0.5	3	2	13.6	42.1	1.6	7.2	44.8	10.5
0.5	3	3	9.6	24.8	0.7	6.9	7.6	2.1
0.5	44	1	62.7	88.4	50.1	438.5	159.4	207.1
0.5	44	2	37.3	73.6	2.3	174.2	41.8	7.3
0.5	44	3	26.9	58.1	0.7	1.8	3.3	1.6
0.5	2000	1	77.5	88.6	114.3	448.1	162.5	202.9
0.5	2000	2	59.2	85.5	6.7	338.6	40.6	6.8
0.5	2000	3	52.0	72.4	1.7	48.2	3.9	1.7

the `GeneralisedLikelihoodRatio` procedure does not exploit sparsity in the vector of mean change.

4.4 | Real data analysis

In this subsection, we illustrate the applicability of the `MissInspect` algorithm on an oceanographic data set covering the Neogene geological period. Oceanographers study historic changes in the global ocean circulation system by examining microfossils that record the isotopic composition of water at the time at which they lived (Wright & Miller, 1996). In particular, large cores are extracted from the ocean floor and a species of microfossils called foraminifera are taken from small slices of sediment at different depths within the core. The ratio of the abundances of ^{13}C to ^{12}C isotopes in their calcium carbonate shells is compared against a standard, to understand the carbon composition within the oceans during their lifetime, and hence to determine the direction in which ocean currents flowed. The depth of the foraminifera within the core is used as a proxy for the geological age of the fossil, measured in millions of years (Ma).

Our data, which are available in a GitHub repository², and were previously analysed by Samworth and Poore (2005) and Poore et al. (2006), consist of measurements from 16 cores extracted from the North Atlantic, Pacific and Southern Oceans and are displayed in Figure 4. In total, there are 7369 observations at 6295 distinct time points, but Figure 4 makes clear that the heterogeneous nature of the data collection process means that it is appropriate to think of the data as containing missingness. The figure also indicates the 10 most prominent changepoints identified by applying the `MissInspect` algorithm in combination with binary segmentation, as discussed in Section 5 below. It is notable that the first changepoint, found by applying the algorithm to the full data set, occurs at 6.13Ma, a time that has previously been identified as a time of rapid change in oceanographic current flow (Poore et al., 2006, p. 13).

5 | EXTENSIONS

As mentioned in the introduction, one of our main theoretical goals in this work is to understand the way in which the missingness and the signal interact in changepoint problems to determine the difficulty of the problem. This is particularly challenging when we seek to handle both high dimensionality and different levels of missingness in different coordinates. For the purposes of our theoretical analysis, then, it is natural to impose stronger assumptions elsewhere, so as to best expose the interesting phenomena at play. Nevertheless, it remains of interest to consider the extent to which the methodology could be generalised, and the assumptions could be relaxed, to cover a wider range of scenarios and problems one might see in practice.

As we saw in analysing the oceanography data in Section 4.4, it may be that we wish to identify multiple changepoints. There are several standard techniques for extending single changepoint procedures to such settings, including binary segmentation and different versions of wild binary segmentation (Fryzlewicz, 2014; Kovács et al., 2020). Any of these approaches can be used in conjunction with the `MissInspect` algorithm to identify multiple changepoints in high-dimensional data streams in the presence of missingness. The theoretical analysis of such a procedure would be technically involved, but would proceed along similar lines to that of Wang and Samworth (2018) for the case of fully observed data.

As always when handling missing data, the situation becomes much more complicated when the missingness and the data are not independent, that is, the missing completely at random (MCAR) assumption does not hold. In the worst case, the missingness may render the changepoint estimation problem impossible, for instance if no signal coordinate has observed data on both sides of the changepoint. A less adversarial setting would be one in which all observations exceeding 1 are censored. Thus, if the vector of mean change had positive entries in signal coordinates, we would expect to see fewer observations in these coordinates after the change. The censoring would lead to (different) truncated Gaussian distributions before and after the change, but a difference in mean of these distributions would persist, so changepoint estimation may still be possible. In general, careful and problem-specific modelling of the dependence of the data and the missingness mechanism is recommended.

In this paper, we have focused on the problem of estimating a changepoint location in settings where such a change is known to exist. An interesting complementary problem is that of detecting whether or not there is a change; in other words, testing $H_0 : \theta = 0$ against $H_1 : \theta \neq 0$. To address

²https://github.com/wangtengyao/MissInspect/tree/main/real_data




FIGURE 4 Ratios of carbon isotope measurements taken from foraminifera in 16 different cores from the North Atlantic, Pacific and Southern Oceans. The label of each panel indicates both the ocean and the number of the core, while the horizontal axis measures geological time (0–23 Ma). The red dashed lines indicate the 10 most prominent changepoints identified by applying the *MissInspect* algorithm in combination with binary segmentation, with the most significant change plotted with a solid line.

this challenge, first, for $j \in [p]$, let $T_{j,1} := \min\{t \in [n] : \omega_{j,t} = 1\}$ denote the first observation time in the j th coordinate, and for $\ell = 2, \dots, N_j$, let $T_{j,\ell} := \min\{t > T_{j,\ell-1} : \omega_{j,t} = 1\}$ denote the ℓ th observation time in the j th coordinate. Then, under the null hypothesis and conditional on Ω ,

$$(\mathcal{T}^{\text{Miss}}(X \circ \Omega, \Omega)_{j, T_{j,\ell}} : \ell \in [N_j - 1]) \stackrel{d}{=} \left(\frac{B_t}{\sqrt{t(1-t)}} : t = 1/N_j, \dots, (N_j - 1)/N_j \right), \quad (5)$$

where $(B_t)_{t \in [0,1]}$ denotes a standard Brownian bridge. Under the alternative hypothesis, the conditional distribution of the absolute value of the peak of the left-hand side in (5), given Ω , is stochastically larger than under the null. This leads to various possible test statistics, including

$$T^{(1)} := \max_{j \in [p]} \max_{\ell \in [N_j - 1]} |\mathcal{T}^{\text{Miss}}(X \circ \Omega, \Omega)_{j, T_{j,\ell}}| \quad \text{and} \quad T^{(2)} := \max_{\ell \in [N_j - 1]} \sum_{j=1}^p \mathcal{T}^{\text{Miss}}(X \circ \Omega, \Omega)_{j, T_{j,\ell}}^2,$$

with pivotal null distributions, conditional on Ω .

Finally, we discuss settings of temporal and spatial dependence in the data. In the former case, a natural model is to replace (3) with

$$X_t = \mu_t + W_t, \quad \text{for } t = 1, \dots, n,$$

where μ_1, \dots, μ_n satisfy (1) and where the noise vectors (W_1, \dots, W_n) form a mean-zero, stationary Gaussian process. In this case, with row-homogeneous missingness independent of the data, the oracle projection direction remains $(\theta \circ \sqrt{\mathbf{q}}) / \|\theta \circ \sqrt{\mathbf{q}}\|_2$, and the `MissInspect` methodology does not need to be altered. On the other hand, if spatial dependence is introduced into Equation (1) by replacing the identity covariance matrix there with a general covariance matrix Σ , then the oracle projection direction becomes proportional to $\Sigma^{-1}(\theta \circ \sqrt{\mathbf{q}})$. If Σ is unknown, then estimating Σ^{-1} may represent a significant challenge, but it may be considerably simplified if our data stream satisfies additional structural assumptions. For instance, if $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$, where $\sigma_1, \dots, \sigma_p$ are unknown, then we can estimate these quantities robustly using, for example, the median absolute deviation of the marginal one-dimensional series (Hampel, 1974). As another example, if Σ is Toeplitz with $\Sigma = (\rho^{|j-k|})_{j,k \in [p]}$ for some $\rho \in (-1, 1)$, then Σ^{-1} is tridiagonal, and its form can again be used to estimate ρ (Wang & Samworth, 2018, lemma 12).

6 | PROOF OF MAIN RESULTS

6.1 | Proof of Proposition 1

The proof of Proposition 1 requires the following result.

Proposition 2 For $n \geq 2$, suppose $(X, \Omega) \sim P_{n,p,z,\theta,\sigma,\mathbf{q}}$ and let $A = \mathcal{T}(\mathbb{E}(X)) \in \mathbb{R}^{p \times (n-1)}$ and $A_\Omega = \mathcal{T}^{\text{Miss}}(\mathbb{E}(X) \circ \Omega, \Omega) \in \mathbb{R}^{p \times (n-1)}$. We have that

$$\mathbb{P}(|(A_\Omega)_{j,t} - \sqrt{q_j} A_{j,t}| > 7\sqrt{6}|\theta_j| \sqrt{\log(kn)}) \leq \frac{4}{k^2 n^2}$$

for all $j \in [p]$ and $t \in [n - 1]$. Consequently,

$$\mathbb{P}(\|A_\Omega - (\text{diag } \sqrt{\mathbf{q}})A\|_F > 7\sqrt{6}\|\theta\|_2\sqrt{n \log(kn)}) \leq \frac{4}{kn}.$$

Proof. Let $\Delta := A_\Omega - (\text{diag } \sqrt{\mathbf{q}})A \in \mathbb{R}^{p \times (n-1)}$ have (j, t) th entry $\Delta_{j,t}$ and let $S := \{j \in [p] : \theta_j \neq 0\}$. Since $\Delta_{j,t} = 0$ for $j \notin S$, it suffices to bound $|\Delta_{j,t}|$ for each $j \in S$ and $t \in [n - 1]$. Without loss of generality, we may assume that $\theta_j > 0$.

