

L_1 -Regularized Least Squares for Support Recovery of High Dimensional Single Index Models with Gaussian Designs

Matey Neykov

*Department of Operations Research and Financial Engineering
Princeton University
Princeton, NJ 08544, USA*

MNEYKOV@PRINCETON.EDU

Jun S. Liu

*Department of Statistics
Harvard University
Cambridge, MA 02138, USA*

JLIU@STAT.HARVARD.EDU

Tianxi Cai

*Department of Biostatistics
Harvard University
Boston, MA 02115, USA*

TCAI@HSPH.HARVARD.EDU

Editor: Xiaotong Shen

Abstract

It is known that for a certain class of single index models (SIMs) $Y = f(\mathbf{X}_{p \times 1}^\top \boldsymbol{\beta}_0, \varepsilon)$, support recovery is impossible when $\mathbf{X} \sim \mathcal{N}(0, \mathbb{I}_{p \times p})$ and a *model complexity adjusted sample size* is below a critical threshold. Recently, optimal algorithms based on Sliced Inverse Regression (SIR) were suggested. These algorithms work provably under the assumption that the design \mathbf{X} comes from an i.i.d. Gaussian distribution. In the present paper we analyze algorithms based on covariance screening and least squares with L_1 penalization (i.e. LASSO) and demonstrate that they can also enjoy optimal (up to a scalar) rescaled sample size in terms of support recovery, albeit under slightly different assumptions on f and ε compared to the SIR based algorithms. Furthermore, we show more generally, that LASSO succeeds in recovering the signed support of $\boldsymbol{\beta}_0$ if $\mathbf{X} \sim \mathcal{N}(0, \boldsymbol{\Sigma})$, and the covariance $\boldsymbol{\Sigma}$ satisfies the irrepresentable condition. Our work extends existing results on the support recovery of LASSO for the linear model, to a more general class of SIMs.

Keywords: Single index models, Sparsity, Support recovery, High-dimensional statistics, LASSO

1. Introduction

Modern data applications often require scientists to deal with high-dimensional problems in which the sample size n could be much less than the dimensionality of the covariates p . To handle such challenging problems, structural assumptions on the data generating mechanism are often imposed. Such assumptions are motivated by the fact that classical procedures such as linear regression provably fail, unless the ratio p/n converges to 0. However, modern procedures based on regularization may work well under the high dimensional setting with additional sparsity assumptions. In the sparse high dimensional setting, it is

often of interest to uncover the sparsity pattern, or in other words to select the relevant variables for the model. Under generalized linear models, support recovery can be achieved by fitting these models with penalized optimization procedures such as LASSO (Tibshirani, 1996) or Dantzig Selector (Candes and Tao, 2007), which are computationally inexpensive compared to exhaustive search approaches. The LASSO algorithm’s variable selection/support recovery capabilities under generalized linear models, have been extensively studied (Meinshausen and Bühlmann, 2006; Zhao and Yu, 2006; Wainwright, 2009; Lee et al., 2013, e.g. among others). However, much less is known under potential mis-specification of these commonly used models or under more general models.

In this paper, we focus on recovering the support of the regression coefficients β_0 under a *single index model* (SIM):

$$Y = f(\mathbf{X}^\top \beta_0, \varepsilon), \quad \varepsilon \perp \mathbf{X} \quad (1)$$

where both the link function f and the distribution of ε are left unspecified. Throughout, we assume that $\beta_0^\top \Sigma \beta_0 = 1$ for identifiability and $E(\mathbf{X}) = 0$, where $\Sigma = E(\mathbf{X} \mathbf{X}^\top)$. We are specifically interested in the case where $\mathbf{X} \in \mathbb{R}^p \sim \mathcal{N}(0, \Sigma)$ and β_0 is s -sparse with $s < p$. Obviously, the SIM includes many commonly used parametric or semi-parametric regression models such as the linear regression model as special cases. Under (1) and the sparsity assumption, we aim to show that the standard least squares LASSO algorithm,

$$\hat{\beta} = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{2n} \sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \beta)^2 + \lambda \|\beta\|_1 \right\}, \quad (2)$$

can successfully recover the support of β_0 provided standard regularity conditions and that the *model complexity adjusted effective sample size*,

$$n_{p,s} = n / \{s \log(p - s)\},$$

is sufficiently large. Obviously, for most choices of f , fitting (2) is essentially making inference under the mis-specified linear regression model.

The least squares LASSO algorithm has been used frequently in practice to perform variable selection for analyzing genomic data (Cantor et al., 2010; Zhao et al., 2011; Wang et al., 2013, e.g.). However, the linear model with Gaussian error is unlikely to be the true model in many such cases. Hence it is of practical importance to theoretically establish that the LASSO’s support recovery capabilities are in fact robust to mis-specification. In addition to arguing that LASSO is robust, and perhaps even more surprisingly, we demonstrate that fitting the mis-specified linear model with LASSO penalty can optimally (up to a scalar) achieve support recovery with respect to the effective sample size $n_{p,s}$, for certain classes of Σ and SIMs in a minimax sense. In the special case when $\Sigma = \mathbb{I}_{p \times p}$, the LASSO algorithm can be slightly modified into a simple covariance screening procedure which possesses similar properties as the LASSO procedure.

1.1 Overview of Related Work

When the dimension p is small, inference under a SIM has been studied extensively in the literature (Xia and Li, 1999; Horowitz, 2009; Peng and Huang, 2011; McCullagh and Nelder,

1989, e.g.) among many others. In the highly relevant line of work on sufficient dimension reduction, many seminal insights can be found in Li and Duan (1989); Li (1991); Cook and Ni (2005). When $\mathbf{X} \sim \mathcal{N}(0, \Sigma)$, results given in Li and Duan (1989) can be used to show that $\operatorname{argmin}_{\beta} \{\sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \beta)^2\}$ consistently estimates β_0 up to a scalar. When β_0 is sparse, the sparse sliced inverse regression procedure given in Li and Nachtsheim (2006) can be used to effectively recover β_0 under model (1) although their procedure requires a consistent estimator of $\Sigma^{-1/2}$.

In the high dimensional setting with diverging p , Alquier and Biau (2013) were the first to consider the sparse SIM, and proposed an estimation framework using a PAC-Bayesian approach. Wang and Zhu (2015) and Wang et al. (2012) demonstrated that when support recovery can be achieved when $p = O(n^k)$ under a SIM via optimizations in the form of $\hat{\beta} = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2n} \sum_{i=1}^n (F_n(Y_i) - 1/2 - \mathbf{X}_i^\top \beta)^2 + \sum_{j=1}^p J_\lambda(\beta_j)$, where J_λ is a penalty function and $F_n(x) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}(Y_i \leq x)$. Regularized procedures have also been proposed for specific choices of f and Y . For example, Yi et al. (2015) study consistent estimation under the model $\mathbb{P}(Y = 1 | \mathbf{X}) = \{f(\beta^\top \mathbf{X}) + 1\}/2$ with binary Y , where $f : \mathbb{R} \mapsto [-1, 1]$. Yang et al. (2015) consider the model $Y = f(\mathbf{X}^\top \beta) + \varepsilon$ with known f , and develop estimation and inferential procedures based on the L_1 regularized least squares loss.

With p potentially growing with n exponentially and under a general SIM, Radchenko (2015) proposed a non-parametric least squares with an equality L_1 constraint to handle simultaneous estimation of β as well as f . The support recovery properties of this procedure are not investigated, and in addition the results do not exhibit the optimal scaling of the triple (n, p, s) . Han and Wang (2015) suggest a penalized approach, in which they use a loss function related to Kendall’s tau correlation coefficient. They also establish the L_2 consistency for the coefficient β but do not consider support recovery. Neykov et al. (2015) analyzed two algorithms based on Sliced Inverse Regression (Li, 1991) under the assumption that $\mathbf{X} \sim \mathcal{N}(0, \mathbb{I}_{p \times p})$, and demonstrated that they can uncover the support optimally in terms of the rescaled sample size. Plan and Vershynin (2015) and Thrampoulidis et al. (2015) demonstrated that a constrained version of LASSO can be used to obtain an L_2 consistent estimator of β_0 . None of these procedures provide results on the performance of the LASSO algorithm in support recovery, which relates to L_2 consistency but is a fundamentally different theoretical aspect. In addition, no existing work on the SIM estimation procedures demonstrates that the performance depends on (n, p, s) only through the effective sample size $n_{p,s}$.

1.2 Organization

The rest of the paper is organized as follows. Our main results are formulated in section 2. In particular, we show results on the covariance screening algorithm when $\Sigma = \mathbb{I}_{p \times p}$ in section 2.2 and our main result on the LASSO support recovery in section 2.3. Proof for the main results are given in section 3. In addition we demonstrate that for a class of SIMs, any algorithm provably fails to recover the support, unless the rescaled sample size $n_{p,s}$ is large enough. Numerical studies, confirming our main result are shown in section 4. We discuss potential future directions in section 5. Technical proofs are deferred to the appendixes.

2. Main Results

In this section we formulate our main results, which include the analysis of a simple covariance screening algorithm, and the LASSO algorithm for SIMs. Before we move on to the algorithms we summarize notation that we use throughout the paper and discuss several useful definitions and preliminary results.

2.1 Preliminary and Notation

For a (sparse) vector $\mathbf{v} = (v_1, \dots, v_p)^\top$, we let $S(\mathbf{v}) := \text{supp}(\mathbf{v}) = \{j : v_j \neq 0\}$ denote its support, $S_\pm(\mathbf{v}) := \{(\text{sign}(v_j), j) : v_j \neq 0\}$ be its signed support, $\|\mathbf{v}\|_p$ denote the L_p norm, $\|\mathbf{v}\|_0 = |\text{supp}(\mathbf{v})|$, and $\mathbf{v}^{\min} = \min_{i \in \text{supp}(\mathbf{v})} |v_i|$. For a real random variable X , define

$$\|X\|_{\psi_2} = \sup_{p \geq 1} p^{-1/2} (\mathbb{E}|X|^p)^{1/p}, \quad \|X\|_{\psi_1} = \sup_{p \geq 1} p^{-1} (\mathbb{E}|X|^p)^{1/p}.^1$$

Recall that a random variable is called *sub-Gaussian* if $\|X\|_{\psi_2} < \infty$ and *sub-exponential* if $\|X\|_{\psi_1} < \infty$. For any integer $k \in \mathbb{N}$ we use the shorthand notation $[k] = \{1, \dots, k\}$. For a matrix $\mathbf{M} \in \mathbb{R}^{d_1 \times d_2}$, sets $S_1, S_2 \subseteq [d]$, we let $\mathbf{M}_{,S_2} = [M_{ij}]_{i \in [d_1]}^{j \in S_2}$ and $\mathbf{M}_{S_1, S_2} = [M_{ij}]_{i \in S_1}^{j \in S_2}$. For a vector $\mathbf{Z} = (Z_1, \dots, Z_d)^\top$ and set $S \subseteq [d]$, \mathbf{Z}_S denotes the subvector corresponding to the set S . Furthermore, let $\|\mathbf{M}\|_{p,q} = \sup_{\|\mathbf{v}\|_p=1} \|\mathbf{M}\mathbf{v}\|_q$. In particular, we have $\|\mathbf{M}\|_{2,2} = \max_{i \in [\max(d_1, d_2)]} \{s_i(\mathbf{M})\}$, where $s_i(\mathbf{M})$ is the i^{th} singular value of \mathbf{M} , and $\|\mathbf{M}\|_{\infty, \infty} = \max_{i \in [d_1]} \sum_{j=1}^{d_2} |M_{ij}|$. For a matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$, we put $D_{\max}(\mathbf{M}) = \max_{i \in [d]} |M_{ii}|$ for its maximal diagonal element, and $\text{diag}(\mathbf{M}) = [M_{ii}]_{i \in [d]}$ for the collection of diagonal entries of \mathbf{M} . We also use standard asymptotic notations. Given two sequences $\{a_n\}, \{b_n\}$ we write $a_n = O(b_n)$ if there exists a constant $C < \infty$ such that $a_n \leq Cb_n$; $a_n = \Omega(b_n)$ if there exists a positive constant $c > 0$ such that $a_n \geq cb_n$, $a_n = o(b_n)$ if $a_n/b_n \rightarrow 0$, and $a_n \asymp b_n$ if there exists positive constants c and C such that $c < a_n/b_n < C$. Throughout, we also assume that there exists a constant $0 < \iota < 1$ such that $s < p - p^\iota$, which implies that $\frac{\log(p-s)}{\log(p)} \geq \iota$.

We assume that data for analysis consists of n independent and identically distributed (i.i.d.) random vectors $\mathcal{D} = \{(Y_i, \mathbf{X}_i^\top)^\top, i = 1, \dots, n\}$ and we focus primarily on $\mathbf{X} \sim \mathcal{N}(0, \Sigma)$. In matrix form, we let $\mathbb{X} = [\mathbf{X}_1, \dots, \mathbf{X}_n]_{n \times p}^\top = [X_{ij}]_{i \in [n]}^{j \in [p]}$, $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$, and $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_n)^\top$. The recovery of $\boldsymbol{\beta}_0$ under SIM often relies on the linearity of expectation assumption given in [Li and Duan \(1989\)](#) and [Li \(1991\)](#):

Definition 2.1.1 (Linearity of Expectation) *A p -dimensional random variable \mathbf{X} is said to satisfy linearity of expectation in the direction $\boldsymbol{\beta}$ if for any direction $\mathbf{b} \in \mathbb{R}^p$:*

$$\mathbb{E}[\mathbf{X}^\top \mathbf{b} | \mathbf{X}^\top \boldsymbol{\beta}] = c_b \mathbf{X}^\top \boldsymbol{\beta} + a_b,$$

where $a_b, c_b \in \mathbb{R}$ are some real constants which might depend on the direction \mathbf{b} .

