



HAL
open science

Adaptive estimation of high-dimensional signal-to-noise ratios

Nicolas Verzelen, Elisabeth Gassiat

► **To cite this version:**

Nicolas Verzelen, Elisabeth Gassiat. Adaptive estimation of high-dimensional signal-to-noise ratios. *Bernoulli*, 2018, 24 (4B), pp.3683-3710. 10.3150/17-BEJ975 . hal-02622458

HAL Id: hal-02622458

<https://hal.inrae.fr/hal-02622458>

Submitted on 26 May 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Adaptive estimation of High-Dimensional Signal-to-Noise Ratios

Nicolas Verzelen* and Elisabeth Gassiat†

Abstract

We consider the equivalent problems of estimating the residual variance, the proportion of explained variance η and the signal strength in a high-dimensional linear regression model with Gaussian random design. Our aim is to understand the impact of not knowing the sparsity of the regression parameter and not knowing the distribution of the design on minimax estimation rates of η . Depending on the sparsity k of the regression parameter, optimal estimators of η either rely on estimating the regression parameter or are based on U -type statistics, and have minimax rates depending on k . In the important situation where k is unknown, we build an adaptive procedure whose convergence rate simultaneously achieves the minimax risk over all k up to a logarithmic loss which we prove to be non avoidable. Finally, the knowledge of the design distribution is shown to play a critical role. When the distribution of the design is unknown, consistent estimation of explained variance is indeed possible in much narrower regimes than for known design distribution.

1 Introduction

1.1 Motivations

In this paper, we investigate the estimation of the proportion of explained variation in high-dimensional linear models with random design, that is the ratio of the variance of the signal to the total amount of variance of the observation. Although this question is of great importance in many applications where the aim is to quantify to what extent covariates explain the variation of the response variable, our analysis is mainly motivated by problems of heritability estimation. In such studies, the response variable is a phenotype measured on n individuals and the predictors are genetic markers on each of these individuals. Then, heritability corresponds to the proportion of phenotypic variance which can be explained by genetic factors. Usually, the number of predictors p greatly exceeds the number n of individuals. When the phenotype under investigation can be explained by a small number of genetic factors, the corresponding regression parameter is sparse, and methods exploiting sparsity are of utmost interest. It appeared recently in biological studies that, for some complex human traits, there was a huge gap (which has been called the “dark matter” of the genome) between the genetic variance explained by populations studies and the one obtained by genome wide associations studies (GWAS), see [29], [33] or [21]. To explain this gap, it has been hypothesized that some traits might be “highly polygenic”, meaning that genetic factors explaining the phenotype could be so numerous that the corresponding regression parameter may no anymore considered to be sparse. This may be the case for instance when psychiatric disorders are associated to neuroanatomical changes as in [2] or [32], see also [35]. As a consequence, sparsity-based methods would be questionable in this situation. When the researcher faces the data, she does not know in

* (corresponding author) INRA, UMR 729 MISTEA, F-34060 Montpellier, FRANCE

† Laboratoire de Mathématiques d’Orsay, Univ. Paris-Sud, CNRS, Université Paris-Saclay, 91405 Orsay, FRANCE.

general the proportion of relevant predictors, that is the level of sparsity of the parameter. In this work, our first aim is to understand the impact of the ignorance of the sparsity level on heritability estimation. Another important feature of the model when estimating proportion of explained variation is the covariance matrix of the predictors. There is a long standing gap between estimation procedures that assume the knowledge of this covariance (e.g. [8, 24]) (which mathematically is the same as assuming that the covariance is the identity matrix) and practical situations where it is generally unknown. Our second aim is to evaluate the impact of the ignorance of the covariance matrix on heritability estimation.

To be more specific, consider the random design high-dimensional linear model

$$y_i = \mathbf{x}_i \beta^* + \epsilon_i, \quad i = 1, \dots, n \quad (1)$$

where $y_i, \epsilon_i \in \mathbb{R}$, $\beta^* \in \mathbb{R}^p$, $i = 1, \dots, n$, and $\mathbf{X} = \begin{pmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_n \end{pmatrix} \in \mathbb{R}^{n \times p}$. We assume that the noise

$\epsilon = (\epsilon_1, \dots, \epsilon_n)^T$ and the lines \mathbf{x}_i , $i = 1, \dots, n$, of \mathbf{X} are independent random variables. We also assume that the ϵ_i , $i = 1, \dots, n$, are independent and identically distributed (i.i.d.) with distribution $\mathcal{N}(0, \sigma^2)$, and that the lines \mathbf{x}_i , $i = 1, \dots, n$, of \mathbf{X} are also i.i.d. with distribution $\mathcal{N}(0, \Sigma)$. Throughout the paper, the covariance matrix Σ is assumed to be invertible and the noise level σ is unknown (the case of known noise level is evoked in the discussion section). Our general objective is the optimal estimation of the signal-to-noise ratio

$$\theta := \frac{\mathbb{E}[\|\mathbf{x}_1^T \beta^*\|_2^2]}{\sigma^2} = \frac{\|\Sigma^{1/2} \beta^*\|_2^2}{\sigma^2}, \quad (2)$$

or equivalently the proportion of explained variation

$$\eta = \eta(\beta^*, \sigma) := \frac{\mathbb{E}[\|\mathbf{x}_1^T \beta^*\|_2^2]}{\text{Var}(y_1)} = \frac{\theta}{1 + \theta} \quad (3)$$

when the vector β^* is unknown and possibly sparse. In the sequel, β^* is said to be k -sparse, when at most k coordinates of β^* are non-zero.

Note that estimating η amounts to decipher the signal strength from the noise level in $\text{Var}(y_1) = \sigma^2 + \|\Sigma^{1/2} \beta^*\|_2^2$. Since $\|Y\|_2^2 / \text{Var}(y_1)$ follows a χ^2 distribution with n degrees of freedom, it follows that $\|Y\|_2^2 / n = \text{Var}(y_1) [1 + O_P(n^{-1/2})]$ and it is therefore almost equivalent (up to a parametric $n^{-1/2}$ loss) to estimate the proportion of explained variation η , the quadratic function $\beta^{*T} \Sigma \beta^*$ or the noise level σ^2 . For the sake of presentation, we mostly express our results in terms of the estimation of η , but they can be easily extended to the signal strength or to the noise estimation problems.

1.2 Main results

There are two main lines of research for estimating σ or η in a high-dimensional setting. Under the assumption that β^* is k -sparse with some small k , it has been established that β^* can be estimated at a fast rate (roughly $\sqrt{k \log p/n}$) using for instance Lasso-type procedures, so that using an adequate plug-in method one could hope to estimate η well. Following this general approach, some authors have obtained $k \log(p)/n$ -consistent [34] and $\sqrt{1/n}$ -consistent [5, 20] estimators of σ in some specific regimes. When β^* is dense (that is when many coordinates of β^* are nonzero), such

approaches fail. In this regime, a U -type estimator [17] has been proved to achieve consistency at the rate \sqrt{p}/n . However, its optimality has never been assessed.

Our first main contribution is the proof that the adaptation to unknown sparsity is indeed possible when Σ is known, but at the price of a $\sqrt{\log(p)}$ loss factor in the convergence rate when β^* is dense. The idea is the following. Let $\hat{\eta}^D(\Sigma^{-1})$ be a U -type estimator which is \sqrt{p}/n -consistent, the true parameter β^* being sparse or not. We shall denote it the dense estimator. Let also $\hat{\eta}^{SL}$ be a $k \log(p)/n$ -consistent estimator when β^* is k -sparse for some small k . Then, if the real β^* is sparse, both estimators should be fairly accurate and should give similar answers, and if the real β^* is dense, or not sparse enough, then $\hat{\eta}^{SL}$ will be quite wrong and will give an answer slightly different from the dense estimator. Therefore, the idea is to choose the sparse estimator $\hat{\eta}^{SL}$ when both estimators are close enough, so that the quickly convergence rate is obtained when the unknown sparsity k is small, and to choose the dense estimator when both estimators are not close, in which case the slower rate is attained which is appropriate in the dense regime. Such a procedure should adapt well to unknown sparsity. Now, to be able to give a precise definition of the estimator, that is to set what “close enough” quantitatively means, one needs a precise understanding of the behavior of the dense and of the sparse estimators. Thus as a first and preliminary step, we obtain a deviation inequalities for the dense estimator, see Theorem 2.1. We also establish the minimax estimation risk of η as a function of (k, n, p) when the parameter β^* is k -sparse (see Table 1 below) and when Σ is known, thereby assessing that Dicker’s procedure [17] is optimal in the dense regime ($k \geq \sqrt{p}$) and an estimator based on the square-root Lasso [34] is near optimal in the sparse regime ($k \leq \sqrt{p}$). Again for known Σ , we finally construct a data-driven combination of $\hat{\eta}^D(\Sigma^{-1})$ (the dense estimator) and $\hat{\eta}^{SL}$ (the sparse estimator) following the idea explained before. We prove that such a procedure is indeed adaptive to unknown sparsity, see Theorem 3.2, and that it achieves the minimax adaptive rate with a $\sqrt{\log(p)}$ loss factor compared to the non adaptive minimax rate. This logarithmic term is proved to be unavoidable, see Proposition 3.1.

Our second main contribution is an analysis of the proportion of explained variance estimation problem under unknown Σ . The construction of dense estimators such as $\hat{\eta}^D(\Sigma^{-1})$ requires the knowledge of the covariance matrix Σ . But in many practical situations, the covariance structure of the covariates is unknown. For unknown Σ , there are basically two main situations:

- Under sufficiently strong structural assumptions on Σ so that Σ^{-1} can be estimated at the rate \sqrt{p}/n in operator norm, a simple plug-in method allows to build a minimax and an adaptive minimax procedure with the same rates as when Σ is known, see Corollary 4.4.
- Our main result is that, for a general covariance matrix Σ , it is basically impossible to build a consistent estimator of η when k is much larger than n ; see Theorem 4.5 and its comments for a precise statement. This is in sharp contrast with the situation where Σ is known, for which the problem of estimating η can be handled in regimes where β^* is impossible to estimate (e.g. $k = p$ and $p = n^{1+\kappa}$ with $\kappa \in (0, 1)$ as depicted in Table 1). For unknown and arbitrary Σ , the range of (k, n, p) for which η can be consistently estimated seems to be roughly the same as for estimating β^* , suggesting that signal estimation (β^*) is nearly as difficult as signal strength estimation ($\beta^{*T} \Sigma \beta^*$). This impossibility result unveils that, in the high-dimensional dense case, the knowledge of the covariance matrix is fundamental and one cannot extend known procedures such as [17, 18] or $\hat{\eta}^D(\Sigma^{-1})$ to this unknown variance setting.

Table 1: Optimal estimation risk $\mathbb{E}[(\hat{\eta} - \eta)^2]$ when β^* is k -sparse and Σ is known. Here, $a \in (0, 1/2)$ is any arbitrarily small constant and it is assumed below that $n \leq p \leq n^2$. The results remain valid for $p \geq n^2$ if we replace the quantities $\frac{k^2 \log^2(p)}{n^2}$ and $\frac{p}{n^2}$ by $\frac{k^2 \log^2(p)}{n^2} \wedge 1$ and $\frac{p}{n^2} \wedge 1$, respectively.

SPARSITY REGIMES	MINIMAX RISK	NEAR-OPTIMAL PROCEDURE
$k \leq \frac{\sqrt{n}}{\log(p)}$	$\frac{1}{n}$	SQUARE-ROOT LASSO ESTIMATOR $\hat{\eta}^{SL}$ (12)
$\frac{\sqrt{n}}{\log(p)} \leq k \leq p^{1/2-a}$	$\frac{k^2 \log^2(p)}{n^2}$	SQUARE-ROOT LASSO ESTIMATOR $\hat{\eta}^{SL}$ (12)
$k \geq \sqrt{p}$	$\frac{p}{n^2}$	DENSE ESTIMATOR $\hat{\eta}^D(\Sigma^{-1})$ (8) (see also [17])

1.3 Related work

The literature on minimax estimation of quadratic functionals initiated in [19] is rather extensive (see e.g. [11, 28]). In the Gaussian sequence model, that is $n = p$ and $\mathbf{X} = \mathbf{I}_p$, Collier et al [14] have derived the minimax estimation rate of the functional $\|\beta^*\|_2^2$ for k -sparse vector β^* when the noise level σ is known. However, we are not aware of any minimax result in the high-dimensional linear model even under known noise level.

Another problem related to the estimation of the quadratic functional $\beta^{*T} \Sigma \beta^*$ is signal detection, which aims at testing the null hypothesis $H_0: \beta^* = 0$ versus $H_{1,k}[r]: \|\Sigma^{1/2} \beta^*\|_2^2 \geq r$ and $|\beta^*|_0 \leq k$ (where $|\beta^*|_0$ denotes the number of non nul coordinates of β^*). The minimax separation distance is then the smallest r such that a test of H_0 vs $H_{1,r}$ is able to achieve small type I and type II error probabilities. This minimax separation distance is somewhat analogous to a local minimax estimation risk of $\|\Sigma^{1/2} \beta^*\|_2^2$ around $\beta^* = 0$. In the Gaussian sequence model, minimax separation distances haven been studied in [4, 23]. These results have been extended to the high-dimensional linear model under both known [3, 22] and unknown [22, 39] noise level. Our first minimax lower bound (Proposition 2.4) is largely inspired from these earlier contributions, but the minimax lower bounds for adaptation problems require more elaborate argument. In particular, the proof of Theorem 4.5 is largely based on new ideas.

Recent works have been devoted to the adaptive estimation of sparse parameters β^* in (1) under unknown variance. As a byproduct, one can then obtain estimators of the variance [5, 34]. See also [20] for more direct approaches to variance estimation. In Section 2, we rely on the square-root Lasso estimator to construct the estimator $\hat{\eta}^{SL}$ which turns out to be minimax in the sparse regime.

In the dense regime, we already mentioned the contribution of Dicker [17] that propose method of moments and maximum likelihood based procedures to estimate η when Σ is known. It is shown that the square risk of these estimators goes to 0 at rate \sqrt{p}/n . When p/n converges to a finite non-negative constant, these estimator are asymptotically normally distributed. Dicker also considers the case of unknown Σ when Σ is highly structured (allowing Σ to be estimable in operator norm at the parametric rate $n^{-1/2}$). Janson et al. [24] introduce the procedure EigenPrism for computing confidence intervals of η and study its asymptotic behavior when Σ is known and p/n converges to a constant $c \in (0, \infty)$. Under similar assumptions, Dicker et al. [18] have considered a maximum likelihood based estimator. Bonnet et al. [8] consider a mixed effect model, which is equivalent to assuming that the parameter β^* follows a prior distribution. In the asymptotic where $p/n \rightarrow c$, they also propose a $n^{-1/2}$ -rate consistent estimator of η . To summarize, none of the aforementioned contributions has studied minimax convergence rates, the problem of adaptation

to sparsity or the estimation problem for unknown Σ (to the exception of [17]).

Finally, there has been a recent interest in the adaptive estimation of other functionals in the linear model (1), such as the coordinates β_i^* of β^* or the sum of coordinates $\sum_{i=1}^n \beta_i^*$ [10, 25, 26, 38, 41]. However, both the statistical methods and the regimes are qualitatively different for these functionals.

1.4 Notations and Organization

The set of integers $\{1, \dots, p\}$ is denoted $[p]$. For any subset J of $[p]$, \mathbf{X}_J is the $n \times |J|$ corresponding submatrix of \mathbf{X} . Given a symmetric matrix \mathbf{A} , $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$ respectively stand for the largest and the smallest eigenvalue of \mathbf{A} , $|A|$ denotes the determinant of A . For a vector u , $\|u\|_p$ denotes its l_p norm and $|u|_0$ stands for its l_0 norm (ie number of non-zero components). For any matrix \mathbf{A} , $\|\mathbf{A}\|_p$ denotes the l_p norm of the vectorized version of \mathbf{A} , that is $(\sum |\mathbf{A}_{i,j}|^p)^{1/p}$. The Frobenius norm is also denoted $\|\mathbf{A}\|_F$. Finally, the l_2 operator norm of a matrix \mathbf{A} writes $\|\mathbf{A}\|_{op}$. In what follows, C, C', \dots denote universal constants whose value may vary from line to line whereas C_1, C_2 and C_3 denote numerical constants that will be used in several places of our work.

In Section 2, we introduce the two main procedures and characterize the minimax estimation risk of η when both the covariance matrix Σ and the sparsity are known. Section 3 is devoted to the problem of adaptation to the unknown sparsity, whereas the case of unknown covariance Σ is studied in Section 4. Extensions to fixed design regression and other related problems are discussed in Section 5. All the proofs are postponed to the end of the paper.

2 Minimax rates for known sparsity

In this section, we consider two estimators. In the spirit of [17], the first estimator $\hat{\eta}^D(\Sigma^{-1})$ is designed for the dense regime ($|\beta^*|_0 \geq p^{1/2}$) and it is proved to be consistent with rate $\sqrt{\bar{p}}/n$ irrespectively of the parameter sparsity. When β^* is in fact highly sparse, the estimator $\hat{\eta}^{SL}$ based on the square-root Lasso better exploits the structure of β^* and achieves the estimation rate $\frac{|\beta^*|_0 \log(p)}{n} + n^{-1/2}$. It turns out that these two procedures (almost) achieve the minimax estimation rate when $|\beta^*|_0$ is known.

2.1 Dense regime

In this subsection, we introduce an estimator of η which will turn out to be mostly interesting for dense parameters β^* . Its definition is close to that in [17]. We provide a detailed analysis of this estimator, and our bounds in Theorem 2.1 below will turn out to be useful both for the adaptation problem and for the case of unknown Σ .

Since $\text{Var}(y_1)$ is easily estimated by $\|Y\|_2^2/n$, the main challenge is to estimate $\|\Sigma^{1/2}\beta^*\|^2$. Thus, the question is how to separate in Y the randomness coming from $\mathbf{X}\beta^*$ from that coming from the ϵ_i 's, $i = 1, \dots, n$. The idea is to use the fact that the noise ϵ is isotropic whereas, conditionally on \mathbf{X} , $\mathbf{X}\beta^*$ is not isotropic. Respectively denote (λ_i, u_i) , $i = 1, \dots, n$ the eigenvalues and eigenvectors of $(\mathbf{X}\mathbf{X}^T)/p$. We will prove, that in a high-dimensional setting where $p > n$, $\mathbf{X}\beta^*$ is slightly more aligned with left eigenvectors of \mathbf{X} associated to large eigenvalues than with those associated to small eigenvalues. This subtle phenomenon suggests that the distribution of the random variable

T

$$T := \frac{1}{n^2} \sum_{i=1}^n (\lambda_i - \bar{\lambda})(Y^T u_i)^2, \quad \text{where} \quad \bar{\lambda} := \sum_{i=1}^n \lambda_i / n,$$

(almost) does not depend on the noise level σ and, at the same time, captures some functional of the signal β^* . This functional turns out to be $\beta^{*T} \Sigma^2 \beta^*$. One can rewrite the random variable as a quadratic form of \mathbf{Y}

$$T = \frac{Y^T (\mathbf{X}\mathbf{X}^T - \text{tr}(\mathbf{X}\mathbf{X}^T)\mathbf{I}_n/n) Y}{n^2}. \quad (4)$$

Working with a normalized estimator $\widehat{V} := \frac{Tn^2}{\|Y\|_2^2(n+1)}$, we state in the following theorem that \widehat{V} concentrates exponentially fast around $\beta^{*T} \Sigma^2 \beta^* / \text{Var}(y_1)$.

Theorem 2.1. *Assume that $p \geq n$.*

There exist numerical constants C_1 and C_2 such that for all $t \leq n^{1/3}$,

$$\mathbb{P} \left[\left| \widehat{V} - \frac{\beta^{*T} \Sigma^2 \beta^*}{\text{Var}(y_1)} \right| \leq C_1 \|\Sigma\|_{op} \frac{\sqrt{pt}}{n} \right] \geq 1 - C_2 e^{-t}. \quad (5)$$

There exists a numerical constant C such that

$$\mathbb{E} \left[\left(\widehat{V} - \frac{\beta^{*T} \Sigma^2 \beta^*}{\text{Var}(y_1)} \right)^2 \right] \leq C \|\Sigma\|_{op}^2 \frac{p}{n^2}. \quad (6)$$

Remark 2.1. The proof relies on recent exponential concentration inequalities for Gaussian chaos [1] and a new concentration inequality of the spectrum of $\mathbf{X}\mathbf{X}^T/n$ around $\text{tr}(\Sigma)/n$ (Lemma A.2). The concentration inequality (5) will be the key tool in the construction of adaptive estimators in the next section.

Remark 2.2. When Σ is the identity matrix, the above theorem enforces that \widehat{V} estimates the proportion of explained variation η at the rate \sqrt{p}/n , uniformly over all β^* and $\sigma > 0$. Note that \widehat{V} is only consistent in the regime where n^2 is large compared to p .

For arbitrary Σ (with bounded eigenvalues), the above theorem only implies that \widehat{V} is of the same order as η , that is, there exists positive constant c and C such that $c\lambda_{\min}(\Sigma) \leq \widehat{V}/\eta \leq C\lambda_{\max}(\Sigma)$. Nevertheless, when the covariance Σ is known, it is possible to get a consistent estimator of η . Replace the design matrix \mathbf{X} in the linear regression model by $\tilde{\mathbf{X}} := \mathbf{X}\Sigma^{-1/2}$ in such a way that its rows $\tilde{\mathbf{x}}_i$ follow i.i.d. standard normal distributions and

$$Y = \tilde{\mathbf{X}}\Sigma^{1/2}\beta^* + \epsilon. \quad (7)$$

Then, we define the estimator $\widehat{\eta}^D$ as \widehat{V} where \mathbf{X} is replaced by $\tilde{\mathbf{X}}$, so that $\widehat{\eta}^D$ is a quadratic form of \mathbf{Y} with a matrix involving the precision matrix, that is the inverse covariance matrix Σ^{-1} . Let us denote $\Omega := \Sigma^{-1}$, and define

$$\widehat{\eta}^D(\Omega) := \frac{Y^T (\mathbf{X}\Omega\mathbf{X}^T - \text{tr}(\mathbf{X}\Omega\mathbf{X}^T)\mathbf{I}_n/n) Y}{(n+1)\|Y\|^2} \quad (8)$$

(we could replace $\text{tr}(\mathbf{X}\Omega\mathbf{X}^T)$ by p in the above definition without changing the rate in the corollary below). We straightforwardly derive from Theorem 2.1 that $\widehat{\eta}^D(\Omega)$ estimates η at the rate \sqrt{p}/n .

Corollary 2.2. Assume that $p \geq n$. There exists a numerical constant C such that the estimator $\hat{\eta}^D(\Omega)$ satisfies

$$\mathbb{E} \left[(\hat{\eta}^D(\Omega) - \eta)^2 \right] \leq C \frac{p}{n^2}. \quad (9)$$

Remark 2.3. It turns out that $\hat{\eta}^D(\Omega)$ is consistent for p small compared to n^2 even though consistent estimation of β^* is impossible in this regime. Although developed independently, the estimator $\hat{\eta}^D(\Omega)$ shares some similarities with the method of moment based estimator of Dicker [17], which also achieves the \sqrt{p}/n convergence rate.

2.2 Sparse regime: square-root Lasso estimator

When β^* is highly sparse, the signal to noise ratio estimator is based on a Lasso-type estimator of β^* proposed in [6, 34]. As customary for Lasso-type methods, we shall work with a standardized version \mathbf{W} of the matrix \mathbf{X} , whose columns $\mathbf{W}_{\bullet j}$ satisfy $\|\mathbf{W}_{\bullet j}\|_2 = 1$. Since the noise-level σ is unknown, we cannot readily use the classical Lasso estimator whose optimal value of the tuning parameter depends on σ . Instead, we rely on the square-root Lasso [6] defined by

$$\tilde{\beta}_{SL} := \arg \min_{\beta \in \mathbb{R}^p} \sqrt{\|Y - \mathbf{W}\beta\|_2^2} + \frac{\lambda_0}{\sqrt{n}} \|\beta\|_1, \quad (\hat{\beta}_{SL})_j := (\tilde{\beta}_{SL})_j / \|\mathbf{x}_j\|_2. \quad (10)$$

In the sequel, the tuning parameter λ_0 is set to $\lambda_0 := 13\sqrt{\log(p)}$ (there is nothing specific with this particular choice). In the proof, we will also use an equivalent definition of the square-root estimator introduced in [34]

$$(\tilde{\beta}_{SL}, \tilde{\sigma}_{SL}) = \arg \min_{\beta \in \mathbb{R}^p, \sigma' > 0} \left[\frac{n\sigma'}{2} + \frac{\|Y - \mathbf{W}\beta\|_2^2}{2\sigma'} \right] + \lambda_0 \|\beta\|_1. \quad (11)$$

(To prove the equivalence between the two definitions, minimize (11) with respect to σ' .) Notice that $\tilde{\sigma}_{SL} = \|Y - \mathbf{W}\hat{\beta}_{SL}\|_2 / \sqrt{n}$. Then, we define the estimator

$$\hat{\eta}^{SL} := 1 - \frac{n\tilde{\sigma}_{SL}^2}{\|Y\|_2^2} = 1 - \frac{\|Y - \mathbf{W}\hat{\beta}_{SL}\|_2^2}{\|Y\|_2^2}, \quad (12)$$

The following proposition is a consequence of Theorem 2 in [34].

Proposition 2.3. There exist two numerical constants C and C' such that the following holds. Assume that β^* is k -sparse, that $p \geq n$ and

$$k \log(p) \frac{\lambda_{\max}(\Sigma)}{\lambda_{\min}(\Sigma)} \leq Cn. \quad (13)$$

Then the square-root Lasso based estimator $\hat{\eta}^{SL}$ satisfies

$$\mathbb{E} [(\hat{\eta}^{SL} - \eta)^2] \leq C' \left[\frac{1}{n} + \frac{k^2 \log^2(p)}{n^2} \frac{\lambda_{\max}^2(\Sigma)}{\lambda_{\min}^2(\Sigma)} \right]. \quad (14)$$

Remark 2.4. Condition (13) is unavoidable, as the minimax risk of proportion of explained variation estimation is bounded away from zero when $k \log(p)$ is large compared to n (see Proposition 2.4 later). To ease the presentation, we have expressed Condition (13) in terms of largest and smallest eigenvalues of Σ . One could in fact replace these quantities by local ones such as compatibility constants (see the proof for more details).