First assume that j satisfies $n\tau q_j \geq 24 \log(kn)$. Let $a_j := \left\lceil \frac{8 \log(kn)}{3q_j} \right\rceil$. It follows that $a_j \leq n\tau/6$. Define

$$\delta_{j,t} := \begin{cases} \frac{4 \log(kn)}{tq_j} & \text{if } 1 \leq t \leq a_j - 1 \\ \sqrt{\frac{6 \log(kn)}{tq_j}} & \text{if } a_j \leq t \leq n. \end{cases}$$

Recall that $L_{j,t} = \sum_{r=1}^t \omega_{j,r}$ and $R_{j,t} = \sum_{r=n-t+1}^n \omega_{j,r}$ for $j \in [p]$ and $t \in [n]$. We consider the event

$$\mathcal{A}_{j,t} := \left\{ \max \left(\left| \frac{L_{j,t}}{tq_j} - 1 \right|, \left| \frac{R_{j,t}}{tq_j} - 1 \right| \right) \leq \delta_{j,t} \right\}.$$

By Bernstein’s inequality (Lemma 3 in the supplementary material), we obtain that

$$\begin{aligned} \mathbb{P}(\mathcal{A}_{j,t}^c) &\leq 4 \exp \left\{ -\frac{\delta_{j,t}^2 tq_j}{2(1 + \delta_{j,t}/3)} \right\} \\ &= 4 \exp \left\{ -\frac{8 \log^2(kn)}{tq_j + 4 \log(kn)/3} \right\} \mathbb{1}_{\{t < a_j\}} + 4 \exp \left\{ -\frac{3 \log(kn)}{1 + \sqrt{2 \log(kn)/(3tq_j)}} \right\} \mathbb{1}_{\{t \geq a_j\}} \\ &\leq 4e^{-2 \log(kn)} \leq \frac{4}{k^2 n^2}, \end{aligned}$$

where we have used the facts that $tq_j \leq 8 \log(kn)/3$ for $t \leq a_j - 1$ and $\sqrt{2 \log(kn)} \leq \sqrt{3tq_j}/2$ for $t \geq a_j$ in the penultimate inequality.

We will now bound $|\Delta_{j,t}|$ for different values of t on $\mathcal{A}_{j,t}$. First, consider $a_j \leq t \leq z$. Since $n - z \geq n\tau > a_j$, we have $\delta_{j,n-z} \leq \delta_{j,n\tau} \leq 1/2$. We deduce from the definition of A_Ω that

$$\begin{aligned} (A_\Omega)_{j,t} &= \sqrt{\frac{L_{j,t}}{R_{j,n-t}(L_{j,t} + R_{j,n-t})}} R_{j,n-z} \theta_j \leq \sqrt{\frac{(1 + \delta_{j,t})tq_j}{(1 - \delta_{j,n-t})^2(n - t)nq_j^2}} (1 + \delta_{j,n-z})(n - z)q_j \theta_j \\ &\leq \sqrt{q_j} A_{j,t} \frac{\sqrt{1 + \delta_{j,t}(1 + \delta_{j,n-z})}}{1 - \delta_{j,n-z}}. \end{aligned} \tag{6}$$

It follows that for $j \in S$ and $a_j \leq t \leq z$,

$$\begin{aligned} \Delta_{j,t} &\leq \sqrt{q_j} A_{j,t} \left\{ \frac{\sqrt{1 + \delta_{j,t}(1 + \delta_{j,n-z})}}{1 - \delta_{j,n-z}} - 1 \right\} \leq \sqrt{q_j} A_{j,t} \{(1 + \delta_{j,t})(1 + 4\delta_{j,n-z}) - 1\} \\ &\leq \sqrt{q_j} A_{j,t} (3\delta_{j,t} + 4\delta_{j,n-z}). \end{aligned}$$

By a similar calculation for deviations in the opposite direction, we have

$$\Delta_{j,t} \geq -\sqrt{q_j} A_{j,t} (\delta_{j,t} + 4\delta_{j,n-z}/3).$$

Thus, using the fact that $A_{j,t} \leq \theta_j \min(\sqrt{t}, \sqrt{n-z})$, we deduce that

$$|\Delta_{j,t}| \leq \sqrt{q_j} A_{j,t} (3\delta_{j,t} + 4\delta_{j,n-z}) \leq 7\sqrt{q_j} \theta_j \min(\sqrt{t}, \sqrt{n-z}) \max(\delta_{j,t}, \delta_{j,n-z}) \leq 7\theta_j \sqrt{6 \log(kn)}.$$

By symmetry, if $z < t \leq n - a_j$, we also have $|\Delta_{j,t}| \leq 7\theta_j \sqrt{6 \log(kn)}$.

Next, if $t \leq a_j - 1$, then we necessarily have $t \leq n\tau$. The calculation in (6) still applies, and we have

$$(A_\Omega)_{j,t} \leq \sqrt{q_j} A_{j,t} \frac{\sqrt{1 + \delta_{j,t}}(1 + \delta_{j,n-z})}{1 - \delta_{j,n-z}} \leq 3\sqrt{q_j} A_{j,t} \sqrt{1 + \delta_{j,t}}.$$

Hence, since $\text{sgn}((A_\Omega)_{j,t}) = \text{sgn}(A_{j,t})$, we have

$$\begin{aligned} |\Delta_{j,t}| &\leq \max((A_\Omega)_{j,t}, \sqrt{q_j} A_{j,t}) \leq A_{j,t} \sqrt{q_j} \max(3\sqrt{1 + \delta_{j,t}}, 1) \\ &\leq \theta_j \sqrt{tq_j} \max\left\{3\sqrt{\frac{tq_j + 4 \log(kn)}{tq_j}}, 1\right\} \leq 8\theta_j \sqrt{\log(kn)}. \end{aligned}$$

A symmetric argument shows that $|\Delta_{j,t}| \leq 8\theta_j \sqrt{\log(kn)}$ for $n - a_j \leq t \leq n - 1$. Combining the above bounds on $|\Delta_{j,t}|$, we see that for j satisfying $n\tau q_j \geq 24 \log(kn)$ and all $t \in [n-1]$, we have that

$$\mathbb{P}(|\Delta_{j,t}| > 7\sqrt{6}\theta_j \sqrt{\log(kn)}) \leq \mathbb{P}(\mathcal{A}_{j,t}^c) \leq \frac{4}{k^2 n^2}. \quad (7)$$

We now turn our attention to j satisfying $n\tau q_j < 24 \log(kn)$. If $q_j = 0$, then $\Delta_{j,t} = 0$. So we may assume $q_j > 0$. Define

$$\epsilon_j := \frac{24 \log(kn)}{n\tau q_j},$$

so that $\epsilon_j > 1$. For $j \in S$, consider the event

$$\mathcal{B}_j := \left\{ \max\left(\frac{L_{j,z}}{zq_j}, \frac{R_{j,n-z}}{(n-z)q_j}\right) \leq 1 + \epsilon_j \right\}.$$

By Lemma 3 in the supplement again, we have

$$\mathbb{P}(\mathcal{B}_j^c) \leq 2 \exp\left\{-\frac{\epsilon_j^2 n\tau q_j}{2(1 + \epsilon_j/3)}\right\} \leq 2e^{-9 \log(kn)} = \frac{2}{(kn)^9}.$$

On \mathcal{B}_j , we have

$$(A_\Omega)_{j,t} \leq (A_\Omega)_{j,z} \leq \theta_j \sqrt{\min\{L_{j,z}, R_{j,n-z}\}} \leq \theta_j \sqrt{(1 + \epsilon_j)n\tau q_j} \leq \theta_j \sqrt{48 \log(kn)}.$$

On the other hand,

$$\sqrt{q_j}A_{j,t} \leq \sqrt{q_j}A_{j,z} \leq \sqrt{q_j} \min\{\sqrt{z}, \sqrt{n-z}\}\theta_j = \theta_j\sqrt{n\tau q_j} \leq \theta_j\sqrt{24 \log(kn)}.$$

Consequently, when $n\tau q_j < 24 \log(kn)$, we have

$$\mathbb{P}(|\Delta_{j,t}| > \theta_j\sqrt{48 \log(kn)}) \leq \mathbb{P}((A_\Omega)_{j,t} > \theta_j\sqrt{48 \log(kn)}) \leq \mathbb{P}(\mathcal{B}_j^c) \leq \frac{2}{(kn)^9}. \tag{8}$$

The first claim follows from Equations (7) and (8). It now follows that

$$\mathbb{P}(\|\Delta\|_F > 7\sqrt{6}\|\theta\|_2\sqrt{n \log(kn)}) \leq \sum_{j \in \mathcal{S}} \sum_{t=1}^{n-1} \mathbb{P}(|\Delta_{j,t}| > 7\sqrt{6}\theta_j\sqrt{\log(kn)}) \leq \frac{4}{kn},$$

as desired.