Remark 2.1.2 *Note that if additionally $\mathbb{E}[\mathbf{X}] = 0$, then by taking expectation it is evident that $a_b \equiv 0$. Clearly, linearity of expectation is direction specific by definition. Elliptical*

1. There are multiple equivalent (up to universal constants) definitions of the so-called Orlicz or ψ norms. See [Vershynin \(2010\)](#) Lemma 5.5 for a succinct formal treatment of this.

distributions (Fang et al., 1990) including multivariate normal are known to satisfy the linearity in expectation uniformly in all directions (Cambanis et al., 1981, e.g.).

Next, we record a simple but very useful observation, which forms the basis of our work. From Theorem 2.1 of Li and Duan (1989), we note that under a SIM and normality of \mathbf{X} with covariance $\Sigma > 0$,

$$\operatorname{argmin}_{\mathbf{b}} \mathbb{E}(Y - \mathbf{b}^\top \mathbf{X})^2 = c_0 \boldsymbol{\beta}_0,$$

for some $c_0 \in \mathbb{R}$. More generally, we have

Lemma 2.1.3 *Assume that the SIM (1) holds, $\Sigma = \mathbb{E}(\mathbf{X}\mathbf{X}^\top) > 0$, and \mathbf{X} satisfies the linearity in expectation condition in the direction $\boldsymbol{\beta}_0$ such that $\mathbb{E}[(\mathbf{X}^\top \boldsymbol{\beta}_0)^2] > 0$. Then we have $\Sigma^{-1} \mathbb{E}(Y\mathbf{X}) = c_0 \boldsymbol{\beta}_0$, where $c_0 := \mathbb{E}(Y\mathbf{X}^\top \boldsymbol{\beta}_0)$. Obviously, $\mathbb{E}(Y\mathbf{X}) = c_0 \boldsymbol{\beta}_0$ when $\Sigma = \mathbb{I}_{p \times p}$.*

In view of Lemma 2.1.3, under sparsity assumptions, an L_1 regularized least square estimator can recover $\boldsymbol{\beta}_0$ proportionally and hence the support of $\boldsymbol{\beta}_0$. Furthermore, when $\Sigma = \mathbb{I}_{p \times p}$, the covariance $\mathbb{E}(Y\mathbf{X})$ can be directly used to recover $\boldsymbol{\beta}_0$. It is noteworthy to remark that in the special case $\mathbf{X} \sim \mathcal{N}(0, \Sigma)$, a simple application of Stein's Lemma (Stein, 1981) can help quantify the constant c_0 precisely:

$$c_0 = \mathbb{E}(Y\mathbf{X}^\top \boldsymbol{\beta}_0) = \mathbb{E}(Z \underbrace{[\mathbb{E}f(Z, \varepsilon)|Z]}_{\varphi(Z)}) = \int D\varphi(z) \frac{\exp(-z^2/2)}{\sqrt{2\pi}} dz = \mathbb{E}D\varphi(Z),$$

where $D\varphi$ is the distributional derivative of φ , $Z \sim \mathcal{N}(0, 1)$ and we abused the notation slightly in the last equality for simplicity.

The remaining of this section is structured as follows — in section 2.2 we study a simple covariance thresholding algorithm, which is a manifestation of program (2) under the assumption $\Sigma = \mathbb{I}_{p \times p}$. In section 2.3 we consider the full-fledged least squares LASSO algorithm (2) with a general covariance matrix Σ . Throughout, we assume $E(\mathbf{X}) = 0$ and let $\sigma^2 = E(Y^2)$, $\eta = \operatorname{Var}(Y^2)$,

$$c_0 = \mathbb{E}(Y\mathbf{X}^\top \boldsymbol{\beta}_0), \quad \gamma = \operatorname{Var}(Y\mathbf{X}^\top \boldsymbol{\beta}_0), \quad \xi^2 = \mathbb{E}\{(Y - c_0 \mathbf{X}^\top \boldsymbol{\beta}_0)^2\}, \quad \text{and} \quad \theta^2 = \operatorname{Var}\{(Y - c_0 \mathbf{X}^\top \boldsymbol{\beta}_0)^2\}.$$

In addition, to simplify the presentation we will assume that the above constants are not scaling with (n, p, s) , and belong to a compact set which is bounded away from 0.

2.2 Covariance Screening under $\Sigma = \mathbb{I}_{p \times p}$

In this section, we propose a simple covariance screening procedure for signed support recovery of $\boldsymbol{\beta}_0$ under a SIM with $\Sigma = \mathbb{I}_{p \times p}$, which relates to the sure independence screening procedures (Fan and Lv, 2008; Fan et al., 2010) under the linear model. Note that the LASSO procedure (2) can equivalently be expressed as:

$$\hat{\boldsymbol{\beta}} = \operatorname{argmin}_{\boldsymbol{\beta} \in \mathbb{R}^p} \left\{ -\frac{1}{n} \sum_{i=1}^n Y_i \mathbf{X}_i^\top \boldsymbol{\beta} + \frac{1}{2n} \boldsymbol{\beta}^\top \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top \boldsymbol{\beta} + \lambda \|\boldsymbol{\beta}\|_1 \right\}.$$

Under the assumption $\Sigma = \mathbb{I}_{p \times p}$, we replace $\frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top$ with $\mathbb{I}_{p \times p}$ and instead consider

$$\hat{\boldsymbol{\beta}} = \operatorname{argmin}_{\boldsymbol{\beta} \in \mathbb{R}^p} \left\{ -\frac{1}{n} \sum_{i=1}^n Y_i \mathbf{X}_i^\top \boldsymbol{\beta} + \frac{1}{2} \boldsymbol{\beta}^\top \boldsymbol{\beta} + \lambda \|\boldsymbol{\beta}\|_1 \right\}.$$

It is well known that the solution to the above program takes the form:

$$\hat{\beta}_j = \operatorname{sign} \left(n^{-1} \sum_{i=1}^n Y_i X_{ij} \right) \left(\left| n^{-1} \sum_{i=1}^n Y_i X_{ij} \right| - \lambda \right)_+,$$

where $x_+ = x \mathbb{1}(x \geq 0)$. Hence, in this special case, the regularization parameter can be equivalently interpreted as a thresholding parameter, filtering all small $|n^{-1} \sum_{i=1}^n Y_i X_{ij}|$. Motivated by this, we consider the following simple covariance screening procedure, which acts as a filter taking covariances and their corresponding signs, only if they pass a critical threshold:

Algorithm 1: Covariance Screening Algorithm

input: $(Y_i, \mathbf{X}_i)_{i=1}^n$: data; tuning parameter $\nu > 0$

1. Calculate $\mathbf{V} := \widehat{\operatorname{Cov}}(Y, \mathbf{X}) := n^{-1} \sum_{i=1}^n Y_i \mathbf{X}_i$,
 2. Set $\widehat{S} := \left\{ V_j : |V_j| > \nu \sqrt{\frac{\log p}{n}} \right\}$;
 3. Output the set $\{(\operatorname{sign}(V_j), j) : |V_j| \in \widehat{S}\}$.
-

The following proposition shows that the algorithm recovers the support with probability approaching 1 provided that the effective sample size $n_{p,s}$ is sufficiently large under the normality assumption of $\mathbf{X} \sim \mathcal{N}(0, \mathbb{I}_{p \times p})$.

Proposition 2.2.1 *Assume that $\mathbf{X} \sim \mathcal{N}(0, \mathbb{I}_{p \times p})$, $c_0 \neq 0$, $E(Y^4) < \infty$, and $\beta_{0j} \in \{\pm \frac{1}{\sqrt{s}}, 0\}$ for all $j \in [p]$ and some $s \in \mathbb{N}$. Set $\nu = \frac{1}{\sqrt{s}} + 2\sqrt{2}\sigma$. Then as long as*

$$n_{p,s} \geq \Upsilon,$$

for a large enough $\Upsilon = O(1)$ (depending on c_0, σ^2) Algorithm 1 recovers the signed support of $c_0 \boldsymbol{\beta}_0$ with asymptotic probability 1.

The proof of Proposition 2.2.1 can be found in Appendix D. We note that this result does not require Y to be sub-Gaussian. When sub-Gaussianity is assumed, the convergence rate of the probability approaching 1 can be improved. Furthermore, under sub-Gaussianity of Y , we can relax the normality assumption of \mathbf{X} . Specifically, one may instead consider the following assumptions

Assumption 2.2.2 (Spherical Distribution of \mathbf{X} and sub-Gaussian of Y) *Let \mathbf{X} be a spherically distributed p dimensional random variable with $\mathbb{E}[\mathbf{X}] = 0$, $\operatorname{Var}[\mathbf{X}] = \mathbb{I}_{p \times p}$, whose moment generating function exists, and takes the form $\mathbb{E}\{\exp(\mathbf{t}^\top \mathbf{X})\} \equiv \Psi(\mathbf{t}^\top \mathbf{t})$, $\mathbf{t} \in \mathbb{R}^p$, and in addition $\Psi : \mathbb{R} \mapsto \mathbb{R}$ is such that $\Psi(t) \leq \exp(Ct)$ for some $C > 0$ for all $t \in \mathbb{R}^+$. In addition, we assume that Y is sub-Gaussian, i.e. $\|Y\|_{\psi_2} \leq K_Y$.*

Note that (Vershynin, 2010, see Lemma 5.5. e.g.) assumption 2.2.2 implies $\max_{j \in [p]} \|X_j\|_{\psi_2} < \infty$. Let $K := \max_{j \in [p]} \|X_j\|_{\psi_2}$. In parallel to Proposition 2.2.1, we have

Proposition 2.2.3 *In addition to Assumption 2.2.2, we assume that $s, c_0 \neq 0$ and $\beta_{0j} \in \{\pm \frac{1}{\sqrt{s}}, 0\}$ for all $j \in [p]$ and some $s \in \mathbb{N}$. Set the tuning parameter $\nu = \omega K_Y K$ for some absolute constant $\omega > 0$. Then there exists an absolute constant $\Upsilon \in \mathbb{R}$ depending on c_0, C, K and K_Y such that if:*

$$n_{p,s} \geq \Upsilon,$$

Algorithm 1 recovers the signed support of $c_0 \beta_0$ with asymptotic probability 1.

2.3 LASSO Algorithm with General Σ

In this section we consider the LASSO algorithm (2),

$$\hat{\beta} = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|\mathbf{Y} - \mathbb{X}\beta\|_2^2 + \lambda \|\beta\|_1, \quad (3)$$

under the assumption $\mathbf{X} \sim \mathcal{N}(0, \Sigma)$ with a generic and unknown covariance matrix Σ . Under certain sufficient conditions our goal is to show that (2) recovers the support of β_0 with asymptotic probability 1 in optimal (up to a scalar) effective sample size. In contrast to section 2.2, we also no longer require each of the signals of β to be of the same magnitude.

We first summarize a primal dual witness (PDW) construction which we borrow from Wainwright (2009). The PDW construction lays out steps allowing one to prove sign consistency for L_1 constrained quadratic programming (3). We will only provide the sufficient conditions to show sign-consistency, and the interested reader can check Wainwright (2009) for the necessary conditions. We note that the validity of the PDW construction is generic, in that it does not rely on the distribution of the residual $\mathbf{w} = \mathbf{Y} - c_0 \mathbb{X} \beta_0$, and hence extends to the current framework. Note that unlike the linear regression case, in our setting \mathbf{w} does not necessarily have mean 0 although $\mathbb{E}[\mathbb{X}^\top \mathbf{w}] = 0$.

Recall that a vector \mathbf{z} is a subgradient of the L_1 norm evaluated at a vector $\mathbf{v} \in \mathbb{R}^p$ (i.e. $\mathbf{z} \in \partial \|\mathbf{v}\|_1$) if we have $z_j = \operatorname{sign}(v_j)$, $v_j \neq 0$ and $z^j \in [-1, 1]$ otherwise. It follows from Karush-Kuhn-Tucker's theorem that a vector $\hat{\beta} \in \mathbb{R}^p$ is optimal for the LASSO problem (3) iff there exists a subgradient $\hat{\mathbf{z}} \in \partial \|\hat{\beta}\|_1$ such that:

$$\frac{1}{n} \mathbb{X}^\top \mathbb{X} (\hat{\beta} - c_0 \beta_0) - \frac{1}{n} \mathbb{X}^\top \mathbf{w} + \lambda \hat{\mathbf{z}} = 0. \quad (4)$$

Put $S_0 := S(\beta_0)$ for brevity. In what follows we will assume that the matrix $\mathbb{X}_{S_0}^\top \mathbb{X}_{S_0}$ is invertible (which it is with probability 1), even though this is not required by the PDW. The PDW method constructs a pair $(\check{\beta}, \check{\mathbf{z}}) \in \mathbb{R}^p \times \mathbb{R}^p$ by following the steps:

- Solve:

$$\check{\beta}_{S_0} = \operatorname{argmin}_{\beta_{S_0} \in \mathbb{R}^s} \frac{1}{2n} \|\mathbf{Y} - \mathbb{X}_{S_0} \beta_{S_0}\|_2^2 + \lambda \|\beta_{S_0}\|_1,$$

where $s = |S_0|$. This solution is unique under the invertibility of $\mathbb{X}_{S_0}^\top \mathbb{X}_{S_0}$ since in this case the function is strictly convex. Set $\check{\beta}_{S_0^c} = 0$.