2.3 Minimax lower bound

We shall prove in the sequel that a combination of the estimators $\hat{\eta}^D(\mathbf{\Omega})$ and $\hat{\eta}^{SL}$ essentially achieves the minimax estimation risk. In the following minimax lower bound we assume that the covariance $\mathbf{\Sigma}$ is the identity matrix \mathbf{I}_p .

Define $\mathbb{B}_0[k]$ the collection of k -sparse vectors of size p . Given any estimator $\hat{\eta}$, define the maximal risk $R(\hat{\eta}, k)$ over k -sparse parameters by

$$R(\hat{\eta}, k) := \sup_{\beta \in \mathbb{B}_0[k], \sigma > 0} \mathbb{E}_{\beta, \sigma} [\{\hat{\eta} - \eta(\beta, \sigma)\}^2] ,$$

where $\mathbb{E}_{\beta, \sigma}[\cdot]$ is the expectation with respect to (Y, \mathbf{X}) where $Y = \mathbf{X}\beta + \epsilon$, with $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$ and the covariance matrix of the rows of \mathbf{X} is \mathbf{I}_p . Then, the minimax risk is denoted $R^*(k) := \inf_{\hat{\eta}} R(\hat{\eta}, k)$.

Proposition 2.4 (Minimax lower bound). *There exists a numerical constant $C > 0$ such that for any $1 \leq k \leq p$,*

$$R^*(k) \geq C \left(\left(\left[\frac{k}{n} \log \left(1 + \frac{p}{k^2} \vee \sqrt{\frac{p}{k^2}} \right) \right]^2 \wedge 1 \right) + \frac{1}{n} \right) . \quad (15)$$

The proof of this proposition follows the lines developed to derive minimax lower bounds for the signal detection problem (see e.g. Theorem 4.3 in [39]). Nevertheless, as this proposition is a first step towards more complex settings, we provide a self-contained proof in Section 7.1.

In (15), we recognize three regimes:

- If $k \geq p^{1/2}$, the minimax rate is larger than $(\sqrt{p}/n) \wedge 1$. This optimal risk is achieved by the dense estimator $\hat{\eta}^D(\mathbf{\Omega})$ up to a constant number.
- If $k \leq p^{1/2-\gamma}$ for some arbitrary small $\gamma > 0$, the minimax rate is of order

$$\frac{1}{\sqrt{n}} + \left(\frac{k \log(p)}{n} \right) \wedge 1 .$$

More precisely for $k \leq \lfloor \sqrt{n}/\log(p) \rfloor$, it is of order $n^{-1/2}$, whereas for larger k it is of order $k \log(p)/n \wedge 1$. This bound is achieved by the square-root Lasso estimator $\hat{\eta}^{SL}$, which does not require the knowledge of $\mathbf{\Sigma}$ and k .

- For k close to $p^{1/2}$ (e.g. $k = (p/\log(p))^{1/2}$), the minimax lower bound (15) and the upper bound (14) only match up to some $\log(p)$ factors. Such a logarithmic mismatch has also been obtained in the related work [4] on minimax detection rates for testing the null hypothesis $\beta^* = 0$ when the design matrix is fixed and orthonormal, that is $p = n$ and $\mathbf{X} = \mathbf{I}_p$. In this orthonormal setting, Collier et al. [14] have very recently closed this gap. Transposed in our setting, their results would suggest that the optimal risk is of order $k \log(p/k^2)/n$, suggesting that Proposition 2.4 is sharp. In the specific case where $\mathbf{\Sigma} = \mathbf{I}_p$, it seems possible to extend the estimator of $\|\beta^*\|_2^2$ introduced by [14] to our setting by considering the pairwise correlations $Y^T \mathbf{W}_{\bullet, j}$ for $j = 1, \dots, p$. Such estimator would then presumably be $k \log(p/k^2)/n$ consistent. As this approach does not seem extend easily to arbitrary $\mathbf{\Sigma}$, we did not go further in this direction.

3 Adaptation to unknown sparsity

In practice, the number $|\beta^*|_0$ of non-zero components of β^* is unknown. In this section, our purpose is to build an estimator $\hat{\eta}$ that adapts to the unknown sparsity $|\beta^*|_0$. Although the computation of the estimators $\hat{\eta}^D(\Omega)$ and $\hat{\eta}^{SL}$ does not require the knowledge of $|\beta^*|_0$, the choice of one estimator over the other depends on this quantity. Observe that, when $p \geq n^2$, the dense estimator $\hat{\eta}^D(\Omega)$ is not consistent. Therefore, only the estimator $\hat{\eta}^{SL}$ is useful and $\hat{\eta}^{SL}$ alone is minimax adaptive to the sparsity k (up to a possible log factor when k is of the order of $p^{1/2}$). This is why we focus on the regime where p is large compared to n and where $p \log p \leq n^2$.

It turns out that no estimator $\hat{\eta}$ can simultaneously achieve the minimax risk $R^*(k)$ over all $k = 1, \dots, p$, and that there is an unavoidable loss for adaptation. This may be seen in the following proposition.

Proposition 3.1. *Assume that $p \log p \leq n^2$, and that for some $a \in]0, 1/2[$, $p^{1-a}(\log p)^2 \geq 16n$. Then for any estimator $\hat{\eta}$, for all k such that $\sqrt{p \log p} \leq k \leq p$, one has*

$$\frac{R(\hat{\eta}, 1)}{\frac{1}{n} \sqrt{\frac{p}{n}}} + \frac{R(\hat{\eta}, k)}{\frac{p \log p}{n^2}} \geq \frac{a^2}{4^5}.$$

Recall that $R^*(1)$ is of order $1/n$ and $R^*(k)$ is of order p/n^2 . Proposition 3.1 implies that any estimator $\hat{\eta}$ whose maximal risk over $\mathbf{B}_0[k]$ is smaller than $p \log(p)/n^2$ exhibits a huge maximal risk over $\mathbf{B}_0[1]$. As a consequence, any estimator admitting a reasonable risk bound over $\mathbf{B}_0[1]$ should have a maximal risk at least of order $p \log(p)/n^2$ for all $k \in [\sqrt{p \log(p)}, p]$. Next, we define an estimator $\hat{\eta}^A$ simultaneously achieving the risk $R^*(k)$ for k small compared to \sqrt{p} and achieving the risk $R^*(k) \log p$ in the dense regime where $k \geq \sqrt{p \log p}$.

Define the numerical constant c_0 as two times the constant C_1 arising in the deviation bound (5) of Theorem 2.1. We build an adaptive estimator by combining the estimator $\hat{\eta}^{SL}$ and $\hat{\eta}^D$ as follows

$$\hat{\eta}^A := \begin{cases} \hat{\eta}^{SL} & \text{if } |\hat{\eta}_T^D(\Omega) - \hat{\eta}^{SL}| \leq c_0 \sqrt{p \log(p)}/n \\ \hat{\eta}_T^D(\Omega) & \text{else} \end{cases} \quad (16)$$

where, for technical reasons, we consider $\hat{\eta}_T^D(\Omega) := \min(1, \max(0, \hat{\eta}^D(\Omega)))$ a truncated version of $\hat{\eta}^D(\Omega)$ which lies in $[0, 1]$.

The rationale behind $\hat{\eta}^A$ is the following. Suppose that β^* is k -sparse, with $k \leq \sqrt{p}$, in which case, $\hat{\eta}^{SL}$ achieves the optimal rate. With large probability, $|\hat{\eta}_T^D(\Omega) - \eta|$ is smaller than $c_0 \sqrt{p \log(p)}/(2n)$ (this is true for arbitrary β^*) and $|\hat{\eta}^{SL} - \eta|$ is smaller than $(1/\sqrt{n} + k \log(p)/n)$ which is smaller than $c_0 \sqrt{p \log(p)}/(2n)$. Hence, $\hat{\eta}^A$ equals $\hat{\eta}^{SL}$ with large probability. Now assume that $k \geq \sqrt{p}$, in which case the optimal rate is of order \sqrt{p}/n and is achieved by $\hat{\eta}_T^D(\Omega)$. Observe that $\hat{\eta}^A = \hat{\eta}_T^D(\Omega)$ except if $\hat{\eta}^{SL}$ is at distance less than $c_0 \sqrt{p \log(p)}/n$ from $\hat{\eta}_T^D(\Omega)$. Consequently, $|\hat{\eta}^A - \eta| \leq c_0 \sqrt{p \log(p)}/n + |\hat{\eta}_T^D(\Omega) - \eta|$. Formalizing the above argument, we arrive at the following.

Theorem 3.2. *There exists a numerical constant C such that the following holds. Assume that $p \geq n$. For any integer $k \in [p]$, any k -sparse vector β^* and any $\sigma > 0$, the estimator $\hat{\eta}^A$ satisfies*

$$\mathbb{E} \left[(\hat{\eta}^A - \eta)^2 \right] \leq C \left[\frac{1}{n} + \left(\frac{k^2 \log^2(p)}{n^2} \frac{\lambda_{\max}^2(\Sigma)}{\lambda_{\min}^2(\Sigma)} \right) \wedge \left(\frac{p \log(p)}{n^2} \right) \right].$$

As a consequence of Propositions 2.4, 3.1 and 3.2, and, in the asymptotic regime where $p \log p \leq n^2$ and p^{1-a} is large compared to n for some positive a , $\hat{\eta}^A$ achieves the optimal adaptive risk for all $k \in \{1, \dots, p^{1/2-\gamma}\} \cup \{(p \log(p))^{1/2}, \dots, p\}$ where $\gamma > 0$ is arbitrary small. For k close to \sqrt{p} , there is still a logarithmic gap between the upper and lower bounds as in the non-adaptive section.

Remark 3.1. Theorem 2.1 is the basic stone for the construction of $\widehat{\eta}^A$ by the use of the deviation inequality. The constant c_0 may be quite large, as the constant C_1 in the deviation inequality, making the estimator difficult to use in practice if n and p are not large enough. Theorem 3.2 however allows to understand how adaptation to sparsity is possible.

4 Minimax estimation when Σ is unknown

In this section, we investigate the case where the covariance matrix Σ is unknown. As the computation of the sparse estimator $\widehat{\eta}_{SL}$ does not require the knowledge of Σ , the optimal estimation rate is therefore unchanged when $|\beta^*|_0$ is much smaller than \sqrt{p} . In what follows we therefore focus on the regime where $|\beta^*|_0 \geq \sqrt{p}$.

4.1 Positive results under restrictions on Σ

Here, we prove that a simple plug-in method allows to achieve the minimax rate as long as one can estimate the inverse covariance matrix Ω sufficiently well. This approach has already been considered by Dicker [17] who has proved a result analogous to Proposition 4.1. For the sake of completeness, we provide detailed arguments and also consider the problem of adaptation to the sparsity. Without loss of generality, we may assume that we have at our disposal an independent copy of \mathbf{X} , denoted $\mathbf{X}^{(2)}$ (if it is not the case, simply divide the data set into two subsamples of the same size).

Given an estimator $\widehat{\Omega}$ of $\Omega := \Sigma^{-1}$ based on the matrix $\mathbf{X}^{(2)}$, the proportion of explained variation η is estimated as in Section 2.1, using (8), except that the true inverse covariance matrix is replaced by its estimator:

$$\widehat{\eta}^D(\widehat{\Omega}) := \frac{Y^T (\mathbf{X}\widehat{\Omega}\mathbf{X}^T - \text{tr}(\mathbf{X}\widehat{\Omega}\mathbf{X}^T)\mathbf{I}_n/n) Y}{(n+1)\|Y\|^2}. \quad (17)$$

Proposition 4.1. Assume that $p \geq n$. For any non-singular estimator $\widehat{\Omega}$ based on the sample $\mathbf{X}^{(2)}$,

$$\mathbb{P} \left[|\widehat{\eta}^D(\widehat{\Omega}) - \eta| \geq C_1 \|\Sigma\|_{op} \|\widehat{\Omega}\|_{op} \frac{\sqrt{pt}}{n} + \|\Sigma\|_{op} \|\widehat{\Omega} - \Omega\|_{op} \left| \mathbf{X}^{(2)} \right| \right] \leq C_2 e^{-t}, \quad (18)$$

for all $t < n^{1/3}$. Here, C_1 and C_2 are the numerical constants that appear in Theorem 2.1.

Thus, if one is able to estimate Ω at the rate \sqrt{p}/n , then $\widehat{\eta}^D(\widehat{\Omega})$ achieves the same estimation rate as if Σ was known. To illustrate this qualitative situation, we describe an example of a class \mathcal{U} of precision matrices and an estimator $\widehat{\Omega}$ satisfying this property.

For any square matrix \mathbf{A} , define its matrix l_1 operator norm by $\|\mathbf{A}\|_{1 \rightarrow 1} = \max_{1 \leq j \leq p} \sum_{1 \leq i \leq p} |\mathbf{A}_{i,j}|$. Given any $M > 0$ and $M_1 > 0$, consider the following collection \mathcal{U} of sparse inverse covariance matrices

$$\mathcal{U} := \mathcal{U}(M, M_1) := \left\{ \Omega : \Omega \succ 0 : \begin{array}{l} \frac{1}{M_1} \leq \lambda_{\min}(\Omega) \leq \lambda_{\max}(\Omega) \leq M_1, \quad \|\Omega\|_{1 \rightarrow 1} \leq M, \\ \max_{1 \leq j \leq p} \sum_{i=1}^p \mathbf{1}_{\Omega_{i,j} \neq 0} \leq \sqrt{\frac{p}{n \log(p)}} \end{array} \right\}. \quad (19)$$

Cai et al [13] introduced the CLIME estimator to estimate sparse precision matrices. Let $\lambda_n > 0$ and $\rho > 0$ be two tuning parameters, whose value will be fixed in Lemma 4.2 below. Denote $\widehat{\Sigma}^{(2)} := \mathbf{X}^{(2)T} \mathbf{X}^{(2)}/n$ the empirical covariance matrix based on the observations $\mathbf{X}^{(2)}$.

Let $\widehat{\Omega}_1$ be the solution of the following optimization problem

$$\min \|\Omega'\|_1, \quad \text{subject to} \quad \|\widehat{\Sigma}^{(2)}\Omega' - \mathbf{I}_p\|_\infty \leq \lambda_n, \quad \Omega' \in \mathbb{R}^{p \times p}. \quad (20)$$

Then, the CLIME estimator $\widehat{\Omega}_{CL}$ is obtained by symmetrizing $\widehat{\Omega}_1$: for all i, j , we take $(\widehat{\Omega}_{CL})_{i,j} = (\widehat{\Omega}_1)_{i,j}$ if $|(\widehat{\Omega}_1)_{i,j}| \leq |(\widehat{\Omega}_1)_{j,i}|$ and $(\widehat{\Omega}_{CL})_{i,j} = (\widehat{\Omega}_1)_{j,i}$ in the opposite case. We may now apply Theorem 1.a in [13] to our setting with $\eta = 1/5 \wedge 1/\sqrt{M_1}$, $K = e^{1/2}$ and $\tau = 1$. This way we obtain the following.

Lemma 4.2. *There exists a numerical constant $C_3 > 0$ such that the following holds. Fix $\lambda_n = 2[25 \vee M_1](3 + e^3(5 \vee \sqrt{M_1})^2 M \sqrt{\log(p)/n})$. Assume that $\log(p) \leq n/8$ and that Ω belongs to \mathcal{U} . Then, the CLIME estimator satisfies*

$$\|\widehat{\Omega}_{CL} - \Omega\|_{op} \leq C_3 M^2 M_1^2 \frac{\sqrt{p}}{n}, \quad (21)$$

with probability larger than $1 - 4/p$.

Let us modify the estimator of η so that it effectively lies in $[0, 1]$. Let $\widehat{\eta}_T^D(\widehat{\Omega}) := \min(1, \max(0, \widehat{\eta}^D(\widehat{\Omega})))$.

Corollary 4.3. *Assume that $p \geq n$ and that Ω belongs to the collection \mathcal{U} defined above. Then, there exists a universal constant $C > 0$ such that the following holds. For any β^* and $\sigma > 0$,*

$$\mathbb{E} \left[\left\{ \widehat{\eta}_T^D(\widehat{\Omega}_{CL}) - \eta \right\}^2 \right] \leq \left[C M^4 M_1^6 \frac{p}{n^2} \right] \wedge 1.$$

We shall now define an adaptive estimator $\widehat{\eta}_{CL}^A$ in the same spirit as $\widehat{\eta}^A$ in the previous subsection. Define $c_0(M, M_1)$ by

$$c_0(M, M_1) := 4C_1 M_1^2 + 2C_3 M^2 M_1^3.$$

Here, C_1 is the numerical constant that appears in Theorem 2.1 and C_3 the numerical constant that appears in Lemma 4.2. Define the estimator as:

$$\widehat{\eta}_{CL}^A := \begin{cases} \widehat{\eta}^{SL} & \text{if } |\widehat{\eta}_T^D(\widehat{\Omega}_{CL}) - \widehat{\eta}^{SL}| \leq c_0(M, M_1) \sqrt{p \log(p)/n} \\ \widehat{\eta}_T^D(\widehat{\Omega}_{CL}) & \text{else} \end{cases}. \quad (22)$$

We then obtain that $\widehat{\eta}_{CL}^A$ is asymptotically minimax adaptive to Ω (if it is known that $\Omega \in \mathcal{U}$) and to sparsity, in the same regimes as those in which $\widehat{\eta}^A$ is asymptotically minimax adaptive to sparsity.

Corollary 4.4. *Assume that Ω belongs to the collection \mathcal{U} defined above. Then, there exists a constant $C(M, M_1) > 0$ only depending on M and M_1 such that the following holds. For any integer $k \in [p]$, any k -sparse vector β^* and any $\sigma > 0$,*

$$\mathbb{E} \left[(\widehat{\eta}_{CL}^A - \eta)^2 \right] \leq C(M, M_1) \left[\frac{1}{n} + \left(\frac{k^2 \log^2(p)}{n^2} \right) \wedge \left(\frac{p \log(p)}{n^2} \right) \right]. \quad (23)$$

Remark 4.1. When Ω belongs to \mathcal{U} , the estimator $\widehat{\eta}_T^D(\widehat{\Omega}_{CL})$ achieves a similar risk bound to that of $\widehat{\eta}_T^D(\Omega)$. Also, $\widehat{\eta}_{CL}^A$ performs as well as estimator $\widehat{\eta}^A$ which requires the knowledge of Ω . As a consequence, there does not seem to be a price to pay for the adaptation to Ω under the restriction $\Omega \in \mathcal{U}$.

Remark 4.2. If the quantity $\sqrt{p/(n \log(p))}$ in the sparsity condition $\max_{1 \leq j \leq p} \sum_{i=1}^p \mathbf{1}_{\Omega_{i,j} \neq 0} \leq \sqrt{p/(n \log(p))}$ in the definition (19) of \mathcal{U} is replaced by some $s \geq \sqrt{p/(n \log(p))}$, the CLIME-based estimator $\widehat{\eta}_T^D(\widehat{\Omega}_{CL})$ will only be consistent at the rate $s \sqrt{\log(p)/n}$ which is slower than the desired \sqrt{p}/n . This is not completely unexpected as we prove in the next subsection that a reliable estimation of η becomes almost impossible when the collection of precision matrices is too large.

4.2 Impossibility results

We now turn to the general problem where Σ is only assumed to have bounded eigenvalues. As explained in the beginning of Section 3, the estimator $\hat{\eta}^{SL}$, which does not require the knowledge of Σ , is minimax adaptive to $\mathbb{B}_0[k]$ when $p \geq n^2$. Hence, we focus in the remainder of this section on the regime $n \leq p \leq n^2$.

In this subsection and the corresponding proofs, we denote $\mathbb{P}_{\beta, \sigma, \Sigma}$ the distribution of (Y, \mathbf{X}) , in order to emphasize the dependency of the data distributions with respect to the covariance matrix of \mathbf{X} . For any $M > 1$, let us introduce $\Xi[M]$ the set of positive symmetric matrices of size p whose eigenvalues lie in the compact $[1/M, M]$. The purpose of these bounded eigenvalues in $(1/M, M)$ is to prove that the difficulty in the estimation problem does not simply arise because of poorly invertible covariance matrices.

Denote $\bar{R}^*[p, M]$ the minimax estimation risk of the the proportion of explained variation η when the covariance matrix is unknown

$$\bar{R}^*[p, M] := \inf_{\hat{\eta}} \sup_{\beta \in \mathbb{B}_0[p], \sigma > 0} \sup_{\Sigma \in \Xi[M]} \mathbb{E}_{\beta, \sigma, \Sigma} \left[(\hat{\eta} - \eta(\beta, \sigma))^2 \right]. \quad (24)$$

When the covariance matrix Σ is known, the minimax rate has been shown to be of order \sqrt{p}/n and therefore goes to 0 as soon as p is small compared to n^2 . The following proposition shows that, for unknown Σ , there is no consistent estimators of η when p is large compared to n .

Theorem 4.5. *Consider an asymptotic setting where both n and p go to infinity. Then, there exists a positive numerical constant C and a function $M : x \mapsto M(x)$ mapping $(0, \infty)$ to $(1, \infty)$ such that the following holds. If for some $\varsigma > 0$,*

$$\frac{n^{1+\varsigma}}{p} \rightarrow 0, \quad (25)$$

then the minimax risk $\bar{R}^[p, M(\varsigma)]$ is bounded away from zero, that is $\underline{\lim} \bar{R}^*[p, M(\varsigma)] \geq C$.*

Remark 4.3. Theorem 4.5 tells us that it is impossible to consistently estimate the proportion of explained variation in a high-dimensional setting where p is much larger than n . This lower bound straightforwardly extends to $\bar{R}^*[k, M(\varsigma)]$ when k is much larger than n in the sense $n^{1+\varsigma}/k \rightarrow 0$ for some $\varsigma > 0$.

Remark 4.4 (Dependency of constants with ς). In the proof of Theorem 4.5, the bound $M[\varsigma]$ is increasing when ς gets closer to zero, thereby allowing the spectrum of Σ to be broader. If we want to consider the minimax risk $\bar{R}^*[p, M]$ with a fixed $M > 0$ (independent of ξ), then one can prove that, whenever $n^{1+\varsigma}/p \rightarrow 0$, then $\underline{\lim} \bar{R}^*[p, M] \geq C(\varsigma)$ for some $C(\varsigma) > 0$ only depending on ς . Some details are provided in the proof.

Let us get a glimpse of the proof by trying to build an estimator of $\eta(\beta^*, \sigma)$ in the high-dimensional regime $p \geq n$. As Ω is unknown and cannot be consistently estimated in this regime, a natural candidate would be to consider $\hat{\eta}^D(\mathbf{I}_p) = \hat{V}$ as defined below (4). By Theorem 2.1, one has

$$\hat{\eta}^D(\mathbf{I}_p) = \frac{\beta^{*T} \Sigma^2 \beta^*}{\text{Var}(y_i)} + O_P\left(\frac{\sqrt{p}}{n}\right).$$

Although the signal strength $\beta^{*T} \Sigma \beta^*$ cannot be consistently estimated for unknown Σ (Theorem 4.5), it is interesting to note that some regularized version of the signal strength $\beta^{*T} \Sigma^2 \beta^*$ is estimable at the rate \sqrt{p}/n (this phenomenon was already observed in [17]).

Going one step further, one can consistently estimate $\beta^{*T}\Sigma^3\beta^*$ for $p \leq n^{3/2}$ by considering a quadratic form of Y as in T (4) but with higher-order polynomials of \mathbf{X} . For p of order $n^{1+\varsigma}$ for some small $\varsigma > 0$, it will be possible to consistently estimate all $a_q := \beta^{*T}\Sigma^q\beta^*$ for $q = 2, 3, \dots, r(\varsigma)$ where $r(\varsigma)$ is a positive integer only depending on ς .

Then, one may wonder whether it is possible to reconstruct $a_1 = \beta^{*T}\Sigma\beta^*$ from (a_q) , $q = 2, \dots, r(\varsigma)$. Observe that a_q is the q -th moment of a positive discrete measure μ supported by the spectrum of Σ and whose corresponding weights are the square norms of the projections of β^* on the eigenvectors of Σ . As a consequence, estimating $\beta^{*T}\Sigma\beta^*$ from (a_q) , $q = 2, \dots, r(\varsigma)$ is a partial moment problem where one aims at recovering the first moment of the measure μ given its higher order moments up to $r(\varsigma)$. Following these informal arguments, we build, in the proof of Theorem 4.5, two discrete measures μ_1 and μ_2 supported on $(1/M(\varsigma), M(\varsigma))$ whose q -th moments coincide for $q = 2, \dots, r(\varsigma)$ and whose first moments are far from each other. Define \mathcal{B}_1 (resp. \mathcal{B}_2) the collection of parameter (β^*, Σ) whose corresponding measure is μ_1 (resp. μ_2). Then, we show that no test can consistently distinguish the hypothesis $H_0 : (\beta^*, \Sigma) \in \mathcal{B}_1$ from $H_1 : (\beta^*, \Sigma) \in \mathcal{B}_2$. As the signal strengths $\beta^{*T}\Sigma\beta^*$ of parameters in \mathcal{B}_1 are far from those in \mathcal{B}_2 , this implies that consistent estimation is impossible in this setting.

Remark 4.5. Let us summarize our findings on the minimax estimation risk when Σ is unknown and $n \leq p \leq n^2$:

- if k is small compared to \sqrt{p} , the minimax risk is of order $[k \log(p)/n \wedge 1] + n^{-1/2}$ and is achieved by the square-root Lasso estimator $\hat{\eta}^{SL}$.
- if k is large compared to n (in the sense $n^{1+\varsigma}/k \rightarrow 0$ for some $\varsigma > 0$), then consistent estimation is impossible.
- if k lies between \sqrt{p} and $n/\log(p)$, the square-root Lasso estimator $\hat{\eta}^{SL}$ is consistent at the rate $k \log(p)/n$. We conjecture that this rate is optimal.
- if k lies between $n/\log(p)$ and n , we are not aware of any consistent estimator η and we conjecture that consistent estimation is impossible.

5 Discussion and extensions

We focused in this work on the estimation risk of η in high-dimensional linear models under two major assumptions: the design is random (with possibly unknown covariance matrix) and the level of noise σ is unknown. We first discuss how the difficulty of the problem is modified when the two assumptions are not satisfied: when the design is not random, then consistent estimation of η is impossible in the dense regime, and when the level of noise is known, then the estimation of η becomes much easier in the dense regime. Finally, we mention the problem of constructing optimal confidence intervals.