Proof of Proposition 1. Let $v \in \mathbb{S}^{p-1}$ denote the leading left singular vector of A_Ω and let $\sigma_1 \geq \sigma_2 \geq 0$ denote the two largest singular values of A_Ω . We start by controlling the angle between \hat{v} and v . Write $\Delta := A_\Omega - (\text{diag} \sqrt{\mathbf{q}})A \in \mathbb{R}^{p \times (n-1)}$ as in the proof of Proposition 2. Since $A = \theta\gamma^\top$, we have $(\text{diag} \sqrt{\mathbf{q}})A = (\theta \circ \sqrt{\mathbf{q}})\gamma^\top$. Hence, by Weyl’s inequality (e.g. Stewart & Sun, 1990, Corollary IV.4.9), we obtain

$$\sigma_1 - \sigma_2 \geq \|\theta\|_{2,\mathbf{q}}\|\gamma\|_2 - 2\|\Delta\|_{\text{op}} \geq \frac{n\tau\|\theta\|_{2,\mathbf{q}}}{4} - 2\|\Delta\|_F,$$

where the final bound uses Wang and Samworth (2018 Lemma 3). By Proposition 2, there is an event \mathcal{A} with probability at least $1 - 4/(kn)$ such that $\|\Delta\|_F \leq 7\sqrt{6}\|\theta\|_2\sqrt{n \log(kn)}$. We may assume that $\sqrt{n\tau}\|\theta\|_2 \geq 112\|\theta\|_{2,\mathbf{q}}\sqrt{6 \log(kn)}$, since otherwise, the proposition is trivially true. With this assumption, we have on \mathcal{A} that $\sigma_1 - \sigma_2 \geq n\tau\|\theta\|_{2,\mathbf{q}}/8$. Thus, by Lemma 2 in the supplementary material, on the event $\mathcal{A} \cap \{\|T_\Omega - A_\Omega\|_\infty \leq \lambda n^{-1/2}\}$, we have that

$$\sin \angle(\hat{v}, v) \leq \frac{4\lambda\sqrt{k}}{\sigma_1 - \sigma_2} \leq \frac{32\lambda\sqrt{k}}{n\tau\|\theta\|_{2,\mathbf{q}}}. \tag{9}$$

On the other hand, by Wang (2016, Theorem 1.4) (an extension of Yu et al., 2015, Corollary 1), on \mathcal{A} , we also have that

$$\sin \angle(v, \theta \circ \sqrt{\mathbf{q}}) \leq \frac{4\|\Delta\|_{\text{op}}}{n\tau\|\theta\|_{2,\mathbf{q}}/4} \leq \frac{112\|\theta\|_2}{\tau\|\theta\|_{2,\mathbf{q}}} \sqrt{\frac{6 \log(kn)}{n}}. \tag{10}$$

By the triangle inequality, we deduce from Equations (9) and (10) that on $\mathcal{A} \cap \{\|T_\Omega - A_\Omega\|_\infty \leq \lambda n^{-1/2}\}$,

$$\sin \angle(\hat{v}, \theta \circ \sqrt{\mathbf{q}}) \leq \frac{32\lambda\sqrt{k}}{n\tau\|\theta\|_{2,\mathbf{q}}} + \frac{112\|\theta\|_2}{\tau\|\theta\|_{2,\mathbf{q}}} \sqrt{\frac{6 \log(kn)}{n}}.$$

The proposition follows on observing that

$$\begin{aligned} \mathbb{P}(\mathcal{A}^c \cup \{\|T_\Omega - A_\Omega\|_\infty > \lambda n^{-1/2}\}) &\leq \frac{4}{kn} + \sum_{j=1}^p \sum_{t=1}^{n-1} \mathbb{P}(|(T_\Omega)_{j,t} - (A_\Omega)_{j,t}| > \lambda n^{-1/2}) \\ &\leq \frac{4}{kn} + pne^{-\lambda^2/(2n\sigma^2)} \leq \frac{6}{kn}, \end{aligned}$$

where the penultimate inequality uses the fact that $(A_\Omega)_{j,t} - (T_\Omega)_{j,t} \mid \Omega \sim N(0, \sigma^2)$ for all $t \in [n-1]$ and $j \in [p]$ such that $L_{j,t}R_{j,n-t} \neq 0$, and is equal to 0 when $L_{j,t}R_{j,n-t} = 0$.

6.2 | Proof of Theorem 1

Proof of Theorem 1. Recall from Algorithm 3 that $n_1 = n/2$, and for $\ell \in \{1, 2\}$, let $\Omega^{(\ell)} \in \{0, 1\}^{p \times n_1}$, $X^{(\ell)} \in \mathbb{R}^{p \times n_1}$ and $X_\Omega^{(\ell)} \in \mathbb{R}^{p \times n_1}$ denote the matrices formed from the n_1 odd columns (when $\ell = 1$) and the n_1 even numbered columns (when $\ell = 2$) of Ω , X and $X_\Omega = X \circ \Omega$ respectively. For $\ell \in \{1, 2\}$, let $T_\Omega^{(\ell)} := \mathcal{T}^{\text{Miss}}(X_\Omega^{(\ell)}, \Omega^{(\ell)}) \in \mathbb{R}^{p \times (n_1-1)}$. By Proposition 1, the output \hat{v} of Algorithm 1 with inputs $T_\Omega^{(1)}$ and λ satisfies

$$\mathbb{P} \left\{ \sin \angle(\hat{v}, \theta \circ \sqrt{\mathbf{q}}) > \frac{32\lambda\sqrt{k}}{n_1\tau\|\theta\|_{2,\mathbf{q}}} + \frac{112\|\theta\|_2}{\tau\|\theta\|_{2,\mathbf{q}}} \sqrt{\frac{6 \log(kn_1)}{n_1}} \right\} \leq \frac{6}{kn_1} = \frac{12}{kn}. \tag{11}$$

We can therefore find a universal constant $C' > 0$ such that whenever (4) holds, we have that the event $\mathcal{A} := \{\sin \angle(\hat{v}, \theta \circ \sqrt{\mathbf{q}}) \leq 1/2\}$ has probability at least $1 - 12/(kn)$.

Writing $\mu^{(2)} := \mathbb{E}(X^{(2)}) \in \mathbb{R}^{p \times n}$, let $A^{(2)} = \mathcal{T}(\mu^{(2)})$ and $A_\Omega^{(2)} = \mathcal{T}^{\text{Miss}}(\mu^{(2)} \circ \Omega^{(2)}, \Omega^{(2)})$. Our main decomposition of interest here is

$$T_\Omega^{(2)} = (\text{diag} \sqrt{\mathbf{q}})A^{(2)} + \Delta^{(2)} + E_\Omega^{(2)},$$

where $\Delta^{(2)} := A_\Omega^{(2)} - (\text{diag} \sqrt{\mathbf{q}})A^{(2)}$ and $E_\Omega^{(2)} := T_\Omega^{(2)} - A_\Omega^{(2)}$. Since Algorithm 3 remains the same if we replace \hat{v} in Step 6 with $-\hat{v}$, we may assume without loss of generality that $\hat{v}^\top (\theta \circ \sqrt{\mathbf{q}}) \geq 0$. Since $(\text{diag} \sqrt{\mathbf{q}})A^{(2)} = (\theta \circ \sqrt{\mathbf{q}})\gamma^{(2)\top}$, where $\gamma^{(2)} = (\gamma_1^{(2)}, \dots, \gamma_{n_1-1}^{(2)})^\top \in \mathbb{R}^{n_1-1}$

$$\gamma_t^{(2)} := \begin{cases} \sqrt{\frac{t}{(n_1-t)n_1}}(n_1 - z/2) & \text{if } t \leq z/2, \\ \sqrt{\frac{(n_1-t)}{n_1}}(z/2) & \text{if } t > z/2, \end{cases}$$

we have $(\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}})A^{(2)})_t \geq 0$ for all $t \in [n_1 - 1]$. On the event \mathcal{A} , we have

$$(\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}})A^{(2)})_{z/2} \geq \frac{\sqrt{3}}{2} \|\theta\|_{2,\mathbf{q}} \gamma_{z/2}^{(2)} \geq \frac{\sqrt{3}}{4} \|\theta\|_{2,\mathbf{q}} \sqrt{n\tau}. \tag{12}$$

Observe that for every $t \in [n_1 - 1]$, we have

$$(\hat{v}^\top E_\Omega^{(2)})_t \mid \Omega^{(2)} \sim N(0, \sigma^2 \|\hat{v}_t\|_2^2),$$

where $J_t := \{j \in [p] : \min(\sum_{r=1}^t (\Omega^{(2)})_{j,r}, \sum_{r=t+1}^{n_1} (\Omega^{(2)})_{j,r}) > 0\}$. Since $\|\hat{v}_{J_t}\|_2 \leq 1$, we deduce that $(\hat{v}^\top E_\Omega^{(2)})_t$ is stochastically dominated by $N(0, \sigma^2)$. Hence, together with the first conclusion of Proposition 2 and a union bound, there exists an event \mathcal{B} with probability at least $1 - 4/(kn_1) - 1/n_1$ such that on \mathcal{B} we have

$$\max_{t \in [n_1-1]} |(\hat{v}^\top \Delta^{(2)})_t| \leq 7\sqrt{6}\|\theta\|_2 \sqrt{\log(kn)} \quad \text{and} \quad \max_{t \in [n_1-1]} |(\hat{v}^\top E_\Omega^{(2)})_t| \leq 2\sigma \sqrt{\log n}. \tag{13}$$

Combining (12) and (13), and by increasing the universal constant $C' > 0$ if necessary, we have by Equation (4) that on $\mathcal{A} \cap \mathcal{B}$,

$$\begin{aligned} (\hat{v}^\top T_\Omega^{(2)})_{z/2} &= (\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}}) A^{(2)})_{z/2} + (\hat{v}^\top E_\Omega^{(2)})_{z/2} + (\hat{v}^\top \Delta^{(2)})_{z/2} \\ &\geq \max \left\{ 0, \max_{t \in [n_1-1]} \left\{ -(\hat{v}^\top E_\Omega^{(2)})_t - (\hat{v}^\top \Delta^{(2)})_t \right\} \right\} > \max_{t \in [n_1-1]} \left(-\hat{v}^\top T_\Omega^{(2)} \right)_t. \end{aligned}$$

In particular, on $\mathcal{A} \cap \mathcal{B}$, we have from the definition of \hat{z} that $(\hat{v}^\top T_\Omega^{(2)})_{\hat{z}/2} \geq (\hat{v}^\top T_\Omega^{(2)})_{z/2} \geq 0$, so on this event we have the basic inequality

$$\begin{aligned} (\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}}) A^{(2)})_{z/2} - (\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}}) A^{(2)})_{\hat{z}/2} &\leq |(\hat{v}^\top E_\Omega^{(2)})_{z/2} - (\hat{v}^\top E_\Omega^{(2)})_{\hat{z}/2}| + |(\hat{v}^\top \Delta^{(2)})_{z/2} - (\hat{v}^\top \Delta^{(2)})_{\hat{z}/2}|. \end{aligned} \tag{14}$$