- Choose a $\check{\mathbf{z}}_{S_0} \in \partial \|\check{\boldsymbol{\beta}}_{S_0}\|_1$.
- For $j \in S_0^c$ set $Z_j := \mathbb{X}_{,j}^\top [\mathbb{X}_{,S_0} (\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0})^{-1} \check{\mathbf{z}}_{S_0} + \mathbf{P}_{\mathbb{X}_{,S_0}^\perp} (\frac{\mathbf{w}}{\lambda n})]^2$, where $\mathbf{P}_{\mathbb{X}_{,S_0}^\perp} = \mathbb{I} - \mathbb{X}_{,S_0} (\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0})^{-1} \mathbb{X}_{,S_0}^\top$ is an orthogonal projection. Checking that $|Z_j| < 1$ for all $j \in S_0^c$ ensures that there is a unique solution $\check{\boldsymbol{\beta}} = (\check{\boldsymbol{\beta}}_{S_0}^\top, \check{\boldsymbol{\beta}}_{S_0^c}^\top)^\top$ satisfying $S(\check{\boldsymbol{\beta}}) \subseteq S(c_0 \boldsymbol{\beta}_0)$.
- To check sign consistency we need $\check{\mathbf{z}}_{S_0} = \text{sign}(c_0 \boldsymbol{\beta}_{0,S_0})$. For each $j \in S_0$, define:

$$\Delta_j := \mathbf{e}_j^\top (n^{-1} \mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0})^{-1} \left[n^{-1} \mathbb{X}_{,S_0}^\top \mathbf{w} - \lambda \text{sign}(c_0 \boldsymbol{\beta}_{0,S_0}) \right],^3$$

where $\mathbf{e}_j \in \mathbb{R}^s$ is a canonical unit vector with 1 at the j^{th} position. Checking $\check{\mathbf{z}}_{S_0} = \text{sign}(c_0 \boldsymbol{\beta}_{0,S_0})$ is equivalent to checking:

$$\text{sign}(c_0 \beta_{0j} + \Delta_j) = \text{sign}(c_0 \beta_{0j}), \quad \forall j \in S_0.$$

To this end we require several restrictions on $\boldsymbol{\Sigma}$ and moment conditions on Y . We partition the covariance matrix

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{S_0, S_0} & \boldsymbol{\Sigma}_{S_0, S_0^c} \\ \boldsymbol{\Sigma}_{S_0^c, S_0} & \boldsymbol{\Sigma}_{S_0^c, S_0^c} \end{bmatrix},$$

where $\boldsymbol{\Sigma}_{S_0, S_0}$ corresponds to the covariance of \mathbf{X}_{S_0} . Furthermore, we let

$$\boldsymbol{\Sigma}_{S_0^c | S_0} := \boldsymbol{\Sigma}_{S_0^c, S_0^c} - \boldsymbol{\Sigma}_{S_0^c, S_0} \boldsymbol{\Sigma}_{S_0, S_0}^{-1} \boldsymbol{\Sigma}_{S_0, S_0^c}, \quad (5)$$

$$\rho_\infty(\boldsymbol{\Sigma}_{S_0, S_0}^{1/2}) := \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty} \|\boldsymbol{\Sigma}_{S_0, S_0}^{1/2}\|_{\infty, \infty}, \quad (6)$$

be the conditional covariance matrix of $\mathbf{X}_{S_0^c} | \mathbf{X}_{S_0}$, and the condition number of $\boldsymbol{\Sigma}_{S_0, S_0}^{1/2}$ with respect to $\|\cdot\|_{\infty, \infty}$, respectively. We assume that

Assumption 2.3.1 (Irrepresentable Condition)

$$\|\boldsymbol{\Sigma}_{S_0^c, S_0} \boldsymbol{\Sigma}_{S_0, S_0}^{-1}\|_{\infty, \infty} \leq (1 - \kappa), \quad \text{for some } \kappa > 0.$$

Assumption 2.3.2 (Bounded Spectrum) For some fixed $0 < \lambda_{\min} \leq \lambda_{\max} < \infty$,

$$\lambda_{\min} \leq \boldsymbol{\Sigma}_{S_0, S_0} \leq \lambda_{\max}.$$

Assumption 2.3.3 (Bounded 4th Moment) $\mathbb{E}(Y^4) < \infty$, and does not scale with (n, p, s) .

Recall that the irrepresentable condition is proved to be necessary for successful support recovery (Wainwright, 2009, Theorem 4) in the linear model, and hence one should not expect that Assumption 2.3.1 can be weakened. Assumption 2.3.3 guarantees that σ^2 , η , c_0 , γ , ξ^2 and θ^2 are well defined and finite. Finally, successful support recovery will depend on the strength of the minimal signal of $\boldsymbol{\beta}_0$. Recall that $\boldsymbol{\beta}_0^{\min} := \min_{j \in S_0} |\beta_{0j}|$, is the minimal non-zero signal in the vector $\boldsymbol{\beta}_0$. We are now ready to provide sufficient conditions for the LASSO signed support recovery, in the setting of SIMs

2. Z_j are derived by simply plugging in $\check{\boldsymbol{\beta}}$ and $\check{\mathbf{z}}_{S_0}$ and solving (4) for $\check{\mathbf{z}}_{S_0^c}$.

3. Δ_j can be seen to equal $\check{\beta}_j - c_0 \beta_{0j}$ for $j \in S_0$, when $\check{\mathbf{z}}_{S_0} = \text{sign}(c_0 \boldsymbol{\beta}_0)$.

Theorem 2.3.4 *Let Assumptions 2.3.1 – 2.3.3 hold. Then for the LASSO estimator given in (3) under a SIM (1), we have the following sufficient conditions:*

i. If

$$n_{p,s} \geq \frac{4D_{\max}(\Sigma_{S_0^c|S_0}) \left(\frac{4}{\lambda_{\min}} + \frac{\xi^2+1}{\lambda^2 s} \right)}{\kappa^2},$$

then $S(\hat{\beta}) \subset S(c_0\beta_0)$, with probability at least $1 - O\{p^{-1} + n^{-1} + e^{-\Omega(s)}\}$.

ii. Let further, for some positive constant $\alpha > 0$ we have $n_{p,s} \geq \alpha$. Then there exist some positive constants $\Upsilon_0, \Upsilon_1, \Upsilon_2 > 0$ which may depend⁴ on c_0, σ such that if:

$$\beta_0^{\min} \geq \|\Sigma_{S_0, S_0}^{-1/2}\|_{\infty, \infty}^2 \lambda \Upsilon_0 + \left[\Upsilon_1 \rho_{\infty}(\Sigma_{S_0, S_0}^{1/2}) \sqrt{\frac{s}{n \log(p-s)}} \|\beta_0\|_{\infty} + \frac{\Upsilon_2 \|\Sigma_{S_0, S_0}^{-1/2}\|_{\infty, \infty}}{\sqrt{s}} \right] n_{p,s}^{-1/2} \quad (7)$$

we have $S_{\pm}(\hat{\beta}) = S_{\pm}(c_0\beta_0)$ with probability at least $1 - O\{e^{-\Omega(s \wedge \log(p-s))} + (\log p)^{-1} + n^{-1}\}$.

Before we proceed with the proof of our main result, we would like to mention a few remarks on our sufficient conditions, in particular the ones suggested in ii.

Remark 2.3.5 *Firstly, the slow probability convergence rate $(\log p)^{-1}$ in part ii. is purely due to the fact we are not willing to assume that Y is coming from a sub-Gaussian distribution. If such an assumption is made the rate reduces to the usual p^{-1} . Secondly, observe that $\lambda_{\max}^{-1} \leq \|\beta_0\|_2 \leq \lambda_{\min}^{-1}$. Hence the value of β_0^{\min} is of “largest” order when $\beta_0^{\min} \asymp \|\beta_0\|_{\infty} \asymp \frac{1}{\sqrt{s}}$. Setting*

$$\lambda := \lambda_T = \sqrt{(\xi^2 + 1) \frac{4C_T D_{\max}(\Sigma_{S_0^c|S_0}) \log(p-s)}{\kappa^2 n}},$$

for some $C_T > 1$ gives us that the condition from i. is equivalent to:

$$n_{p,s} \geq \frac{16D_{\max}(\Sigma_{S_0^c|S_0})}{(1 - C_T^{-1})\kappa^2 \lambda_{\min}}.$$

Note that due to positive definiteness: $D_{\max}(\Sigma_{S_0^c|S_0}) \leq D_{\max}(\Sigma_{S_0^c, S_0^c})$, and hence it is reasonable to assume $D_{\max}(\Sigma_{S_0^c|S_0}) = O(1)$. Assume additionally that $\|\Sigma_{S_0, S_0}^{-1/2}\|_{\infty, \infty} = O(1)$, $\rho_{\infty}(\Sigma_{S_0, S_0}^{1/2}) = O(1)$, and let $\beta_0^{\min} \asymp \frac{1}{\sqrt{s}}$. Notice that the space of matrices satisfying assumptions 2.3.1 – 2.3.3 and the latter assumptions is non-empty as one can easily show that Toeplitz matrices with entries $\Sigma_{kj} = \rho^{|k-j|}$ with $|\rho| < 1$ satisfy these conditions e.g. Using the same $\lambda = \lambda_T$ we can clearly achieve the sufficient condition in ii. provided that the model complexity adjusted sample size $n_{p,s}$ is large enough. On the other hand, this

4. The dependency is inversely proportional to $|c_0|$ and proportional to σ ; For more details refer to the proof.

scaling can no longer be guaranteed if $\beta_0^{\min} \asymp \frac{1}{\sqrt{s}}$ fails to hold. In any case, it is clear that when $\Sigma = \mathbb{I}_{p \times p}$ all conditions required are met, and hence Theorem 2.3.4 shows that the LASSO algorithm will work optimally (up to a scalar) in terms of the rescaled sample size. Below we formulate a result which allows us to claim certain optimality for the LASSO algorithm in the case of a more general Σ matrix whenever we have approximately equally sized coefficients.

Proposition 2.3.6 *Consider a special example of a SIM with $Y = G\{h(\mathbf{X}^\top \beta_0) + \varepsilon\}$ for $\mathbf{X} \sim \mathcal{N}(0, \Sigma)$, and G, h are known strictly increasing continuous functions and in addition h is an L -Lipschitz function. Assume that there exists a set $S \subset [p], |S| = s$ such that $\|\Sigma_{S,S}\|_{\infty, \infty} < R$, Assumptions 2.3.1 and 2.3.2 hold on S and $0 < d \leq \text{diag}(\Sigma_{S^c, S^c}) \leq D < \infty$. In addition, assume that ε is a continuously distributed random variable with density p_ε satisfying:*

$$p_\varepsilon(x) \propto \exp(-P(x^2)), \tag{8}$$

where P is any non-zero polynomial with non-negative coefficients such that $P(0) = 0$. We restrict the parameter space to $\left\{ \beta \in \mathbb{R}^p : \beta^\top \Sigma \beta = 1, \|\beta\|_0 = s, \frac{\beta^{\min}}{\|\beta\|_\infty} \geq c_\Sigma \right\}$, where $c_\Sigma > 0$ is a sufficiently small constant depending solely on Σ (see (20) for details). If

$$n_{p,s} < C,$$

for some constant $C > 0$ (depending on P, G, h, Σ) and s is sufficiently large, any algorithm for support recovery makes errors with probability at least $\frac{1}{2}$ asymptotically.

Proposition 2.3.6 justifies that the LASSO is sample size optimal even when more generic covariance matrices than identity are considered, provided that we can show the class of SIMs defined satisfy the assumptions of this section. We do so in the following Remark.

Remark 2.3.7 *We will now argue that if we have a model as described in Proposition 2.3.6, the LASSO will recover the support provided that the the covariance matrix Σ satisfies Assumptions 2.3.1 and 2.3.2, $\mathbb{E}(Y^4) < \infty$ and the minimal signal strength is sufficiently strong as required by Remark 2.3.5. Notice that by Chebyshev's association inequality (Boucheron et al., 2013), we have:*

$$\mathbb{E}(Y \mathbf{X}^\top \beta_0) > \mathbb{E}(Y) \mathbb{E}(\mathbf{X}^\top \beta_0) = 0,$$

where the inequality is strict since G and h are strictly increasing and $\mathbf{X}^\top \beta_0 \sim \mathcal{N}(0, 1)$. In fact, using exactly the same argument, one can show more generally that if $r(z) = \mathbb{E}[Y | \mathbf{X}^\top \beta_0 = z]$ is a strictly monotone function it follows that $\mathbb{E}(Y \mathbf{X}^\top \beta_0) \neq 0$. We close this remark by pointing out that the logistic regression model $P(Y = 1 | \mathbf{X}) = g(\mathbf{X}^\top \beta_0)$ with $g(x) = e^x / (1 + e^x)$ satisfies the condition since $r(z) = g(z)$ is strictly monotone, and hence using the LASSO algorithm one can recover the support correctly. This is an example that even discrete valued Y outcomes can be solved by the least squares LASSO algorithm.

2.3.1 OUTCOME TRANSFORMATIONS

In this subsection we provide brief comments on possible strategies to transform the data in view of the results of Theorem 2.3.4. A crucial condition in order for the signed support recovery of β_0 to hold is $c_0 \neq 0$, which should not be expected to hold in general, but nevertheless naturally occurs in many cases. If this condition does not hold, one can potentially transform the outcome $\tilde{Y} = g(Y)$ by a function g in order to achieve $\mathbb{E}(\tilde{Y} \mathbf{X}^\top \beta_0) \neq 0$ even if $\mathbb{E}(Y \mathbf{X}^\top \beta_0) = 0$. If we use $\tilde{Y}_i = g(Y_i)$ instead of Y_i then clearly LASSO succeeds under the assumptions from Theorem 2.3.4, only with assumptions on Y_i being replaced by \tilde{Y}_i . The following proposition characterizes when one should expect a correlation inducing transformation g to exist.