5.1 Fixed design

If the regression design \mathbf{X} is considered as fixed, then the counterpart of the proportion of explained variation would be

$$\eta[\beta^*, \sigma, \mathbf{X}] := \frac{\|\mathbf{X}\beta^*\|_2^2/n}{\|\mathbf{X}\beta^*\|_2^2/n + \sigma^2} .$$

In this new setting, the square-root Lasso estimator still estimates $\eta[\beta^*, \sigma, \mathbf{X}]$ at the rate $n^{-1/2} + k \log(p)/n$ up to multiplicative constants only depending on the sparse eigenvalues and compatibility constants of \mathbf{X} . In contrast, the construction of \hat{V} relies on the fact that \mathbf{X} is random and is independent of the isotropic noise ϵ . When \mathbf{X} is considered as fixed, \hat{V} does not consistently estimate $\eta[\beta^*, \sigma, \mathbf{X}]$ for p small compared to n^2 . As a simple example, take $\sigma = 1$ and define β^* by $\beta^{*T} v_i = \lambda_i^{-1/2}$ for $i = 1, \dots, n$ where $(v_i)_i$ denote the right eigenvectors of \mathbf{X} and $(\lambda_i^{1/2})_i$ its singular values. Then, the random variables T and \hat{V} (defined in Section 2.1) are concentrated around 0, whereas $\eta[\beta^*, \sigma, \mathbf{X}]$ equals 1/2.

More generally, the next proposition states that it is impossible to consistently estimate $\eta[\beta^*, \sigma, \mathbf{X}]$ in a high-dimensional setting $p \geq n + 1$. The randomness of \mathbf{X} therefore plays a fundamental role in the problem.

Proposition 5.1. *Assume that $p > n$ and consider any fixed design \mathbf{X} such that $\text{Rank}(\mathbf{X}) = n$. Given β^* and σ , denote $\mathbb{P}_{\beta^*, \sigma}$ and $\mathbb{E}_{\beta^*, \sigma}$ the probability and expectation with respect to the distribution $Y = \mathbf{X}\beta^* + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$. Then, the minimax estimation risk satisfies*

$$\inf_{\hat{\eta}} \sup_{\beta^* \in \mathbb{R}^p, \sigma \geq 0} \mathbb{E}_{\beta^*, \sigma} [(\hat{\eta} - \eta[\beta^*, \sigma, \mathbf{X}])^2] \geq \frac{1}{4}. \quad (26)$$

5.2 Knowledge of the noise level

Throughout this manuscript, we assumed that the noise level σ was unknown. As explained in the introduction, the situation is qualitatively different when σ is known. Let us briefly sketch the optimal convergence rates in this setting, still restricting ourselves to $p \geq n$. For any $k = 1, \dots, p$ define the maximal risk and the minimax risks

$$R(\hat{\eta}, k, \sigma) := \sup_{\beta \in \mathbb{B}_0[k]} \mathbb{E}_{\beta, \sigma} [\{\hat{\eta} - \eta(\beta, \sigma)\}^2], \quad R^*(k, \sigma) := \inf_{\hat{\eta}} R(\hat{\eta}, k, \sigma),$$

It follows from the minimax lower bounds for signal detection [3, 22], that for some $C > 0$ (lower bounds in [3, 22] are asymptotic but it is not difficult to adapt the arguments to obtain non-asymptotic bounds to the price of worse multiplicative constants),

$$R^*(k, \sigma) \geq C \left(\frac{k}{n} \log(1 + \frac{p}{k^2}) \right)^2 \wedge \frac{1}{n}, \quad (27)$$

which is of order $[k \log(p)/n]^2 \wedge n^{-1}$ except in the regime where n is of order p and where k is of order $p^{1/2}$ in which case the logarithmic factors do not match. As for the upper bounds, since $\|Y\|_2^2 / [\sigma^2 + \beta^{*T} \Sigma \beta^*]$ follows a χ^2 distribution with n degrees of freedom, the estimator $\hat{\eta}^{D, \sigma} := 1 - \frac{n\sigma^2}{\|Y\|_2^2}$ admits a quadratic risk (up to constants) smaller than $1/n$. This implies that the proportion of explained variation η can be efficiently estimated for arbitrarily large p . For small k , one can use the Gauss-Lasso estimator based on $\tilde{\beta}^{SL}$. Let \hat{J} be the set of integers j such that $\tilde{\beta}^{SL} \neq 0$ and define:

$$\hat{\eta}^{GL, \sigma} := \frac{\|\mathbf{\Pi}_{\hat{J}} Y\|_2^2 / n}{\sigma^2 + \|\mathbf{\Pi}_{\hat{J}} Y\|_2^2 / n}$$

where $\mathbf{\Pi}_{\hat{J}} = \mathbf{X}_{\hat{J}} (\mathbf{X}_{\hat{J}}^T \mathbf{X}_{\hat{J}})^{-1} \mathbf{X}_{\hat{J}}^T$ is the orthogonal projector of \mathbb{R}^n onto the space spanned by the columns of $\mathbf{X}_{\hat{J}}$. The Gauss-Lasso estimator was introduced to get an estimator of heritability in the sparse situation in a first version of this work [40]. Following the proof of Theorem 2.3 in [40] we may obtain that, under Assumption (13) and when $|\beta^*|_0 = k$,

$$\mathbb{E} [(\hat{\eta}^{GL, \sigma} - \eta)^2] \leq C' \frac{k^2 \log^2(p)}{n^2} \frac{\lambda_{\max}^2(\Sigma)}{\lambda_{\min}^2(\Sigma)}.$$

In conclusion, the rate $[k \log(p)/n] \wedge n^{-1/2}$ is (up to a possible logarithmic multiplicative term) optimal. These results contrast with the case of unknown σ in two ways: (i) The optimal rate is order-wise faster when σ is known especially when k is small ($n^{-1/2}$ versus $k \log(p)/n$) and when k, p are larger ($p^{1/2}/n$ versus $n^{-1/2}$). (ii) Since $\hat{\eta}^{D,\sigma}$ and $\hat{\eta}^{GL,\sigma}$ do not use the knowledge of Σ , adaptation to unknown covariance of the covariates is possible.

5.3 Minimax confidence intervals

In practice, one may not only be interested in the estimation of $\eta(\beta^*, \sigma)$, but also on building confidence intervals [24]. In the proof of Theorem 2.1 and in Proposition 2.3, we obtain exponential concentration inequalities of $\hat{\eta}^D(\mathbf{\Omega})$ and $\hat{\eta}^{SL}$ around β^* . This allows to get, for any $\alpha > 0$ and any $k = 1, \dots, p$, confidence intervals

$$\begin{aligned} IC_\alpha^D &:= \left[\hat{\eta}^D(\mathbf{\Omega}) \pm C(\alpha) \frac{\sqrt{p}}{n} \right], \\ IC_{\alpha,k}^{SL} &:= \left[\hat{\eta}^{SL} \pm C'(\alpha) \left(\frac{1}{n^{1/2}} + \frac{k \log(p)}{n} \frac{\lambda_{\max}^2(\mathbf{\Sigma})}{\lambda_{\min}^2(\mathbf{\Sigma})} \right) \right], \end{aligned}$$

where $C(\alpha)$ and $C'(\alpha)$ are universal constants only depending on α . When $p \geq n$, IC_α^D is honest over \mathbb{R}^p in the sense that

$$\inf_{\beta \in \mathbb{B}_0[p], \sigma > 0} \mathbb{P}_{\beta, \sigma} [\eta \in IC_\alpha^D] \geq 1 - \alpha.$$

For $p \geq n$ and if Assumption (13) is satisfied, then the confidence interval $IC_{\alpha,k}^{SL}$ is honest over $\mathbb{B}_0[k]$ in the sense that

$$\inf_{\beta \in \mathbb{B}_0[k], \sigma > 0} \mathbb{P}_{\beta, \sigma} [\eta \in IC_{\alpha,k}^{SL}] \geq 1 - \alpha.$$

In high-dimensional linear regressions, there have been recent advances towards the construction of optimal confidence regions both for the unknown vector β^* [30] or low-dimensional functional of the parameters such as components β_i^* [10, 25, 38, 41] or $\sum_i \beta_i^*$ [10]. Building on this line of work, it seems at hand to prove the minimax optimality of IC_α^D and $IC_{\alpha,k}^{SL}$, proving the existence of such honest confidence intervals. Of course, as already noticed when constructing our adaptive estimator, the choice of the constants $C(\alpha)$ and $C'(\alpha)$ are probably far to be optimal in applications.

A further step would be to study the problem of the construction (if possible) of adaptive confidence intervals. We leave those important questions for future research.

6 Proof of the upper bounds

6.1 Proof of Theorem 2.1

6.1.1 Some preliminary notation and deviation bounds

Consider the spectral decomposition $\mathbf{\Sigma} = \mathbf{O}\mathbf{\Gamma}\mathbf{O}^T$ where $\mathbf{\Gamma}$ is a diagonal matrix and \mathbf{O} is an orthogonal matrix. Define the matrix $\mathbf{Z} = \mathbf{X}\mathbf{O}\mathbf{\Gamma}^{-1/2}$ whose entries are independent standard normal variables. We denote

$$\rho := \frac{\mathbf{\Gamma}^{1/2} \mathbf{O}^T \beta^*}{\|\mathbf{\Sigma}^{1/2} \beta^*\|_2}. \quad (28)$$

In the following, we need to control the eigenvalues of $\mathbf{X}\mathbf{X}^T$. Define $\mathbf{A} := \mathbf{X}\mathbf{X}^T - \text{tr}[\boldsymbol{\Sigma}]\mathbf{I}_n$ and note that

$$\mathbf{A} = \mathbf{Z}\boldsymbol{\Gamma}\mathbf{Z}^T - \text{tr}(\boldsymbol{\Gamma})\mathbf{I}_n = \sum_{j=1}^p \boldsymbol{\Gamma}_{jj} (\mathbf{Z}_{\bullet j}\mathbf{Z}_{\bullet j}^T - \mathbf{I}_n) ,$$

where $\mathbf{Z}_{\bullet j}$ stands for the j -th column of \mathbf{Z} , so that \mathbf{A} is a weighted sum of centered Wishart matrices with parameters $(1, n)$. Extending the deviation inequalities of Davidson and Szarek [16] for Wishart matrices to weighted sums of Wishart matrices, we obtain the following, which is proved in Appendix B.

Lemma 6.1. *For any $t > 0$,*

$$\mathbb{P} \left[\|\mathbf{A}\|_{op} \leq 2\sqrt{\text{tr}(\boldsymbol{\Sigma})}\|\boldsymbol{\Sigma}\|_{op}^{1/2} \left(\sqrt{n} + 10 + \sqrt{2t} \right) + 3\|\boldsymbol{\Sigma}\|_{op} [n + 100 + 2t] \right] \geq 1 - 2e^{-t} . \quad (29)$$

As a consequence, for all $n \geq 20$, we get that

$$\mathbb{P} [\|\mathbf{A}\|_{op} \leq 25\|\boldsymbol{\Sigma}\|_{op}(\sqrt{np} + n)] \geq 1 - 2e^{-n} . \quad (30)$$

To control $\|\mathbf{A}\|_{op}$, we could have applied non-commutative Bernstein inequalities (Theorem 6.1.1 in [36]). However, this approach would have produced additional logarithmic terms.

6.1.2 Analysis of T

We decompose T into four terms, whose deviations will be controlled independently.

$$\begin{aligned} T &:= T_a + T_b + T_c + T_d \\ T_a &:= \frac{\beta^{*T}\mathbf{X}^T (\mathbf{X}\mathbf{X}^T - \text{tr}(\boldsymbol{\Sigma})\mathbf{I}_n) \mathbf{X}\beta^*}{n^2} \\ T_b &:= \frac{Y^T [\{ \text{tr}(\boldsymbol{\Sigma}) - \text{tr}(\mathbf{X}\mathbf{X}^T)/n \} \mathbf{I}_n] Y}{n^2} \\ T_c &:= \frac{\epsilon^T \mathbf{A} \epsilon}{n^2} , \quad T_d := 2 \frac{\epsilon^T \mathbf{A} \mathbf{X} \beta^*}{n^2} . \end{aligned}$$

Control of T_a . The main term in the above decomposition is T_a . Since its control is quite technical, we only state a deviation bound for the time being. Subsections 6.1.4 and 6.1.5 below are devoted to the proof of this lemma.

Lemma 6.2. *For all $t \leq n^{1/3}$, we have*

$$\mathbb{P} \left[\left| T_a - (1 + n^{-1}) \|\boldsymbol{\Sigma}\beta^*\|_2^2 \right| \geq C \|\boldsymbol{\Sigma}^{1/2}\beta^*\|_2^2 \|\boldsymbol{\Sigma}\|_{op} \frac{\sqrt{pt}}{n} \right] \leq 2e^{-t} . \quad (31)$$

Control of T_b , T_c , and T_d . Since $\text{tr}(\mathbf{X}\mathbf{X}^T)\mathbf{I}_n/n$ is a Gaussian quadratic form, we have by Lemma A.1 that

$$\mathbb{P} \left[\left| \text{tr}(\boldsymbol{\Sigma}) - \text{tr}(\mathbf{X}\mathbf{X}^T)/n \right| \geq 8\|\boldsymbol{\Sigma}\|_{op} \sqrt{\frac{pt}{n}} \right] \leq 2e^{-t} , \quad \forall t < np, \quad (32)$$

where we used that $\|\boldsymbol{\Sigma}\|_F \leq \sqrt{p}\|\boldsymbol{\Sigma}\|_{op}$. Also, $\|\mathbf{Y}\|_2^2 / \text{Var}(y_1)$ follows a χ^2 distribution with n degrees of freedom, which implies $\mathbb{P}[\|\mathbf{Y}\|_2^2 \geq \text{Var}(y_1)(n + 4\sqrt{nt})] \leq e^{-t}$ for all $t < n$. We conclude that for all $t < n$,

$$\mathbb{P} \left[|T_b| \geq 40 \text{Var}(y_1) \|\boldsymbol{\Sigma}\|_{op} \frac{\sqrt{pt}}{n^{3/2}} \right] \leq 3e^{-t} . \quad (33)$$

The term T_c is a Gaussian chaos of order 4. We could apply the general deviation bounds from [1], but it is easier to work here conditionally to \mathbf{X} . Conditionally to \mathbf{X} , T_c is a quadratic form with respect to ϵ . By Lemma A.1,

$$\mathbb{P} \left[\frac{n^2 |T_c|}{\sigma^2} \geq |\text{tr}(\mathbf{A})| + 8 \|\mathbf{A}\|_{op} \sqrt{nt} \right] \leq 2e^{-t}, \quad \forall t < n,$$

where we used $\|\mathbf{A}\|_F \leq \sqrt{n} \|\mathbf{A}\|_{op}$. Gathering this bound with the deviation inequality (32) for $\text{tr}(\mathbf{A})$, the deviation inequality (30) for $\|\mathbf{A}\|_{op}$, and using the fact that $p \geq n$, we conclude that, if $n \geq 20$, for all $t < n$,

$$\mathbb{P} \left[|T_c| \geq 208 \sigma^2 \|\Sigma\|_{op} \frac{\sqrt{pt}}{n^{3/2}} \right] \leq 6e^{-t}. \quad (34)$$

Conditionally to \mathbf{X} , $n^2 T_d / (2\sigma)$ follows a centered normal distribution with variance $\|\mathbf{A}\mathbf{X}\beta^*\|_2^2 \leq \|\mathbf{A}\|_{op}^2 \|\mathbf{X}\beta^*\|_2^2$. Hence,

$$\mathbb{P} \left[\frac{n^2 |T_d|}{2\sigma} \geq \|\mathbf{A}\|_{op} \|\mathbf{X}\beta^*\|_2 \sqrt{2t} \right] \leq 2e^{-t}, \quad \forall t > 0.$$

Then, $\|\mathbf{A}\|_{op}$ is controlled by (30) and $\|\mathbf{X}\beta^*\|_2^2 / \|\Sigma^{1/2}\beta^*\|_2^2$ follows a $\chi^2(n)$ distribution so that it can be controlled using Lemma A.1. If $n \geq 20$, for all $t < n$, we arrive at

$$\mathbb{P} \left[T_d \geq 125 \sigma \|\Sigma^{1/2}\beta^*\|_2 \|\Sigma\|_{op} \frac{\sqrt{pt}}{n} \right] \leq 5e^{-t}. \quad (35)$$

Gathering all the deviation inequalities (31–35), we obtain that for some constants $C, C' > 0$, if $n \geq 20$,

$$\mathbb{P} \left[\left| T - \beta^{*T} \Sigma^2 \beta^* (1 + n^{-1}) \right| \geq C \text{Var}(y) \|\Sigma\|_{op} \frac{\sqrt{pt}}{n} \right] \leq C' e^{-t} \quad (36)$$

for all $t \leq n^{1/3}$.

6.1.3 Analysis of \widehat{V}

Since $\|Y\|_2^2 / \text{Var}(y_1)$ follows a χ^2 distribution with n degrees of freedom, we obtain by Lemma A.1 that $\mathbb{P} \left[\left| \|Y\|_2^2 / n - \text{Var}(y_1) \right| \geq 4 \text{Var}(y_1) \sqrt{t/n} \right] \leq 2e^{-t}$ for all $t < n$, so that using (36) and the fact that $p \geq n$, we conclude that, for all $t \leq n^{1/3}$, with probability larger than $1 - (2 + C')e^{-t}$, for some constant $C > 0$,

$$\left| \widehat{V} - \frac{\beta^{*T} \Sigma^2 \beta^*}{\text{Var}(y)} \right| \leq C \|\Sigma\|_{op} \frac{\sqrt{pt}}{n},$$

and the first part of Theorem 2.1 is proved.

Let us now turn to the second moment of $U := \widehat{V} - \frac{\beta^{*T} \Sigma^2 \beta^*}{\text{Var}(y)}$. Define \mathcal{A} the event such that $|U| \leq C \|\Sigma\|_{op} \frac{\sqrt{pn^{1/3}}}{n}$, where C is the same constant as in the above bound. The probability of \mathcal{A} is larger than $1 - C' e^{-n^{1/3}}$ for some $C' > 0$. Then, the square risk decomposes as

$$\begin{aligned} \mathbb{E}[U^2] &\leq \mathbb{E}[\mathbf{1}_{\mathcal{A}} U^2] + 2 \mathbb{E} \left[\mathbf{1}_{\mathcal{A}^c} \left(\frac{\beta^{*T} \Sigma^2 \beta^*}{\text{Var}(y)} \right)^2 \right] + 2 \mathbb{E} \left[\mathbf{1}_{\mathcal{A}^c} (\widehat{V})^2 \right] \\ &\leq C \|\Sigma\|_{op}^2 \frac{p}{n^2} + 2 \mathbb{P}(\mathcal{A}^c) \|\Sigma\|_{op}^2 + 2 [\mathbb{P}(\mathcal{A}^c)]^{1/2} \left[\mathbb{E} \left\{ (\widehat{V})^4 \right\} \right]^{1/2}, \end{aligned}$$

where we have integrated the above deviation inequality in the last line. It remains to control the fourth moment of \widehat{V} . We have

$$\widehat{V} \leq \frac{Y^T (\mathbf{X}\mathbf{X}^T - \text{tr}(\mathbf{X}\mathbf{X}^T)\mathbf{I}_n/n) Y}{n\|Y\|_2^2} \leq \frac{\|\mathbf{A}\|_{op}}{n} + \frac{|\text{tr}(\boldsymbol{\Sigma}) - \text{tr}(\mathbf{X}\mathbf{X}^T)/n|}{n}.$$

Gathering the deviation inequalities (29) and (32), we derive that for some constants \tilde{C} and \tilde{C}' , for any $t > 0$,

$$\mathbb{P} \left[\widehat{V} \geq \tilde{C}\|\boldsymbol{\Sigma}\|_{op} \left(\sqrt{\frac{p}{n}} + \frac{\sqrt{pt}}{n} + t \right) \right] \leq \tilde{C}'e^{-t}.$$

Integrating this deviation inequality, we obtain that $\left[\mathbb{E} \left\{ (\widehat{V})^4 \right\} \right]^{1/2}$ is upper bounded by a constant times $\|\boldsymbol{\Sigma}\|_{op}^2 p/n$. In conclusion, for some numerical constant numbers C and C' ,

$$\mathbb{E} \left[\left| \widehat{V} - \frac{\beta^{*T} \boldsymbol{\Sigma}^2 \beta^*}{\text{Var}(y_i)} \right|^2 \right] \leq C \|\boldsymbol{\Sigma}\|_{op}^2 \frac{p}{n^2} \left[1 + ne^{-\frac{n^{1/3}}{2}} \right] \leq C' \|\boldsymbol{\Sigma}\|_{op}^2 \frac{p}{n^2}.$$

and the second part of Theorem 2.1 is proved.

6.1.4 Deviation inequalities for Gaussian chaos

We shall use deviation inequalities for (non necessarily homogeneous) Gaussian chaos. Let us recall a recent result from Adamczak and Wolff [1]. In order to state this result, we need to introduce some new notation.

Let d and q denote positive integers. Consider a d -indexed matrix $\mathbf{B} = (b_{i_1, \dots, i_d})_{i_1, \dots, i_d=1}^q$. For $\mathbf{i} = (i_1, \dots, i_d) \in [q]^d$ and $I \subset [d]$ we write $\mathbf{i}_I = (i_k)_{k \in I}$. Let P_d be the set of partitions of $[d]$ into non empty disjoint subsets. Given a partition $\mathcal{J} = \{J_1, \dots, J_k\}$, define the norm

$$\|\mathbf{B}\|_{\mathcal{J}} = \sup \left\{ \sum_{\mathbf{i} \in [q]^d} b_{\mathbf{i}} \prod_{l=1}^k x_{\mathbf{i}_{J_l}}^{(l)} : \|x^{(l)}\|_F \leq 1, \quad 1 \leq l \leq k \right\}, \quad (37)$$

where $x^{(l)}$ is a $|J_l|$ -indexed matrix and $\|x^{(l)}\|_F$ is its Frobenius norm.

Note that taking union of subsets in the partition increases the norm: given $\mathcal{J} = \{J_1, \dots, J_k\}$, the partition $\mathcal{J}' = \{\{J_1 \cup J_2\}, J_3, \dots, J_k\}$ satisfies $\|\mathbf{B}\|_{\mathcal{J}} \leq \|\mathbf{B}\|_{\mathcal{J}'}$. Indeed, the $|J_1| + |J_2|$ -dimensional matrix $x^{(1)} \otimes x^{(2)}$ in the definition (37) of $\|\mathbf{B}\|_{\mathcal{J}}$ satisfies $\|x^{(1)} \otimes x^{(2)}\|_F \leq 1$.

Proposition 6.3 (Theorem 2 in [1]). *Let $f : \mathbb{R}^q \mapsto \mathbb{R}$ be a polynomial of q variables of total degree smaller or equal to D . For any integer $d \geq 1$, let $\boldsymbol{\Gamma}^d f$ denote the d -th derivative of f . Let $Z = (Z_1, \dots, Z_q)$ denote a q -dimensional standard Gaussian vector. Then, for any $t > 0$,*

$$\mathbb{P} [|f(Z) - \mathbb{E}(f(Z))| \geq t] \leq 2 \exp \left[- \min_{1 \leq d \leq D} \min_{\mathcal{J} \in \mathcal{P}_d} \left(\frac{C^d t}{\|\boldsymbol{\Gamma}^d f(Z)\|_{\mathcal{J}}} \right)^{\frac{2}{\#\mathcal{J}}} \right], \quad (38)$$

where C is a numerical constant.

6.1.5 Proof of Lemma 6.2

Define the variable

$$V := \frac{n^2 T_a}{\|\Sigma^{1/2} \beta^*\|_2^2} = \rho^T \mathbf{Z}^T \left[(\mathbf{Z} \Gamma \mathbf{Z}^T) - \text{tr}(\Sigma) \mathbf{I}_n \right] \mathbf{Z} \rho ,$$

where we recall that ρ is introduced in (28). First, we compute the expectation of V :

$$\mathbb{E}[V] = (n^2 + n) \rho^T \Gamma \rho = \|\Sigma \beta^*\|_2^2 / \|\Sigma^{1/2} \beta^*\|_2^2 . \quad (39)$$

V is a polynomial $f(\mathbf{Z})$ of degree 4 of the $q = np$ independent standard Gaussian variables $\mathbf{Z} = (Z_{i,j})_{1 \leq i \leq n, 1 \leq j \leq p}$ so that we can apply Proposition 6.3. Since V is the sum of an homogeneous polynomial of degree 4 and an homogeneous polynomial of degree 2, we only have to consider the derivatives of order 2 and of order 4 in (38), all the other terms (of order 1 and 3) being null. Write V as

$$V = \sum_{i,j=1}^n \sum_{k,l,m=1}^p \mathbf{Z}_{ik} \mathbf{Z}_{jl} (\mathbf{Z}_{im} \mathbf{Z}_{jm} - \delta_{i,j}) \rho_k \rho_l \Gamma_{mm} , \quad (40)$$

where $\delta_{i,j} = 1$ is the indicator function of $i = j$. We may express V using the four-indexed matrix \mathbf{B} :

$$\mathbf{B}_{(j_1, k_1), (j_2, k_2), (j_3, k_3), (j_4, k_4)} = \rho_{k_1} \rho_{k_4} \delta_{j_1, j_2} \delta_{j_3, j_4} \delta_{k_2, k_3} \Gamma_{k_2 k_2}$$

as follows:

$$V = f(\mathbf{Z}) := \sum_{j_1, \dots, j_4=1}^n \sum_{k_1, \dots, k_4=1}^p \mathbf{B}_{(j_1, k_1), (j_2, k_2), (j_3, k_3), (j_4, k_4)} \mathbf{Z}_{j_1 k_1} (\mathbf{Z}_{j_2 k_2} \mathbf{Z}_{j_3 k_3} - \delta_{j_2, j_3}) \mathbf{Z}_{j_4 k_4}$$

so that the expectation $\mathbb{E}[\Gamma^4 f(\mathbf{Z})]$ of the fourth derivative of $f(\mathbf{Z})$ is obtained by a symmetrization of \mathbf{B} . More precisely, for any index (i_1, \dots, i_4) in $([n] \times [p])^4$, $\mathbb{E}[\Gamma^4 f(\mathbf{Z})]_{i_1, i_2, i_3, i_4} = \sum_{\sigma} \mathbf{B}_{i_{\sigma(1)}, \dots, i_{\sigma(4)}}$ where the sum runs over all permutations of $\{1, \dots, 4\}$. Using the triangular inequality, we shall obtain a bound on $\|\mathbb{E}[\Gamma^4 f(\mathbf{Z})]\|_{\mathcal{J}}$ from a bound on $\|\mathbf{B}\|_{\mathcal{J}}$. Thus it suffices to bound $\|\mathbf{B}\|_{\mathcal{J}}^2$ for all partitions \mathcal{J} . We start with $\mathcal{J} = \{1, 2, 3, 4\}$.

$$\|\mathbf{B}\|_{\{1,2,3,4\}}^2 = \|\mathbf{B}\|_F^2 \leq n^2 \|\rho\|_2^4 \text{tr}(\Gamma^2) = n^2 \text{tr}(\Sigma^2)$$

Let us now consider any partition $\mathcal{J} = \{J_1, J_2\}$ of size 2. Without loss of generality, there exists $t \in \{1, 2, 3\}$ such that $t \in J_1$ and $t+1 \in J_2$. Since each entry of \mathbf{B} contains a Dirac $\delta_{j_t, j_{t+1}}$ or $\delta_{k_t, k_{t+1}}$, there is a n or p factor less in $\|\mathbf{B}\|_{\mathcal{J}}^2$ in comparison to $\|\mathbf{B}\|_{\{1,2,3,4\}}^2$, and we get $\|\mathbf{B}\|_{J_1, J_2}^2 \leq n \lambda_{\max}^2(\Sigma)(p \vee n)$. Let us illustrate this with $J_1 = \{1\}$, $J_2 = \{2, 3, 4\}$. By symmetry, $\|\mathbf{B}\|_{\mathcal{J}}$ is achieved for $x_{j,k}^{(1)} = \rho_j n^{-1/2}$, and by Cauchy-Schwarz inequality, we obtain $\|\mathbf{B}\|_{\mathcal{J}}^2 = n \text{tr}(\Sigma^2)$.