By (Wang and Samworth 2018 Lemma 7) on the event $\mathcal{A} \cap \mathcal{B}$, for every $t \in [n_1 - 1]$, we have

$$\begin{aligned} (\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}}) A^{(2)})_{z/2} - (\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}}) A^{(2)})_t &= |\hat{v}^\top (\sqrt{\mathbf{q}} \circ \theta)| (\gamma_{z/2} - \gamma_t) \\ &\geq \frac{\sqrt{3}}{2} \|\theta\|_{2,\mathbf{q}} \cdot \frac{2}{3\sqrt{6}} \min \left(\frac{|z/2 - t|}{\sqrt{n_1 \tau}}, \frac{\sqrt{n_1 \tau}}{2} \right) \\ &= \frac{1}{3\sqrt{2}} \|\theta\|_{2,\mathbf{q}} \min \left(\frac{|z/2 - t|}{\sqrt{n_1 \tau}}, \frac{\sqrt{n_1 \tau}}{2} \right). \end{aligned} \tag{15}$$

Combining Equations (13)–(15), we then have on $\mathcal{A} \cap \mathcal{B}$ that,

$$\frac{1}{3\sqrt{2}} \|\theta\|_{2,\mathbf{q}} \min \left(\frac{|z - \hat{z}|}{2\sqrt{n_1 \tau}}, \frac{\sqrt{n_1 \tau}}{2} \right) \leq 4\sigma \sqrt{\log n} + 14\sqrt{6}\|\theta\|_2 \sqrt{\log(kn)}. \tag{16}$$

For $C' \geq 84\sqrt{3}$, we have by Equation (4) that

$$\frac{24\sqrt{2}\sigma}{\|\theta\|_{2,\mathbf{q}}} \sqrt{\frac{\log n}{n_1 \tau}} + \frac{168\sqrt{3}\|\theta\|_2}{\|\theta\|_{2,\mathbf{q}}} \sqrt{\frac{\log(kn)}{n_1 \tau}} < \frac{2C'}{\tau} \sqrt{\frac{\log(pn)}{n}} \left(\frac{\sigma \sqrt{k}}{\|\theta\|_{2,\mathbf{q}}} + \frac{\|\theta\|_2}{\|\theta\|_{2,\mathbf{q}}} \right) \leq 1,$$

which means that the minimum on the left-hand side of (16) must be achieved by the first term. We therefore deduce with probability at least $\mathbb{P}(\mathcal{A} \cap \mathcal{B}) \geq 1 - 22/n$ that,

$$\frac{|\hat{z} - z|}{n\tau} \leq \frac{24\sigma\sqrt{\log n} + 84\sqrt{6}\|\theta\|_2\sqrt{\log(kn)}}{\|\theta\|_{2,\mathbf{q}}\sqrt{n\tau}} \leq 84\sqrt{6} \left(\frac{\sigma}{\|\theta\|_{2,\mathbf{q}}} + \frac{\|\theta\|_2}{\|\theta\|_{2,\mathbf{q}}} \right) \sqrt{\frac{\log(kn)}{n\tau}},$$

as desired.

6.3 | Proof of Theorem 2

The proof of Theorem 2 will make use of the following propositions.

Proposition 3 *Let $(X, \Omega) \sim P_{n,p,z,\theta,\sigma,\mathbf{q}}$, let $A = (A_{j,t}) = \mathcal{T}(\mathbb{E}(X)) \in \mathbb{R}^{p \times (n-1)}$ and let $A_\Omega = (A_\Omega)_{j,t} = \mathcal{T}^{\text{Miss}}(\mathbb{E}(X) \circ \Omega, \Omega) \in \mathbb{R}^{p \times (n-1)}$. Write $\Delta = (\Delta_{j,t}) = A_\Omega - (\text{diag} \sqrt{\mathbf{q}})A \in \mathbb{R}^{p \times (n-1)}$, fix $v = (v_1, \dots, v_p)^\top \in \mathbb{S}^{p-1}$ and let $\tau := n^{-1} \min\{z, n-z\}$. For any given $\delta \in (0, 1]$, if $n\tau \min_{j \in [p]} q_j \geq 60k \log(12p/\delta)$, then for t satisfying $|z - t| \leq n\tau/50$, we have with probability at least $1 - \delta$ that*

$$|(v^\top \Delta)_z - (v^\top \Delta)_t| \leq \frac{2\|\theta\|_{2,\mathbf{q}}|z-t|}{9\sqrt{n\tau}} + \sqrt{\frac{2|z-t| \sum_{j \in [p]} v_j^2 \theta_j^2 \log(12/\delta)}{n\tau}} + \frac{4 \log(12p/\delta)}{3\sqrt{n\tau}} \max_{j \in [p]} \frac{|v_j \theta_j|}{q_j^{1/2}}.$$

Proof. Without loss of generality, we may assume that $t < z$. For each $j \in [p]$, by two Taylor expansions, there exist $\xi_j, \tilde{\xi}_j \in [t, z]$ such that

$$\begin{aligned} A_{j,z} - A_{j,t} &= \frac{(n-z)\theta_j}{\sqrt{n}} \left(\sqrt{\frac{z}{n-z}} - \sqrt{\frac{t}{n-t}} \right) \\ &= \theta_j(z-t) \frac{\sqrt{z^{-1} + (n-z)^{-1}}}{2} + \theta_j(z-t)^2 \frac{n^{1/2}(n-z)(n-4\xi_j)}{8\xi_j^{3/2}(n-\xi_j)^{5/2}} \\ &= \theta_j(z-t) \frac{\sqrt{t^{-1} + (n-z)^{-1}}}{2} + \theta_j(z-t)^2 \left\{ \frac{n^{1/2}(n-z)(n-4\xi_j)}{8\xi_j^{3/2}(n-\xi_j)^{5/2}} - \frac{1}{4\xi_j^2 \sqrt{\xi_j^{-1} + (n-z)^{-1}}} \right\}. \end{aligned}$$

Similarly, by another two Taylor expansions, there exist random variables $\Xi_j, \tilde{\Xi}_j \in [L_{j,t}, L_{j,z}]$ such that

$$\begin{aligned} (A_\Omega)_{j,z} - (A_\Omega)_{j,t} &= \theta_j(L_{j,z} - L_{j,t}) \frac{\sqrt{L_{j,t}^{-1} + R_{j,n-z}^{-1}}}{2} \\ &\quad + \theta_j(L_{j,z} - L_{j,t})^2 \left\{ \frac{N_j^{1/2} R_{j,n-z} (N_j - 4\Xi_j)}{8\Xi_j^{3/2} (N_j - \Xi_j)^{5/2}} - \frac{1}{4\tilde{\Xi}_j^2 \sqrt{\tilde{\Xi}_j^{-1} + R_{j,n-z}^{-1}}} \right\}. \end{aligned}$$

We write

$$\begin{aligned} D_{1,j} &:= \frac{\theta_j}{2} \sqrt{\frac{t^{-1} + (n-z)^{-1}}{q_j}} \{(L_{j,z} - L_{j,t}) - q_j(z-t)\} \\ D_{2,j} &:= \frac{\theta_j(L_{j,z} - L_{j,t})}{2} \left\{ \sqrt{L_{j,t}^{-1} + R_{j,n-z}^{-1}} - \sqrt{\frac{t^{-1} + (n-z)^{-1}}{q_j}} \right\} \end{aligned}$$

$$D_{3,j} := |\theta_j| q_j^{1/2} (z - t)^2 \left\{ \frac{1}{2} \left(\frac{n}{\xi_j(n - \xi_j)} \right)^{3/2} + \frac{1}{4\xi_j^{3/2}} \right\}$$

$$D_{4,j} := |\theta_j| (L_{j,z} - L_{j,t})^2 \left\{ \frac{1}{2} \left(\frac{N_j}{\Xi_j(N_j - \Xi_j)} \right)^{3/2} + \frac{1}{4\Xi_j^{3/2}} \right\}.$$

We then have the bound

$$|(v^\top \Delta)_z - (v^\top \Delta)_t| \leq \left| \sum_{j=1}^p v_j D_{1,j} \right| + \left| \sum_{j=1}^p v_j D_{2,j} \right| + \sum_{j=1}^p |v_j| D_{3,j} + \sum_{j=1}^p |v_j| D_{4,j}. \tag{17}$$

We control the four terms on the right-hand side of (17) separately. For the first term, setting $y := \sqrt{2(z - t) \sum_{j \in [p]} v_j^2 \theta_j^2 \log(12/\delta)} + (1/3) \max_{j \in [p]} |v_j \theta_j| q_j^{-1/2} \log(12/\delta)$, we consider the event

$$B_t := \left\{ \left| \sum_{j=1}^p \frac{v_j \theta_j}{q_j^{1/2}} \sum_{r=t+1}^z (\omega_{j,r} - q_j) \right| \leq y \right\}.$$

Since $(\omega_{j,r})_{j \in [p], r \in (t, z]}$ are independent $\text{Bern}(q_j)$ random variables, we have by Lemma 3 in the supplementary material that $\mathbb{P}(B_t^c) \leq \delta/6$. On B_t , we have that

$$\left| \sum_{j=1}^p v_j D_{1,j} \right| = \frac{1}{2} \sqrt{t^{-1} + (n - z)^{-1}} \left| \sum_{j=1}^p \frac{v_j \theta_j}{q_j^{1/2}} \sum_{r=t+1}^z (\omega_{j,r} - q_j) \right| \leq \frac{0.8y}{\sqrt{n\tau}}. \tag{18}$$