Proposition 2.3.8 *There exists a measurable function $g : \mathbb{R} \mapsto \mathbb{R}$ such that $\mathbb{E}\{g(Y) \mathbf{X}^\top \beta_0\} \neq 0$ if and only if $\text{Var}\{\mathbb{E}(\mathbf{X}^\top \beta_0 | Y)\} > 0$.*

Another potential advantage of performing a transformation is to ensure that $\tilde{Y} = g(Y)$ is sub-Gaussian. For example, if we let $g(y) = F(y) = P(Y \leq y)$, then the sub-Gaussianity of \tilde{Y} is guaranteed, which would improve the rate of support recovery. For many choices of g such as F , the transformation may be defined at the population level and is unknown a priori. Thus it would be desirable to employ data dependent estimate \hat{g} of g . In other words we consider fitting the following LASSO to recover the support:

$$\hat{\beta} = \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{2n} \|\hat{g}(\mathbf{Y}) - \mathbb{X}\beta\|_2^2 + \lambda \|\beta\|_1, \quad (9)$$

where $\hat{g}(\mathbf{Y})$ should be understood as element-wise application of \hat{g} . The following Corollary extends Theorem 2.3.4 to allow for data dependent transformations.

Corollary 2.3.9 *Let the assumptions of Theorem 2.3.4 hold for $\tilde{Y}_i = g(Y_i)$ in place of Y_i and assume additionally that:*

$$\|\hat{g}(\mathbf{Y}) - g(\mathbf{Y})\|_2 \leq O(\sqrt{\log p}), \quad (10)$$

with probability at least $1 - O(p^{-1})$, then if (7) holds we have $S_\pm(\hat{\beta}) = S_\pm(c_0 \beta_0)$ with probability at least $1 - O\{e^{-\Omega(s \wedge \log(p-s))} + (\log p)^{-1} + n^{-1}\}$.

Remark 2.3.10 *Akin to (2), we do not require an intercept in the model after doing a transformation in (9). This is possible since \mathbf{X} is assumed to have mean 0. In practice if this were not the case, one would have to either center \mathbf{X} or include an intercept which is not penalized.*

3. Proof of Theorem 2.3.4

Our proof follows similar steps as Theorem 3 of Wainwright (2009) although many critical modifications are needed, since the error term $\mathbf{w} = \mathbf{Y} - c_0 \mathbb{X} \beta_0$ is no longer independent of \mathbb{X} and is not mean 0.

3.1 Verifying Strict Dual Feasibility

For $j \in S_0^c$ decompose

$$\mathbb{X}_{\cdot j}^\top = \boldsymbol{\Sigma}_{\{j\}, S_0} \boldsymbol{\Sigma}_{S_0, S_0}^{-1} \mathbb{X}_{\cdot, S_0}^\top + \mathbf{E}_j^\top, \quad (11)$$

where the entries of the prediction error vector $\mathbf{E}_j = (E_{1j}, \dots, E_{nj})^\top \in \mathbb{R}^n$ are i.i.d. with $E_{ij} \sim \mathcal{N}(0, [\boldsymbol{\Sigma}_{S_0^c | S_0}]_{jj})$, $i \in [n]$. In addition, observe that by this construction we have that \mathbf{E}_j is independent of \mathbb{X}_{\cdot, S_0} which can be verified upon multiplication by \mathbb{X}_{\cdot, S_0} in (11) and taking expectation. Following the definition of Z_j gives us that $Z_j = A_j + B_j$, where:

$$A_j := \mathbf{E}_j^\top \left[\mathbb{X}_{\cdot, S_0} (\mathbb{X}_{\cdot, S_0}^\top \mathbb{X}_{\cdot, S_0})^{-1} \check{\mathbf{z}}_{S_0} + \mathbf{P}_{\mathbb{X}_{\cdot, S_0}^\perp} \left(\frac{\mathbf{w}}{\lambda n} \right) \right], \quad (12)$$

$$B_j := \boldsymbol{\Sigma}_{\{j\}, S_0} (\boldsymbol{\Sigma}_{S_0, S_0})^{-1} \check{\mathbf{z}}_{S_0}. \quad (13)$$

Under the irrerepresentable condition, we have that $\max_{j \in S_0^c} |B_j| \leq (1 - \kappa)$. Conditional on \mathbb{X}_{\cdot, S_0} and $\boldsymbol{\varepsilon}$ (which determine $\mathbf{w} = \mathbf{Y} - c_0 \mathbb{X} \boldsymbol{\beta}_0$) we have that the gradient $\check{\mathbf{z}}_{S_0}$ is independent of the vector \mathbf{E}_j because the gradient is deterministic after conditioning on these quantities. We have that $\text{Var}(E_{ij}) \leq D_{\max}(\boldsymbol{\Sigma}_{S_0^c | S_0})$, and thus conditionally on \mathbb{X}_{\cdot, S_0} and $\boldsymbol{\varepsilon}$ we get:

$$\begin{aligned} \text{Var}(A_j) &\leq D_{\max}(\boldsymbol{\Sigma}_{S_0^c | S_0}) \left\| \mathbb{X}_{\cdot, S_0} (\mathbb{X}_{\cdot, S_0}^\top \mathbb{X}_{\cdot, S_0})^{-1} \check{\mathbf{z}}_{S_0} + \mathbf{P}_{\mathbb{X}_{\cdot, S_0}^\perp} \left(\frac{\mathbf{w}}{\lambda n} \right) \right\|_2^2 \\ &= D_{\max}(\boldsymbol{\Sigma}_{S_0^c | S_0}) \left[\check{\mathbf{z}}_{S_0}^\top (\mathbb{X}_{\cdot, S_0}^\top \mathbb{X}_{\cdot, S_0})^{-1} \check{\mathbf{z}}_{S_0} + \left\| \mathbf{P}_{\mathbb{X}_{\cdot, S_0}^\perp} \left(\frac{\mathbf{w}}{\lambda n} \right) \right\|_2^2 \right]. \end{aligned}$$

Next we need a lemma, which is a slight modification of Lemma 4 in [Wainwright \(2009\)](#). The reason for this modification is that in our case \mathbf{w} is no longer $\sim \mathcal{N}(0, \sigma^2 \mathbb{I})$.

Lemma 3.1.1 *Assume that $\frac{s}{n} \leq \frac{1}{16}$. Then we have:*

$$\max_{j \in S_0^c} \text{Var}(A_j) \leq \underbrace{D_{\max}(\boldsymbol{\Sigma}_{S_0^c | S_0}) \left(\frac{4s}{\lambda_{\min} n} + \frac{\xi^2 + 1}{\lambda^2 n} \right)}_M,$$

with probability at least $1 - 2e^{-s/2} - n^{-1}\theta^2$.

Now since conditionally on \mathbb{X}_{\cdot, S_0} and $\boldsymbol{\varepsilon}$ we have $A_j \sim \mathcal{N}(0, \text{Var}(A_j))$, using a standard normal tail bound and the union bound we conclude:

$$\mathbb{P}(\max_{j \in S_0^c} |A_j| \geq \kappa) \leq 2(p - s)e^{-\frac{\kappa^2}{2M}} + 2e^{-\frac{s}{2}} + n^{-1}\theta^2.$$

We need to select M so that the exponential term is decaying in the above display. A sufficient condition for this is $\kappa^2/(2M) \geq 2 \log(p - s)$. The last is equivalent to:

$$n_{p,s} \geq \frac{4D_{\max}(\boldsymbol{\Sigma}_{S_0^c | S_0}) \left(\frac{4}{\lambda_{\min}} + \frac{\xi^2 + 1}{\lambda^2 s} \right)}{\kappa^2}.$$

3.2 Verifying Sign Consistency

The first part of the proof shows that the LASSO has a unique solution $\widehat{\boldsymbol{\beta}}$ which satisfies $S(\widehat{\boldsymbol{\beta}}) \subseteq S(c_0\boldsymbol{\beta}_0)$ with high probability. Now we need to verify the sign-consistency, in order to show that the supports coincide. We have the following:

$$\max_{j \in S_0} |\Delta_j| \leq \lambda \underbrace{\left\| (n^{-1} \mathbb{X}_{S_0}^\top \mathbb{X}_{S_0})^{-1} \text{sign}(c_0 \boldsymbol{\beta}_{0S_0}) \right\|_\infty}_{\mathbf{I}_1} + \underbrace{\left\| (\mathbb{X}_{S_0}^\top \mathbb{X}_{S_0})^{-1} \mathbb{X}_{S_0}^\top \mathbf{w} \right\|_\infty}_{\mathbf{I}_2}.$$

To deal with the first term we need the following:

Lemma 3.2.1 *There exist positive constants $K_1, C_2 > 0$, such that the following holds:*

$$\mathbb{P}(\mathbf{I}_1 \geq \lambda K_1 \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty}^2) \leq 4 \exp(-C_2(s \wedge \log(p-s))),$$

The proof of this lemma is part of the proof of Theorem 3 in [Wainwright \(2009\)](#) and we omit the details. Next we turn to bounding the term \mathbf{I}_2 . Here our proof departs substantially from the proof in [Wainwright \(2009\)](#), as \mathbf{I}_2 no longer has a simple structure required in the original argument. In our case \mathbf{w} depends on \mathbb{X}_{S_0} , and it is not mean 0. We will make usage of the following result, whose proof is provided in the appendix

Lemma 3.2.2 *Let $\|\boldsymbol{\beta}_0\|_2 = 1$. We have n i.i.d. observations $Y = f(\mathbf{X}_{S_0}^\top \boldsymbol{\beta}_{0S_0}, \varepsilon)$ from a SIM, where $\mathbf{X}_{S_0} \sim \mathcal{N}(0, \mathbb{I}_{s \times s})$, with $s < n$ and $n_{p,s} \geq \alpha > 0$ for some positive constant α . Then there exist some positive constants $\Upsilon_1, \Upsilon_2 > 0$ (depending on σ and $|c_0|$), such that:*

$$\|[\mathbb{X}_{S_0}^\top \mathbb{X}_{S_0}]^{-1} \mathbb{X}_{S_0}^\top \mathbf{Y} - c_0 \boldsymbol{\beta}_{0S_0}\|_\infty \leq \left(\Upsilon_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + \Upsilon_2 \sqrt{\frac{\log(p-s)}{n}} \right).$$

with probability at least $1 - O\{e^{-s/2} + (\log p)^{-1} + n^{-1} + p^{-1}\}$. Denote for brevity the RHS of the inequality as $\delta(\|\boldsymbol{\beta}_{0S_0}\|_\infty, n, s, p)$.

While Lemma 3.2.2 is stated in terms of standard multivariate normal distribution $\mathcal{N}(0, \mathbb{I}_{s \times s})$, we can easily adapt it to more general situations where we observe non-standard normal random variables $\mathcal{N}(0, \boldsymbol{\Sigma}_{S_0, S_0})$. Recall that the rows of \mathbb{X}_{S_0} are distributed as $\mathcal{N}(0, \boldsymbol{\Sigma}_{S_0, S_0})$, $Y_i = f(\mathbf{X}_i^\top \boldsymbol{\beta}_0, \varepsilon)$, and $\boldsymbol{\beta}_{0S_0}^\top \boldsymbol{\Sigma}_{S_0, S_0} \boldsymbol{\beta}_{0S_0} = 1$. Denote with $\mathbf{Z} = \mathbb{X}_{S_0} \boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}$. Then we have the following inequality, with high probability:

$$\begin{aligned} \mathbf{I}_2 &= \|[\mathbb{X}_{S_0}^\top \mathbb{X}_{S_0}]^{-1} \mathbb{X}_{S_0}^\top \mathbf{Y} - c_0 \boldsymbol{\beta}_{0S_0}\|_\infty = \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2} [\mathbf{Z}^\top \mathbf{Z}]^{-1} \mathbf{Z}^\top \mathbf{Y} - c_0 \boldsymbol{\beta}_{0S_0}\|_\infty \\ &\leq \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty} \|[\mathbf{Z}^\top \mathbf{Z}]^{-1} \mathbf{Z}^\top \mathbf{Y} - c_0 \boldsymbol{\Sigma}_{S_0, S_0}^{1/2} \boldsymbol{\beta}_{0S_0}\|_\infty \\ &\leq \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty} \delta(\|\boldsymbol{\Sigma}_{S_0, S_0}^{1/2} \boldsymbol{\beta}_{0S_0}\|_\infty, s, n, p). \end{aligned}$$

The last two inequalities imply that:

$$\max_{j \in S} |\Delta_j| \leq \lambda K_1 \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty}^2 + \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty} \delta(\|\boldsymbol{\Sigma}_{S_0, S_0}^{1/2} \boldsymbol{\beta}_{0S_0}\|_\infty, s, n, p).$$

Hence as long as for $\boldsymbol{\beta}_0^{\min} = \min\{|\beta_{0j}| : j \in S_0\}$ we have:

$$|c_0| \boldsymbol{\beta}_0^{\min} \geq \lambda K_1 \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty}^2 + \|\boldsymbol{\Sigma}_{S_0, S_0}^{-1/2}\|_{\infty, \infty} \delta(\|\boldsymbol{\Sigma}_{S_0, S_0}^{1/2}\|_{\infty, \infty} \|\boldsymbol{\beta}_0\|_\infty, s, n, p),$$

the LASSO will recover the support with high-probability. This concludes the proof.