If now the partition $\mathcal{J} = \{J_1, \dots, J_r\}$ has cardinality larger than 2, it was observed in the previous subsection that $\|\mathbf{B}\|_{\mathcal{J}}^2 \leq \|\mathbf{B}\|_{J_1, \cup_{s>2} J_s}^2$. We have thus proved that, for all $t > 0$,

$$\min_{\mathcal{J} \in \mathcal{P}_4} \left(\frac{t}{\|\mathbb{E}[\Gamma^4 f(\mathbf{Z})]\|_{\mathcal{J}}} \right)^{\frac{2}{\#\mathcal{J}}} \geq C \frac{t^2}{n^2 \text{tr}(\Sigma^2)} \wedge \left[\bigwedge_{k=2,3,4} \left(\frac{t^2}{n \lambda_{\max}^2(\Sigma)(p \vee n)} \right)^{1/k} \right] . \quad (41)$$

Let us now turn to the second derivative of $f(\mathbf{Z})$. Denote $\mathbf{B}' = \mathbb{E}[\Gamma^2 f(\mathbf{Z})]$. Coming back to the definition (40) of V , observe that $\mathbf{B}'_{(j_1, k_1), (j_2, k_2)}$ is zero when $(j_1 \neq j_2)$ because any term involving

j_1 and j_2 in V contains exactly two terms with index j_1 and two terms with index j_2 . Now, if $j_1 = j_2$, the entries of \mathbf{B}' are bounded in absolute values by

$$|\mathbf{B}'_{(j,k_1),(j,k_2)}| \leq Cn|\rho_{k_1}||\rho_{k_2}|(\Gamma_{k_1k_1} + \Gamma_{k_2k_2}) + \Gamma_{k_1k_1}.$$

As a consequence,

$$\begin{aligned} \|\mathbf{B}'\|_{\{1,2\}}^2 &= \|\mathbf{B}'\|_F^2 \leq Cn \sum_{k_1,k_2} n^2 |\rho_{k_1}|^2 |\rho_{k_2}|^2 \lambda_{\max}^2(\Sigma) + n \sum_{k_1} \Gamma_{k_1k_1}^2 \\ &\leq [n^3 \lambda_{\max}^2(\Sigma) + ntr(\Sigma^2)] , \end{aligned}$$

since $\|\rho\|_2^2 = 1$ and $\sum \Gamma_{k,k}^2 = tr(\Sigma^2)$. For $\mathcal{J} = \{\{1\}, \{2\}\}$, $\|\mathbf{B}'\|_{\mathcal{J}}$ is the spectral norm of \mathbf{B}' when considered as 2-dimensional $np \times np$ matrix. Since \mathbf{B}' can be seen as a block diagonal matrix, we obtain

$$\|\mathbf{B}'\|_{\{1\},\{2\}} \leq Cn\lambda_{\max}(\Sigma).$$

We arrive at

$$\min_{\mathcal{J} \in \mathcal{P}_2} \left(\frac{t}{\|\mathbb{E}[\Gamma^2 f(\mathbf{Z})]\|_{\mathcal{J}}} \right)^{\frac{2}{\#\mathcal{J}}} \geq C \frac{t^2}{n^3 \lambda_{\max}^2(\Sigma) + ntr(\Sigma^2)} \wedge \frac{t}{n\lambda_{\max}(\Sigma)}. \quad (42)$$

Proposition 6.3 together with (39), (41), and (42) allows us to conclude.

6.2 Proof of Proposition 2.3

This proposition is a consequence of the analysis of the square-root Lasso in [34]. We start with the decomposition

$$\widehat{\eta}^{SL} - \eta_* = \widetilde{\sigma}_{SL}^2 \left(\frac{n}{\|Y\|_2^2} - \frac{1}{\text{Var}(y_1)} \right) + \frac{\widetilde{\sigma}_{SL}^2 - \|\epsilon\|_2^2/n}{\text{Var}(y_1)} + \frac{\|\epsilon\|_2^2/n - \sigma^2}{\text{Var}(y_1)}. \quad (43)$$

By definition of the Lasso estimator, we have $n\widetilde{\sigma}_{SL}^2 = \|Y - \mathbf{W}\widehat{\beta}_{SL}\|_2^2 \leq \|Y\|_2^2$. As a consequence, the first term in the above equation is smaller in absolute value than $|1 - \|Y\|_2^2/(n \text{Var}(y_1))|$. Since $\|Y\|_2^2/\text{Var}(y_1)$ and $\|\epsilon\|_2^2/\sigma^2$ each follow a χ^2 distribution with n degrees of freedom, we have

$$\mathbb{E} \left[\left| \frac{\|Y\|_2^2}{n \text{Var}(y_1)} - 1 \right|^2 \right] = \frac{2}{n} \quad \text{and} \quad \mathbb{E} \left[\left(\frac{\|\epsilon\|_2^2/n - \sigma^2}{\text{Var}(y_1)} \right)^2 \right] = \frac{2\sigma^4}{n(\text{Var}(y_1))^2} \leq \frac{2}{n},$$

where we used $\text{Var}(y_1) = \|\Sigma^{1/2}\beta^*\|_2^2 + \sigma^2$. Let \mathcal{A} be an event of large probability to be defined below. Since $|\widehat{\eta}^{SL} - \eta_*| \leq 1$, we deduce from (43) that

$$\mathbb{E} \left[|\widehat{\eta}^{SL} - \eta_*|^2 \right] \leq \mathbb{P}(\mathcal{A}^c) + \frac{12}{n} + \frac{3 \mathbb{E} \left[(\widetilde{\sigma}_{SL}^2 - \|\epsilon\|_2^2/n)^2 \mathbf{1}_{\mathcal{A}} \right]}{\text{Var}^2(y_1)}, \quad (44)$$

so that we only have to focus on $\mathbb{P}(\mathcal{A}^c)$ and the difference $\widetilde{\sigma}_{SL}^2 - \|\epsilon\|_2^2/n$. We need a few more notation. In the sequel, J_* denotes the support of β^* , that is the set of indices i such that $\beta_i^* \neq 0$. For $T \subset [p]$ and $\xi > 0$, the compatibility constant $\kappa[\xi, T; \mathbf{W}]$ is defined by

$$\kappa[\xi, T; \mathbf{W}] = \min_{u \in \mathcal{C}(\xi, T)} \left\{ \frac{|T|^{1/2} \|\mathbf{W}u\|_2}{\|u_T\|_1} \right\}, \quad \text{where } \mathcal{C}(\xi, T) = \{u : \|u_{T^c}\|_1 < \xi \|u_T\|_1\}.$$

The compatibility constant, which quantifies how the design acts on the cone $\mathcal{C}(\xi, T)$, arises in state of the art results for the Lasso estimator [7, 27, 37]. We now define \mathcal{A} as the event on which the following conditions are satisfied:

$$\|\mathbf{W}^T \epsilon\|_\infty \leq 2\sigma\sqrt{\log(p)}, \quad (45)$$

$$\frac{8}{9}n\sigma^2 \leq \|\epsilon\|_2^2 \leq \frac{6}{5}n\sigma^2, \quad (46)$$

$$\kappa[5, J_*, \mathbf{W}] \geq 16^{-1}\lambda_{\min}^{1/2}(\boldsymbol{\Sigma})/\lambda_{\max}^{1/2}(\boldsymbol{\Sigma}). \quad (47)$$

The first lemma provides a deterministic prediction error for the square-root estimator. It is a simplified version of Theorem 2 in [34] (the notation and normalizations are slightly different).

Lemma 6.4 ([34]). *On the event \mathcal{A} , the design \mathbf{W} and the noise ϵ are such that*

$$12\lambda_0^2|\beta^*|_0 \log(p) \leq n\kappa^2[5, J_*; \mathbf{W}] \quad \text{and} \quad \|\mathbf{W}^T \epsilon\|_\infty \leq \frac{\lambda_0 \|\epsilon\|_2}{4\sqrt{n}}, \quad (48)$$

and the square root Lasso estimator satisfies

$$\max \left[1 - \frac{\sqrt{n}\tilde{\sigma}_{SL}}{\|\epsilon\|_2}, 1 - \frac{\|\epsilon\|_2}{\sqrt{n}\tilde{\sigma}_{SL}} \right] \leq 3\lambda_0^2 \frac{|\beta^*|_0}{n\kappa^2[5, J_*; \mathbf{W}]} \leq 1/2. \quad (49)$$

Proof. First, the second part of (48) is enforced by Conditions (45) and (46) together with the definition of λ_0 . The first part in (48) is a consequence of (47) and hypothesis (13). Then, we apply Theorem 2 of [34] to the estimator $\tilde{\sigma}_{SL}$. Notice that the choice of λ_0 and (48) in the above lemma differs by a factor \sqrt{n} from Theorem 2 in [34] because the design is normalized differently. Using the notation of [34], we fix $\xi = 2$ so that Condition (48) implies that $\tau_*^2 \leq 1/4$ (we fix $\nu = 1/2$ in [34, Eq.(16)]). Then, the condition on z^* in [34, Th.2] is a consequence of the second part of (48). The result follows. \square

It follows from (49) that, under \mathcal{A} ,

$$\begin{aligned} (\tilde{\sigma}_{SL}^2 - \|\epsilon\|_2^2/n)^2 &\leq 9 \frac{\|\epsilon\|_2^2}{n} (\tilde{\sigma}_{SL} - \|\epsilon\|_2/\sqrt{n})^2 \\ &\leq C \frac{\|\epsilon\|_2^4}{n^2} \lambda_0^4 \frac{|\beta^*|_0^2}{n^2 \kappa^4[5, J_*; \mathbf{W}]} \\ &\leq C' \sigma^4 \frac{|\beta^*|_0^2 \log^2(p)}{n^2} \frac{\lambda_{\max}^2(\boldsymbol{\Sigma})}{\lambda_{\min}^2(\boldsymbol{\Sigma})}, \end{aligned}$$

where we used the conditions (46) and (47) in the last line. In view of (44), Proposition 2.3 follows finally from the following lemma.

Lemma 6.5. *Under Assumption (13), we have for some positive constants C , C' , and C'' that*

$$\mathbb{P}[\mathcal{A}^c] \leq C[pe^{-C'(n \wedge p)} + p^{-1}] \leq \frac{C''}{n}.$$

Proof of Lemma 6.5. We control the probability of each event defined by (45), (46), and (47). Conditionally to \mathbf{W} , $\|\mathbf{W}^T \epsilon\|_\infty/\sigma$ is distributed as a supremum of p independent standard Gaussian variables. Applying an union bound over all variables $(\mathbf{W}^T \epsilon)_i$, we derive that

$$\mathbb{P} \left[\|\mathbf{W}^T \epsilon\|_\infty \leq \sigma\sqrt{4\log(p)} \right] \geq 1 - p^{-1}.$$

Turning to (46), we see that $\|\epsilon\|_2^2/\sigma^2$ follows a χ^2 distribution with n degrees of freedom. By Lemma A.1, we obtain

$$\mathbb{P} \left[\frac{8}{9}n\sigma^2 \leq \|\epsilon\|_2^2 \leq \frac{6}{5}n\sigma^2 \right] \geq 1 - 2e^{-Cn} ,$$

for some positive constant $C > 0$.

Finally, we need to control the compatibility constant $\kappa[5, J_*; \mathbf{W}]$. As the compatibility constant is larger than restricted eigenvalues, we can readily apply the results of [31]. In particular, their Corollary 1 entails that, with probability larger than $1 - c_1 \exp[-c_2 n]$ for some $c_1 > 0$ and $c_2 > 0$,

$$\kappa[5, J_*; \mathbf{X}/(\sqrt{n}\lambda_{\min}^{1/2}(\boldsymbol{\Sigma}))] \geq 1/8 ,$$

as long as $|J_*| \log(p) < c_3 n$. The latter condition is satisfied by hypothesis (13). Coming back to the definition of \mathbf{W} and of the compatibility constant, we have

$$\kappa[5, J_*; \mathbf{W}] \geq \kappa[5, J_*; \mathbf{X}/(\sqrt{n}\lambda_{\min}^{1/2}(\boldsymbol{\Sigma}))] \frac{\sqrt{n}\lambda_{\min}^{1/2}(\boldsymbol{\Sigma})}{\max \|\mathbf{X}_{\bullet i}\|_2} .$$

Since, for all i , $\boldsymbol{\Sigma}_{ii}$ is larger than $\lambda_{\min}(\boldsymbol{\Sigma})$, we can apply Lemma A.1 to get

$$\mathbb{P} \left[\min_{i=1, \dots, p} \|\mathbf{X}_{\bullet i}\|_2^2 \geq n\lambda_{\min}(\boldsymbol{\Sigma})/2 \right] \geq 1 - pe^{-Cn} \quad (50)$$

for some positive constant $C > 0$. Finally, Assumption (13) enforces that $\log(p)$ is small compared to n so that pe^{-Cn} is smaller than C'/n for some positive constant C' . \square

6.3 Proof of Theorem 3.2

Notice first that we always have $|\hat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \leq |\hat{\eta}^D(\boldsymbol{\Omega}) - \eta|$ and $|\hat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \leq 1$.

We first consider the case where (β^*, σ) is arbitrary. The difference $\hat{\eta}^A - \eta$ decomposes as

$$\hat{\eta}^A - \eta = (\hat{\eta}^{SL} - \hat{\eta}^D(\boldsymbol{\Omega})) \mathbf{1}_{\hat{\eta}^A = \hat{\eta}^{SL}} + (\hat{\eta}^D(\boldsymbol{\Omega}) - \eta) \mathbf{1}_{\hat{\eta}^A = \hat{\eta}^{SL}} + (\hat{\eta}_T^D(\boldsymbol{\Omega}) - \eta) \mathbf{1}_{\hat{\eta}^A = \hat{\eta}_T^D(\boldsymbol{\Omega})} .$$

The difference $\hat{\eta}^D - \eta$ is controlled thanks to Corollary 2.2, whereas the difference $\hat{\eta}^{SL} - \hat{\eta}^D(\boldsymbol{\Omega})$ is small when $\hat{\eta}^A = \hat{\eta}^{SL}$ by definition of $\hat{\eta}^A$.

$$\begin{aligned} \mathbb{E} \left[(\hat{\eta}^A - \eta)^2 \right] &\leq 3 \mathbb{E} \left[(\hat{\eta}^D(\boldsymbol{\Omega}) - \eta)^2 \right] + 3 \mathbb{E} \left[(\hat{\eta}_T^D(\boldsymbol{\Omega}) - \eta)^2 \right] + 3 \mathbb{E} \left[(\hat{\eta}^{SL} - \hat{\eta}^D(\boldsymbol{\Omega}))^2 \mathbf{1}_{\hat{\eta}^A = \hat{\eta}^{SL}} \right] \\ &\leq 6 \mathbb{E} \left[(\hat{\eta}^D(\boldsymbol{\Omega}) - \eta)^2 \right] + 3 \mathbb{E} \left[(\hat{\eta}^{SL} - \hat{\eta}^D(\boldsymbol{\Omega}))^2 \mathbf{1}_{\hat{\eta}^A = \hat{\eta}^{SL}} \right] \\ &\leq C \frac{p}{n^2} + 2c_0^2 \frac{p \log(p)}{n^2} , \end{aligned}$$

where we used the definition of $\hat{\eta}^A = \hat{\eta}^{SL}$ in the last line. Thus, considering that the risk of the estimator is bounded by 1, it is possible to choose the numerical constant C such that Theorem 3.2 holds true if β^* is k -sparse with k such that (13) does not hold.

Assume now that β^* is k -sparse with k such that (13) holds. We start from the decomposition

$$\mathbb{E} \left[(\hat{\eta}^A - \eta)^2 \right] = \mathbb{E} \left[(\hat{\eta}^{SL} - \eta)^2 \mathbf{1}_{\hat{\eta}^A = \hat{\eta}^{SL}} \right] + \mathbb{E} \left[(\hat{\eta}_T^D(\boldsymbol{\Omega}) - \eta)^2 \mathbf{1}_{\hat{\eta}^A = \hat{\eta}_T^D(\boldsymbol{\Omega})} \right] .$$

In this sparse setting, the risk of $\hat{\eta}^{SL}$ is minimax optimal but the risk of $\hat{\eta}_T^D(\boldsymbol{\Omega})$ is possibly quite large. We have to work around the event $\hat{\eta}^A = \hat{\eta}_T^D(\boldsymbol{\Omega})$. This event can only be achieved if either we

have $|\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \geq c_0 \sqrt{p \log(p)}/(2n)$ or if we have simultaneously $|\widehat{\eta}^{SL} - \eta| \geq c_0 \sqrt{p \log(p)}/(2n)$ and $|\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \leq c_0 \sqrt{p \log(p)}/(2n)$. Under this last possibility, observe that $|\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \leq |\widehat{\eta}^{SL} - \eta|$. Thus, we obtain

$$\begin{aligned} \mathbb{E} \left[(\widehat{\eta}^A - \eta)^2 \right] &\leq \mathbb{E} \left[(\widehat{\eta}^{SL} - \eta)^2 \mathbf{1}_{\widehat{\eta}^A = \widehat{\eta}^{SL}} \right] + \mathbb{E} \left[(\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta)^2 \mathbf{1}_{|\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \geq c_0 \sqrt{p \log(p)}/(2n)} \right] \\ &\quad + \mathbb{E} \left[(\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta)^2 \mathbf{1}_{|\widehat{\eta}^{SL} - \eta| \geq c_0 \sqrt{p \log(p)}/(2n)} \mathbf{1}_{|\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \leq c_0 \sqrt{p \log(p)}/(2n)} \right] \\ &\leq 2 \mathbb{E} \left[(\widehat{\eta}^{SL} - \eta)^2 \right] + \mathbb{P} \left[|\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \geq c_0 \sqrt{p \log(p)}/(2n) \right]. \end{aligned}$$

The risk $\mathbb{E}[(\widehat{\eta}^{SL} - \eta)^2]$ is bounded thanks to Proposition 2.3 whereas the deviation inequality

$$\mathbb{P} \left[|\widehat{\eta}_T^D(\boldsymbol{\Omega}) - \eta| \geq c_0 \sqrt{p \log(p)}/(2n) \right]$$

is smaller than C_2/p by Theorem 2.1. Together with the fact that $p \geq n$, we have proved that when β^* is k -sparse with k such that 13 holds,

$$\mathbb{E} \left[(\widehat{\eta}^A - \eta)^2 \right] \leq C \left[\frac{1}{n} + \frac{k^2 \log^2(p)}{n^2} \frac{\lambda_{\max}^2(\boldsymbol{\Sigma})}{\lambda_{\min}^2(\boldsymbol{\Sigma})} \right].$$

Theorem 3.2 follows.

6.4 Analysis of the plug-in method

Proof of Proposition 4.1. We first note that the estimator $\widehat{\eta}^D(\widehat{\boldsymbol{\Omega}})$ is built using the following linear regression model

$$Y = \left[\mathbf{X} \widehat{\boldsymbol{\Omega}}^{1/2} \right] \left[\widehat{\boldsymbol{\Omega}}^{-1/2} \beta^* \right] + \epsilon.$$

It then follows from Theorem 2.1 in Section 2.1 that $\widehat{\eta}^D(\widehat{\boldsymbol{\Omega}})$ is an estimator of $\|\widehat{\boldsymbol{\Omega}}^{1/2} \boldsymbol{\Sigma} \beta^*\|_2^2 / \text{Var}(y_1)$. More precisely, we have

$$\mathbb{P} \left[\left| \widehat{\eta}^D(\widehat{\boldsymbol{\Omega}}) - \frac{\|\widehat{\boldsymbol{\Omega}}^{1/2} \boldsymbol{\Sigma} \beta^*\|_2^2}{\text{Var}(y_1)} \right| \geq C_1 \|\boldsymbol{\Sigma}\|_{op} \|\widehat{\boldsymbol{\Omega}}\|_{op} \frac{\sqrt{pt}}{n} \mid \mathbf{X}^{(2)} \right] \leq C_2 e^{-t}$$

for all $t < n^{1/3}$. Decomposing the difference $\widehat{\eta}^D(\widehat{\boldsymbol{\Omega}}) - \eta$ into

$$\widehat{\eta}^D(\widehat{\boldsymbol{\Omega}}) - \eta = \widehat{\eta}^D(\widehat{\boldsymbol{\Omega}}) - \frac{\|\widehat{\boldsymbol{\Omega}}^{1/2} \boldsymbol{\Sigma} \beta^*\|_2^2}{\text{Var}(y_1)} + \frac{\|\widehat{\boldsymbol{\Omega}}^{1/2} \boldsymbol{\Sigma} \beta^*\|_2^2}{\text{Var}(y_1)} - \eta,$$

we only have to consider the second term

$$\begin{aligned} \left| \frac{\|\widehat{\boldsymbol{\Omega}}^{1/2} \boldsymbol{\Sigma} \beta^*\|_2^2}{\text{Var}(y_1)} - \eta \right| &= \left| \frac{\beta^{*T} \boldsymbol{\Sigma}^{1/2} (\boldsymbol{\Sigma}^{1/2} \widehat{\boldsymbol{\Omega}} \boldsymbol{\Sigma}^{1/2} - \mathbf{I}_p) \boldsymbol{\Sigma}^{1/2} \beta^*}{\text{Var}(y_1)} \right| \\ &\leq \|\boldsymbol{\Sigma}^{1/2} \widehat{\boldsymbol{\Omega}} \boldsymbol{\Sigma}^{1/2} - \mathbf{I}_p\|_{op} \\ &\leq \|\boldsymbol{\Sigma}\|_{op} \|\widehat{\boldsymbol{\Omega}} - \boldsymbol{\Omega}\|_{op} \end{aligned}$$

where we used in the second line that $\beta^{*T} \boldsymbol{\Sigma} \beta^* / \text{Var}(y_1) \leq 1$. □

Proof of Corollary 4.3. Define the event \mathcal{B} such that inequality (21) is true. Assume first that

$$CM^2M_1^2\frac{\sqrt{p}}{n} \leq (2M_1)^{-1}, \quad (51)$$

where C is the numerical constant in (21).

$$\begin{aligned} \mathbb{E} \left[\left\{ \widehat{\eta}_T^D(\widehat{\Omega}_{CL}) - \eta \right\}^2 \right] &= \mathbb{E} \left[\left\{ \widehat{\eta}_T^D(\widehat{\Omega}_{CL}) - \eta \right\}^2 \mathbf{1}_{\mathcal{B}} \right] + \mathbb{E} \left[\left\{ \widehat{\eta}_T^D(\widehat{\Omega}_{CL}) - \eta \right\}^2 \mathbf{1}_{\mathcal{B}^c} \right] \\ &\leq \mathbb{E} \left[\left\{ \widehat{\eta}_T(\widehat{\Omega}_{CL}) - \eta \right\}^2 \mathbf{1}_{\mathcal{B}} \right] + \mathbb{P}[\mathcal{B}^c], \end{aligned}$$

where we used that $\widehat{\eta}_T^D(\widehat{\Omega}_{CL})$ belongs to $[0, 1]$. By Lemma 4.2, $\mathbb{P}[\mathcal{B}^c] \leq 4/p$. Under event \mathcal{B} , $\lambda_{\min}(\widehat{\Omega}_{CL}) \geq \lambda_{\min}(\Omega) - \|\widehat{\Omega}_{CL} - \Omega\|_{op} \geq (2M_1)^{-1}$ and $\widehat{\Omega}_{CL}$ is therefore non-singular. Plugging (21) in Proposition 4.1 and integrating the deviation bound with respect to $t > 0$, we get that for some numerical constant C ,

$$\mathbb{E} \left[\left\{ \widehat{\eta}_T^D(\widehat{\Omega}_{CL}) - \eta \right\}^2 \right] \leq CM^4M_1^6\frac{p}{n^2}.$$

If now, (51) is not satisfied, we just use that since the thresholded estimator is $\widehat{\eta}_T^D(\widehat{\Omega}_{CL})$ belongs to $[0, 1]$, the risk is always smaller than 1, which is smaller than $CM^2M_1^3\frac{\sqrt{p}}{n}$. \square