For the second term on the right-hand side of (17), let \mathcal{F} denote the σ -algebra generated by $(\omega_{j,r} : j \in [p], r \in [n], r \notin [t + 1, z])$, then $L_{j,z} - L_{j,t}$ is independent of \mathcal{F} , whereas

$$G_j := \frac{v_j \theta_j}{2} \left\{ \sqrt{L_{j,t}^{-1} + R_{j,n-z}^{-1}} - q_j^{-1/2} \sqrt{t^{-1} + (n - z)^{-1}} \right\}$$

is measurable with respect to \mathcal{F} . We can therefore apply Lemma 3 in the supplement conditional on \mathcal{F} to obtain that there is an event C_t with $\mathbb{P}(C_t^c | \mathcal{F}) \leq \delta/6$ on which

$$\left| \sum_{j=1}^p v_j D_{2,j} \right| = \left| \sum_{j=1}^p G_j \sum_{r=t+1}^z \omega_{j,r} \right|$$

$$\leq (z - t) \left| \sum_{j=1}^p G_j q_j \right| + \sqrt{2(z - t) \log(12/\delta) \sum_{j=1}^p G_j^2 q_j} + \frac{1}{3} \max_{j \in [p]} |G_j| \log(12/\delta). \tag{19}$$

For $0 \leq a < b \leq n$, define

$$H_{j,(a,b)} := \left| \frac{L_{j,b} - L_{j,a}}{(b - a)q_j} - 1 \right|. \tag{20}$$

We consider the event

$$\mathcal{A}_{j,t} := \left\{ \max(H_{j,(0,t)}, H_{j,(0,z)}, H_{j,(z,n)}) \leq \frac{1}{5} \right\} \cap \left\{ H_{j,(t,z)} \leq \sqrt{\frac{2 \log(12p/\delta)}{(z-t)q_j}} + \frac{\log(12p/\delta)}{3(z-t)q_j} \right\}.$$

By Lemma 3 in the supplement again, we obtain that

$$\mathbb{P}(\mathcal{A}_{j,t}^c) \leq 6 \exp \left\{ -\frac{(1/5)^2(1-1/50)n\tau q_j}{2(1+1/15)} \right\} + \frac{\delta}{6p} \leq 6e^{-n\tau q_j/60} + \frac{\delta}{6p} \leq \frac{2\delta}{3p},$$

where we used the assumption $n\tau q_j \geq 60 \log(12p/\delta)$ in the final inequality. On $\mathcal{A}_{j,t}$, we have

$$|G_j| \leq \frac{|v_j \theta_j|}{2q_j^{1/2}} (\sqrt{5/4} - 1) \sqrt{t^{-1} + (n-z)^{-1}} \leq \frac{0.084 |v_j \theta_j|}{(n\tau q_j)^{1/2}},$$

where we have used the fact that $t \geq (49/50)n\tau$. Combining the above inequality with (17), on $\cap_{j \in [p]} \mathcal{A}_{j,t} \cap C_t$, we have by the Cauchy–Schwarz inequality that

$$\left| \sum_{j=1}^p v_j D_{2,j} \right| \leq 0.084 \frac{\|\theta\|_{2,\mathbf{q}}(z-t) + y}{(n\tau)^{1/2}}. \tag{21}$$

For the third and fourth terms on the right-hand side of (17), since $|z-t| \leq n\tau/50$, we have

$$D_{3,j} \leq |\theta_j| q_j^{1/2} (z-t)^2 \left\{ \frac{1}{2} \left(\frac{2}{\min(t, n-z)} \right)^{3/2} + \frac{1}{4t^{3/2}} \right\} \leq \frac{1.8 |\theta_j| q_j^{1/2} (z-t)^2}{(n\tau)^{3/2}}.$$

Moreover, on $\mathcal{A}_{j,t}$,

$$\begin{aligned} D_{4,j} &\leq |\theta_j| q_j^2 (z-t)^2 (1 + H_{j,(t,z)})^2 \left\{ \frac{1}{2} \left(\frac{2}{\min(L_{j,t}, R_{j,n-z})} \right)^{3/2} + \frac{1}{4L_{j,t}^{3/2}} \right\} \\ &\leq \frac{2.4 |\theta_j| q_j^{1/2} (z-t)^2}{(n\tau)^{3/2}} (1 + H_{j,(t,z)})^2 \leq \frac{4.8 |\theta_j| q_j^{1/2} (z-t)^2}{(n\tau)^{3/2}} + \frac{60 |\theta_j| \log^2(12p/\delta)}{(n\tau q_j)^{3/2}}, \end{aligned}$$

where the final step uses the fact that $(1 + \sqrt{2a + a/3})^2 \leq 2 + 25a^2$ for any $a > 0$. Therefore, on $\cap_{j \in [p]} \mathcal{A}_{j,t}$, since $|z-t| \leq n\tau/50$ and $n\tau \min_{j \in [p]} q_j \geq 60k \log(12p/\delta)$, we have by the Cauchy–Schwarz inequality again that

$$\begin{aligned} \sum_{j=1}^p |v_j| (D_{3,j} + D_{4,j}) &\leq \frac{6.6 \|\theta\|_{2,\mathbf{q}}(z-t)^2}{(n\tau)^{3/2}} + \frac{60 \log^2(12p/\delta)}{(n\tau)^{3/2}} \sum_{j: \theta_j \neq 0} \frac{|v_j \theta_j|}{q_j^{3/2}} \\ &\leq \frac{0.132 \|\theta\|_{2,\mathbf{q}}(z-t)}{(n\tau)^{1/2}} + \frac{\log(12p/\delta)}{(n\tau)^{1/2}} \max_{j: \theta_j \neq 0} \frac{|v_j \theta_j|}{q_j^{1/2}}. \end{aligned} \tag{22}$$

Therefore, combining (17), (18), (21) and (22), we have on $\cap_{j \in [p]} \mathcal{A}_{j,t} \cap \mathcal{B}_t \cap C_t$ that

$$\begin{aligned}
 |(v^\top \Delta)_z - (v^\top \Delta)_t| &\leq \frac{y}{\sqrt{n\tau}} + \frac{2\|\theta\|_{2,q}(z-t)}{9\sqrt{n\tau}} + \frac{\log(12p/\delta)}{(n\tau)^{1/2}} \max_{j \in [p]} \frac{|v_j \theta_j|}{q_j^{1/2}} \\
 &\leq \frac{2\|\theta\|_{2,q}(z-t)}{9\sqrt{n\tau}} + \sqrt{\frac{2(z-t) \sum_{j \in [p]} v_j^2 \theta_j^2 \log(12/\delta)}{n\tau}} + \frac{4 \log(12p/\delta)}{3(n\tau)^{1/2}} \max_{j \in [p]} \frac{|v_j \theta_j|}{q_j^{1/2}}.
 \end{aligned}$$

Since $\sum_{j=1}^p \mathbb{P}(\mathcal{A}_{j,t}^c) + \mathbb{P}(\mathcal{B}_t^c) + \mathbb{P}(\mathcal{C}_t^c) \leq \delta$, the proof is complete.

Proposition 4 *Suppose that $\Omega = (\omega_{j,t})_{j \in [p], t \in [n]}$ and $W = (W_{j,t})_{j \in [p], t \in [n]}$ are independent, with $\omega_{j,t} \sim \text{Bern}(q_j)$ independently and $q_j \in (0, 1]$, and with $W_{j,t} \stackrel{\text{iid}}{\sim} N(0, \sigma^2)$. Let $E_\Omega := \mathcal{T}^{\text{Miss}}(W \circ \Omega, \Omega)$, let $z \in [n-1]$ and let $\tau := n^{-1} \min\{z, n-z\}$. Suppose that $t \in [n-1]$ satisfies $|z-t| \leq n\tau/2$. For a fixed $v \in \mathbb{S}^{p-1}$, if $n\tau \min_{j \in [p]} q_j \geq 20 \log(11p/\delta)$, then we have for any $\delta \in (0, 1]$ that*

$$\mathbb{P} \left\{ |(v^\top E_\Omega)_z - (v^\top E_\Omega)_t| > 70\sigma \sqrt{\frac{|z-t| \log(11/\delta) + \log^2(11/\delta) \max_{j \in [p]} v_j^2 / q_j}{n\tau}} \right\} \leq \delta.$$

Proof. By symmetry, we may assume without loss of generality that $t < z$. We note that $(E_\Omega)_{j,z} - (E_\Omega)_{j,t}$ is a centred normal random variable conditional on Ω , so we start by looking at its conditional variance. By definition of $\mathcal{T}^{\text{Miss}}$, we have

$$\begin{aligned}
 (E_\Omega)_{j,z} - (E_\Omega)_{j,t} &= \sqrt{\frac{N_j}{L_{j,z}R_{j,n-z}}} \left(\frac{L_{j,z}}{N_j} \sum_{r=1}^n W_{j,r} \omega_{j,r} - \sum_{r=1}^z W_{j,r} \omega_{j,r} \right) \\
 &\quad - \sqrt{\frac{N_j}{L_{j,t}R_{j,n-t}}} \left(\frac{L_{j,t}}{N_j} \sum_{r=1}^n W_{j,r} \omega_{j,r} - \sum_{r=1}^t W_{j,r} \omega_{j,r} \right) \\
 &= \sqrt{\frac{N_j}{L_{j,z}R_{j,n-z}}} \left(\frac{L_{j,z} - L_{j,t}}{N_j} \sum_{r=1}^n W_{j,r} \omega_{j,r} - \sum_{r=t+1}^z W_{j,r} \omega_{j,r} \right) \\
 &\quad + \left(\sqrt{\frac{N_j}{L_{j,z}R_{j,n-z}}} - \sqrt{\frac{N_j}{L_{j,t}R_{j,n-t}}} \right) \left(\frac{L_{j,t}}{N_j} \sum_{r=1}^n W_{j,r} \omega_{j,r} - \sum_{r=1}^t W_{j,r} \omega_{j,r} \right). \tag{23}
 \end{aligned}$$