4. Numerical Studies

To support our theoretical claims, and in particular Theorem 2.3.4 we provide brief numeric analysis in this section. We consider the following models:

$$Y = \mathbf{X}^\top \boldsymbol{\beta}_0 + \sin(\mathbf{X}^\top \boldsymbol{\beta}_0) + \mathcal{N}(0, 1), \quad (14)$$


$$Y = 2 \operatorname{atan}(\mathbf{X}^\top \boldsymbol{\beta}_0) + \mathcal{N}(0, 1), \quad (15)$$

$$Y = (\mathbf{X}^\top \boldsymbol{\beta}_0)^3 + \mathcal{N}(0, 1), \quad (16)$$


$$Y = \sinh(\mathbf{X}^\top \boldsymbol{\beta}_0) + \mathcal{N}(0, 1). \quad (17)$$

We use a Toeplitz covariance matrix for the simulations $\boldsymbol{\Sigma} = \mathbb{I}$ and $\Sigma_{kj} = 2^{-|k-j|}$. The vector $\boldsymbol{\beta}_0$ is selected so that $\boldsymbol{\beta}_0^\top \boldsymbol{\Sigma} \boldsymbol{\beta}_0 = 1$, and its entries have equal magnitude, with the first one having a negative sign, and the remaining being positive. We check whether the solution path of the LASSO contains an s -sparse vector $\hat{\boldsymbol{\beta}}_0$ whose support coincides with the support of $\boldsymbol{\beta}_0$. This verifies the validity of one implication of our theory, as it shows that the solution path indeed contains the true signed support of $\boldsymbol{\beta}_0$.


Figure 1: LASSO, $s = \sqrt{p}$, $\boldsymbol{\Sigma} = \mathbb{I}_{p \times p}$




(a) Model (14)



(b) Model (15)



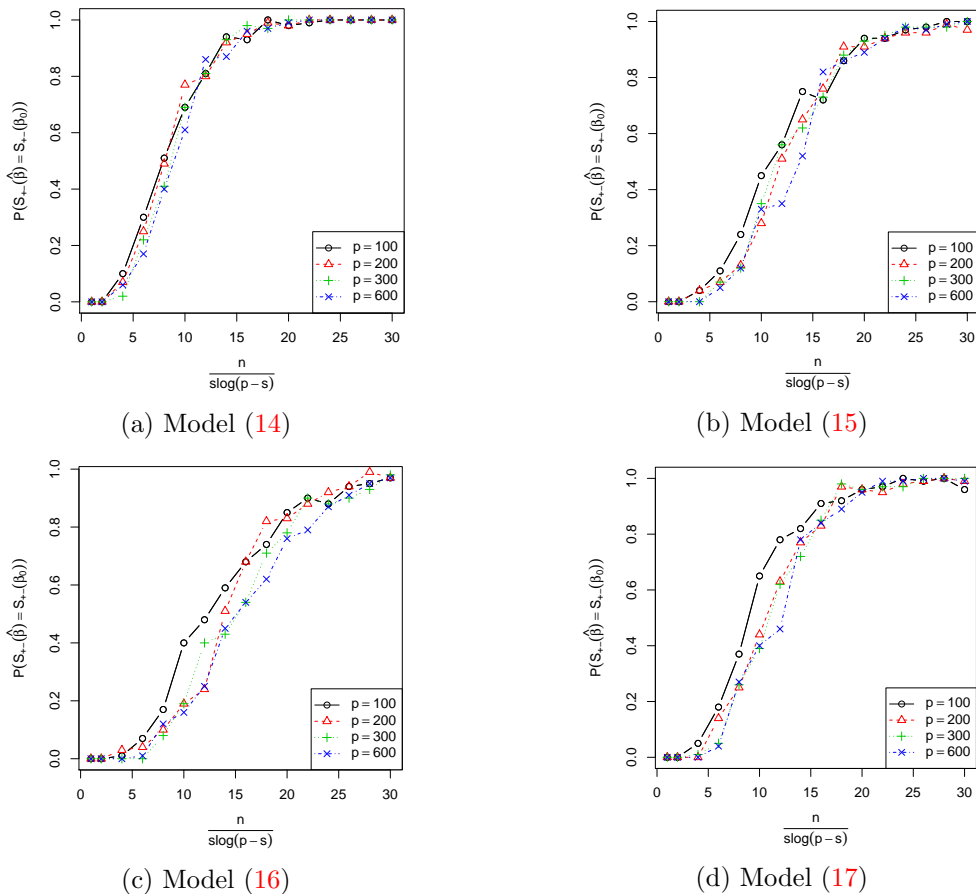
(c) Model (16)



(d) Model (17)

In Figure 1, we show results of signed support recovery for different values of p , in the regime $s = \sqrt{p}$ in the case of an identity covariance matrix $\Sigma = \mathbb{I}_{p \times p}$. Similarly, Figure 2 shows results for the case of Toeplitz covariance matrix $\Sigma_{kj} = 2^{-|k-j|}$. As expected, the support recovery is harder in the presence of correlation between the variables. These results illustrate different phase transitions occurring for the four different models. We observe empirically that values of $n_{p,s}$ achieving reasonable success probability can be large in some cases. It could be the case that using a transformed version of Y might lead to better results for the model complexity adjusted sample size, as suggested in Section 2.3.1. Figure 2 supports the result of Theorem 2.3.4 as all curves merge when the effective sample size $n_{p,s}$ becomes sufficiently large. In addition, these results suggest that the performance of the support recovery is largely determined by (n, p, s) through the magnitude of $n_{p,s}$.

Figure 2: LASSO, $s = \sqrt{p}$, $\Sigma_{kj} = 2^{-|k-j|}$



In addition to the verification of our theory, we also compare the vanilla least squares LASSO to a *version* of the sparse sliced inverse regression (SSIR) algorithm suggested by Li and Nachtsheim (2006). The SSIR algorithm is also based on a LASSO estimation. In its original form however, this algorithm is not applicable for high-dimensional settings such as ours, since it needs an estimate of the matrix $\Sigma^{-1/2}$. To make use of the SSIR under the

high-dimensional setting, we estimate $\Sigma^{-1/2}$ by the CLIME procedure (Cai et al., 2011) to $\widehat{\Sigma}^{1/2}$ under sparsity assumptions. Due to space considerations we only show comparisons for models (14) and (17), and the plots are attached in Appendix F. In the majority of cases LASSO outperforms the SSIR algorithm substantially for small values of $n_{p,s}$, although it seems that both approaches reach perfect support recovery at similar rescaled sample sizes. We would like to emphasize that the SSIR algorithm requires solving an extra optimization problem, and that furthermore, there are no theoretical results ensuring that SSIR in general recovers the support. On an important note, the SSIR algorithm is designed to estimate the central space of the more general class of multi-index models, which we do not discuss in the present paper. For a brief discussion on how our work can be related to multi-index model please refer to section 5.

5. Discussion

In this paper, we demonstrate that under a high dimensional SIM, L_1 -regularized least squares, including a simplified covariance screening procedure under orthogonal design, is robust in terms of model selection consistency, in that it correctly recovers the support of the true regression parameter β_0 provided that $c_0 = \mathbb{E}(Y\mathbf{X}^\top\beta_0) \neq 0$, the minimal signal strength is sufficiently large and $\mathbf{X} \sim \mathcal{N}(0, \Sigma)$ under standard assumptions on Σ which are necessary even in the linear regression case. Thus, our results extend known results on the support recovery performance of LASSO under linear models to a much broader class of SIMs. We furthermore demonstrate that the support recovery is achieved in a sample size optimal $n_{p,s}$ manner within a certain class of SIMs.

As we indicated in section 2.3.1, the assumption $c_0 \neq 0$ does not always hold, and in addition it cannot be easily verified. A potential remedy for this approach will be to transform the Y variable. From theoretical point of view it is of interest to develop procedures which can adaptively estimate a “good” outcome transformation. Additionally, a downside of the L_1 -regularization is the fact that the irrepresentable condition on the covariance matrix is unavoidable. This could potentially be remedied by using more general and non-convex penalties such as the SCAD penalty (Fan and Li, 2001). We focused on the setting with $\mathbf{X} \sim \mathcal{N}(0, \Sigma)$, however we suspect that the support recovery holds in more general cases where \mathbf{X} comes from an elliptical distribution. It is less clear, however, whether sample size optimality continues to hold in such situations, as we crucially rely on the normality of \mathbf{X} , in particular when using Lemma A.0.1 which follows by Gordon’s comparison theorem, and through numerous projection-independence properties which are characteristic of the Gaussian distribution.

The proposed method focuses on SIMs as the true underlying model. Extensions to incorporate general multi-index models are not straightforward. For the special case of multi-index models of the form

$$Y = \sum_{j=1}^k f_j(\mathbf{X}^\top\beta_j, \varepsilon_j),$$

our method should also be able to recover the support, assuming that k is fixed and the vectors β_j have disjoint supports, and $\mathbb{E}(Y\mathbf{X}^\top\beta_j) \neq 0$ for all $j \in [k]$. When applied to such a model, the LASSO estimate $\widehat{\beta}$ will include the union of the supports of β_j , provided that

sufficiently strong minimal signal (or sufficiently large sample size) and an irrepressantable condition are present. How to apply the LASSO algorithm for support recovery under the more general class of multi-index models warrants further research.

A note on the choice of the tuning parameter λ in practice is also in order. According to Remark 2.3.5 and optimal choice of λ may depend on the unknowable parameter ξ^2 . A procedure which we found to work well in practice is based on simple cross validation. To find a good tuning parameter λ , from a grid of ℓ values of λ : $\{\lambda_1, \dots, \lambda_\ell\}$ which are of $\sqrt{\frac{\log p}{n}}$ magnitude, we recommend using K -fold cross validation, and setting λ to the value minimizing the average least squares loss across (i.e. mean squared-error) the K folds. Notice that according to Lemma 2.1.3, this criteria is very sensible. An important question is whether one can arrive at a procedure with theoretical guarantees for λ selection, and we hope to address this problem in our future work.

Finally, under a SIM and proper distributional assumptions on \mathbf{X} , one may also recover β_0 proportionally using other convex loss functions. For example, when Y is binary, the logistic log-likelihood loss may be more efficient than the L_2 loss. Thus, it is of interest to investigate the support recovery properties of the LASSO (or more general penalization procedures) with other convex losses — such as the logistic/hinge losses, which could be less susceptible to outliers.

Acknowledgments

The authors would like to thank the Associate Editor and anonymous reviewers for their insightful remarks which led to the improvement of this manuscript. We are also grateful to Professor Nouredine El Karoui, who among other observations, brought to our attention that outcome transformations may be performed to facilitate the usage of LASSO even when $c_0 = 0$ for the original outcome Y . This research was partially supported by Research Grants NSF DMS1208771, NIH R01 GM113242-01, NIH U54 HG007963 and NIH RO1 HL089778.

Appendix A. Auxiliary Lemmas

In this section, for convenience of the reader, we state couple of lemmas that we use often in our analysis.

Lemma A.0.1 (Corollary 5.35 Vershynin (2010)) *Let $\mathbf{A}_{n \times s}$ matrix whose entries are i.i.d. standard normal random variables. Then for every $t \geq 0$, with probability at least $1 - 2 \exp(-t^2/2)$ one has:*

$$\sqrt{n} - \sqrt{s} - t \leq s_{\min}(\mathbf{A}) \leq s_{\max}(\mathbf{A}) \leq \sqrt{n} + \sqrt{s} + t,$$

where $s_{\min}(\mathbf{A})$ and $s_{\max}(\mathbf{A})$ are the smallest and largest singular values of \mathbf{A} correspondingly.