Proof of Corollary 4.4. In order to show (23), we follow the same steps as for proving Proposition 3.2, the only difference being that we need to prove that $\mathbb{P}[|\widehat{\eta}_T^D(\widehat{\Omega}_{CL}) - \eta| \geq c_0(M, M_1)\sqrt{p \log p}/(2n)]$ is larger than C/p for some $C > 0$. As above, we consider two cases whether (51) is satisfied or not. If Condition (51) is satisfied, we use Proposition 4.1 with $t = \log(p)$ and the event \mathcal{B} to prove that

$$\mathbb{P} \left[\left| \widehat{\eta}_T^D(\widehat{\Omega}) - \eta \right| \geq 2C_1M_1^2\frac{\sqrt{p \log(p)}}{n} + C_3M^2M_1^3\frac{\sqrt{p}}{n} \right] \leq \frac{C_2 + 4}{p}.$$

If Condition (51) is not satisfied, we again use that

$$|\widehat{\eta}_T^D(\widehat{\Omega}) - \eta| \leq 1 \leq 2CM^2M_1^3\frac{\sqrt{p}}{n}.$$

\square

7 Proofs of the minimax lower bounds

7.1 Proof of Proposition 2.4

7.1.1 Proof of the parametric rate $R^*(k) \geq R^*(1) \geq Cn^{-1}$

First, we prove that η cannot be estimated faster than the parametric rate $n^{-1/2}$. Fix $\sigma = 1$, $\beta_1^* = (1, 0, \dots, 0)^T$ and $\beta_2^* = (1 + n^{-1/2}, 0, \dots, 0)^T$. Then $\eta_1 = \eta(\beta_1^*, \sigma) = 1/2$ and $\eta_2 = \eta(\beta_2^*, \sigma) \geq 1/2 + n^{-1/2}/4$. Denoting $\mathbb{K}(\mathbb{P}_{\beta_1^*, \sigma}; \mathbb{P}_{\beta_2^*, \sigma})$ the Kullback-Leibler divergence between $\mathbb{P}_{\beta_1^*, \sigma}$ and $\mathbb{P}_{\beta_2^*, \sigma}$, we have

$$\mathbb{K}(\mathbb{P}_{\beta_1^*, \sigma}; \mathbb{P}_{\beta_2^*, \sigma}) = \mathbb{E} \left[\frac{\|\mathbf{X}(\beta_1^* - \beta_2^*)\|_2^2}{2} \right] = \frac{1}{2}.$$

Using Pinsker's inequality, we provide a lower bound of $R^*(1)$ in terms of $\mathbb{K}(\mathbb{P}_{\beta_1^*, \sigma}; \mathbb{P}_{\beta_2^*, \sigma})$ and $(\eta_1 - \eta_2)^2$ as follows:

$$\begin{aligned}
R^*(1) &\geq \inf_{\hat{\eta}} \mathbb{E}_{\beta_1^*, \sigma} [(\hat{\eta} - \eta_1)^2] \vee \mathbb{E}_{\beta_2^*, \sigma} [(\hat{\eta} - \eta_2)^2] \\
&\geq \frac{(\eta_2 - \eta_1)^2}{4} \inf_{\hat{\eta}} \mathbb{P}_{\beta_1^*, \sigma} [\hat{\eta} \geq (\eta_1 + \eta_2)/2] \vee \mathbb{P}_{\beta_2^*, \sigma} [\hat{\eta} \leq (\eta_1 + \eta_2)/2] \\
&\geq \frac{(\eta_2 - \eta_1)^2}{8} \inf_{\mathcal{A}} \mathbb{P}_{\beta_1^*, \sigma}(\mathcal{A}) + \mathbb{P}_{\beta_2^*, \sigma}(\mathcal{A}^c), \quad \text{where } \mathcal{A} \text{ is any measurable event} \\
&\geq \frac{(\eta_2 - \eta_1)^2}{8} [1 - \|\mathbb{P}_{\beta_1^*, \sigma} - \mathbb{P}_{\beta_2^*, \sigma}\|_{TV}] \\
&\geq \frac{(\eta_2 - \eta_1)^2}{8} \left[1 - 2^{-1/2} \mathbb{K}^{1/2}(\mathbb{P}_{\beta_1^*, \sigma}; \mathbb{P}_{\beta_2^*, \sigma})\right], \quad \text{by Pinsker's inequality} \\
&\geq \frac{(\eta_2 - \eta_1)^2}{16} \geq \frac{1}{16^2 n},
\end{aligned}$$

which concludes the proof.

7.1.2 Proof of $R^*(k) \geq C \left\{ \left[\frac{k}{n} \log \left(1 + \frac{p}{k^2} \vee \sqrt{\frac{p}{k^2}} \right) \right]^2 \wedge 1 \right\}$

In this proof, we follow the standard strategy of reducing the heritability estimation problem to a detection problem, thereby taking advantage on available bounds of [39]. We could simply derive Proposition 2.4 from Theorem 4.3 in [39], but we prefer to detail the arguments as a first step towards the minimax lower bounds for adaptation problems.

Denote \mathbb{P}_0 the distribution of (Y, \mathbf{X}) when $\beta^* = 0$ and $\sigma = 1$. Let $\rho > 0$ be a positive quantity that will be fixed later. Also, denote \mathcal{B} the collection of all vectors $\beta \in \mathbb{R}^p$ with exactly k non-zero components that are either equal to $\frac{\rho}{[(1+\rho^2)k]^{1/2}}$ or $-\frac{\rho}{[(1+\rho^2)k]^{1/2}}$. Defining $\sigma_\rho^2 := (1 + \rho^2)^{-1}$, we obtain, for all $\beta \in \mathcal{B}$, $\eta(\beta, \sigma_\rho) = \rho^2 / (1 + \rho^2)$. Following the beaten path of Le Cam's approach, we consider μ the uniform measure on \mathcal{B} and denote \mathbf{P}_μ the mixture probability measure

$$\mathbf{P}_\mu = \int_{\mathcal{B}} \mathbb{P}_{\beta, \sigma_\rho} \mu(d\beta) \tag{52}$$

Let $\hat{\eta}$ be any estimator of η . The minimax risk $R^*(k)$ is obviously lower bounded as follows:

$$\begin{aligned}
R^*(k) &\geq \mathbb{E}_0 [\hat{\eta}^2] \vee \vee_{\beta \in \mathcal{B}} \mathbb{E}_{\beta, \sigma_\rho} \left[\left(\hat{\eta} - \frac{\rho^2}{1 + \rho^2} \right)^2 \right] \\
&\geq \frac{1}{2} \left[\mathbb{E}_0 [\hat{\eta}^2] + \mathbf{E}_\mu \left[\left(\hat{\eta} - \frac{\rho^2}{1 + \rho^2} \right)^2 \right] \right] \\
&\geq \frac{\rho^4}{8(1 + \rho^2)^2} \left[\mathbb{P}_0 \left[\hat{\eta} > \frac{\rho^2}{2(1 + \rho^2)} \right] + \mathbf{P}_\mu \left[\hat{\eta} \leq \frac{\rho^2}{2(1 + \rho^2)} \right] \right].
\end{aligned}$$

Defining the test statistic $\hat{T} := \mathbf{1}\{\hat{\eta} > \rho^2 / [2(1 + \rho^2)]\}$, one recognizes in the bound above the sum

of type I and type II errors of the test \mathbb{P}_0 versus \mathbf{P}_μ . We arrive at

$$\begin{aligned}
R^*(k) &\geq \frac{\rho^4}{8(1+\rho^2)^2} \left[\mathbb{P}_0[\widehat{T} = 1] + \mathbf{P}_\mu[\widehat{T} = 0] \right] \\
&\geq \frac{\rho^4}{8(1+\rho^2)^2} \left[1 - |\mathbb{P}_0(\widehat{T} = 0) - \mathbf{P}_\mu(\widehat{T} = 0)| \right] \\
&\geq \frac{\rho^4}{8(1+\rho^2)^2} [1 - \mathbb{E}_0 |L_\mu - 1|] \quad \text{where } \mathbb{L}_\mu = \frac{d\mathbf{P}_\mu}{d\mathbb{P}_0} \\
&\geq \frac{\rho^4}{8(1+\rho^2)^2} \left[1 - (\chi^2(\mathbf{P}_\mu, \mathbb{P}_0))^{1/2} \right], \quad (\text{by Cauchy-Schwarz inequality}) \quad (53)
\end{aligned}$$

where $\chi^2(\mathbf{P}_\mu, \mathbb{P}_0) = \mathbb{E}_0[(L_\mu - 1)^2]$ stands for the χ^2 distance between probability distributions. As a consequence, we only need to bound the χ^2 distance between \mathbf{P}_μ and \mathbb{P}_0 . Fortunately, this distance has been controlled in [39] (take $v = 1$, $\text{Var}(y) = 1$ in [39, p.741, line 14] and note that $k\lambda^2 = \rho^2/(1+\rho^2)$).

Lemma 7.1 ([39]). *We have*

$$\chi^2(\mathbf{P}_\mu, \mathbb{P}_0) \leq \exp \left[k \log \left(1 + \frac{k}{p} \left(\cosh \left(\frac{n\rho^2}{k} \right) - 1 \right) \right) \right] - \frac{1}{2}. \quad (54)$$

Let us fix ρ^2 in such a way that

$$\frac{n\rho^2}{k} = \log \left[1 + \frac{p}{k^2} \log(5/4) + \sqrt{(1 + \frac{p}{k^2} \log(5/4))^2 - 1} \right]. \quad (55)$$

Using the classical equality $\cosh(\log(1+x+\sqrt{x^2+2x})) = 1+x$ for $x \geq 0$, we arrive at

$$\chi^2(\mathbf{P}_\mu, \mathbb{P}_0) \leq \exp [k \log(1 + \log(5/4)/k)] - 1/2 \leq 3/4,$$

which, together with (53), implies

$$R^*(k) \geq \frac{\rho^4}{8(1+\rho^2)^2} (1 - (3/4)^{1/2}).$$

Since $\log(1+ux) \geq u \log(1+x)$ for any $u \in (0,1)$ and $x > 0$, we derive from (55) that

$$\rho^2 \geq \log \left(\frac{5}{4} \right) \left[\frac{k}{n} \log \left(1 + \frac{p}{k^2} \vee \sqrt{\frac{p}{k^2}} \right) \right]^2.$$

But

$$\frac{\rho^2}{1+\rho^2} \geq \frac{\rho^2}{2} \wedge 1,$$

which concludes the proof.

7.2 Proof of Proposition 3.1

Define the quantity $\rho > 0$ by

$$\rho^2 := \frac{a\sqrt{p \log p}}{4n}. \quad (56)$$

We consider $\mu, \mathbf{P}_\mu, \mathbf{E}_\mu$ as introduced in the proof of Proposition 2.4. Let $\hat{\eta}$ be a given estimator. Define

$$R := n\sqrt{\frac{n}{p}}\mathbf{E}_0[\hat{\eta}^2] + \frac{n^2}{p\log p}\mathbf{E}_\mu\left[\left(\hat{\eta} - \frac{\rho^2}{1+\rho^2}\right)^2\right].$$

Then,

$$\frac{R(\hat{\eta}, 1)}{\frac{1}{n}\sqrt{\frac{p}{n}}} + \frac{R(\hat{\eta}, k)}{\frac{p\log p}{n^2}} \geq R.$$

Now define the event $\mathcal{A}(\hat{\eta}) := \{\hat{\eta} \geq \rho^2/[2(1+\rho^2)]\}$. Then, one has

$$n\sqrt{\frac{n}{p}}\mathbf{E}_0[\hat{\eta}^2] \geq n\sqrt{\frac{n}{p}}\mathbb{P}_0[\mathcal{A}(\hat{\eta})]\frac{\rho^4}{4(1+\rho^2)^2} \geq \frac{a^2}{4^4}\sqrt{\frac{p}{n}}\log p\mathbb{P}_0[\mathcal{A}(\hat{\eta})].$$

Similarly, $\mathbf{E}_\mu(\hat{\eta} - \rho^2/(1+\rho^2))^2 \geq \mathbf{P}_\mu(\mathcal{A}^c(\hat{\eta}))\rho^4/4(1+\rho^2)^2 \geq \frac{a^2}{4^4}\frac{p\log p}{n^2}\mathbf{P}_\mu(\mathcal{A}^c(\hat{\eta}))$ so that

$$R \geq \frac{a^2}{4^4} \inf_{\mathcal{A}} \left\{ \mathbb{P}_0[\mathcal{A}] \sqrt{\frac{p}{n}} \log p + \mathbf{P}_\mu[\mathcal{A}^c] \right\},$$

were the infimum is taken over all measurable events \mathcal{A} . Restricting the events \mathcal{A} to have small probability, we arrive at

$$R \geq \frac{a^2}{4^4} \left[1 \wedge_{\mathcal{A}, \mathbb{P}_0[\mathcal{A}] \leq \sqrt{n}/(\sqrt{p}\log p)} \inf_{\mathcal{A}} \mathbf{P}_\mu[\mathcal{A}^c] \right], \quad (57)$$

so that it suffices to obtain a uniform lower bound $\mathbf{P}_\mu[\mathcal{A}^c]$ over events \mathcal{A} of small \mathbb{P}_0 -probability.

$$\begin{aligned} \mathbf{P}_\mu(\mathcal{A}^c) &\geq 1 - \mathbb{P}_0(\mathcal{A}) - |\mathbf{P}_\mu(\mathcal{A}) - \mathbb{P}_0(\mathcal{A})| \\ &\geq 1 - \mathbb{P}_0(\mathcal{A}) - |\mathbf{E}_0[(\mathbb{L}_\mu - 1)\mathbf{1}_{\mathcal{A}}]| \quad \text{where } \mathbb{L}_\mu = \frac{d\mathbf{P}_\mu}{d\mathbb{P}_0} \\ &\geq 1 - \mathbb{P}_0(\mathcal{A}) - (\mathbb{P}_0[\mathcal{A}]\chi^2(\mathbf{P}_\mu, \mathbb{P}_0))^{1/2}. \quad (\text{by Cauchy-Schwarz inequality}) \end{aligned} \quad (58)$$

Define $x = \frac{ap}{2k^2}\log(p)$. Since $\sqrt{p\log p} \leq k$, and since $\log(1+ux) \geq u\log(1+x)$ for any $u \in (0, 1)$ and $x > 0$, we have

$$\frac{n\rho^2}{k} \leq \log[1+x\vee\sqrt{x}] \leq \log\left[1+x+\sqrt{2x+x^2}\right].$$

Together with Lemma 7.1 and the classical identity $\cosh[\log(1+u+\sqrt{2u+u^2})] = 1+u$ for all $u > 0$, we arrive at

$$\frac{\sqrt{n}}{\sqrt{p}\log p}\chi^2(\mathbf{P}_\mu, \mathbb{P}_0) \leq \frac{\sqrt{n}}{\sqrt{p}\log p} \exp\left[\frac{k^2}{p}x\right] \leq \frac{\sqrt{n}}{\sqrt{p}\log p} p^{a/2} = \sqrt{\frac{n}{p^{1-a}\log p}}. \quad (59)$$

Coming back to the lower bound (58), we conclude that, for any event \mathcal{A} satisfying $\mathbb{P}_0(\mathcal{A}) \leq \sqrt{n}/(\sqrt{p}\log p)$, we have

$$\mathbf{P}_\mu(\mathcal{A}^c) \geq 1 - \sqrt{\frac{n}{p}\frac{1}{\log p}} - \left(\sqrt{\frac{n}{p^{1-a}\log p}}\right)^{1/2}.$$

Plugging this result in (57) and using the fact that $p^{1-a}(\log p)^2 \geq 16n$ leads to the desired result.

7.3 Proof of Theorem 4.5

7.3.1 General arguments

Suppose that Condition (25) is satisfied for some $\varsigma > 0$. Define r be the smallest integer such that $\varsigma \geq 1/(2r)$ so that we can assume henceforth that $n^{1+1/(2r)}/p \rightarrow 0$.

In this proof, we follow the same general approach as in the other minimax lower bounds, that is we define two mixture distributions \mathbf{P}_0 and \mathbf{P}_1

$$\mathbf{P}_0 := \int \mathbb{P}_{\beta, \sigma_0, \Sigma} \mu_0(d\beta, d\Sigma), \quad \mathbf{P}_1 := \int \mathbb{P}_{\beta, \sigma_1, \Sigma} \mu_1(d\beta, d\Sigma),$$

in such a way that \mathbf{P}_0 and \mathbf{P}_1 are almost indistinguishable and at the same time the function $\eta(\beta, \sigma)$ takes different values for parameters in the support of the prior distribution μ_0 and parameters in the support of the prior distribution μ_1 . The main difference with previous proofs lies in the fact that μ_0 and μ_1 are now prior probabilities on both the regression coefficient β and the covariance matrix Σ .

Let $\alpha_0 = (\alpha_{i,0})$, $\gamma_0 = (\gamma_{i,0})$, $i = 1, \dots, r$ and $\alpha_1 = (\alpha_{i,1})$, $\gamma_1 = (\gamma_{i,1})$, $i = 0, \dots, r$ be positive parameters whose exact values will be fixed later. We emphasize that the values of these parameters will only depend on r and not on n and p . Given a positive integer q and $\alpha = (\alpha_1, \dots, \alpha_q)$ whose coordinates α_j are positive, define the probability distribution π_α on vectors of $\mathbb{R}^{q \times p}$ whose density is proportional to $(\mathbf{I}_p + \sum_{i=1}^q \alpha_i x_i x_i^T)^{-n/2} e^{-\sum_{i=1}^q p \|x_i\|_2^2 / 2}$ for $x_i \in \mathbb{R}^p$, $i = 1, \dots, q$.

The distribution μ_0 is defined as follows. Let $(v_{i,0})$, $i = 1, \dots, r$ be independently sampled according to the distribution π_{α_0} . Then, conditionally to $(v_{1,0}, \dots, v_{r,0})$, β and Σ are fixed to the following values

$$\beta = \sum_{i=1}^r \gamma_{i,0} v_{i,0}; \quad \Sigma^{-1} = \mathbf{I}_p + \sum_{i=1}^r \alpha_{i,0} v_{i,0} v_{i,0}^T. \quad (60)$$

Similarly, under μ_1 ,

$$\beta = \sum_{i=0}^r \gamma_{i,1} v_{i,1}; \quad \Sigma^{-1} = \mathbf{I}_p + \sum_{i=0}^r \alpha_{i,1} v_{i,1} v_{i,1}^T \quad (61)$$

where the vectors $(v_{i,1})$, $i = 1, \dots, r$ are independently sampled according to the distribution π_{α_1} . Finally, the noise variances are fixed to the following values.

$$\sigma_0^2 = 3/2, \quad \sigma_1^2 = 1/2. \quad (62)$$

To prove that \mathbf{P}_0 and \mathbf{P}_1 are almost indistinguishable we will consider separately the marginal distribution of \mathbf{X} and the conditional distribution of Y given \mathbf{X} . We will see that the centered Gaussian distribution of \mathbf{X} under both \mathbf{P}_0 and \mathbf{P}_1 are indistinguishable from the standard normal distribution when $n = o(p)$, see Lemma 7.4 below.

Let us now choose the parameters $\gamma_{i,j}$ and $\alpha_{i,j}$ in such a way that the conditional distribution of Y given \mathbf{X} under \mathbf{P}_0 is indistinguishable from that under \mathbf{P}_1 when $n^{1+1/(2r)} = o(p)$. We first consider a truncated moment problem.

Lemma 7.2. *There exist two discrete positive measures $\rho_0 = \sum_{i=1}^r \xi_{i,0} \delta_{\tau_{i,0}}$ on $\rho_1 = \sum_{i=0}^r \xi_{i,1} \delta_{\tau_{i,1}}$ supported on $(0, 1)$ such that*

1. *The atoms $\tau_{i,j}$ for $j = 0, 1$ and $i = 1, \dots, r$ lie in $[1/5, 4/5]$, whereas the first atom $\tau_{0,1}$ of ρ_1 is allowed to be smaller.*

2. The total mass of ρ_0 equals $1/2$, whereas the total mass of ρ_1 is $3/2$.
3. For all $q = 1, \dots, 2r - 1$, the q -th moment of ρ_0 and ρ_1 coincide

$$\int x^q d\rho_0 = \int x^q d\rho_1 = \int_{1/4}^{3/4} x^q dx := m_q .$$

For $j = 0, 1$, we set the values $\gamma_{i,j} = [\xi_{i,j}/\tau_{i,j}]^{1/2}$ and $\alpha_{i,j} = \tau_{i,j}^{-1} - 1$.

Let us give a hint why such a choice leads to what we need. As a consequence of our parameter choices, the following identities are satisfied

$$\sum_{i=1}^r \frac{\gamma_{i,0}^2}{1 + \alpha_{i,0}} + \sigma_0^2 = \sum_{i=0}^r \frac{\gamma_{i,1}^2}{1 + \alpha_{i,1}} + \sigma_1^2 = 2 \quad (63)$$

$$\sum_{i=1}^r \frac{\gamma_{i,0}^2}{(1 + \alpha_{i,0})^q} = \sum_{i=0}^r \frac{\gamma_{i,1}^2}{(1 + \alpha_{i,1})^q} = \frac{3^q - 1}{q4^q} = m_{q-1}, \quad \forall q = 2, \dots, 2r \quad (64)$$

Had the random vectors $(v_{i,j})$ introduced in μ_j formed an orthonormal family, then we would have had $\beta^T \Sigma^q \beta = \sum_i \frac{\gamma_{i,j}^2}{(1 + \alpha_{i,j})^q}$ for any positive integer q . We shall prove later that, under the distribution μ_j , the vectors $v_{i,j}$ have a norm close to one and are almost orthogonal with large probability. Hence, identities (64) imply that the moments $\beta^T \Sigma^q \beta$ concentrate around the same value under μ_0 and μ_1 , this for all $q = 2, \dots, 2r$. This will lead to the fact that the conditional distribution of Y given \mathbf{X} under \mathbf{P}_0 is indistinguishable from that under \mathbf{P}_1 when $n^{1+1/(2r)} = o(p)$ as proved in Lemma 7.5 below. In the same way, (63) will imply that $\beta^T \Sigma \beta + \sigma_j^2$ concentrate around 2 under μ_j for $j = 0, 1$ so that η will concentrate around different values under \mathbf{P}_0 and \mathbf{P}_1 since $\sigma_0^2 \neq \sigma_1^2$. This is stated in Lemma 7.3 below.

Remark. As, with large probability, the random vectors $v_{i,j}$ will be proved to be almost orthonormal, the spectrum of Σ almost lies in $(1/5, 1)$ with high probability under μ_0 . Under μ_1 all the eigenvalues of Σ , except the smallest one, almost lie in $(1/5, 1)$ with high probability, whereas the smallest eigenvalue of Σ , which is of order $1/(1 + \alpha_{0,1})$, will be closer to zero. If we had wanted to restrict ourselves to covariance matrices with uniformly bounded eigenvalues (in say $[M, 1/M]$) as suggested in the discussion below Theorem 4.5, we would have defined the parameters thanks to discrete measures ρ_0 and ρ_1 with support in $[1/M, 1]$. However, to constrain the q -th moment of ρ_0 and ρ_1 to coincide for $q = 1, \dots, 2r - 1$, the difference in total mass between ρ_0 and ρ_1 would now depend on r . The remainder of the proof would be unchanged except that the quantities η_0 and η_1 in (65) would depend on r and the ultimate conclusion would be that $\underline{\lim} \bar{R}^*[p, M] \geq C(r)$.

Let us now define the quantities

$$\eta_0 := 1 - \frac{\sigma_0^2}{\sum_{i=1}^r \frac{\gamma_{i,0}^2}{1 + \alpha_{i,0}} + \sigma_0^2} = 1/4, \quad \eta_1 := 1 - \frac{\sigma_0^2}{\sum_{i=0}^r \frac{\gamma_{i,1}^2}{1 + \alpha_{i,1}} + \sigma_1^2} = 3/4. \quad (65)$$

The next lemma states that, for $j = 0, 1$, $\eta(\beta, \sigma)$ is close to η_j under μ_j .

Lemma 7.3. *There exists three positive constants $C_1(r)$, $C_2(r)$ and $C_3(r)$ such that the following holds. If p is larger than $nC_1(r)$, then*

$$\mu_j \left[|\eta(\beta, \sigma) - \eta_j| \geq C_2(r) \left(p^{-1/4} + \left(\frac{n}{p} \right)^{1/2} \right) \right] \leq e^{-C_3(r)p^{1/2}}, \quad (66)$$

for $j = 0, 1$. Also, the spectrum of Σ is bounded away from zero with large probability, that is for $j = 0, 1$,

$$\mu_j \left[\lambda_{\min}(\Sigma) \geq \frac{1}{2} \min_i \frac{1}{1 + \alpha_{i,j}} \right] \leq e^{-C_3(r)p^{1/2}}. \quad (67)$$

Define $M(r) := [2 \max_{i,j} (1 + \alpha_{i,j})^{-1}]$. By definition of μ_0 and μ_1 , the largest eigenvalue of Σ is always equal to one. By Lemma 7.3, with μ_0 and μ_1 probability going to one, the spectrum of Σ lies in $[1/M(r), M(r)]$.

Let us now bound the minimax risk $\overline{R}^*[p, M(r)]$. Contrary to the prior distributions chosen in the proof of Proposition 3.1, the proportion of explained variation $\eta(\beta, \sigma)$ is not constant either on μ_0 or on μ_1 , so that we cannot directly relate the minimax estimation rate to the total variation distance as done before. Nevertheless, these proportions of explained variation concentrate around η_0 and η_1 so that it will be possible to work around this difficulty. This slight refinement of Le Cam's method has already been applied for other functional estimation problems (see e.g. [12]). Also to circumvent the issue that some eigenvalues of Σ are smaller than $M[r]$ with positive (but very small) probability, we consider a thresholded version of the risk $\mathbf{E}_1^*[\cdot] := \mathbf{E}_1[\cdot \mathbf{1}_{\lambda_{\min}(\Sigma) \geq M^{-1}(r)}]$ and $\mathbf{E}_0^*[\cdot] := \mathbf{E}_0[\cdot \mathbf{1}_{\lambda_{\min}(\Sigma) \geq M^{-1}(r)}]$.