Now, by the mean value theorem, there exists a random variable $\Xi_j \in [L_{j,t}, L_{j,z}]$ such that

$$\left| \sqrt{\frac{N_j}{L_{j,z}R_{j,n-z}}} - \sqrt{\frac{N_j}{L_{j,t}R_{j,n-t}}} \right| \leq (L_{j,z} - L_{j,t}) \left| \frac{\Xi_j}{N_j} - \frac{1}{2} \right| \left(\frac{N_j}{\Xi_j(N_j - \Xi_j)} \right)^{3/2} \leq \frac{\sqrt{2}(L_{j,z} - L_{j,t})}{\min(\Xi_j, N_j - \Xi_j)^{3/2}}. \tag{24}$$

Also, observe that

$$\frac{L_{j,t}}{N_j} \sum_{r=1}^n W_{j,r} \omega_{j,r} - \sum_{r=1}^t W_{j,r} \omega_{j,r} = \sum_{r=t+1}^n W_{j,r} \omega_{j,r} - \frac{R_{j,n-t}}{N_j} \sum_{r=1}^n W_{j,r} \omega_{j,r}. \tag{25}$$

Substituting (24) and (25) into Equation (23), and observing that $\sum_{r=1}^n W_{j,r} \omega_{j,r}$ is positively correlated with each of $\sum_{r=t+1}^z W_{j,r} \omega_{j,r}$, $\sum_{r=1}^t W_{j,r} \omega_{j,r}$ and $\sum_{r=t+1}^n W_{j,r} \omega_{j,r}$, we have that

$$\begin{aligned} \text{Var}((E_\Omega)_{j,z} - (E_\Omega)_{j,t} \mid \Omega) &\leq \frac{2\sigma^2 N_j}{L_{j,z} R_{j,n-z}} \left(\frac{(L_{j,z} - L_{j,t})^2}{N_j} + L_{j,z} - L_{j,t} \right) \\ &\quad + \frac{4\sigma^2 (L_{j,z} - L_{j,t})^2}{\min(\Xi_j, N_j - \Xi_j)^3} \min \left(\frac{L_{j,t}^2}{N_j} + L_{j,t}, R_{j,n-t} + \frac{R_{j,n-t}^2}{N_j} \right) \\ &\leq 4\sigma^2 (L_{j,z} - L_{j,t}) \left(\frac{1}{L_{j,z}} + \frac{1}{R_{j,n-z}} \right) \\ &\quad + \frac{8\sigma^2 (L_{j,z} - L_{j,t})^2}{\min\{L_{j,t}, R_{j,n-z}\}^2} \max \left(\frac{L_{j,z}}{L_{j,t}}, \frac{R_{j,n-t}}{R_{j,n-z}} \right). \end{aligned} \tag{26}$$

Recalling the definition of $H_{j,(a,b)}$ from Equation (20) in the proof of Proposition 3, we consider the event

$$\mathcal{A}_{j,t} := \left\{ \max\{H_{j,(0,z)}, H_{j,(0,t)}, H_{j,(z,n)}, H_{j,(t,n)}\} \leq \frac{1}{2} \right\} \cap \left\{ H_{j,(t,z)} \leq \frac{n\tau}{z-t} \right\}.$$

By Bernstein’s inequality (Lemma 3 in the supplementary material), we obtain that

$$\mathbb{P}(\mathcal{A}_{j,t}^c) \leq 8 \exp \left\{ -\frac{(1/2)^2(n\tau/2)q_j}{2(1+1/6)} \right\} + \exp \left\{ -\frac{(n\tau/(z-t))^2(z-t)q_j}{2(1+n\tau/\{3(z-t)\})} \right\} \leq 9e^{-n\tau q_j/20},$$

where we used the fact that $z - t \leq n\tau/2$ in the final inequality. It therefore follows from Equation (26) that on the event $\mathcal{A}_{j,t}$,

$$\text{Var}((E_\Omega)_{j,z} - (E_\Omega)_{j,t} \mid \Omega) \leq \frac{16\sigma^2}{n\tau q_j} (L_{j,z} - L_{j,t}) + \frac{768\sigma^2}{(n\tau q_j)^2} (L_{j,z} - L_{j,t})^2 \leq \frac{1168\sigma^2}{n\tau q_j} (L_{j,z} - L_{j,t}).$$

Hence, on $\cap_{j \in [p]} \mathcal{A}_{j,t}$, we have

$$\text{Var} \left\{ (v^\top E_\Omega)_z - (v^\top E_\Omega)_t \mid \Omega \right\} \leq \frac{1168\sigma^2}{n\tau} \sum_{j=1}^p \frac{v_j^2}{q_j} (L_{j,z} - L_{j,t}) = \frac{1168\sigma^2}{n\tau} \sum_{j=1}^p \frac{v_j^2}{q_j} \sum_{r=t+1}^z \omega_{j,r}.$$

Now, setting $y := \sqrt{2(z-t) \sum_{j \in [p]} v_j^4 q_j^{-1} \log(11/\delta)} + (1/3) \max_{j \in [p]} v_j^2 q_j^{-1} \log(11/\delta)$, consider the event

$$\mathcal{B}_t := \left\{ \sum_{j=1}^p \frac{v_j^2}{q_j} \sum_{r=t+1}^z \omega_{j,r} \leq z - t + y \right\}.$$

We have by Bernstein’s inequality (Lemma 3 in the supplement) that $\mathbb{P}(\mathcal{B}_t^c) \leq \delta/11$. Noting that $\sum_{j \in [p]} v_j^4 q_j^{-1} \leq \max_{j \in [p]} v_j^2 q_j^{-1}$, and using the fact that $a + \sqrt{2ab} + b/3 \leq 2(a+b)$ for any $a, b > 0$, we have from the Gaussian tail bound that for every $u > 0$,

$$\begin{aligned} &\mathbb{P} \left\{ |(v^\top E_\Omega)_z - (v^\top E_\Omega)_t| > 49u\sigma \sqrt{\frac{z-t + \log(11/\delta) \max_{j \in [p]} v_j^2/q_j}{n\tau}} \right\} \\ &\leq e^{-u^2/2} + \sum_{j=1}^p \mathbb{P}(\mathcal{A}_{j,t}^c) + \mathbb{P}(\mathcal{B}_t^c) \leq e^{-u^2/2} + 9 \sum_{j=1}^p e^{-n\tau q_j/20} + \frac{\delta}{11}. \end{aligned}$$

The result follows by taking $u := \sqrt{2 \log(11/\delta)}$ and using the fact that $n\tau \min_{j \in [p]} q_j \geq 20 \log(11p/\delta)$.

Proof of Theorem 2. We write $z_1 := z/2$ and $n_1 := n/2$. Taking $C, C' > 0$ from Theorem 1, we may assume that $c \in (0, 1/50]$ is small enough that the hypothesis (4) of Theorem 1 is satisfied when $\rho \leq c$. Hence, by Theorem 1, there is an event \mathcal{E} with probability at least $1 - 22/n$ such that

$$\frac{|\hat{z} - z|}{n\tau} \leq C \sqrt{\frac{\log(kn)}{n\tau}} \left(\frac{\sigma}{\|\theta\|_{2,\mathbf{q}}} + \frac{\|\theta\|_2}{\|\theta\|_{2,\mathbf{q}}} \right) \leq C\rho.$$

By further reducing $c > 0$ if necessary, we may assume that on \mathcal{E} , and when $\rho \leq c$, we have $|\hat{z} - z| \leq n\tau/50$.

Let $A^{(2)}, \Delta^{(2)}$ and $E_\Omega^{(2)}$ be defined as in the proof of Theorem 1. With $\hat{v} = (\hat{v}_1, \dots, \hat{v}_p)^\top \in \mathbb{S}^{p-1}$ as defined in Algorithm 3, an inspection of the proof of Theorem 1 reveals that on \mathcal{E} , we also have for all $t \in [z_1 - n_1\tau\rho, z_1 + n_1\tau\rho]$ and $\rho \leq c$ that

$$(\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}}) A^{(2)})_{z_1} - (\hat{v}^\top (\text{diag} \sqrt{\mathbf{q}}) A^{(2)})_t \geq \frac{|z_1 - t| \|\theta\|_{2,\mathbf{q}}}{3\sqrt{2n_1\tau}}. \tag{27}$$

Recall that \hat{v} is measurable with respect to the σ -algebra generated by the odd-numbered time points, and that $\Delta^{(2)}$ and $E_\Omega^{(2)}$ are measurable with respect to the σ -algebra generated by the even-numbered time points. By taking the universal constant $C_1 > 0$ in the statement of the theorem to be sufficiently large, we can ensure that the lower bounds on $n\tau \min_{j \in [p]} q_j$ in Propositions 3 and 4 are satisfied. It follows by these propositions that when $\rho \leq c$, for each $t \in [z_1 - n_1\tau\rho, z_1 + n_1\tau\rho]$, there is an event \mathcal{A}_t of probability at least $1 - n^{-2}$ on which both

$$|(\hat{v}^\top \Delta^{(2)})_{z_1} - (\hat{v}^\top \Delta^{(2)})_t| - \frac{2\|\theta\|_{2,\mathbf{q}}|z_1 - t|}{9\sqrt{n_1\tau}} \lesssim \sqrt{\frac{|z_1 - t| \sum_{j \in [p]} \hat{v}_j^2 \theta_j^2 \log n}{n\tau}} + \frac{\log(pn)}{\sqrt{n\tau}} \max_{j \in [p]} \frac{|\hat{v}_j \theta_j|}{q_j^{1/2}}, \tag{28}$$

$$|(\hat{v}^\top E_\Omega^{(2)})_{z_1} - (\hat{v}^\top E_\Omega^{(2)})_t| \lesssim \sqrt{\frac{|z_1 - t| \sigma^2 \log n}{n\tau}} + \frac{\sigma \log n}{\sqrt{n\tau}} \max_{j \in [p]} \frac{|\hat{v}_j|}{q_j^{1/2}}. \tag{29}$$