Lemma A.0.2 *Consider a fixed nonzero vector $\mathbf{z} \in \mathbb{R}^s$ and a random matrix $\mathbf{A}_{n \times s}$, whose entries are i.i.d. standard normal random variables. If p, s, n are such that $s \geq 2$, $\frac{s}{n} \leq \frac{1}{64}$ and $\frac{\log p}{n-s+1} \leq \frac{1}{32}$, there are positive absolute constants C_1 and C_2 satisfying:*

$$\mathbb{P} \left(\left\| (n^{-1} \mathbf{A}^\top \mathbf{A})^{-1} - \mathbb{I}_{s \times s} \right\| \|\mathbf{z}\|_\infty \geq C_1 \frac{s}{n} \|\mathbf{z}\|_\infty + C_2 \|\mathbf{z}\|_2 \sqrt{\frac{\log p}{n}} \right) \leq 4p^{-1}.$$

Proof [Proof of Lemma A.0.2] This Lemma is a generalization/modification of Lemma 5 of Wainwright (2009), allowing us to make usage of the L_2 norm $\|\mathbf{z}\|_2$ to obtain more precise bounds. For self-content we spell out the full details of the proof. Using the spectral theorem, decompose the matrix $(n^{-1} \mathbf{A}^\top \mathbf{A})^{-1} - \mathbb{I} = \mathbf{V} \mathbf{D} \mathbf{V}^\top$, where \mathbf{D} is a diagonal matrix and \mathbf{V} is an independent of \mathbf{D} unitary matrix. Define the random variables:

$$U_i = \mathbf{e}_i^\top \mathbf{V} \mathbf{D} \mathbf{V}^\top \mathbf{z} = z_i \mathbf{v}_i^\top \mathbf{D} \mathbf{v}_i + \mathbf{v}_i^\top \mathbf{D} \sum_{j \neq i} z_j \mathbf{v}_j,$$

where \mathbf{v}_i^\top is the i^{th} row vector of the matrix \mathbf{V} . To bound $\max_i |U_i|$ we deal with these two terms in turn. First notice that $\mathbf{v}_i^\top \mathbf{D} \mathbf{v}_i$ is the i^{th} diagonal entry of the matrix $(n^{-1} \mathbf{A}^\top \mathbf{A})^{-1} - \mathbb{I}$. By the assumption $(\mathbf{A}^\top \mathbf{A})^{-1} \sim \mathcal{W}^{-1}(\mathbb{I}_{s \times s}, n)$ where \mathcal{W}^{-1} is an inverse Wishart distribution. By the properties of the inverse Wishart distribution we conclude that $\mathbf{v}_i^\top \mathbf{D} \mathbf{v}_i \sim n \chi^{-2}(n-s+1) - 1$, where χ^{-2} is the inverse χ^2 distribution. Hence using Lemma 1 of Laurent and Massart (2000) and the union bound we have:

$$\mathbb{P}(\max_i |1 - ((\mathbf{v}_i^\top \mathbf{D} \mathbf{v}_i + 1)(n-s+1)/n)^{-1}| \geq 2\sqrt{y} + 2y) \leq 2s \exp(-(n-s+1)y).$$

Selecting $y = \frac{2 \log p}{n-s+1}$ bounds the above probability by $2s/p^2 \leq 2/p$. For values of $y < 1/32$ we can lump $2\sqrt{y} + 2y < \left(2 + \frac{\sqrt{2}}{4}\right) \sqrt{y}$. Thus inverting the inequality inside the probability we conclude that for each $i \in [s]$:

$$\frac{n}{(n-s+1) \left(1 + \left(2 + \frac{\sqrt{2}}{4}\right) \sqrt{\frac{2 \log p}{n-s+1}}\right)} \leq \mathbf{v}_i^\top \mathbf{D} \mathbf{v}_i + 1 \leq \frac{n}{(n-s+1) \left(1 - \left(2 + \frac{\sqrt{2}}{4}\right) \sqrt{\frac{2 \log p}{n-s+1}}\right)}$$

It is simple to see that when $\frac{\log(p)}{n-s+1} \leq \frac{1}{32}$ the above implies:

$$\max_i |z_i \mathbf{v}_i^\top \mathbf{D} \mathbf{v}_i| < \|\mathbf{z}\|_\infty \tilde{c}_1 \frac{s}{n-s} + \|\mathbf{z}\|_2 \tilde{c}_2 \sqrt{\frac{\log p}{n-s+1}}, \quad (18)$$

where $\tilde{c}_1 = (1 - (1/2 + \sqrt{2}/16))^{-1}$ and $\tilde{c}_2 = (2 + \sqrt{2}/4)\tilde{c}_1$. Next we show that the function $F(\mathbf{v}_i) = \mathbf{v}_i^\top \mathbf{D} \sum_{j \neq i} z_j \mathbf{v}_j$ is Lipschitz with a constant $8\sqrt{s/n}\|\mathbf{z}\|_2$. We have:

$$\|\nabla F\|_2 \leq \|\mathbf{D}\|_{2,2} \left\| \sum_{j \neq i} z_j \mathbf{v}_j \right\|_2 \leq 9\sqrt{\frac{s}{n}} \|\mathbf{z}\|_2,$$

with the last inequality holding with probability at least $1 - 2\exp(-s/2)$ when $\frac{s}{n} \leq \frac{1}{64}$. We used that \mathbf{v}_j are orthonormal, and we bounded the maximum eigenvalue of \mathbf{D} , which follows just as in the proof of Lemma E.0.2 so we omit the details. Since the variables $\{\mathbf{v}_j\}_{j \neq i}$ are uniformly distributed on a $(s-1)$ -dimensional sphere, the proof is completed by invoking the concentration of Lipschitz functions on the sphere to bound $\max_{i \in [s]} |F(\mathbf{v}_i)|$:

$$\mathbb{P}(\max_{i \in [s]} |F(\mathbf{v}_i)| \geq t) \leq 2s \exp\left(-\tilde{c}(s-1) \frac{t^2}{81 \frac{s}{n} \|\mathbf{z}\|_2^2}\right),$$

for an absolute constant \tilde{c} . Under the assumption $s \geq 2$, we can select $t = 18\|\mathbf{z}\|_2 \sqrt{\frac{\log p}{\tilde{c}n}}$ and taking into account that $p > s$ completes the proof, after noticing that we can absorb the second term of (18) to the above expression. \blacksquare

Appendix B. Preliminary Results

Proof [Proof of Lemma 2.1.3] First let $\Sigma = \mathbb{I}_{p \times p}$ (hence by assumption $\|\beta_0\|_2 = 1$) and take any $\mathbf{b} \perp \beta_0$. Note that by the linearity of expectation:

$$\mathbb{E}[\mathbf{X}^\top \beta_0 \mathbf{X}^\top \mathbf{b} | \mathbf{X}^\top \beta_0] = c_b (\mathbf{X}^\top \beta_0)^2.$$

Taking another expectation above we have $\mathbb{E}[\mathbf{X}^\top \beta_0 \mathbf{X}^\top \mathbf{b}] = c_b \mathbb{E}[(\mathbf{X}^\top \beta_0)^2]$. However

$$\mathbb{E}[\mathbf{X}^\top \beta_0 \mathbf{X}^\top \mathbf{b}] = \mathbf{b}^\top \beta_0 = 0,$$

and hence $c_b = 0$. Thus if $\mathbf{b} \perp \beta_0$, $\mathbb{E}[\mathbf{X}^\top \mathbf{b} | \mathbf{X}^\top \beta_0] = 0$. Next, for any $\mathbf{b} \perp \beta_0$ we have:

$$\mathbb{E}[Y \mathbf{X}^\top \mathbf{b}] = \mathbb{E}[\mathbb{E}[Y \mathbf{X}^\top \mathbf{b} | \mathbf{X}^\top \beta_0]] = \mathbb{E}[\mathbb{E}[Y | \mathbf{X}^\top \beta_0] \mathbb{E}[\mathbf{X}^\top \mathbf{b} | \mathbf{X}^\top \beta_0]] = 0.$$

Hence $\mathbb{E}[Y \mathbf{X}] \propto \beta_0$. Finally, a projection on β_0 yields

$$c_0 \|\beta_0\|_2^2 = \mathbb{E}[Y \mathbf{X}^\top \beta_0].$$

This completes the proof in the case when $\Sigma = \mathbb{I}_{p \times p}$. For the more general case observe that $Y = f(\mathbf{X}^\top \beta_0, \varepsilon) = f(\mathbf{X}^\top \Sigma^{-1/2} \Sigma^{1/2} \beta_0, \varepsilon)$, and thus by what we just saw we have:

$$\mathbb{E}[Y \Sigma^{-1/2} \mathbf{X}] = c_0 \Sigma^{1/2} \beta_0,$$

which becomes what we wanted to show after multiplying by $\Sigma^{-1/2}$ on the left. \blacksquare

Appendix C. Lower Bound

For two probability measures P and Q , which are absolutely continuous with respect to a third probability measure μ (i.e. $P, Q \ll \mu$), define their KL divergence by $D_{\text{KL}}(P\|Q) = \int p \log \frac{p}{q} d\mu$, where $p = \frac{dP}{d\mu}$, $q = \frac{dQ}{d\mu}$.

Lemma C.0.1 *Assume conditions required in Proposition 2.3.6. In addition, let for any fixed $u, v \in \mathbb{R}$ and some positive constant Ξ , $f(u, \varepsilon)$ and $f(v, \varepsilon)$ satisfy*

$$D_{\text{KL}}(p(f(u, \varepsilon))\|p(f(v, \varepsilon))) \leq \exp(\Xi(u - v)^2) - 1, \quad (19)$$

Then if

$$n_{p,s} < \frac{1}{8\Xi}, \quad \text{and} \quad s \geq 8\Xi,$$

any algorithm recovering the support of β_0 under (1) will have errors with probability at least $\frac{1}{2}$ asymptotically.

Proof [Proof of Lemma C.0.1] We start by constructing a set of $p - s$ vectors $\mathbf{B} = \{\beta_1, \dots, \beta_{p-s}\}$, belonging to the parameter space $\{\beta \in \mathbb{R}^p : \beta^\top \Sigma \beta = 1, \|\beta\|_0 = s, \frac{\beta^{\min}}{\|\beta\|_\infty} \geq c_\Sigma\}$, for a sufficiently small (to be chosen) $c_\Sigma > 0$, such that $\text{Var}(\mathbf{X}^\top(\beta_k - \beta_j)) \leq \frac{4}{s}$ for all $k, j \in [p - s]$. Once this set is constructed we will use standard Fano type of argument to finish the proof.

Without loss of generality let us assume that $S_0 = [s]$. To construct the set \mathbf{B} , first focus on the sub-matrix Σ_{S_0, S_0} . We take the s dimensional vector $\gamma = a(1/\sqrt{s}, \dots, 1/\sqrt{s}, 0)^\top$ where $a > 0$ is selected so that $\gamma^\top \Sigma_{S_0, S_0} \gamma = \frac{s-1}{s}$. Since Σ_{S_0, S_0} is assumed to have bounded eigenvalues we know that such a indeed exists, and can be chosen in the interval $a \in [\frac{1}{2} \frac{1}{\sqrt{\lambda_{\max}}}, \frac{1}{\sqrt{\lambda_{\min}}}]$ for $s \geq 2$. Next to construct β_k we use:

$$\beta_{kr} = \gamma_r \mathbf{1}(r \in S_0) + b_k \mathbf{1}(r = k + s) / \sqrt{s},$$

where b_k is chosen so that $\beta_k^\top \Sigma \beta_k = 1$. Below we argue that such b_k indeed exists. By Hölder's inequality we have $\|\Sigma_{S_0, S_0} \gamma\|_\infty \leq a \|\Sigma_{S_0, S_0}\|_{\infty, \infty} / \sqrt{s} \leq aR / \sqrt{s}$. Hence using Assumption 2.3.1 we have:

$$\|\Sigma_{S_0^c, S_0} \Sigma_{S_0, S_0}^{-1} \Sigma_{S_0, S_0} \gamma\|_\infty \leq \|\Sigma_{S_0^c, S_0} \Sigma_{S_0, S_0}^{-1}\|_{\infty, \infty} aR / \sqrt{s} \leq (1 - \kappa) aR / \sqrt{s}.$$

Note that due to the last inequality, for any $k \in [p - s]$ we have:

$$\beta_k^\top \Sigma \beta_k = \gamma^\top \Sigma_{S_0, S_0} \gamma + 2 \Sigma_{\{k+s\}, S_0} \gamma \frac{b_k}{\sqrt{s}} + \frac{b_k^2 \Sigma_{k+s, k+s}}{s} \leq \frac{s-1}{s} + 2 \frac{(1-\kappa)a|b_k|R}{s} + \frac{b_k^2 \Sigma_{k+s, k+s}}{s},$$

where we remind the reader that $k+s \in S_0^c$. We chose b_k such that $\text{sign}(b_k) = \text{sign}(\Sigma_{\{k+s\}, S_0} \gamma)$. Hence we also have:

$$\frac{s-1}{s} + \frac{b_k^2 \Sigma_{k+s, k+s}}{s} \leq \beta_k^\top \Sigma \beta_k.$$

Combining the last two inequalities we conclude that there exists:

$$|b_k| \in \left[\frac{\sqrt{(1-\kappa)^2 a^2 R^2 + \Sigma_{k+s, k+s}} - (1-\kappa)aR}{\Sigma_{k+s, k+s}}, \Sigma_{k+s, k+s}^{-1/2} \right],$$

with the desired properties. One can easily check that when:

$$c_{\Sigma} \leq \min \left(\frac{1}{2} \sqrt{\frac{d}{\lambda_{\max}}}, \frac{\sqrt{(1-\kappa)^2 R^2 + d \lambda_{\min}} - (1-\kappa)R}{D} \right), \quad (20)$$

we have that $\min_{j \in [p-s]} \frac{\beta_j^{\min}}{\|\beta_j\|_{\infty}} \geq c_{\Sigma}$, and in addition as we promised we have:

$$\text{Var}(\mathbf{X}^{\top}(\beta_k - \beta_j)) = \frac{\Sigma_{k+s, k+s} b_k^2 - 2\Sigma_{k+s, j+s} b_k b_j + \Sigma_{j+s, j+s} b_j^2}{s} \leq \frac{4}{s}.$$