Without loss of generality, we may assume that all the estimators $\hat{\eta}$ below only take values in $[0, 1]$.

$$\begin{aligned} \overline{R}^*[p, M(r)] &\geq \inf_{\hat{\eta}} \mathbf{E}_0^*[\{\hat{\eta} - \eta(\beta, \sigma)\}^2] \vee \mathbf{E}_1^*[\{\hat{\eta} - \eta(\beta, \sigma)\}^2] \\ &\geq \inf_{\hat{\eta}} \mathbf{E}_0[\{\hat{\eta} - \eta(\beta, \sigma)\}^2] \vee \mathbf{E}_1[\{\hat{\eta} - \eta(\beta, \sigma)\}^2] - \bigvee_{i=0,1} \mu_i [\lambda_{\min}(\Sigma) \leq M^{-1}(r)] \\ &\geq \inf_{\hat{\eta}} \frac{1}{2} \bigvee_{i=1,2} \mathbf{E}_i[\{\hat{\eta} - \eta_i\}^2] - \bigvee_{i=1,2} \mathbf{E}_i[\{\eta(\beta, \sigma) - \eta_i\}^2] - \bigvee_{i=0,1} \mu_i [\lambda_{\min}(\Sigma) \leq M^{-1}(r)], \end{aligned}$$

where we used $(x - y)^2 \geq (x - z)^2/2 - (y - z)^2$. From (66) and the fact that $\eta(\beta, \sigma)$ belongs to $[0, 1]$, we derive that

$$\bigvee_{i=1,2} \mathbf{E}_i[\{\eta(\beta, \sigma) - \eta_i\}^2] \leq C(r)(p^{-1/2} + n/p),$$

when p is large enough. Besides, the probabilities $\mu_i[\lambda_{\min}(\Sigma) \leq M^{-1}(r)]$ are smaller than $e^{-Cp^{1/2}}$ by (67). Then, we control the maximum $\bigvee_{i=1,2} \mathbf{E}_i[\{(\hat{\eta} - \eta_i)^2\}]$ using the total variation distance between \mathbf{P}_0 and \mathbf{P}_1 as we did in the proof of Proposition 2.4. More precisely,

$$\begin{aligned} \overline{R}^*[p, M(r)] + C(r)[p^{-1/2} + (n/p)] &\geq \frac{(\eta_1 - \eta_0)^2}{8} \inf_{\hat{\eta}} \mathbf{P}_0(\hat{\eta} \geq \frac{\eta_1 + \eta_0}{2}) \vee \mathbf{P}_1(\hat{\eta} \leq \frac{\eta_1 + \eta_0}{2}) \\ &\geq \frac{(\eta_1 - \eta_0)^2}{16} \inf_{\mathcal{A}} \mathbf{P}_0(\mathcal{A}) + \mathbf{P}_1(\mathcal{A}^c) \\ &\geq \frac{(\eta_1 - \eta_0)^2}{16} [1 - \|\mathbf{P}_1 - \mathbf{P}_0\|_{TV}], \end{aligned}$$

so that we only have to focus on $\|\mathbf{P}_1 - \mathbf{P}_0\|_{TV}$. Let us decompose the total variation distance between \mathbf{P}_0 and \mathbf{P}_1 in a way enabling to consider separately the marginal distribution of \mathbf{X} and the conditional distributions of Y given \mathbf{X} . Since the total variation distance is, up to a multiplicative

constant, the l_1 distance between the density functions, we obtain

$$\begin{aligned}
2\|\mathbf{P}_1 - \mathbf{P}_0\|_{TV} &= \int |f_0(y, \mathbf{x}) - f_1(y, \mathbf{x})| dy d\mathbf{x} \\
&= \int |f_0(y|\mathbf{x})f_0(\mathbf{x}) - f_1(y|\mathbf{x})f_1(\mathbf{x})| dy d\mathbf{x} \\
&\leq \int f_1(y|\mathbf{x})|f_0(\mathbf{x}) - f_1(\mathbf{x})| dy d\mathbf{x} + \int f_0(x)|f_0(y|\mathbf{x}) - f_1(y|\mathbf{x})| dy d\mathbf{x} \\
&\leq \int |f_0(\mathbf{x}) - f_1(\mathbf{x})| d\mathbf{x} + \int f_0(x)|f_0(y|\mathbf{x}) - f_1(y|\mathbf{x})| dy d\mathbf{x} \\
&\leq 2\|\mathbf{P}_0^{\mathbf{X}} - \mathbf{P}_1^{\mathbf{X}}\|_{TV} + 2\mathbf{E}_0^{\mathbf{X}} \left[\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV} \right], \tag{68}
\end{aligned}$$

where, for $i = 0, 1$, $\mathbf{P}_i^{\mathbf{X}}$ (resp. f_i) denotes the marginal probability distribution (resp. density) of \mathbf{X} under \mathbf{P}_i , $\mathbf{P}_i^{Y|\mathbf{X}}$ (resp. $f_i(\cdot|\mathbf{x})$) is the conditional distribution (resp. density) of Y given \mathbf{X} and $\mathbf{E}_0^{\mathbf{X}}$ stands for the expectation with respect to $\mathbf{P}_0^{\mathbf{X}}$. The main difficulty in the proof lies in controlling these two total deviation distances $\|\mathbf{P}_0^{\mathbf{X}} - \mathbf{P}_1^{\mathbf{X}}\|_{TV}$ and $\mathbf{E}_0^{\mathbf{X}}[\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV}]$.

The marginal distribution of \mathbf{X} under \mathbf{P}_0 and \mathbf{P}_1 is that of a n sample of p -dimensional normal distribution whose precision matrix is a rank r perturbation of the identity matrix and whose r principal directions are sampled nearly uniformly. In a high-dimensional setting, such perturbations are indistinguishable from the standard normal distribution as shown in the next lemma.

Lemma 7.4. *There exist two positive constants $C(r)$ and $C'(r)$ only depending on r such that the following holds. If $p \geq C(r)n$, then*

$$\|\mathbf{P}_0^{\mathbf{X}} - \mathbf{P}_1^{\mathbf{X}}\|_{TV} \leq C'(r) \sqrt{\frac{n}{p}}.$$

The intricate construction of μ_0 and μ_1 (and especially the choices of the parameters $\alpha_{i,j}$ and $\gamma_{i,j}$) has been made to force the conditional $\mathbf{P}_0^{Y|\mathbf{X}}$ and $\mathbf{P}_1^{Y|\mathbf{X}}$ to be close to each other. Informally, the fact that the quantities $\beta^T \Sigma^q \beta$ almost coincide under μ_0 and μ_1 , this for all $q = 2, \dots, 2r$, will translate into the total distance $\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV}$ as illustrated by the next lemma.

Lemma 7.5. *There exist two positive constants $C(r)$ and $C'(r)$ only depending on r such that the following holds. If $p \geq C(r)n$, then*

$$\mathbf{E}_0^{\mathbf{X}} \left[\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV} \right] \leq C'(r) \left(\frac{n^{1+1/(2r)}}{p} \right)^r.$$

Under assumption (25), the distance $\|\mathbf{P}_1 - \mathbf{P}_0\|_{TV}$ goes to 0, and the minimax risk $\overline{R}^*[p, M(r)]$ is therefore bounded away from zero:

$$\underline{\lim} \overline{R}^*[p, M(r)] \geq \frac{(\eta_1 - \eta_0)^2}{32}.$$

7.3.2 Proof of the truncated moment problem (Lemma 7.2)

Define $\overline{\rho}$ the uniform measure over the interval $[1/4, 3/4]$. First, we want to construct ρ_0 an r -atomic measure whose support is in $[1/4, 3/4]$ and whose moments up to order $2r - 1$ coincide with those of $\overline{\rho}$. This truncated moment problem has received a lot of attention in the literature. For

instance, Theorem 4.1 (equivalence between (i) and (ii)) in [15] ensures the existence of ρ_0 . Define the Hankel matrix \mathbf{A} of order $r - 1$ and the matrix \mathbf{B} by

$$\mathbf{A}_{i,j} := m_{i+j}, \quad \mathbf{B}_{i,j} := m_{i+j+1}, \quad \forall 0 \leq i, j \leq r - 1 .$$

(Here, $m_0 = \int_{1/4}^{3/4} dx = 1/2$). The same theorem 4.1 in [15] then ensures that the symmetric Hankel matrix \mathbf{A} is positive semidefinite, $\mathbf{A} \geq 0$, and $3/4\mathbf{A} \geq \mathbf{B} \geq 1/4\mathbf{A}$ where $\mathbf{B} \geq 1/4\mathbf{A}$ implies that $\mathbf{B} - 1/4\mathbf{A}$ is positive semidefinite. Since the representation of the truncated moment problem (m_0, \dots, m_{2r-1}) is not unique ($\bar{\rho}$ and ρ_0 are admissible) Theorem 3.8 in [15] ensures that \mathbf{A} is non-singular. Hence, up to modifying the constants, we can obtain strict inequalities in the bounds

$$\mathbf{A} > 0, \quad \frac{4}{5}\mathbf{A} > \mathbf{B} > \frac{1}{5}\mathbf{A} . \quad (69)$$

Given $\varepsilon > 0$, define the modified matrices \mathbf{A}^ε and \mathbf{B}^ε by

$$\mathbf{A}_{i,j}^\varepsilon := \begin{cases} m_0 & \text{if } i = j = 0 \\ m_{i+j-1} - \varepsilon^{i+j} & \text{else} \end{cases}, \quad \mathbf{A}_{i,j}^\varepsilon = m_{i+j} - \varepsilon^{i+j} .$$

Since the set of positive matrices is open, there exists some $\varepsilon_0 > 0$ such that the Hankel matrix $\mathbf{A}^{\varepsilon_0}$ is positive and $\frac{4}{5}\mathbf{A}^{\varepsilon_0} > \mathbf{B}^{\varepsilon_0} > \frac{1}{5}\mathbf{A}^{\varepsilon_0}$. As a consequence of Theorem 4.1 in [15], there exists an r -atomic measure ρ^{ε_0} with support in $[1/5, 4/5]$ whose q -th moment is $m_q - \varepsilon_0^q$ for $q = 1, \dots, 2r - 1$ and m_0 for $q = 0$.

Finally, the measure $\rho_1 := \delta_{\varepsilon_0} + \rho^{\varepsilon_0}$ satisfies all the desired moment conditions, that is $\int d\rho_1 = m_0 + 1$ and $\int x^q d\rho_1 = m_q$ for $q = 1, \dots, 2r - 1$.

7.3.3 Additional lemma

Lemma 7.6. *There exists three positive constants $C_4(r)$, $C_5(r)$ and $C_6(r)$ such that the following holds. Assuming that $p \geq nC_4(r)$, we have, for both $j = 0$ and $j = 1$,*

$$\pi_{\alpha_j} \left[\max_i \|v_i\|_2^2 - 1 \geq C_5(r) \left(\frac{n}{p}\right)^{1/2} + p^{-1/4} \right] \leq e^{-C_6(r)p^{1/2}} . \quad (70)$$

Proof. We only prove the lemma for $j = 0$, since for $j = 1$ the proof is similar. To ease the notation, we simply write v and α for $(v_{i,0})_i$ and $(\alpha_{i,0})_i$ and μ for μ_0 . Recall that the density of $v = (v_1, \dots, v_r)$ is proportional to $e^{-p \sum_i \|v_i\|_2^2 / 2} |\mathbf{I}_p + \sum_i \alpha_i v_i v_i^T|^{-n/2}$. Denote ω the uniform probability measure on the p -dimensional sphere. Let us first change the coordinate system. The density g of $t = (\|v_i\|_2^2)$ and $w = (v_i / \|v_i\|_2)$ satisfies $g(t, w) = \phi(t, w) [\int \phi(x, w') \prod dx_i d\omega(dw'_i)]^{-1}$, where

$$\phi(t, w) := (\prod_i t_i^{-1}) e^{-\frac{p}{2} \sum_{i=1}^r (t_i - 1 - \log(t_i))} |\mathbf{I}_p + \sum_i \alpha_i t_i w_i w_i^T|^{-n/2} .$$

In order to control the density g , we first provide a lower bound on the normalizing constant. For $t \in (1 - (rp)^{-1/2}, 1)^r$, $|\mathbf{I}_p + \sum_i \alpha_i t_i w_i w_i^T| \leq \|\mathbf{I}_p + \sum_i \alpha_i t_i w_i w_i^T\|_{op}^r \leq (1 + \|\alpha\|_\infty r)^r$. As a consequence,

$$\begin{aligned} \int \phi(x, w') \prod_i dx_i d\omega(w'_i) &\geq \left[\int_0^{(rp)^{-1/2}} \frac{1}{1-u} e^{-\frac{p}{2}(-u - \log(1-u))} du \right]^r (1 + \|\alpha\|_\infty r)^{-nr/2} \\ &\geq (rp)^{-r/2} e^{-r^2/4} (1 + \|\alpha\|_\infty r)^{-nr/2} . \end{aligned}$$

Since the determinant $|\mathbf{I}_p + \sum_i \alpha_i t_i w_i w_i^T|$ is always larger than one, we obtain for some constant $C(r)$ depending only on r and some universal constant C

$$\begin{aligned} g(t, w) &\leq (rp)^{r/2} e^{-r^2/4} (1 + \|\alpha\|_\infty r)^{nr/2} (\prod_i t_i^{-1}) e^{-\frac{p}{2} \sum_{i=1}^r (t_i - 1 - \log(t_i))} \\ &\leq C(r) p^{r/2} (1 + \|\alpha\|_\infty r)^{nr/2} \exp \left[-Cp \sum_i \left\{ (t_i - 1)^2 \mathbf{1}_{t_i \in (1/2, 3/2)} + \mathbf{1}_{t_i \leq 1/2} + t \mathbf{1}_{t_i \geq 3/2} \right\} \right]. \end{aligned} \quad (71)$$

If $\|t - 1\|_\infty$ belongs to $((\frac{nr}{2Cp} \log(1 + \|\alpha\|_\infty r))^{1/2} + p^{-1/4}, 1/2)$, then for some constant $C'(r)$ depending only on r

$$g(t, w) \leq C'(r) p^{r/2} \exp(-Cp^{1/2}).$$

If $\|t - 1\|_\infty \geq 1/2$, then $g(t, w) \leq C'(r) p^{r/2} \exp(-C'p\|t - 1\|_\infty)$ for some universal constant C' . Integrating these bounds with respect to w and t , we conclude that for some constant $C_5(r)$ depending only on r and some universal constant C''

$$\pi \left[\max_i \|v_i\|_2^2 - 1 \geq (C_5(r) \frac{n}{p})^{1/2} + p^{-1/4} \right] \leq C'(r) p^{r/2} e^{-C''p^{1/2}},$$

which is smaller than $e^{-C_6(r)p^{1/2}}$ for some constant $C_6(r)$ for p is large compared to r . □

7.3.4 Control of $\|\mathbf{P}_1^{\mathbf{X}} - \mathbf{P}_0^{\mathbf{X}}\|_{TV}$ (Proof of Lemma 7.4)

Define the probability distribution $\bar{\mathbf{P}}^{\mathbf{X}}$, such that, under $\bar{\mathbf{P}}^{\mathbf{X}}$, the entries of \mathbf{X} follow independent standard normal distributions. By the triangular inequality, we have

$$\|\mathbf{P}_0^{\mathbf{X}} - \mathbf{P}_1^{\mathbf{X}}\|_{TV} \leq \|\bar{\mathbf{P}}^{\mathbf{X}} - \mathbf{P}_0^{\mathbf{X}}\|_{TV} + \|\bar{\mathbf{P}}^{\mathbf{X}} - \mathbf{P}_1^{\mathbf{X}}\|_{TV}, \quad (72)$$

We will prove that $\|\bar{\mathbf{P}}^{\mathbf{X}} - \mathbf{P}_0^{\mathbf{X}}\|_{TV}$ is small, the distance $\|\bar{\mathbf{P}}^{\mathbf{X}} - \mathbf{P}_1^{\mathbf{X}}\|_{TV}$ being handled similarly. In order to simplify the notation in the remainder of this proof, we drop the subscript 0 in the vector α_0 and v_0 . Let ω denote the Haar measure on dimension p orthogonal matrices. We work out the marginal density $f_0(\mathbf{X})$ as follows

$$\begin{aligned} f_0(\mathbf{X}) &= \int (2\pi)^{-(np)/2} |\mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T|^{n/2} \exp \left[-tr(\mathbf{X}\mathbf{X}^T)/2 - \sum_{i=1}^r \alpha_i \|\mathbf{X}v_i\|_2^2/2 \right] \pi_\alpha(dv) \\ &= \int_{\mathbb{R}^+} h_0(v, \mathbf{X}) \pi_\alpha(dv), \end{aligned}$$

$$\text{where } h_0(v, \mathbf{X}) := |\mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T|^{n/2} e^{-tr(\mathbf{X}\mathbf{X}^T)/2} \int e^{-\sum_{i=1}^r \alpha_i \|\mathbf{X}\mathbf{O}v_i\|_2^2/2} \omega(d\mathbf{O}),$$

since the values of $|\mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T|$ and $\sum_{i=1}^r \|v_i\|^2$ are rotation invariant.

For fixed $v = (v_1, \dots, v_r)$, $h_0(v, \mathbf{X})$ stands for the density of \mathbf{X} when the corresponding precision matrix of the rows of \mathbf{X} is a rank r perturbation of the identity matrix whose directions are sampled uniformly on the unit sphere. We shall prove below that it is impossible to distinguish this distribution from $\bar{\mathbf{P}}^{\mathbf{X}}$ (i.e. no perturbation) when p is large compared to n . Denote $\bar{f}(\mathbf{X})$ the

density of \mathbf{X} under $\bar{\mathbf{P}}^{\mathbf{X}}$. In the following equations, $\|f_0(\mathbf{X}) - \bar{f}(\mathbf{X})\|_1$ denotes the l_1 distance (in $\mathbb{R}^{n \times p}$) between the densities. Using Fubini's Theorem, we obtain

$$\begin{aligned} 2\|\mathbf{P}_0^{\mathbf{X}} - \bar{\mathbf{P}}^{\mathbf{X}}\|_{TV} &= \|f_0(\mathbf{X}) - \bar{f}(\mathbf{X})\|_1 \\ &\leq \int \|h_0(\mathbf{X}, v) - \bar{f}(\mathbf{X})\|_1 \pi_\alpha(dv), \end{aligned} \quad (73)$$

so that we will bound the l_1 distance $\|h_0(\mathbf{X}, v) - \bar{f}(\mathbf{X})\|_1$ for all $v = (v_1, \dots, v_r)$. Denote $L_v(\mathbf{X})$ the likelihood ratio $h_0(v, \mathbf{X})/\bar{f}(\mathbf{X})$.

$$\begin{aligned} \|h(\mathbf{X}, v) - \bar{f}(\mathbf{X})\|_1 &= \bar{\mathbf{E}}^{\mathbf{X}} [|L_v(\mathbf{X}) - 1|] \\ &\leq \sqrt{\bar{\mathbf{E}}^{\mathbf{X}} [(L_v(\mathbf{X}) - 1)^2]} = \sqrt{\bar{\mathbf{E}}^{\mathbf{X}} [L_v^2(\mathbf{X})]} - 1, \end{aligned} \quad (74)$$

so that we have to compute the second moment of the likelihood $L_v(\mathbf{X})$. As the proof of the following lemma is a bit tedious, it is postponed to the end of the subsection.

Lemma 7.7. *There exist three positive constants $C_7(r)$, $C_8(r)$, $C_9(r)$, only depending on r such that the following holds. Assuming $p \geq nC_7(r)$, we have*

$$\bar{\mathbf{E}}^{\mathbf{X}} [L_v^2(\mathbf{X})] \leq 1 + C_8(r) \frac{n}{p} + 4r \exp(-C_9(r)p),$$

simultaneously for all $v = (v_1, \dots, v_r)$ satisfying $\max_i \|v_i\|_2 \leq 2$.

This lemma, together with (74), gives us a uniform bound of $\|h(\mathbf{X}, v) - \bar{f}(\mathbf{X})\|_1$ over all v satisfying the above condition

$$\|h(\mathbf{X}, v) - \bar{f}(\mathbf{X})\|_1^2 \leq C_8(r) \frac{n}{p} + 4r \exp(-C_9(r)p) \leq C_{10}(r) \frac{n}{p}.$$

Denote \mathcal{V} the collection of v such that $v = (v_1, \dots, v_r)$ satisfying $\max_i \|v_i\|_2 \leq 2$. Coming back to the decomposition (73), we conclude that

$$\begin{aligned} \|\mathbf{P}_0^{\mathbf{X}} - \bar{\mathbf{P}}^{\mathbf{X}}\|_{TV} &\leq \int_{\mathcal{V}} \|h(\mathbf{X}, v) - \bar{f}(\mathbf{X})\|_1 \pi_\alpha(dv) + \int_{\mathcal{V}^c} \|h(\mathbf{X}, v) - \bar{f}(\mathbf{X})\|_1 \pi_\alpha(dv) \\ &\leq C(r) \sqrt{\frac{n}{p}} + \mu_0 \left[\max_{i=1}^r \|v_i\|_2 \geq 2 \max_i |\alpha_i| \right] \leq C'(r) \sqrt{\frac{n}{p}}, \end{aligned}$$

when p is large compared to n using Lemma 7.6. Handling analogously the difference $\|\mathbf{P}_1^{\mathbf{X}} - \bar{\mathbf{P}}^{\mathbf{X}}\|_{TV}$, we conclude that

$$\|\mathbf{P}_1^{\mathbf{X}} - \bar{\mathbf{P}}^{\mathbf{X}}\|_{TV} \leq C'(r) \sqrt{\frac{n}{p}}. \quad (75)$$

Proof of Lemma 7.7. Relying on Fubini identity and the fact that \mathbf{X} follows a normal distribution, we have

$$\begin{aligned} \bar{\mathbf{E}}^{\mathbf{X}} [L_v^2(\mathbf{X})] &= \bar{\mathbf{E}}^{\mathbf{X}} \left[\left| \mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T \right|^n \int e^{-\sum_{i=1}^r \alpha_i (\|\mathbf{X}\mathbf{O}v_i\|_2^2 + \|\mathbf{X}\mathbf{O}'v_i\|_2^2)/2} \omega(d\mathbf{O}) \omega(d\mathbf{O}') \right] \\ &= \int \left[\frac{|\mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T|^2}{|\mathbf{I}_p + \sum_{i=1}^r \alpha_i (\mathbf{O}v_i)(\mathbf{O}v_i)^T + \alpha_i (\mathbf{O}'v_i)(\mathbf{O}'v_i)^T|^2} \right]^{n/2} \omega(d\mathbf{O}) \omega(d\mathbf{O}') \\ &= \int \left[\frac{|\mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T|^2}{|\mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T + \alpha_i (\mathbf{O}v_i)(\mathbf{O}v_i)^T|^2} \right]^{n/2} \omega(d\mathbf{O}). \end{aligned}$$

Diagonalizing the matrix $\mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T$, we find $\tilde{\alpha}_1, \dots, \tilde{\alpha}_r > 0$ and an orthonormal family w_1, \dots, w_r such that $\sum_{i=1}^r \alpha_i v_i v_i^T = \sum_{i=1}^r \tilde{\alpha}_i w_i w_i^T$. Note that $\|\tilde{\alpha}\|_\infty \leq \|\alpha\|_\infty \sum_{i=1}^r \|v_i\|_2^2 \leq 2r \|\alpha\|_\infty$. We arrive at the following representation

$$\bar{\mathbf{E}}^{\mathbf{X}} [L_v^2(\mathbf{X})] = \int \left[\frac{|\mathbf{I}_p + \sum_{i=1}^r \tilde{\alpha}_i w_i w_i^T|^2}{|\mathbf{I}_p + \sum_{i=1}^r \tilde{\alpha}_i w_i w_i^T + \tilde{\alpha}_i (\mathbf{O}w_i)(\mathbf{O}w_i)^T|^2} \right]^{n/2} \omega(d\mathbf{O}) . \quad (76)$$

We see that $\bar{\mathbf{E}}^{\mathbf{X}} [L_v^2(\mathbf{X})]$ expresses as the $n/2$ -moment of a ratio of determinants. In order to ease the notation, we extend the vector $\tilde{\alpha}$ in a $2r$ -dimensional vector by concatenating it with itself. Define the size diagonal matrix \mathbf{D} by $\mathbf{D}_{i,i} = \tilde{\alpha}_i$ for $1 \leq i \leq 2r$. Extend the orthonormal family (w_1, \dots, w_r) into (w_1, \dots, w_{2r}) by a Gram-Schmidt orthonormalization of $(w_1, \dots, w_r, \mathbf{O}w_1, \dots, \mathbf{O}w_r)$. Since the determinant of $\mathbf{I}_p + \sum_{i=1}^r \tilde{\alpha}_i w_i w_i^T + \sum_{i=1}^r \tilde{\alpha}_i (\mathbf{O}w_i)(\mathbf{O}w_i)^T$ is only determined by its restriction of the basis (w_1, \dots, w_{2r}) , we introduce the matrix \mathbf{A} of this linear application into the space spanned by (w_1, \dots, w_{2r}) . We arrive at

$$\bar{\mathbf{E}}^{\mathbf{X}} [L_v^2(\mathbf{X})] = \int \left(\frac{|\mathbf{I}_{2r} + \mathbf{D}|}{|\mathbf{A}|} \right)^{n/2} \omega(d\mathbf{O}) . \quad (77)$$

In order to prove that this quantity is close to one, we shall show that the matrix close \mathbf{A} is close (in entry-wise supremum norm) to $\mathbf{I}_{2r} + \mathbf{D}$. The difference matrix $\mathbf{V} := \mathbf{A} - \mathbf{I}_{2r} - \mathbf{D}$ writes as

$$\mathbf{V}_{l,m} = \sum_{i=1}^r \tilde{\alpha}_i \langle \mathbf{O}w_i, w_l \rangle \langle \mathbf{O}w_i, w_m \rangle - \tilde{\alpha}_l \mathbf{1}_{l=m} \mathbf{1}_{l>r}$$

Given $i = 1, \dots, r$, define the space $S_i = \text{Vect}\{w_1, \dots, w_r, \mathbf{O}w_1, \dots, \mathbf{O}w_{i-1}\}$. By definition of w_l , observe that $\langle w_i, w_l \rangle = 0$ for all $i < l$. As a consequence, for any $l < m$,

$$|\mathbf{V}_{l,m}| \leq \sum_{i=m}^r \tilde{\alpha}_i |\langle \mathbf{O}w_i, w_l \rangle| \leq \sum_{i=1}^r \tilde{\alpha}_i \|\Pi_{S_i} \mathbf{O}w_i\|_2 ,$$

where Π_S denote the orthogonal projection onto the vector space S . The diagonal terms of \mathbf{V} satisfy

$$|\mathbf{V}_{l,l}| \leq \sum_{i=1}^r \tilde{\alpha}_i \|\Pi_{S_i} \mathbf{O}w_i\|_2^2 .$$

For $i = 1, \dots, r$, denote $W_i = \|\Pi_{S_i} \mathbf{O}w_i\|_2$. Define the matrix $\mathbf{V}_2 := (\mathbf{I}_{2r} + \mathbf{D})^{-1/2} \mathbf{V} (\mathbf{I}_{2r} + \mathbf{D})^{-1/2}$. As a consequence of the previous inequalities, we obtain

$$|\text{tr}(\mathbf{V}_2)| \leq 2r^2 \|\tilde{\alpha}\|_\infty \max_i W_i^2 , \quad \|\mathbf{V}_2\|_\infty \leq r \|\tilde{\alpha}\|_\infty \max_i W_i .$$

Assume that $2r \|\mathbf{V}_2\|_\infty \leq 1/2$ so that $\|\mathbf{V}_2\|_{op} \leq 1/2$. Also denote $\lambda_i(\mathbf{V}_2)$ the ordered eigenvalues of \mathbf{V}_2 .

$$\begin{aligned} \log \left[\frac{|\mathbf{A}|}{|\mathbf{I}_{2r} + \mathbf{D}|} \right] &= \log [|\mathbf{I}_{2r} + \mathbf{V}_2|] = \sum_{i=1}^{2r} \log(1 + \lambda_i(\mathbf{V}_2)) \\ &\geq \text{tr}(\mathbf{V}_2) - \sum_i \lambda_i^2(\mathbf{V}_2) \quad (\text{since } \log(1+x) \geq x - x^2 \text{ for } x \in [-1/2, 1/2]) \\ &\geq -6r^4 (\|\tilde{\alpha}\|_\infty^2 \vee 1) \max_i W_i^2 . \end{aligned}$$

Define the event $\mathcal{A} := \{4r^2\|\tilde{\alpha}\|_\infty \max_i W_i < 1\}$, so that, under \mathcal{A} , we have $2r\|\mathbf{V}_2\|_\infty \leq 1/2$. Under \mathcal{A} , we bound $\log[|\mathbf{A}|]$ as above, whereas, under $|\mathcal{A}|^c$, we simply use that $|\mathbf{A}| > 1$. We also write \mathbb{E}^ω for the expectation with respect to the Haar measure ω .

$$\begin{aligned} \overline{\mathbb{E}}^{\mathbf{X}} [L_v^2(\mathbf{X})] &\leq \mathbb{E}^\omega \left[\exp \left(4nr^4 (\|\tilde{\alpha}\|_\infty^2 \vee 1) \max_i W_i^2 \right) \right] + \omega[\mathcal{A}^c] |\mathbf{I}_{2r} + \mathbf{D}|^n \\ &\leq \mathbb{E}^\omega \left[\exp \left(\frac{nC(r)}{p} (p \max_i W_i^2) \right) \right] + \omega[\mathcal{A}^c] \exp(nC'(r)) , \end{aligned} \quad (78)$$

where $C(r)$ and $C'(r)$ only depend on r . We used in the last line that $\|\tilde{\alpha}\|_\infty \leq 2r\|\alpha\|_\infty$ and that the choice of α only depends on r .