Combining Equations (27)–(29) and the basic inequality as in (14) in the proof of Theorem 1, we have on the event $\mathcal{E} \cap \bigcap_{t \in [z/2 - n_1\tau\rho, z/2 + n_1\tau\rho]} \mathcal{A}_t$ and with $\rho \leq c$ that

$$\frac{|\hat{z} - z| \|\theta\|_{2,\mathbf{q}}}{\sqrt{n\tau}} \lesssim \sqrt{\frac{|z - z|(\sigma^2 + \sum_{j \in [p]} \hat{v}_j^2 \theta_j^2) \log n}{n\tau}} + \frac{\log(pn)}{\sqrt{n\tau}} \max_{j \in [p]} \frac{|\hat{v}_j \theta_j|}{q_j^{1/2}} + \frac{\sigma \log n}{\sqrt{n\tau}} \max_{j \in [p]} \frac{|\hat{v}_j|}{q_j^{1/2}}. \tag{30}$$

Define $v = (v_j)_{j \in [p]} \in \mathbb{R}^p$ such that $v_j := \theta_j q_j^{1/2} / \|\theta\|_{2,\mathbf{q}}$. Then we can write

$$\hat{v} = \alpha v + \beta w,$$

for some unit-length (random) vector $w = (w_j)_{j \in [p]}$ that is orthogonal to v and some $\alpha, \beta \in \mathbb{R}$ such that $\alpha^2 + \beta^2 = 1$. Moreover, by inspecting the proof of Theorem 1, we see that on \mathcal{E} , we have $|\beta| = \sin \angle(\hat{v}, v) \leq \rho$. Then from Equation (30), we have on $\mathcal{E} \cap \bigcap_{t \in [z/2 - n_1\tau\rho, z/2 + n_1\tau\rho]} \mathcal{A}_t$ that

$$\begin{aligned}
 |\hat{z} - z| &\lesssim \frac{\sigma^2 + \alpha^2 \sum_{j \in [p]} v_j^2 \theta_j^2 + \beta^2 \sum_{j \in [p]} w_j^2 \theta_j^2}{\|\theta\|_{2,\mathbf{q}}^2} \log n \\
 &\quad + \frac{|\alpha| \max_{j \in [p]} |v_j \theta_j| q_j^{-1/2} + |\beta| \max_{j \in [p]} |w_j \theta_j| q_j^{-1/2}}{\|\theta\|_{2,\mathbf{q}}} \log(pn) \\
 &\quad + \frac{\sigma (|\alpha| \max_{j \in [p]} |v_j| q_j^{-1/2} + |\beta| \max_{j \in [p]} |w_j| q_j^{-1/2})}{\|\theta\|_{2,\mathbf{q}}} \log n \\
 &\lesssim \frac{\sigma^2 \log n}{\|\theta\|_{2,\mathbf{q}}^2} + \frac{\|\theta\|_{4,\mathbf{q}}^4 \log n}{\|\theta\|_{2,\mathbf{q}}^4} + \frac{\rho^2 \|\theta\|_{\infty}^2 \log n}{\|\theta\|_{2,\mathbf{q}}^2} \\
 &\quad + \frac{\|\theta\|_{\infty}^2 \log(pn)}{\|\theta\|_{2,\mathbf{q}}^2} + \frac{\rho \max_{j \in [p]} |\theta_j| q_j^{-1/2} \log(pn)}{\|\theta\|_{2,\mathbf{q}}} \\
 &\quad + \frac{\|\theta\|_{\infty} \sigma \log n}{\|\theta\|_{2,\mathbf{q}}^2} + \frac{\rho \sigma \log n}{\|\theta\|_{2,\mathbf{q}} \min_{j \in [p]} q_j^{1/2}} \\
 &\lesssim \frac{(\sigma^2 + \|\theta\|_{\infty}^2) \log(pn)}{\|\theta\|_{2,\mathbf{q}}^2} + \frac{\rho(\sigma + \|\theta\|_{\infty}) \log(pn)}{\|\theta\|_{2,\mathbf{q}} \min_{j \in [p]} q_j^{1/2}} \lesssim \frac{(\sigma^2 + \|\theta\|_{\infty}^2) \log(pn)}{\|\theta\|_{2,\mathbf{q}}^2},
 \end{aligned}$$

where the final bound uses the definition of ρ and the fact that $n\tau^2 \min_{j \in [p]} q_j \geq C_1 k \log(pn)$. The desired result follows since $\mathbb{P}\left(\mathcal{E} \cap \bigcap_{t \in [z/2 - n_1 \tau \rho, z/2 + n_1 \tau \rho]} \mathcal{A}_t\right) \geq 1 - 22/n - (2n_1 \tau \rho + 1)/n^2 \geq 1 - 23/n$.

6.4 | Proof of Theorem 3

Proof of Theorem 3. For notational simplicity, we abbreviate $P_{n,\rho,z,\theta,\sigma,\mathbf{q}}$ as P_z in this proof, with corresponding expectation operator E_z . For any $1 \leq z_1 < z_2 \leq n - 1$, by Le Cam’s two point testing lemma (e.g. Yu, 1997, Lemma 1), we have that

$$\inf_{z \in \bar{z}} \max_{z \in [n-1]} E_z |\bar{z} - z| \geq \frac{1}{2} |z_1 - z_2| \{1 - d_{\text{TV}}(P_{z_1}, P_{z_2})\}. \tag{31}$$

By Pinsker’s inequality (e.g. Wainwright, 2019, Lemma 15.2), we have

$$\begin{aligned}
 2d_{\text{TV}}^2(P_{z_1}, P_{z_2}) &\leq \text{KL}(P_{z_1} \| P_{z_2}) = \mathbb{E}_{P_{z_1}} \left[\mathbb{E}_{P_{z_1}} \left\{ \log \left(\frac{dP_{z_1}}{dP_{z_2}}(X, \Omega) \right) \middle| \Omega \right\} \right] \\
 &= \sum_{j=1}^p \sum_{t=z_1+1}^{z_2} \mathbb{E}_{P_{z_1}} \frac{\theta_j^2 \omega_{j,t}}{2\sigma^2} = \frac{(z_2 - z_1) \|\theta\|_{2,\mathbf{q}}^2}{2\sigma^2}.
 \end{aligned}$$

Choosing $z_2 - z_1 = \min\{\lfloor \sigma^2 / \|\theta\|_{2,\mathbf{q}}^2 \rfloor, n - 2\}$, we have $d_{\text{TV}}(P_{z_1}, P_{z_2}) \leq 1/2$ and consequently if $\sigma^2 \geq \|\theta\|_{2,\mathbf{q}}^2$, then by Equation (31),

$$\inf_{\tilde{z} \in \tilde{\mathcal{Z}}_{z \in [n-1]}} \max_{z \in [n-1]} E_z |\tilde{z} - z| \geq \frac{1}{4} \min \left\{ \left\lfloor \frac{\sigma^2}{\|\theta\|_{2,\mathbf{q}}^2} \right\rfloor, n-2 \right\} \geq \frac{1}{12} \min \left\{ \frac{\sigma^2}{\|\theta\|_{2,\mathbf{q}}^2}, n \right\}. \tag{32}$$

On the other hand, if $\|\theta\|_\infty^2 \geq 2M^2\|\theta\|_{2,\mathbf{q}}^2$, then

$$\sum_{j:\theta_j \neq 0} q_j \leq \frac{\|\theta\|_{2,\mathbf{q}}^2}{\min_{j:\theta_j \neq 0} \theta_j^2} \leq \frac{M^2\|\theta\|_{2,\mathbf{q}}^2}{\|\theta\|_\infty^2} \leq 1/2.$$

Define $S := \{j \in [p] : \theta_j \neq 0\}$ and

$$\mathcal{A} := \{((x_{j,t})_{j \in [p], t \in [n]}, (\omega_{j,t})_{j \in [p], t \in [n]}) : \omega_{j,t} = 0 \text{ whenever } j \in S \text{ and } z_1 + 1 \leq t \leq z_2\}.$$

Then the distributions of P_{z_1} given \mathcal{A} and P_{z_2} given \mathcal{A} are identical. Moreover, $P_{z_1}(\mathcal{A}) = P_{z_2}(\mathcal{A})$. Thus, for any Borel measurable subset B of $\mathbb{R}^{n \times p} \times \{0, 1\}^{n \times p}$, we have

$$|P_{z_1}(B) - P_{z_2}(B)| = |P_{z_1}(B|\mathcal{A}^c) - P_{z_2}(B|\mathcal{A}^c)|P_{z_1}(\mathcal{A}^c) \leq P_{z_1}(\mathcal{A}^c).$$

Hence, using the fact that $1 - x \geq e^{-2x \log 2}$ for $x \in [0, 1/2]$, we have

$$1 - d_{\text{TV}}(P_{z_1}, P_{z_2}) \geq P_{z_1}(\mathcal{A}) = \prod_{j:\theta_j \neq 0} (1 - q_j)^{z_2 - z_1} \geq \exp \left\{ -2(\log 2)(z_2 - z_1) \sum_{j:\theta_j \neq 0} q_j \right\}.$$

Choosing $z_2 - z_1 = \min\{[(2 \sum_{j:\theta_j \neq 0} q_j)^{-1}, n-2]\}$, we have $1 - d_{\text{TV}}(P_{z_1}, P_{z_2}) \geq 1/4$, and consequently, on combining with (31) we obtain that

$$\inf_{\tilde{z} \in \tilde{\mathcal{Z}}_{z \in [n-1]}} \max_{z \in [n-1]} E_z |\tilde{z} - z| \geq \frac{1}{8} \min \left\{ \frac{1}{2 \sum_{j:\theta_j \neq 0} q_j}, n-2 \right\} \geq \frac{1}{24} \min \left\{ \frac{\|\theta\|_\infty^2}{M^2\|\theta\|_{2,\mathbf{q}}^2}, n \right\}. \tag{33}$$