Next, let \mathcal{J} be a uniform distribution on \mathbf{B} . Under the 0 – 1 loss the risk equals the probability of error:

$$\frac{1}{p-s} \sum_{j \in [p-s]} P_{\beta_j} \{\widehat{S} \neq S(\beta_j)\}, \quad (21)$$

where by P_{β_j} we are measuring the probability under a dataset generated with β_j , and \widehat{S} is an estimate of the true support produced by any (possibly randomized) algorithm. By Fano's inequality that:

$$\mathbb{P}(\text{error}) \geq 1 - \frac{\mathcal{I}(\mathcal{J}; \mathcal{D}) + \log(2)}{\log |\mathbf{B}|}, \quad (22)$$

where $\mathcal{I}(\mathcal{J}; \mathcal{D})$ is the mutual information between the sample \mathcal{J} and the sample \mathcal{D} . Note now that for the mutual information we have

$$\begin{aligned} \mathcal{I}(\mathcal{J}; \mathcal{D}) &= \mathcal{I}(\mathcal{J}; [\{f(\mathbf{X}_j^{\top} \beta_0, \varepsilon_j), \mathbf{X}_j\}, j = 1, \dots, n]) \\ &\leq n\mathcal{H}[\{f(\mathbf{X}^{\top} \beta_0, \varepsilon), \mathbf{X}\}] - n\mathcal{H}[\{f(\mathbf{X}^{\top} \beta_0, \varepsilon), \mathbf{X}\} | \mathcal{J}] \\ &\leq n \max_{k, j \in [p-s]} D_{\text{KL}}(p[\{f(\mathbf{X}^{\top} \beta_k, \varepsilon), \mathbf{X}\}] \| p[\{f(\mathbf{X}^{\top} \beta_j, \varepsilon), \mathbf{X}\}]), \end{aligned}$$

where $\mathcal{H}(\cdot)$ denotes the marginal entropy, $\mathcal{H}(\cdot | \cdot)$ denotes the conditional entropy, D_{KL} denotes the KL divergence, the first inequality follows from the chain inequality of entropy and the second inequality follows from a standard bound. Since the KL divergence is invariant under change of variables, we let $U_k = \mathbf{X}^{\top} \beta_k$ and $\mathbf{W}_{kj} = \mathbf{P}_{\Sigma, \{\beta_k, \beta_j\}^{\perp}} \mathbf{X}$, where $\mathbf{P}_{\Sigma, \{\beta_k, \beta_j\}^{\perp}} \in \mathbb{R}^{(p-2) \times p}$ is chosen such that $\mathbf{P}_{\Sigma, \{\beta_k, \beta_j\}^{\perp}} \Sigma \beta_k = 0$ and $\mathbf{P}_{\Sigma, \{\beta_k, \beta_j\}^{\perp}} \Sigma \beta_j = 0$. Noting that \mathbf{W}_{kj} is independent of U_k, U_j, ε , $U_k - U_j \sim \mathcal{N}(0, V)$ with $V \leq 4/s$, and applying assumption (19), we have

$$\begin{aligned} n^{-1} \mathcal{I}(\mathcal{J}; \mathcal{D}) &\leq \max_{k, j \in [p-s]} D_{\text{KL}}(p[\{f(U_k, \varepsilon), U_k, U_j\}] \| p[\{f(U_j, \varepsilon), U_k, U_j\}]) \\ &\leq \max_{k, j \in [p-s]} \mathbb{E} \exp(\Xi(U_k - U_j)^2) - 1 \leq \sqrt{\frac{s}{s - 8\Xi}} - 1 \leq \frac{8\Xi}{s}. \end{aligned}$$

where the second inequality can be obtained by conditioning on U_k, U_j , and we assume that the value of s is large enough so that $s \geq 16\Xi$. We conclude that:

$$\mathcal{I}(\mathcal{J}; \mathcal{D}) \leq \frac{8\Xi n}{s}.$$

Consequently by (22) if $n_{p,s} < 1/(8\Xi)$ we will have errors with probability at least $\frac{1}{2}$, asymptotically. This is what we wanted to show. \blacksquare

Proof [Proof of Proposition 2.3.6] Note that all moments of the random variable ε exist. Next we verify that condition (19) of Lemma C.0.1 holds in this setup. Since G is 1-1 and KL divergence is invariant under changes of variables WLOG we can assume our model is simply $Y = h(\mathbf{X}^\top \beta_0) + \varepsilon$ or in other words $f(u, \varepsilon) = h(u) + \varepsilon$. This is a location family for $u \in \mathbb{R}$ and thus the normalizing constant of the densities will stay the same regardless of the value of u . Direct calculation yields:

$$D_{\text{KL}}[p\{f(u, \varepsilon)\} \| p\{f(v, \varepsilon)\}] = \mathbb{E}[P((\xi + h(u) - h(v))^2) - P(\xi^2)] = \tilde{P}((h(u) - h(v))^2),$$

where ξ has a density $p_\xi(x) \propto \exp(-P(x^2))$, and \tilde{P} is another non-zero polynomial with nonnegative coefficients, with $\tilde{P}(0) = 0$ of the same degree as P . The last observation follows from the fact that all odd moments of ξ are 0, since ξ is a symmetric about 0 distribution. Since h is L -Lipschitz we conclude that:

$$D_{\text{KL}}[p\{f(u, \varepsilon)\} \| p\{f(v, \varepsilon)\}] \leq \tilde{P}(L^2(u - v)^2).$$

The last can be clearly dominated by $\exp(C(u - v)^2) - 1$ for a large enough constant C . \blacksquare

Appendix D. Covariance Thresholding

Proof [Proof of Proposition 2.2.3] Using the fact that for any two random variables R, T we have $\|RT\|_{\psi_1} \leq 2\|R\|_{\psi_2}\|T\|_{\psi_2}$ we can conclude that the random vector $Y\mathbf{X}$ is coordinate-wise sub-exponentially distributed since $\sup_{j \in [p]} \|YX_j\|_{\psi_1} \leq \mathcal{K} := 2K_Y K$. An application of Proposition 5.16 of Vershynin (2010) and the union bound then gives us that:

$$\mathbb{P}\left(\left\|\frac{1}{n} \sum_{i=1}^n Y_i \mathbf{X}_i - \mathbb{E}[Y\mathbf{X}]\right\|_{\infty} \geq t\right) \leq 2p \exp\left[-\tilde{c} \min\left(\frac{nt^2}{\mathcal{K}^2}, \frac{nt}{\mathcal{K}}\right)\right],$$

where $\tilde{c} > 0$ is some absolute constant. This inequality then implies:

$$\sup_{j \in [p]} \left| \frac{1}{n} \sum_{i=1}^n Y_i X_{ij} - \mathbb{E}(YX_j) \right| \leq \mathcal{K} \sqrt{\frac{2 \log p}{\tilde{c}n}},$$

with probability at least $1 - 2p^{-1}$ for values of n, p such that $\frac{\log p}{n} \leq \frac{\tilde{c}}{2}$. Note that the inequality in the preceding display implies that if:

$$\frac{|c_0|}{\sqrt{s}} > R\mathcal{K} \sqrt{\frac{2 \log p}{\tilde{c}n}},$$

for any $R > 2$ there will be a gap in the absolute values of the coefficients of $U_j = |n^{-1} \sum_{i=1}^n Y_i X_{ij}|$ for $j \in S_0$ and $j \notin S_0$. The latter happens because:

$$\frac{|c_0|}{\sqrt{s}} - \mathcal{K} \sqrt{\frac{2 \log p}{\tilde{c}n}} \geq (R - 1)\mathcal{K} \sqrt{\frac{2 \log p}{\tilde{c}n}} > \mathcal{K} \sqrt{\frac{2 \log p}{\tilde{c}n}}.$$

This also shows that the coefficients will achieve the correct sign. Thus, as long as $\frac{n}{s \log p} \geq \Upsilon$, for $\Upsilon = \frac{2R^2\mathcal{K}^2}{c_0^2c^2}$ signed support recovery happens with asymptotic probability 1. Under our assumption the latter is implied by $n_{p,s} > \Upsilon/\iota$ which completes the proof. \blacksquare

Proof [Proof of Proposition 2.2.1] We follow the same steps as the proof of Proposition 2.2.3. We will use the following Lemma which we justify after the proof:

Lemma D.0.1 *Let us observe n data points from the model described in Proposition 2.2.1 with $\boldsymbol{\beta}$ being an arbitrary unit vector. Then with probability at least $1 - \frac{\eta}{n\sigma^4} - \frac{\gamma}{\log p} - \frac{2}{p}$ the following event holds:*

$$\left\| \frac{1}{n} \sum_{i=1}^n Y_i \mathbf{X}_i - \mathbb{E}(Y \mathbf{X}) \right\|_{\infty} \leq \left(\|\boldsymbol{\beta}_0\|_{\infty} + 2\sqrt{2\sigma^2} \right) \sqrt{\frac{\log p}{n}}.$$

Using the fact that in our case $\mathbb{E}(Y \mathbf{X}) = c_0 \boldsymbol{\beta}_0$, and that $\|\boldsymbol{\beta}_0\|_{\infty} = \frac{1}{\sqrt{s}}$, we have that if:

$$\left(\|\boldsymbol{\beta}_0\|_{\infty} + 2\sqrt{2\sigma^2} \right) \sqrt{\frac{\log p}{n}} \leq (1 + 2\sqrt{2}\sigma) \frac{1}{\sqrt{s}} \sqrt{\frac{s \log p}{n}} < \frac{|c_0|}{2} \frac{1}{\sqrt{s}}$$

there will be a gap between the coefficients corresponding to $j \in S_0 := S(\boldsymbol{\beta}_0)$ and $j \notin S_0$. Note that the last inequality holds if $\frac{s \log p}{n} < \frac{c_0^2}{4(1+2\sqrt{2}\sigma)^2}$. \blacksquare

Remark D.0.2 *The slow convergence in probability rate $(\log p)^{-1}$ observed in Lemma D.0.1 is due to the fact that we are not requiring that Y is sub-Gaussian. If we do require it, the convergence rate of the probability can be seen to reduce to the usual p^{-1} level.*

Proof [Proof of Lemma D.0.1] Note that, sub-exponential concentration bounds do not apply in this case. However, observe that by the properties of the multivariate normal distribution the random variable $(\mathbb{I} - \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^{\top}) \mathbf{X}$ is independent of $\mathbf{X}^{\top} \boldsymbol{\beta}_0$ and hence is independent of Y . Furthermore it is clear that the random variable $Y(\mathbb{I} - \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^{\top}) \mathbf{X}$ has mean 0. Note that conditional on $Y_i, i \in [n]$ we have that $\frac{1}{n} \sum_{i=1}^n Y_i (\mathbb{I} - \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^{\top}) \mathbf{X}_i \sim \mathcal{N}(0, n^{-2} \sum_{i=1}^n Y_i^2 (\mathbb{I} - \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^{\top}))$. Thus by a standard Gaussian tail bound:

$$\mathbb{P} \left(\left\| \frac{1}{n} \sum_{i=1}^n Y_i (\mathbb{I} - \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^{\top}) \mathbf{X}_i \right\|_{\infty} \geq t \mid \mathbf{Y} \right) \leq 2p \exp \left[-\frac{nt^2}{2\bar{Y}^2} \right],$$

where $\bar{Y}^2 = n^{-1} \sum_{i=1}^n Y_i^2$, and we used that $\|\mathbb{I} - \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^{\top}\|_{2,2} \leq 1$. By Chebyshev's inequality $\mathbb{P}(|\bar{Y}^2 - \sigma^2| \geq r) \leq \frac{\eta}{nr^2}$. Hence selecting $r = \sigma^2$ will keep the above probability going to 1 at rate $\frac{1}{n}$ and moreover for large n we have $\bar{Y}^2 \leq 2\sigma^2$. Using this bound in the tail bound above yields that for a choice of $t = 2\sqrt{2\sigma^2 \frac{\log p}{n}}$ the tail bound will go to 0 at rate $2p^{-1}$ as claimed.

Next consider controlling:

$$\mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^n Y_i \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^\top \mathbf{X}_i - c_0 \boldsymbol{\beta}_0\right\|_\infty \geq t\right) = \mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^n Y_i \mathbf{X}_i^\top \boldsymbol{\beta}_0 - c_0\right| \geq t/\|\boldsymbol{\beta}_0\|_\infty\right),$$

where recall that $\mathbb{E}(Y \mathbf{X}) = c_0 \boldsymbol{\beta}_0$, and c_0 is defined in the main text. Applying Chebyshev's inequality once again we get that $t = \|\boldsymbol{\beta}_0\|_\infty \sqrt{\frac{\log p}{n}}$ suffices to keep the above probability going to 0. By the triangle inequality we conclude that, with probability going to 1:

$$\left\|\frac{1}{n}\sum_{i=1}^n Y_i \mathbf{X}_i - \mathbb{E}(Y \mathbf{X})\right\|_\infty \leq \|\boldsymbol{\beta}_0\|_\infty \sqrt{\frac{\log p}{n}} + 2\sqrt{2\sigma^2 \frac{\log p}{n}}.$$

This is what we claimed. ■

Appendix E. LASSO Support Recovery

Proof [Proof of Lemma 3.1.1] Note that since $\mathbf{P}_{\mathbb{X}_{\cdot, S_0}^\perp}$ is an orthogonal projection matrix it contracts length and hence:

$$\left\|\mathbf{P}_{\mathbb{X}_{\cdot, S_0}^\perp} \left(\frac{\mathbf{w}}{\lambda n}\right)\right\|_2 \leq \frac{\|\mathbf{w}\|_2^2}{\lambda^2 n^2}.$$

Next observe that $\mathbf{w} = \mathbf{Y} - c_0 \mathbf{X} \boldsymbol{\beta}_0$ is a vector with non-zero mean. However, by Chebyshev's inequality we have:

$$\mathbb{P}\left(\left|\frac{\|\mathbf{w}\|_2^2}{n} - \xi^2\right| \geq t\right) \leq \frac{\theta^2}{nt^2}.$$

Then setting $t = 1$ brings the above probability to 0 at a rate $\frac{\theta^2}{n}$. Next:

$$n^{-1} \tilde{\mathbf{z}}_{S_0}^\top (n^{-1} \mathbb{X}_{\cdot, S_0}^\top \mathbb{X}_{\cdot, S_0})^{-1} \tilde{\mathbf{z}}_{S_0} \leq \frac{1}{\lambda_{\min}(1 - 2\sqrt{\frac{s}{n}})^2} \frac{\|\tilde{\mathbf{z}}_{S_0}\|_2^2}{n} \leq \frac{1}{\lambda_{\min}(1 - 2\sqrt{\frac{s}{n}})^2} \frac{s}{n},$$

with probability at least $1 - 2\exp(-s/2)$, where we used Lemma A.0.1. This completes the proof. ■

Proof [Proof of Lemma 3.2.2] First, we note the following decomposition:

$$[\mathbb{X}_{\cdot, S_0}^\top \mathbb{X}_{\cdot, S_0}]^{-1} \mathbb{X}_{\cdot, S_0}^\top \mathbf{Y} - c_0 \boldsymbol{\beta}_{0S_0} = (n[\mathbb{X}_{\cdot, S_0}^\top \mathbb{X}_{\cdot, S_0}]^{-1} - \mathbb{I}) n^{-1} \mathbb{X}_{\cdot, S_0}^\top \mathbf{Y} + (n^{-1} \mathbb{X}_{\cdot, S_0}^\top \mathbf{Y} - c_0 \boldsymbol{\beta}_{0S_0}).$$

Note that the second term is mean 0. Applying Lemma D.0.1 gives us a bound on the second term. We next move on to consider the first term.