In order to work out this quantity, we need to control the deviations of $p \max_i W_i^2$. Recall that (w_1, \dots, w_r) form an orthonormal family. Hence, conditionally to $(\mathbf{O}w_1, \dots, \mathbf{O}w_{i-1})$, $\mathbf{O}w_i$ follows a uniform distribution on the unit sphere intersected with the orthogonal space of $\text{Vect}(\mathbf{O}w_1, \dots, \mathbf{O}w_{i-1})$. As a consequence, W_i^2 follows the same distribution as $\sum_{i=1}^r Z_i^2 / \|Z\|_2^2$ where $Z = (Z_1, \dots, Z_{p-i+1}) \sim \mathcal{N}(0, \mathbf{I}_{p-i+1})$ (since the Gaussian distribution is isotropic). Noting that W_i^2 is always smaller than one 1, we consider for any $t \in (0, p)$

$$\begin{aligned} \omega [pW_i^2 \geq t + 2r] &= \mathbb{P} \left[\frac{p \sum_{i=1}^r Z_i^2}{\|Z\|_2^2} \geq t + 2r \right] \leq \mathbb{P} [\|Z\|_2^2 \leq p/2] + \mathbb{P} \left[\sum_{i=1}^r (Z_i^2 - 1) \geq \frac{t}{2} \right] \\ &\leq e^{-Cp} + e^{-\frac{t}{8} \wedge \frac{t^2}{64r}} \leq 2e^{-C(r)t} , \end{aligned}$$

where we used Lemma A.1 and $r \leq p/8$ in the last line. Taking an union bound and integrating this deviation bound, we derive that

$$\mathbb{E}^\omega \left[\exp \left(\frac{nC(r)}{p} (p \max_i W_i^2) \right) \right] \leq 1 + C'(r) \frac{n}{p} ,$$

where we used that p is large compared to n . We also derive from the above deviation inequality that $\omega[\mathcal{A}^c] \leq 4re^{-C(r)p}$. Together with (78), this concludes the proof. \square

7.3.5 Control of $\mathbb{E}_0^{\mathbf{X}} [\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV}]$ (Proof of Lemma 7.5)

Let us first characterize the conditional distributions $\mathbf{P}_0^{Y|\mathbf{X}}$ and $\mathbf{P}_1^{Y|\mathbf{X}}$.

Lemma 7.8 (Distribution of Y conditionally to \mathbf{X} under \mathbf{P}_0 and \mathbf{P}_1). *Define the matrices*

$$\mathbf{B}_0 = \sum_{i=1}^r \gamma_{i,0}^2 (p\mathbf{I}_p + \alpha_{i,0} \mathbf{X}^T \mathbf{X})^{-1} , \quad \mathbf{B}_1 = \sum_{i=0}^r \gamma_{i,1}^2 (p\mathbf{I}_p + \alpha_{i,1} \mathbf{X}^T \mathbf{X})^{-1} . \quad (79)$$

and for $j = 0, 1$,

$$\mathbf{\Gamma}_j := \frac{1}{\sigma_j^2} \left[\mathbf{I}_n - \frac{1}{\sigma_j^2} \mathbf{X} (\mathbf{B}_j^{-1} + \mathbf{X}^T \mathbf{X} / \sigma_j^2)^{-1} \mathbf{X}^T \right] \quad (80)$$

Under \mathbf{P}_0 (resp. \mathbf{P}_1), Y follows, conditionally to \mathbf{X} , a centered normal distribution with precision matrix $\mathbf{\Gamma}_0$ (resp. $\mathbf{\Gamma}_1$).

The precision matrices $\mathbf{\Gamma}_0$ and $\mathbf{\Gamma}_1$ are both diagonalizable in the same basis that diagonalizes $\mathbf{X}\mathbf{X}^T$. Denoting $\lambda_i, i = 1, \dots, n$, the ordered eigenvalues of $\mathbf{X}\mathbf{X}^T/p$, we define

$$h_j(\lambda_i) = \left(\sigma_j^2 + \sum_l \frac{\gamma_{l,j}^2 \lambda_i}{1 + \alpha_{l,j} \lambda_i} \right)^{-1}, \quad j = 0, 1 \quad (81)$$

the corresponding eigenvalues of $\mathbf{\Gamma}_0$ and $\mathbf{\Gamma}_1$.

Suppose that λ_i lies in $(1/2, 3/2)$ (this occurs with high probability). Since, by (63), $\sigma_j^2 + \sum_l \frac{\gamma_{l,j}^2}{1 + \alpha_{l,j}} = 2$, we have

$$1/4 \leq h_j(\lambda_i) \leq 4 \quad (82)$$

Let us develop the Taylor's expansion of $1/h_j(\lambda_i)$ with respect to $(\lambda_i - 1)$:

$$\frac{1}{h_j(\lambda_i)} = \sigma_j^2 + \sum_l \frac{\gamma_{l,j}^2}{1 + \alpha_{l,j}} + \sum_{m=1}^{+\infty} (\lambda_i - 1)^m (-1)^{m-1} \sum_l \frac{\gamma_{l,j}^2 \alpha_{l,j}^{m-1}}{(1 + \alpha_{l,j})^{m+1}}. \quad (83)$$

By the definition (62) of σ_j^2 and the property (63), the constant term $\sigma_j^2 + \sum_l \frac{\gamma_{l,j}^2}{1 + \alpha_{l,j}}$ equals 2. Now consider any $m \in \{1, \dots, 2r - 1\}$. The rational function $x^{m-1}/(1+x)^{m+1}$ decomposes as a linear combination of $\frac{1}{(1+x)^2}, \dots, \frac{1}{(1+x)^{m+1}}$. As a consequence, the term of order $(\lambda_i - 1)^m$ in (83) is a linear combination of the moments m_q (defined in (64)) for $q = 1, \dots, m - 1$. Thus, the choice of the parameters $\gamma_{l,j}$ and $\alpha_{l,j}$ makes all the terms of the development of $1/h_0(\lambda_i)$ and $1/h_1(\lambda_i)$ coincide up to order $2r - 1$.

$$\begin{aligned} \left| \frac{1}{h_1(\lambda_i)} - \frac{1}{h_0(\lambda_i)} \right| &= \left| \sum_{m=2r}^{\infty} (\lambda_i - 1)^m (-1)^m \left[\sum_l \frac{\gamma_{l,0}^2 \alpha_{l,0}^{m-1}}{(1 + \alpha_{l,0})^{m+1}} - \frac{\gamma_{l,1}^2 \alpha_{l,1}^{m-1}}{(1 + \alpha_{l,1})^{m+1}} \right] \right| \\ &\leq \sum_{m=2r}^{\infty} |\lambda_i - 1|^m \left[\sum_l \frac{\gamma_{l,0}^2}{(1 + \alpha_{l,0})^2} + \sum_i \frac{\gamma_{l,1}^2}{(1 + \alpha_{l,1})^2} \right] \\ &\leq |\lambda_i - 1|^{2r}, \end{aligned}$$

where we used in the last line that $|\lambda_i - 1| \leq 1/2$ and $m_1 = 1/4$ (with all m_1 defined in (64)). Suppose that all the eigenvalues λ_i lie in $(1/2, 3/2)$ so that $\frac{h_0(\lambda_i)}{h_1(\lambda_i)}$ belongs to $(1/16, 16)$ by (82).

$$\begin{aligned} 2\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV}^2 &\leq \mathbb{K}[\mathbf{P}_0^{Y|\mathbf{X}}; \mathbf{P}_1^{Y|\mathbf{X}}] \quad (\text{by Pinsker's inequality}) \\ &= \frac{1}{2} \sum_{i=1}^n \left(\frac{h_1(\lambda_i)}{h_0(\lambda_i)} - 1 - \log \left[\frac{h_1(\lambda_i)}{h_0(\lambda_i)} \right] \right) \\ &\leq C \sum_{i=1}^n h_1^2(\lambda_i) \left(\frac{1}{h_0(\lambda_i)} - \frac{1}{h_1(\lambda_i)} \right)^2 \\ &\leq C' \sum_{i=1}^n (\lambda_i - 1)^{4r} \leq C'n \|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\|_{op}^{4r} \end{aligned}$$

since, for some fixed $C > 0$, $x - 1 - \log(x) \leq C(x - 1)^2$ for $x \in (1/16, 16)$. The total variation distance is always smaller than one, so that

$$\begin{aligned} \mathbf{E}_0^{\mathbf{X}} \left[\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV} \right] &\leq \mathbf{P}_0^{\mathbf{X}} [\|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\|_{op} \geq 1/2] + 2^{-1/2} \mathbf{E}_0^{\mathbf{X}} \left[\sqrt{\mathbb{K}[\mathbf{P}_0^{Y|\mathbf{X}}; \mathbf{P}_1^{Y|\mathbf{X}}]} \mathbf{1}_{\|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\|_{op} \leq 1/2} \right] \\ &\leq \mathbf{P}_0^{\mathbf{X}} [\|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\|_{op} \geq 1/2] + Cn^{1/2} \mathbf{E}_0^{\mathbf{X}} [\|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\|_{op}^{2r}] \end{aligned}$$

Conditionally to $v = (v_1, \dots, v_r)$, \mathbf{X} follows a Gaussian distribution with inverse covariance $\Sigma^{-1}(v) = \mathbf{I}_p + \sum_{i=1}^r \alpha_{i,0} v_i v_i^T$ with $\alpha_{i,0} > 0$. As a consequence, $\|\Sigma(v)\|_{op} = 1$ (since $p > r$) and $tr(\Sigma(v))$ belongs to $[p - r, p]$. As a consequence, we may apply the deviation inequality for Wishart matrices with non-identity covariances (Lemma 6.1) to \mathbf{X} and reintegrate with respect to v .

$$\mathbf{P}_0^{\mathbf{X}} \left[\|\mathbf{X}\mathbf{X}^T - tr(\Sigma(v))\mathbf{I}_n\|_{op} \leq 2\sqrt{p} \left(\sqrt{n} + 10 + \sqrt{2t} \right) + 3[n + 100 + 2t] \right] \geq 1 - 2e^{-t}, \quad (84)$$

for any $t > 0$. Since $n \leq p$, this simplifies in

$$\mathbf{P}_0^{\mathbf{X}} \left[\|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\|_{op} \leq \frac{r}{p} + C \left\{ \sqrt{\frac{n}{p}} + \sqrt{\frac{t}{p}} + \frac{t}{p} \right\} \right] \geq 1 - 2e^{-t}.$$

Integrating this deviation bound, we obtain $\mathbf{E}_0^{\mathbf{X}} [\|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\|_{op}^{2r}] \leq C(r)(n/p)^r$ where the constant $C(r)$ only depends on the integer r . Also, from (84), we derive that the probability that $\|\mathbf{X}\mathbf{X}^T/p - \mathbf{I}_n\| \geq 1/2$ is smaller than $2e^{-C'p}$ for some $C' > 0$. We have proved that

$$\mathbf{E}_0^{\mathbf{X}} \left[\|\mathbf{P}_0^{Y|\mathbf{X}} - \mathbf{P}_1^{Y|\mathbf{X}}\|_{TV} \right] \leq 2e^{-C'p} + C(r) \left(\frac{n^{1+1/(2r)}}{p} \right)^r,$$

where the constant $C(r)$ only depends on $r > 0$.

Proof of Lemma 7.8. We only prove the result for \mathbf{P}_1 , the result for \mathbf{P}_0 being handled similarly. Since \mathbf{P}_1 is a mixture distribution, we introduce $f_1(Y, \mathbf{X}, v_0, \dots, v_r)$ the total density of (Y, \mathbf{X}, v) where $v = (v_0, \dots, v_r)$.

$$f_1(Y, \mathbf{X}, v) \propto \exp \left[-\frac{\|Y - \mathbf{X} \sum_{i=0}^r \gamma_{i,1} v_i\|^2}{2\sigma_1^2} - \frac{tr(\mathbf{X}\mathbf{X}^T)}{2} - \sum_{i=0}^r \frac{\alpha_{i,1}}{2} \|\mathbf{X}v_i\|_2^2 - \sum_{i=0}^r \frac{p}{2} \|v_i\|_2^2 \right]$$

Denote $z = \sum_{i=0}^r \gamma_{i,1} v_i$. If the density of v is proportional to $\prod_i \exp \left[-\sum_{i=0}^r \frac{\alpha_{i,1}}{2} \|\mathbf{X}v_i\|_2^2 - \sum_{i=0}^r \frac{p}{2} \|v_i\|_2^2 \right]$, then z follows a centered normal distribution with covariance matrix $\mathbf{B}_1 = \sum_{i=0}^r \gamma_{i,1}^2 (p\mathbf{I}_p + \alpha_{i,1}\mathbf{X}^T\mathbf{X})^{-1}$. As a consequence, integrating $f_1(Y, \mathbf{X}, v)$ with respect to v leads to the marginal density

$$\begin{aligned} f_1(Y, \mathbf{X}) &\propto \int \frac{1}{(2\pi)^{n/2} |\mathbf{B}_1|^{1/2}} e^{-\|Y - \mathbf{X}z\|^2/(2\sigma_1^2) - z^T \mathbf{B}_1^{-1} z/2} e^{-tr(\mathbf{X}\mathbf{X}^T)/2} dz \\ &\propto e^{-\frac{1}{2\sigma_1^2} Y^T \left(\mathbf{I}_n - \frac{1}{\sigma_1^2} \mathbf{X}(\mathbf{B}_1^{-1} + \mathbf{X}^T\mathbf{X}/\sigma_1^2)^{-1} \mathbf{X}^T \right) Y} e^{-tr(\mathbf{X}\mathbf{X}^T)/2} \left[|\mathbf{B}_1| |\mathbf{B}_1^{-1} + \frac{\mathbf{X}^T\mathbf{X}}{\sigma_1^2}| \right]^{-1/2} \end{aligned}$$

Hence, conditionally to \mathbf{X} , Y follows a centered normal distribution with precision matrix $\mathbf{\Gamma}_1$. \square

7.3.6 Proof of Lemma 7.3

We prove the result for μ_0 , the proof for μ_1 being handled similarly. In order to ease the notation we respectively write α_i , γ_i , v_i and σ for $\alpha_{i,0}$, $\gamma_{i,0}$, $v_{i,0}$ and σ_0 . From (63) and the definition (65), we have the following decomposition

$$\eta_0 = \frac{\sum_{i=1}^r \frac{\gamma_i^2}{1+\alpha_i}}{\sigma^2 + \sum_{i=1}^r \frac{\gamma_i^2}{1+\alpha_i}}.$$

Define $s_0 = \sum_{i=1}^r \gamma_i^2 / (1 + \alpha_i)$ so that $\eta_0 = s_0 / (\sigma^2 + s_0)$. Since σ^2 is fixed to $3/2$, it suffices to prove that $\beta^T \Sigma \beta$ is concentrated around s_0 to obtain a concentration bound for $\eta(\beta, \sigma)$ around η_0 .

We use a similar approach to that of the proof of Lemma 7.7. Conditionally to (v_1, \dots, v_r) , $\Sigma^{-1} = \mathbf{I}_p + \sum_{i=1}^r \alpha_i v_i v_i^T$. Define $w_i := v_i / \|v_i\|_2$ the standardized version of v_i and $t_i := \|v_i\|_2^2$. Also define the Gram-Schmidt orthonormalization basis (ν_1, \dots, ν_r) obtained from (w_1, \dots, w_r) . Define the $r \times r$ matrix $\tilde{\Sigma}$ which represents the restriction of Σ in the orthonormal basis (ν_1, \dots, ν_r) . Similarly, define the r -dimensional vector $\tilde{\beta}$ so that $s_0 = \tilde{\beta}^T \tilde{\Sigma} \tilde{\beta}$. Define the diagonal matrix $\bar{\Sigma}$ by $\bar{\Sigma}_{l,l} = 1 / (1 + \alpha_l t_l)$ for all $l = 1, \dots, r$. and the vector $\bar{\beta}$ by $\bar{\beta}_l = \gamma_l$ for $l = 1, \dots, r$. Then, $\beta^T \Sigma \beta - s_0$ decomposes as

$$\begin{aligned} \beta^T \Sigma \beta - s_0 &= \left(\bar{\beta}^T \bar{\Sigma} \bar{\beta} - s_0 \right) + \bar{\beta}^T \left[\bar{\Sigma} - \tilde{\Sigma} \right] \bar{\beta} + \left(\text{tr} \left[\tilde{\Sigma} \left\{ \tilde{\beta} \tilde{\beta}^T - \bar{\beta} \bar{\beta}^T \right\} \right] \right) \\ &= (I) + (II) + (III). \end{aligned} \quad (85)$$

We shall prove that each of these three terms is small in absolute value. Recall that the α_i and γ_i are positive constants only depending on r .

$$\begin{aligned} |(I)| &= \left| \sum_{i=1}^r \frac{\gamma_i^2}{1 + \alpha_i t_i} - \frac{\gamma_i^2}{1 + \alpha_i} \right| \leq r \|\gamma\|_\infty^2 \|\alpha\|_\infty \max_i |t_i - 1| \\ |(II)| &\leq \|\bar{\beta}\|_2^2 \|\bar{\Sigma} - \tilde{\Sigma}\|_{op} \leq \|\gamma\|_2^2 \|\bar{\Sigma} - \tilde{\Sigma}\|_{op} \\ |(III)| &\leq \|\tilde{\Sigma}\|_F \|\tilde{\beta} \tilde{\beta}^T - \bar{\beta} \bar{\beta}^T\|_F \leq \|\tilde{\Sigma}\|_F (\|\tilde{\beta}\|_2 + \|\bar{\beta}\|_2) \|\tilde{\beta} - \bar{\beta}\|_2 \leq C(r) \|\tilde{\beta} - \bar{\beta}\|_2, \end{aligned}$$

where we used in the last line that the eigenvalues of $\tilde{\Sigma}$ are all smaller than one. Coming back to (85), we have proved that

$$|\beta^T \Sigma \beta - s_0| \leq C(r) \left[\max_i |t_i - 1| + \|\bar{\Sigma} - \tilde{\Sigma}\|_{op} + \|\tilde{\beta} - \bar{\beta}\|_2 \right]. \quad (86)$$

Let us bound $\|\bar{\Sigma} - \tilde{\Sigma}\|_{op}$ and $\|\tilde{\beta} - \bar{\beta}\|_2$ in terms of t and w . By definition of the basis (ν_1, \dots, ν_r) ,

$$\begin{aligned} |[\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}]_{l,m}| &= \left| \sum_{i=1}^r \alpha_i t_i \langle w_i, \nu_l \rangle \langle w_i, \nu_m \rangle - \alpha_l t_l \mathbf{1}_{l=m} \right| \\ &\leq r \|\alpha\|_\infty \|t\|_\infty \max_{i=1, \dots, r} \max_{l \neq i} |\langle w_i, \nu_l \rangle| \\ &\leq r \|\alpha\|_\infty \|t\|_\infty \max_{i=1, \dots, r} W_i, \end{aligned}$$

where $W_i := \|\Pi_{\text{Vect}(w_1, \dots, w_{i-1}, w_{i+1}, \dots, w_r)} w_i\|_2$ and Π_S the orthogonal projection onto the space S . Here we used that for all i, l , $|\langle w_i, \nu_l \rangle| \leq 1$ and that $1 - \langle w_l, \nu_l \rangle^2 = \sum_{i \neq l} |\langle w_i, \nu_l \rangle|^2$.

Assuming that $2r \|\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}\|_\infty \leq 1$, we have $\|\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}\|_{op} \leq 1/2$, so that also $\|\bar{\Sigma}^{1/2} (\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}) \bar{\Sigma}^{1/2}\|_{op} \leq 1/2$ since all the eigenvalues of $\bar{\Sigma}$ are smaller than one. Then

$$\begin{aligned} \|\tilde{\Sigma} - \bar{\Sigma}\|_{op} &= \|\bar{\Sigma}^{1/2} [(\mathbf{I}_r + \bar{\Sigma}^{1/2} (\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}) \bar{\Sigma}^{1/2})^{-1} - \mathbf{I}_r] \bar{\Sigma}^{1/2}\|_{op} \\ &\leq \|(\mathbf{I}_r + \bar{\Sigma}^{1/2} (\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}) \bar{\Sigma}^{1/2})^{-1} - \mathbf{I}_r\|_{op} \\ &\leq \frac{\|\bar{\Sigma}^{1/2} (\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}) \bar{\Sigma}^{1/2}\|_{op}}{1 - \|\bar{\Sigma}^{1/2} (\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}) \bar{\Sigma}^{1/2}\|_{op}} \\ &\leq 2 \|\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}\|_{op} \\ &\leq 2r \|\tilde{\Sigma}^{-1} - \bar{\Sigma}^{-1}\|_\infty, \end{aligned}$$

where we used that all the eigenvalues of $\bar{\Sigma}$ are smaller than one. Turning to the difference $\tilde{\beta} - \bar{\beta}$, we have, for any $l = 1, \dots, r$,

$$|(\tilde{\beta} - \bar{\beta})_l| = \left| \sum_{i=1}^r \gamma_i \langle w_i, \nu_l \rangle - \gamma_l \right| \leq r \|\gamma\|_\infty \max_i W_i .$$

Thus, we obtain $\|\tilde{\beta} - \bar{\beta}\|_2 \leq C(r) \max_i W_i$. Together with (86), this gives us

$$|\beta^T \Sigma \beta - s_0| \leq C(r) \left[\max_i |t_i - 1| + \|t\|_\infty \max_i W_i \right] , \quad (87)$$

as soon as $\|t\|_\infty \max_i W_i \leq 1/2r^2\|\alpha\|_\infty$. The deviations of $\max_i |t_i - 1|$ are given by Lemma 7.6 so that it only remains to control the deviations of W_i . Let \mathcal{A} be any event on $w = (w_1, \dots, w_r)$. From (71), we derive that

$$\mu_0(\mathcal{A}) = \int \mathbf{1}_{\mathcal{A}} g(t, w) \Pi_i dt_i d\omega(w_i) \leq C'(r) p^{r/2} (1 + \|\alpha\|_\infty r)^{nr/2} \int \mathbf{1}_{\mathcal{A}} \Pi_i d\omega(w_i) .$$

As a consequence, the probability $\mu_0(\mathcal{A})$ is always smaller than $C'(r) p^{r/2} (1 + \|\alpha\|_\infty r)^{nr/2}$ than the probability of \mathcal{A} , when w_1, \dots, w_r are independently and uniformly distributed on the unit sphere. When w_1, \dots, w_r are independently and uniformly distributed on the unit sphere, W_i^2 follows the same distribution as $\sum_{i=1}^{r-1} Z_i^2 / \|Z\|_2^2$ where $Z = (Z_1, \dots, Z_p) \sim \mathcal{N}(0, \mathbf{I}_p)$ (since the Gaussian distribution is isotropic). Arguing as in the proof of Lemma 7.7, we derive that for any $t \in (0, p)$

$$\begin{aligned} \omega [pW_i^2 \geq t + 2r] &= \mathbb{P} \left[\frac{p \sum_{i=1}^{r-1} Z_i^2}{\|Z\|_2^2} \geq t + 2r \right] \leq \mathbb{P} [\|Z\|_2^2 \leq p/2] + \mathbb{P} \left[\sum_{i=1}^{r-1} (Z_i^2 - 1) \geq \frac{t}{2} \right] \\ &\leq e^{-p/16} + e^{-\frac{t}{8} \wedge \frac{t^2}{64r}} \leq 2e^{-C(r)t} , \end{aligned}$$

where we used Lemma A.1 in the last line. Taking an union bound, we derive that

$$\mu_0 \left[p \max_i W_i^2 \geq t + 2r \right] \leq C'(r) p^{1+r/2} (1 + \|\alpha\|_\infty r)^{nr/2} e^{-C(r)t} .$$

Thus,

$$\mu_0 \left[\|t\|_\infty \max_i W_i \geq p^{-1/4} \right] \leq \mu_0 [\|t\|_\infty \geq 2] + C'(r) p^{1+r/2} (1 + \|\alpha\|_\infty r)^{nr/2} e^{-C'(r)p^{1/2}} \quad (88)$$

since p is large compared to r . Together with (87) and Lemma 7.6, this gives us

$$\mu_0 \left[|\beta^T \Sigma \beta - s_0| \geq \tilde{C}(r) (p^{-1/4} + (\frac{n}{p})^{1/2}) \right] \leq e^{-\tilde{C}(r)p^{1/2}} .$$

for p large enough compared to n and r . This last deviation inequality easily transfers to that of η_0 .