By combining (32) and (33), and considering the three possible cases of (i) $\sigma^2 \geq \|\theta\|_{2,\mathbf{q}}^2 > \|\theta\|_\infty^2/(2M^2)$, (ii) $\|\theta\|_\infty^2/(2M^2) \geq \|\theta\|_{2,\mathbf{q}}^2 > \sigma^2$ and (iii) $\min\{\sigma^2, \|\theta\|_\infty^2/(2M^2)\} \geq \|\theta\|_{2,\mathbf{q}}^2$, we have

$$\begin{aligned} \inf_{\tilde{z} \in \tilde{\mathcal{Z}}_{z \in [n-1]}} \max_{z \in [n-1]} E_z |\tilde{z} - z| &\geq \frac{1}{12} \max \left\{ \frac{\sigma^2}{\|\theta\|_{2,\mathbf{q}}^2} \wedge \frac{n}{2}, \frac{\|\theta\|_\infty^2}{2M^2\|\theta\|_{2,\mathbf{q}}^2} \wedge \frac{n}{2} \right\} \\ &\geq \frac{1}{12} \min \left(\max \left\{ \frac{\sigma^2}{\|\theta\|_{2,\mathbf{q}}^2}, \frac{\|\theta\|_\infty^2}{2M^2\|\theta\|_{2,\mathbf{q}}^2} \right\}, \frac{n}{2} \right) \\ &\geq \frac{1}{24} \min \left(\frac{\sigma^2}{\|\theta\|_{2,\mathbf{q}}^2} + \frac{\|\theta\|_\infty^2}{2M^2\|\theta\|_{2,\mathbf{q}}^2}, n \right), \end{aligned}$$

as required.

ACKNOWLEDGEMENTS

BF was supported by a Knox studentship from Trinity College, Cambridge and EOX funding from Ecole polytechnique; TW was supported by EPSRC grant EP/T02772X/1; RJS was supported by

EPSRC grants EP/P031447/1 and EP/N031938/1, as well as ERC Advanced Grant 101019498. We thank the anonymous reviewers for constructive comments that helped to improve the paper.

REFERENCES

- Aston, J. & Kirch, C. (2013) Evaluating stationarity via change-point alternatives with applications to fMRI data. *Annals of Applied Statistics*, 6, 1906–1948.
- Bai, J. (2010) Common breaks in means and variances for panel data. *Journal of Economics*, 157, 78–92.
- Cao, Y., Thompson, A., Wang, M. & Xie, Y. (2019) Sketching for sequential change-point detection. *EURASIP Journal on Advances in Signal Processing*, 2019, 1–22.
- Chan, H.P. (2017) Optimal sequential detection in multi-stream data. *The Annals of Statistics*, 45, 2736–2763.
- Chan, H.P. & Walther, G. (2015) Optimal detection of multi-sample aligned sparse signals. *The Annals of Statistics*, 43, 1865–1895.
- Chen, J. & Gupta, A.K. (1997) Testing and locating variance change-points with application to stock prices. *Journal of the American Statistical Association*, 92, 739–747.
- Chen, Y., Wang, T. & Samworth, R.J. (2022) High-dimensional, multiscale online change-point detection. *Journal of the Royal Statistical Society: Series B*, 84, 234–266.
- Cho, H. (2016) Change-point detection in panel data via double CUSUM statistic. *Electronic Journal of Statistics*, 10, 2000–2038.
- Cho, H. & Fryzlewicz, P. (2014) Multiple-change-point detection for high dimensional time series via sparsified binary segmentation. *Journal of the Royal Statistical Society: Series B*, 77, 475–507.
- Csörgő, M. & Horváth, L. (1997) *Limit theorems in change-point analysis*. New York: John Wiley and Sons.
- Cule, M., Samworth, R. & Stewart, M. (2010) Maximum likelihood estimation of a multi-dimensional log-concave density. *Journal of the Royal Statistical Society: Series B (with discussion)*, 72, 545–607.
- Dümbgen, L. & Rufibach, K. (2009) Maximum likelihood estimation of a log-concave density and its distribution function: basic properties and uniform consistency. *Bernoulli*, 15, 40–68.
- Enikeeva, F. & Harchaoui, Z. (2019) High-dimensional change-point detection under sparse alternatives. *The Annals of Statistics*, 47, 2051–2079.
- Enikeeva, F. & Klopp, O. (2021) Change-point detection in dynamic networks with missing links. *arXiv preprint*, arXiv:2106.14470.
- Fryzlewicz, P. (2014) Wild binary segmentation for multiple change-point detection. *The Annals of Statistics*, 42, 2243–2281.
- Hampel, F.R. (1974) The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69, 383–393.
- Henry, D., Simani, S. & Patton, R. (2010) Fault detection and diagnosis for aeronautic and aerospace missions. *Fault Tolerant Flight Control*, 399, 91–128.
- Horváth, L. & Hušková, M. (2012) Change-point detection in panel data. *Journal of Time Series Analysis*, 33, 631–648.
- Horváth, L. & Rice, G. (2014) Extensions of some classical methods in change point analysis. *TEST*, 23, 219–255.
- Huopaniemi, I., Nadkarni, G., Nadukuru, R., Lotay, V., Ellis, S., Gottesman, O. et al. (2014) Disease progression subtype discovery from longitudinal EMR data with a majority of missing values and unknown initial time points. In: *AMIA Annual Symposium Proceedings*, vol. 2014, 709, American Medical Informatics Association.
- Jirak, M. (2015) Uniform change point tests in high dimension. *The Annals of Statistics*, 43, 2451–2483.
- Kovács, S., Li, H., Haubner, L., Munk, A. & Bühlmann, P. (2020) Optimistic search strategy: change point detection for large-scale data via adaptive logarithmic queries. *arXiv preprint*, arxiv:2010.10194.
- Liu, H., Gao, C. & Samworth, R.J. (2021) Minimax rates in sparse, high-dimensional change point detection. *The Annals of Statistics*, 49, 1081–1112.
- Londschien, M., Kovács, S. & Bühlmann, P. (2021) Change-point detection for graphical models in the presence of missing values. *Journal of Computational and Graphical Statistics*, 30, 768–779.
- Mazumder, R., Hastie, T. & Tibshirani, R. (2010) Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11, 2287–2322.
- Mei, Y. (2010) Efficient scalable schemes for monitoring a large number of data streams. *Biometrika*, 97, 419–433.

- Olshen, A., Venkatraman, E., Lucito, R. & Wigler, M. (2004) Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics*, 5, 557–572.
- Padilla, O.H.M., Yu, Y., Wang, D. & Rinaldo, A. (2022) Optimal nonparametric multivariate change point detection and localization. *IEEE Transactions on Information Theory*, to appear.
- Page, E.S. (1955) A test for a change in a parameter occurring at an unknown point. *Biometrika*, 42, 523–527.
- Peng, T., Leckie, C. & Ramamohanarao, K. (2004) Proactively detecting distributed denial of service attacks using source IP address monitoring. In: *International Conference on Research in Networking*, pp. 771–782, Springer.
- Poore, H., Samworth, R., White, N., Jones, S. & McCave, I. (2006) Neogene over flow of northern component water at the Greenland-Scotland Ridge. *Geochemistry, Geophysics, Geosystems*, 7, Q06010.
- Samworth, R. & Poore, H. (2005) Understanding past ocean circulations: a nonparametric regression case study. *Statistical model*, 5, 289–307.
- Soh, Y.S. & Chandrasekaran, V. (2017) High-dimensional change-point estimation: combining filtering with convex optimization. *Applied and Computational Harmonic Analysis*, 43, 122–147.
- Sparks, R., Keighley, T. & Muscatello, D. (2010) Early warning CUSUM plans for surveillance of negative binomial daily disease counts. *Journal of Applied Statistics*, 37, 1911–1929.
- Stewart, G.W. & Sun, J. (1990) *Matrix perturbation theory*. San Diego: Academic Press.
- Wainwright, M.J. (2019) *High-dimensional statistics: a non-asymptotic viewpoint*. Cambridge: Cambridge University Press.
- Wang, T. (2016) Spectral methods and computational trade-offs in high-dimensional statistical inference. PhD Thesis, University of Cambridge.
- Wang, T. & Samworth, R.J. (2018) High dimensional change point estimation via sparse projection. *Journal of the Royal Statistical Society: Series B*, 80, 57–83.
- Wang, T., Follain, B. & Samworth, R.J. (2022) InspectChangepoint: high-dimensional changepoint estimation via sparse projection. *R package version 1.2*.
- Wright, J.D. & Miller, K.G. (1996) Control of North Atlantic deep water circulation by the greenland-scotland ridge. *Paleoceanography*, 11, 157–170.
- Xie, Y. & Siegmund, D. (2013) Sequential multi-sensor change-point detection. *The Annals of Statistics*, 41, 670–692.
- Xie, Y., Huang, J. & Willett, R. (2013) Change-point detection for high-dimensional time series with missing data. *IEEE Journal of Selected Topics in Signal Processing*, 7, 12–27.
- Yu, B. (1997) Assouad, Fano and Le Cam. In: Pollard, D., Torgersen, E. & Yang G.L. (Eds.). *Festschrift for Lucien Le Cam: research papers in probability and statistics*. New York: Springer, pp. 423–435.
- Yu, Y., Wang, T. & Samworth, R.J. (2015) A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102, 315–323.
- Zhang, N.R., Siegmund, D.O., Ji, H. & Li, J.Z. (2010) Detecting simultaneous changepoints in multiple sequences. *Biometrika*, 97, 631–645.
- Zhu, Z., Wang, T. & Samworth, R.J. (2019) High-dimensional principal component analysis with heterogeneous missingness. *arXiv preprint*, arxiv:1906.12125.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Follain, B., Wang, T. & Samworth, R.J. (2022) High-dimensional changepoint estimation with heterogeneous missingness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(3), 1023–1055. Available from: <https://doi.org/10.1111/rssb.12540>