Let us now turn to the spectrum of Σ . By definition of $\tilde{\Sigma}$,

$$\lambda_{\min}(\Sigma) = \lambda_{\min}(\tilde{\Sigma}) \geq \lambda_{\min}(\bar{\Sigma}) - \|\tilde{\Sigma} - \bar{\Sigma}\|_{op} \geq \min \frac{1}{1 + \alpha_i} - \|\tilde{\Sigma} - \bar{\Sigma}\|_{op}$$

so that $\lambda_{\min}(\Sigma) \geq \frac{1}{2} \min \frac{1}{1 + \alpha_i}$ as soon as $2r^2\|\alpha\|_\infty\|t\|_\infty \max_{i=1, \dots, r} W_i \leq \frac{1}{2} \min \frac{1}{1 + \alpha_i}$ and the end of the lemma follows easily from (88) and Lemma 7.6.

7.4 Proof of Proposition 5.1

As in the previous minimax lower bounds, we use Le Cam's approach and build two mixture measures. Denote $\underline{\mathbb{P}}_0$ the distribution of Y when $\beta^* = 0$ and $\sigma = 1$. Under $\underline{\mathbb{P}}_0$, Y follows a standard normal distribution and $\eta[0, 1, \mathbf{X}] = 0$. Given μ a continuous prior measure on \mathbb{R}^p , we take

$$\underline{\mathbf{P}}_\mu = \int_{\mathcal{B}} \mathbb{P}_{\beta,0} \mu(d\beta)$$

Note that μ -almost surely, $\eta[\beta, 0, \mathbf{X}] = 1$. Recall that $\lambda_i^{1/2}$, $i = 1, \dots, n$ denote the singular values of \mathbf{X} and v_i , $i = 1, \dots, n$ its right eigenvectors. Let us choose μ such that, under μ , $(\beta^T v_i) \lambda_i^{1/2}$ follow independent standard normal distributions. Obviously, under $\underline{\mathbf{P}}_\mu$, Y also follows a standard normal distribution, that is $\underline{\mathbf{P}}_\mu = \underline{\mathbb{P}}_0$.

Consider any estimator $\hat{\eta}$. Then,

$$\begin{aligned} \sup_{\beta \in \mathbb{R}^p, \sigma \geq 0} \mathbb{E}_{\beta^*, \sigma} [(\hat{\eta} - \eta[\beta^*, \sigma, \mathbf{X}])^2] &\geq \mathbb{E}_{0,1} [\hat{\eta}^2] \sqrt{\vee_{\beta \in \mathbb{R}^p} \mathbb{E}_{\beta,0} [(\hat{\eta} - 1)]} \\ &\geq \mathbb{E}_0 [\hat{\eta}^2] \sqrt{\mathbb{E}_\mu [(\hat{\eta} - 1)^2]} \\ &\geq \mathbb{E}_0 [\hat{\eta}^2] \sqrt{\mathbb{E}_0 [\hat{\eta}^2 - 1]} \quad (\text{since } \underline{\mathbb{P}}_0 = \underline{\mathbf{P}}_\mu) \\ &\geq \frac{1}{2} [\mathbb{E}_0 [\hat{\eta}^2] + \mathbb{E}_0 [(\hat{\eta} - 1)^2]] \\ &\geq \frac{1}{2} + \mathbb{E}_0 [\hat{\eta}]^2 - \mathbb{E}_0 [\hat{\eta}] \quad (\text{by Cauchy-Schwarz inequality}) \\ &\geq 1/4. \end{aligned}$$

A Auxiliary lemmas

Lemma A.1 (χ^2 distributions [28]). *Let Z stands for a standard Gaussian vector of size k and let \mathbf{A} be a symmetric matrix of size k . For any $t > 0$,*

$$\mathbb{P} \left[Z^T \mathbf{A} Z \geq \text{tr}(\mathbf{A}) + 2\|\mathbf{A}\|_F \sqrt{t} + 2\|\mathbf{A}\|_{\text{opt}} t \right] \leq e^{-t}.$$

When \mathbf{A} is the identity matrix, the above bound simplifies as

$$\mathbb{P} \left[\chi^2(k) \geq k + 2\sqrt{kt} + 2t \right] \leq e^{-t},$$

where $\chi^2(k)$ stand for a χ^2 -distributed random variable with k degrees of freedom. We also have

$$\mathbb{P} \left[\chi^2(k) \leq k - 2\sqrt{kt} \right] \leq e^{-t},$$

for any $t > 0$.

Laurent and Massart [28] have only stated a specific version of Lemma A.1 for positive matrices \mathbf{A} , but their argument straightforwardly extend to general symmetric matrices \mathbf{A} .

Lemma A.2 (Wishart distributions [16]). *Let \mathbf{Z} be a $n \times d$ matrix whose entries follow independent standard normal distributions. For any positive number x ,*

$$\begin{aligned} \mathbb{P} \left[\lambda_{\max} [\mathbf{Z}^T \mathbf{Z}] \geq n \left(1 + \sqrt{d/n} + \sqrt{2x/n} \right)^2 \right] &\leq \exp(-x) \\ \mathbb{P} \left[\lambda_{\min} [\mathbf{Z}^T \mathbf{Z}] \leq n \left(1 - \sqrt{d/n} - \sqrt{2x/n} \right)^2 \right] &\leq \exp(-x) \end{aligned}$$

B Proof of Lemma 6.1

Recall that $\mathbf{X}\mathbf{X}^T$ a weighted sum of Wishart matrices with parameters $(1, n)$

$$\mathbf{X}\mathbf{X}^T = \mathbf{Z}\mathbf{\Gamma}\mathbf{Z}^T = \sum_{j=1}^p \mathbf{\Gamma}_{jj} (\mathbf{z}_{\bullet j} \mathbf{z}_{\bullet j}^T)$$

Define the matrix \mathbf{U} by

$$\mathbf{U} := \mathbf{\Gamma}^{1/2} \mathbf{Z}^T \quad (89)$$

The singular values of \mathbf{U} are the same as those of \mathbf{X}^T . Denote $s_1(\mathbf{U}) \geq s_2(\mathbf{U}) \geq \dots \geq s_n(\mathbf{U})$ the ordered singular values of \mathbf{U} . From the previous remark, the following decomposition holds

$$\lambda_{\max}(\mathbf{A}) = s_1^2(\mathbf{U}) - \text{tr}(\mathbf{\Sigma}) \quad \text{and} \quad \lambda_{\min}(\mathbf{A}) = s_n^2(\mathbf{U}) - \text{tr}(\mathbf{\Sigma}) .$$

Hence, it will suffice to derive deviation inequalities for both $s_1(\mathbf{U})$ and $s_n(\mathbf{U})$ to get the result (29).

Denote \mathbb{S}^{p-1} the p dimensional unit sphere. Since $s_1(\mathbf{U}) = \sup_{x \in \mathbb{S}^{p-1}} \|\mathbf{U}x\|_2$ and $s_n(\mathbf{U}) = \inf_{x \in \mathbb{S}^{p-1}} \|\mathbf{U}x\|_2$, both $s_1(\mathbf{U})$ and $s_n(\mathbf{U})$ are Lipschitz (with respect to the Frobenius norm) functions with constant 1 of the entries of U . As a consequence, $s_1(\mathbf{U})$ and $s_n(\mathbf{U})$ are Lipschitz functions with constant $\max_i(\mathbf{\Gamma}_{i,i}^{1/2}) = \|\mathbf{\Sigma}\|_{op}^{1/2}$ of the entries of \mathbf{Z} . Applying the Gaussian isoperimetric inequality [9], it follows that

$$\begin{aligned} \mathbb{P} \left[s_1(\mathbf{U}) \geq \mathbb{E}[s_1(\mathbf{U})] + \|\mathbf{\Sigma}\|_{op}^{1/2} \sqrt{2t} \right] &\leq \exp[-t] \\ \mathbb{P} \left[s_n(\mathbf{U}) \leq \mathbb{E}[s_n(\mathbf{U})] - \|\mathbf{\Sigma}\|_{op}^{1/2} \sqrt{2t} \right] &\leq \exp[-t] . \end{aligned} \quad (90)$$

In order to control $\mathbb{E}[s_1(\mathbf{U})]$ and $\mathbb{E}[s_n(\mathbf{U})]$, we apply Gordon-Slepian lemma following the approach of Davidson and Szarek [16, Appendix IIc]. First, recall Gordon's extension of Slepian lemma.

Lemma B.1 (Gordon-Slepian lemma). *Let $(X_t)_{t \in T}$ and $(Y_t)_{t \in T}$ be two finite families of jointly Gaussian mean zero random variables such that $\text{Var}(X_t - X_{t'}) \leq \text{Var}(Y_t - Y_{t'})$ for all $t, t' \in T$. Then $\mathbb{E}[\max_{t \in T} X_t] \leq \mathbb{E}[\max_{t \in T} Y_t]$. Similarly, if $T = \cup_{s \in S} T_s$ and*

$$\text{Var}(X_t - X_{t'}) \leq \text{Var}(Y_t - Y_{t'}) \quad \text{if } t \in T_s, t' \in T_{s'} \text{ with } s \neq s' \quad (91)$$

$$\text{Var}(X_t - X_{t'}) \geq \text{Var}(Y_t - Y_{t'}) \quad \text{if } t, t' \in T_s \text{ for some } s \quad (92)$$

one has $\mathbb{E}[\max_{s \in S} \min_{t \in T_s} X_t] \leq \mathbb{E}[\max_{s \in S} \min_{t \in T_s} Y_t]$

Define the Gaussian process $P_{(u,v)}$ indexed by $(u, v) \in \mathbb{S}^{p-1} \times \mathbb{S}^{n-1}$,

$$P_{(u,v)} := \langle u, \mathbf{U}v \rangle = \text{tr}(\mathbf{Z}(\mathbf{\Gamma}^{1/2} u v^T))$$

For any (u, v) and (u', v') , this process satisfies

$$\begin{aligned} \text{Var}(P_{(u,v)} - P_{(u',v')}) &= \|(\mathbf{\Gamma}^{1/2} u)^T v - (\mathbf{\Gamma}^{1/2} u')^T v'\|_F^2 \\ &= \|\mathbf{\Gamma}^{1/2}(u - u')\|_2^2 + u^T \mathbf{\Gamma} u' \|v - v'\|_2^2 \\ &\leq \|\mathbf{\Gamma}^{1/2}(u - u')\|_2^2 + \|\mathbf{\Sigma}\|_{op} \|v - v'\|_2^2 . \end{aligned}$$

Let Z_1 and Z_2 be two independent standard Gaussian vectors of respective size p and n . For any $u \in \mathbb{R}^p$ and any $v \in \mathbb{R}^n$, define

$$Q_{(u,v)} := (\mathbf{\Gamma}^{1/2}u)^T Z_1 + \|\boldsymbol{\Sigma}\|_{op}^{1/2} v^T Z_2$$

Hence,

$$\text{Var} (Q_{(u,v)} - Q_{(u',v')}) = \|\mathbf{\Gamma}^{1/2}(u - u')\|_2^2 + \|\boldsymbol{\Sigma}\|_{op} \|v - v'\|_2^2 .$$

We are therefore in position to apply Slepian lemma to the processes $P_{u,v}$ and $Q_{u,v}$ (although the set $\mathbb{S}^{p-1} \times \mathbb{S}^{n-1}$ is not finite, the result is still true). Observe that $\max_{(u,v) \in \mathbb{S}^{p-1} \times \mathbb{S}^{n-1}} P_{u,v} = s_1(\mathbf{U})$. It follows that

$$\begin{aligned} \mathbb{E} [s_1(\mathbf{U})] &\leq \mathbb{E} \left[\max_{(u,v) \in \mathbb{S}^{p-1} \times \mathbb{S}^{n-1}} Q_{(u,v)} \right] = \mathbb{E} \left[\|\mathbf{\Gamma}^{1/2} Z_1\|_2 \right] + \|\boldsymbol{\Sigma}\|_{op}^{1/2} \mathbb{E} [\|Z_2\|_2] \\ &\leq \sqrt{\text{tr}(\mathbf{\Gamma})} + \|\boldsymbol{\Sigma}\|_{op}^{1/2} \sqrt{n} , \end{aligned} \quad (93)$$

by Cauchy-Schwarz inequality.

For any $v \in \mathbb{S}^{n-1}$, define $T_v := \{(u, v), u \in \mathbb{S}^{p-1}\}$. Hypothesis (91) is still satisfied for $P_{(u,v)}$ and $Q_{(u,v)}$. For (u, v) and $(u', v) \in T_v$,

$$\text{Var} (P_{(u,v)}(\mathbf{Z}) - P_{(u',v)}(\mathbf{Z})) = \|\mathbf{\Gamma}^{1/2}(u - u')\|_2^2 = \text{Var} (Q_{(u,v)}(\mathbf{Z}) - Q_{(u',v)}(\mathbf{Z})) ,$$

and Hypothesis (92) is also satisfied. Applying Gordon-Slepian lemma, we obtain

$$\begin{aligned} -\mathbb{E} [s_n(\mathbf{U})] &= \mathbb{E} \left[\max_{v \in \mathbb{S}^{n-1}} \min_{u \in T_v} P_{(u,v)} \right] \\ &\leq \mathbb{E} \left[\max_{v \in \mathbb{S}^{n-1}} \min_{u \in T_v} Q_{(u,v)} \right] \\ &\leq \mathbb{E} \left[\|\boldsymbol{\Sigma}\|_{op}^{1/2} \|Z_2\|_2 - \|\mathbf{\Gamma}^{1/2} Z_1\|_2 \right] \end{aligned} \quad (94)$$

By Cauchy-Schwarz inequality $\mathbb{E} [\|Z_2\|_2] \leq \sqrt{n}$. It remains to lower bound $\mathbb{E} [\|\mathbf{\Gamma}^{1/2} Z_1\|_2]$. Denote $V = \|\mathbf{\Gamma}^{1/2} Z_1\|_2$. By isoperimetric Gaussian inequality $\mathbb{P} [V \geq \mathbb{E}[V] + \|\boldsymbol{\Sigma}\|_{op}^{1/2} \sqrt{2t}] \leq e^{-t}$ for any $t > 0$. Squaring the above inequality, it follows that for any $t > 0$,

$$\mathbb{P} \left[\frac{V^2 - \mathbb{E}^2[V]}{2\|\boldsymbol{\Sigma}\|_{op}} \geq t \right] \leq \exp \left[-\frac{\|\boldsymbol{\Sigma}\|_{op} t^2}{2\mathbb{E}^2[V]} \wedge t \right] .$$

Integrating this bound with respect to $t > 0$, we obtain

$$\frac{\mathbb{E}[V^2] - \mathbb{E}^2[V]}{2\|\boldsymbol{\Sigma}\|_{op}} \leq \int_0^\infty e^{-\frac{\|\boldsymbol{\Sigma}\|_{op} t^2}{4\mathbb{E}^2[V]}} dt + \int_0^\infty e^{-t} dt \leq \sqrt{\frac{\pi \mathbb{E}^2[V]}{2\|\boldsymbol{\Sigma}\|_{op}}} + 1 \leq \sqrt{\frac{\pi \mathbb{E}[V^2]}{2\|\boldsymbol{\Sigma}\|_{op}}} + 1$$

by Cauchy-Schwarz inequality, which implies

$$\begin{aligned} \mathbb{E}[V] &\geq \sqrt{\left(\mathbb{E}[V^2] - \sqrt{8\pi\|\boldsymbol{\Sigma}\|_{op}} \mathbb{E}[V^2] - 4\|\boldsymbol{\Sigma}\|_{op} \right)_+} \\ &\geq \sqrt{\mathbb{E}[V^2]} \left[1 - \sqrt{\frac{8\pi\|\boldsymbol{\Sigma}\|_{op}}{\mathbb{E}[V^2]}} - 4\frac{\|\boldsymbol{\Sigma}\|_{op}}{\mathbb{E}[V^2]} \right] \\ &\geq \sqrt{\mathbb{E}[V^2]} - \sqrt{8\pi\|\boldsymbol{\Sigma}\|_{op}} - 4\|\boldsymbol{\Sigma}\|_{op} / \sqrt{\mathbb{E}[V^2]} , \end{aligned}$$

where we used $\sqrt{1-x} \geq 1-x$ for all $x \in (0, 1)$ in the second line. Since $\mathbb{E}[V^2] = \text{tr}(\Sigma)$, we conclude that

$$\mathbb{E}[V] \geq \sqrt{\text{tr}(\Sigma)} - \|\Sigma\|_{op}^{1/2} \pi^{1/2} 8^{1/2} - 4 \frac{\|\Sigma\|_{op}}{\sqrt{\text{tr}(\Sigma)}} \geq \sqrt{\text{tr}(\Sigma)} - \|\Sigma\|_{op}^{1/2} (\sqrt{8\pi} + 4).$$

Gathering this bound together with (90), (93), and (94), we obtain

$$\begin{aligned} \mathbb{P} \left[s_1(\mathbf{U}) \geq \sqrt{\text{tr}(\Sigma)} + \|\Sigma\|_{op}^{1/2} \sqrt{n} + \|\Sigma\|_{op}^{1/2} \sqrt{2t} \right] &\leq e^{-t} \\ \mathbb{P} \left[s_n(\mathbf{U}) \leq \sqrt{\text{tr}(\Sigma)} - \|\Sigma\|_{op}^{1/2} \sqrt{n} - \|\Sigma\|_{op}^{1/2} (10 + \sqrt{2t}) \right] &\leq e^{-t} \end{aligned}$$

Recalling that $s_1^2(\mathbf{U}) = \lambda_{\max}(\mathbf{X}\mathbf{X}^T)$ and $s_n^2(\mathbf{U}) = \lambda_{\min}(\mathbf{X}\mathbf{X}^T)$ concludes the proof.

References

- [1] Radosław Adamczak and Paweł Wolff. Concentration inequalities for non-Lipschitz functions with bounded derivatives of higher order. *Probab. Theory Related Fields*, 162(3-4):531–586, 2015.
- [2] D.G. Amaral, C.M. Schumann, and C.W. Nordahl. Neuroanatomy of autism. *Trends in neurosciences*, 31(3):137–45, 2008.
- [3] E. Arias-Castro, E. Candes, and Y. Plan. Global testing under sparse alternatives: Anova, multiple comparisons and the higher criticism. *Annals of Statistics*, 39:2533–2556, 2011.
- [4] Yannick Baraud. Non-asymptotic minimax rates of testing in signal detection. *Bernoulli*, 8(5):577–606, 2002.
- [5] Mohsen Bayati, Murat A Erdogdu, and Andrea Montanari. Estimating lasso risk and noise level. In *Advances in Neural Information Processing Systems*, pages 944–952, 2013.
- [6] A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.
- [7] P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.*, 37(4):1705–1732, 2009.
- [8] Anna Bonnet, Elisabeth Gassiat, and Céline Lévy-Leduc. Heritability estimation in high dimensional sparse linear mixed models. *Electron. J. Stat.*, 9(2):2099–2129, 2015.
- [9] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013. A nonasymptotic theory of independence, With a foreword by Michel Ledoux.
- [10] T Tony Cai and Zijian Guo. Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity. *arXiv preprint arXiv:1506.05539*, 2015.
- [11] T. Tony Cai and Mark G. Low. Nonquadratic estimators of a quadratic functional. *Ann. Statist.*, 33(6):2930–2956, 2005.
- [12] T. Tony Cai and Mark G. Low. Testing composite hypotheses, Hermite polynomials and optimal estimation of a nonsmooth functional. *Ann. Statist.*, 39(2):1012–1041, 2011.
- [13] Tony Cai, Weidong Liu, and Xi Luo. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *J. Amer. Statist. Assoc.*, 106(494):594–607, 2011.
- [14] Olivier Collier, Laëticia Comminges, and Alexandre B Tsybakov. Minimax estimation of linear and quadratic functionals on sparsity classes. *arXiv preprint arXiv:1502.00665*, 2015.
- [15] Raúl E. Curto and Lawrence A. Fialkow. Recursiveness, positivity, and truncated moment problems. *Houston J. Math.*, 17(4):603–635, 1991.

- [16] K. R. Davidson and S. J. Szarek. Local operator theory, random matrices and Banach spaces. In *Handbook of the geometry of Banach spaces, Vol. I*, pages 317–366. North-Holland, Amsterdam, 2001.
- [17] Lee H. Dicker. Variance estimation in high-dimensional linear models. *Biometrika*, 101(2):269–284, 2014.
- [18] Lee H Dicker and Murat A Erdogdu. Maximum likelihood for variance estimation in high-dimensional linear models. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 159–167, 2016.
- [19] David L. Donoho and Michael Nussbaum. Minimax quadratic estimation of a quadratic functional. *J. Complexity*, 6(3):290–323, 1990.
- [20] Jianqing Fan, Shaojun Guo, and Ning Hao. Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 74(1):37–65, 2012.
- [21] D. Goldstein. Common genetic variation and human traits. *New England Journal of Medicine*, (360):1696–1698, 2009.
- [22] Y. Ingster, A. Tsybakov, and N. Verzelen. Detection boundary in sparse regression. *Electronic Journal of Statistics*, 4:1476–1526, 2010.
- [23] Yu. I. Ingster and I. A. Suslina. *Nonparametric goodness-of-fit testing under Gaussian models*, volume 169 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 2003.
- [24] Lucas Janson, Rina Foygel Barber, and Emmanuel Candès. Eigenprism: Inference for high-dimensional signal-to-noise ratios. *arXiv preprint arXiv:1505.02097*, 2015.
- [25] Adel Javanmard and Andrea Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *J. Mach. Learn. Res.*, 15:2869–2909, 2014.
- [26] Adel Javanmard and Andrea Montanari. De-biasing the lasso: Optimal sample size for gaussian designs. *arXiv preprint arXiv:1508.02757*, 2015.
- [27] V. Koltchinski, K. Lounici, and A. Tsybakov. Nuclear norm penalization and optimal rates for noisy low rank matrix completion. *Annals of Statistics*, 39(5):2302–2329, 2011.
- [28] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28(5):1302–1338, 2000.
- [29] Brendan Maher. Personal genomes: The case of the missing heritability. *Nature*, 456(7218):1821, Nov. 6 2008.
- [30] Richard Nickl and Sara van de Geer. Confidence sets in sparse regression. *Ann. Statist.*, 41(6):2852–2876, 2013.
- [31] G. Raskutti, M. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- [32] R.G. Steen, C. Mull, R. McClure, R.M. Hamer, and J.A. Lieberman. Brain volume in first-episode schizophrenia: systematic review and meta-analysis of magnetic resonance imaging studies. *Br J Psychiatry*, (188):510–8, 2006.
- [33] J.L. Stein, S.E. Medland, A.A. Vasquez, D.P. Hibar, R.E. Senstad, M.A. Winkler, R. Toro, K. Appel, R. Bartecek, O. Bergmann, M. Bernard, A.A. Brown, D.M. Cannon, M.M. Chakravarty, A. Christoforou, M. Domin, O. Grimm, M. Hollinshead, A.J. Holmes, G. Homuth, J.-J. Hottenga, C. Langan, L.M. Lopez, N.K. Hansell, K.S. Hwang, and et al. Identification of common variants associated with human hippocampal and intracranial volumes. *Nature Genetics*, (44(5)):552–61, 2012.
- [34] Tingni Sun and Cun-Hui Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- [35] Roberto Toro, Jean-Baptiste Poline, Guillaume Hugué, Eva Loth, Vincent Frouin, Tobias Banaschewski, Gareth J Barker, Arun Bokde, Christian Büchel, Fabiana Carvalho, Patricia Conrod, Mira Fauth-Bühler, Herta Flor, Jürgen Gallinat, Hugh Garavan, Penny Gowloan,

Andreas Heinz, Bernd Ittermann, Claire Lawrence, Hervé Lemaître, Karl Mann, Frauke Nees, Tomá Paus, Zdenka Pausova, Marcella Rietschel, Trevor Robbins, Michael Smolka, Andreas Ströhle, Gunter Schumann, and Thomas Bourgeron. Genomic architecture of human neuroanatomical diversity. *Molecular Psychiatry*, 2014.

- [36] J. A. Tropp. An introduction to matrix concentration inequalities. *ArXiv e-prints*, Jan 2015.
- [37] S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.*, 3:1360–1392, 2009.
- [38] Sara van de Geer, Peter Bühlmann, Ya’acov Ritov, and Ruben Dezeure. On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Statist.*, 42(3):1166–1202, 2014.
- [39] N. Verzelen and F. Villers. Goodness-of-fit tests for high-dimensional gaussian linear models. *Annals of Statistics*, 38:704–752, 2010.
- [40] Nicolas Verzelen and Elisabeth Gassiat. Adaptive estimation of high-dimensional signal-to-noise ratios (version 1). *arXiv preprint arXiv:1602.08006v1*, 2016.
- [41] Cun-Hui Zhang and Stephanie S. Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 76(1):217–242, 2014.