Consider a ‘‘symmetrization’’ transformation of the predictor matrix $\tilde{\mathbb{X}}_{\cdot, S_0}^\top = (\mathbb{I} - \boldsymbol{\beta}_{0S_0} \boldsymbol{\beta}_{0S_0}^\top) \mathbb{X}_{\cdot, S_0}^\top + \boldsymbol{\beta}_{0S_0} \boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{\cdot, S_0}^*$, where $[\mathbb{X}_{\cdot, S_0}^*]_{n \times s}$ is an i.i.d. copy of \mathbb{X}_{\cdot, S_0} , or in other words the columns of $\mathbb{X}_{\cdot, S_0}^*$: $\mathbf{X}_j^* \sim \mathcal{N}(0, \mathbb{I}_{n \times n})$, $j = 1, \dots, s$ and are independent of \mathbb{X}_{\cdot, S_0} . Note that in doing this

construction, we guarantee that $\tilde{\mathbb{X}}_{,S_0}$ is independent of $\mathbb{X}_{,S_0}^\top \boldsymbol{\beta}_{0S_0}$. Now we further decompose the first term as follows:

$$\begin{aligned}
 (n[\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} - \mathbb{I})n^{-1}\mathbb{X}_{,S_0}^\top \mathbf{Y} &= \underbrace{(n[\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} - \mathbb{I})\boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top n^{-1}\mathbb{X}_{,S_0}^\top \mathbf{Y}}_{\mathbf{I}_1} \\
 &+ \underbrace{n([\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} - [\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1})(\mathbb{I} - \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top)n^{-1}\mathbb{X}_{,S_0}^\top \mathbf{Y}}_{\mathbf{I}_2} \\
 &+ \underbrace{(n[\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} - \mathbb{I})n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbf{Y}}_{\mathbf{I}_3} \\
 &- \underbrace{(n[\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} - \mathbb{I})\boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top n^{-1}\mathbb{X}_{,S_0}^{*\top} \mathbf{Y}}_{\mathbf{I}_4}.
 \end{aligned}$$

We next deal with each of these terms separately. For the first and last terms we can directly apply Lemma A.0.2. Under the same event as in Lemma E.0.1, taking into account that $\|\boldsymbol{\beta}_{0S_0}\|_2 = 1$ we have that $\|([n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} - \mathbb{I})\boldsymbol{\beta}_{0S_0}\|_\infty \leq C_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + C_2 \sqrt{\frac{\log p}{n}}$ and $\|([n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} - \mathbb{I})\boldsymbol{\beta}_{0S_0}\|_\infty \leq C_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + C_2 \sqrt{\frac{\log p}{n}}$. Furthermore, $\boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{,S_0}^\top \mathbf{Y}/n$ is a mean c_0 random variable. Just as in the proof of Lemma D.0.1 by Chebyshev's inequality we have that with probability at least $1 - \frac{\gamma}{\log p}$ we have $|\boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{,S_0}^\top \mathbf{Y}/n| \leq |c_0| + \sqrt{\frac{\log p}{n}}$.

Furthermore, notice that $n^{-1}\boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{,S_0}^{*\top} \mathbf{Y}$ is a mean 0 random variable. Conditionally on \mathbf{Y} it has a $\mathcal{N}(0, n^{-2} \sum Y_i^2)$ distribution. With exactly the same argument as in the proof of Lemma D.0.1 we conclude that with probability at least $1 - \frac{\eta}{n\sigma^4} - \frac{2}{p}$:

$$|n^{-1}\boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{,S_0}^{*\top} \mathbf{Y}| \leq 2\sqrt{\sigma^2 \frac{\log p}{n}}.$$

Hence, combining the above results we obtain:

$$\|\mathbf{I}_1\|_\infty + \|\mathbf{I}_4\|_\infty \leq \left(C_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + C_2 \sqrt{\frac{\log p}{n}}\right) \left(|c_0| + \sqrt{\frac{\log p}{n}} + 2\sqrt{\sigma^2 \frac{\log p}{n}}\right). \quad (23)$$

To deal with the term \mathbf{I}_2 first note that by Hölder's inequality we have:

$$\|\mathbf{I}_2\|_\infty \leq \|n([\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} - [\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1})\|_{\infty, \infty} \|(\mathbb{I} - \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top)n^{-1}\mathbb{X}_{,S_0}^\top \mathbf{Y}\|_\infty. \quad (24)$$

To deal with the first term we make usage of the following result:

Lemma E.0.1 *Suppose that s, n satisfy $\frac{s}{n} \leq \frac{1}{16}$. The following bound holds:*

$$\|([n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} - [n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1})\|_{\infty, \infty} \leq 40\sqrt{s} \left(C_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + \tilde{C}_2 \sqrt{\frac{\log p}{n}}\right),$$

with probability at least $1 - 4 \exp(-s/2) - \frac{12}{p} - \frac{4}{n}$, where $C_1 > 0$ and $\tilde{C}_2 = C_2 + 4$ are the same constants as in (23).

Note that the second term is a mean 0 random variable since $(\mathbb{I} - \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top)\mathbb{X}_{,S_0}$ is independent of \mathbf{Y} . Just as in Lemma D.0.1 we can show that $\|(\mathbb{I} - \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top)n^{-1}\mathbb{X}_{,S_0}^\top\mathbf{Y}\|_\infty \leq 2\sqrt{2\sigma^2\frac{\log p}{n}}$ with probability at least $1 - \frac{2s}{p^2} \geq 1 - \frac{2}{p}$ (this event is in fact a sub-event of the bounds of the first term $n^{-1}\mathbb{X}_{,S_0}^\top\mathbf{Y} - c_0\boldsymbol{\beta}_{0S_0}$). Lemma E.0.1 gives us a bound on $\|n([\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1} - [\tilde{\mathbb{X}}_{,S_0}^\top\tilde{\mathbb{X}}_{,S_0}]^{-1})\|_{\infty,\infty}$ which in conjunction with the previous inequality suffices to control the term \mathbf{I}_2 .

Finally, to deal with the term \mathbf{I}_4 we will make use of the following:

Lemma E.0.2 *Let $\frac{s}{n} \leq \frac{1}{64}$. Then there exists a constant $\Upsilon \asymp \sigma > 0$, such that the term:*

$$\|(n[\tilde{\mathbb{X}}_{,S_0}^\top\tilde{\mathbb{X}}_{,S_0}]^{-1} - \mathbb{I})n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top\mathbf{Y}\|_\infty \leq \Upsilon\sqrt{\frac{\log p}{n}},$$

with probability at least $1 - \frac{2}{p} - \frac{\eta}{n\sigma^4} - 2\exp(-s/2)$.

Applying Lemma E.0.2 we have in conjunction with our previous bounds (23) and (24) we get:

$$\begin{aligned} \|[\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}\mathbb{X}_{,S_0}^\top\mathbf{Y} - c_0\boldsymbol{\beta}_{0S_0}\|_\infty &\leq \left(C_1\frac{s}{n}\|\boldsymbol{\beta}_{0S_0}\|_\infty + C_2\sqrt{\frac{\log p}{n}}\right) \left(|c_0| + \sqrt{\frac{\log p}{n}} + 2\sqrt{\sigma^2\frac{\log p}{n}}\right) \\ &\quad + 80\sqrt{s} \left(C_1\frac{s}{n}\|\boldsymbol{\beta}_{0S_0}\|_\infty + \tilde{C}_2\sqrt{\frac{\log p}{n}}\right) \sqrt{2\sigma^2\frac{\log p}{n}} \\ &\quad + \Upsilon\sqrt{\frac{\log p}{n}} + \|\boldsymbol{\beta}_{0S_0}\|_\infty\sqrt{\frac{\log p}{n}} + 2\sqrt{\sigma^2\frac{\log p}{n}}, \end{aligned}$$

with probability at least $1 - 4\exp(-s/2) - \frac{4}{n} - \frac{18}{p} - 2\frac{\eta}{n\sigma^4} - 2\frac{\gamma}{\log p}$ ⁵, which finishes the proof, after grouping terms and recalling the fact that $\log(p-s) \asymp \log p$. \blacksquare

Proof [Proof of Lemma E.0.1] We first compare $[n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}$ to $[n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top\mathbb{X}_{,S_0} + \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top]^{-1}$. The latter matrix might happen to be non-invertible but this is irrelevant for our proof as we argue below. Using Woodbury's matrix identity we have:

$$[n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top\mathbb{X}_{,S_0} + \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top]^{-1} - [n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1} = \frac{[n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}\boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top\mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}}{1 - \boldsymbol{\beta}_{0S_0}^\top\mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}\boldsymbol{\beta}_{0S_0}},$$

where $\mathbf{M} = n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0} - \mathbb{I} - n^{-1}\mathbb{X}_{,S_0}^{*\top}\mathbb{X}_{,S_0}$. Note that whenever the right hand side of Woodbury's identity is well defined, the matrix $n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top\mathbb{X}_{,S_0} + \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top$ is indeed invertible, and the inverse satisfies the above identity. As we argue below the right hand side is well defined (i.e. the denominator is non-zero) with high probability hence the proof goes through. Next we handle the term $\boldsymbol{\beta}_{0S_0}^\top\mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}$. By the triangle inequality have:

$$\|\boldsymbol{\beta}_{0S_0}^\top\mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}\|_\infty \leq \|\boldsymbol{\beta}_{0S_0}^\top([n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1} - \mathbb{I})\|_\infty + \|\boldsymbol{\beta}_{0S_0}^\top n^{-1}\mathbb{X}_{,S_0}^{*\top}\mathbb{X}_{,S_0}[n^{-1}\mathbb{X}_{,S_0}^\top\mathbb{X}_{,S_0}]^{-1}\|_\infty.$$

5. Here we are recognizing the fact that the events of some probability bounds we derived above, in fact coincide.

For the first term Lemma A.0.2 is directly applicable. Applying this lemma gives us the existence of constants C_1 and C_2 such that:

$$\|([n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} - \mathbb{I})\boldsymbol{\beta}_{0S_0}\|_\infty \leq C_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + C_2 \sqrt{\frac{\log p}{n}},$$

with probability at least $1 - 4p^{-1}$. For the second term, we have that conditionally on $\mathbb{X}_{,S_0}$ it has a normal distribution: $\mathcal{N}(0, n^{-1}(n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0})^{-1})$. Since $\mathbb{X}_{,S_0}$ is standard normal, we can apply Lemma A.0.1 to claim that $\|n(\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0})^{-1}\|_{2,2} \leq \left(\frac{1}{1-\sqrt{\frac{s}{n}}-t}\right)^2$ with probability at least $1 - 2\exp(-nt^2/2)$. Taking $t = \sqrt{\frac{s}{n}}$ gives us that $\|n(\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0})^{-1}\|_{2,2} \leq \frac{1}{(1-2\sqrt{\frac{s}{n}})^2}$ with probability at least $1 - 2\exp(-s/2)$. Thus conditioning on this event, by a standard normal tail bound and a union bound we have:

$$\mathbb{P}(\|\boldsymbol{\beta}_{0S_0}^\top n^{-1}\mathbb{X}_{,S_0}^{*\top} \mathbb{X}_{,S_0} [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}\|_\infty \geq t) \leq 2s \exp\left(-t^2 n \left(1 - 2\sqrt{\frac{s}{n}}\right)^2 / 2\right).$$

Selecting $t = 4\sqrt{\frac{\log p}{n}}$, we get the probability above is bounded by $\frac{2s}{p^2} \leq \frac{2}{p}$ (where we used the assumption $\sqrt{\frac{s}{n}} \leq \frac{1}{4}$). So finally on the intersection event we have:

$$\|\boldsymbol{\beta}_{0S_0}^\top \mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}\|_\infty \leq C_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + \tilde{C}_2 \sqrt{\frac{\log p}{n}},$$

with probability at least $1 - 6p^{-1} - 2\exp(-s/2)$ where $\tilde{C}_2 = C_2 + 4$. Let us now consider the denominator:

$$\begin{aligned} & 1 - \boldsymbol{\beta}_{0S_0}^\top \mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0} \\ &= 1 - \boldsymbol{\beta}_{0S_0}^\top (\mathbb{I} - [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}) \boldsymbol{\beta}_{0S_0} + n^{-1} \boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{,S_0}^{*\top} \mathbb{X}_{,S_0} [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0} \\ &= \boldsymbol{\beta}_{0S_0}^\top [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0} + n^{-1} \boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{,S_0}^{*\top} \mathbb{X}_{,S_0} [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0}. \end{aligned}$$

Using Lemma A.0.1 we have $\lambda_{\min}([n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}) \geq \frac{1}{(1+2\sqrt{\frac{s}{n}})^2} > \frac{1}{4}$ with the last bound holding since $\frac{s}{n} < \frac{1}{4}$. Hence $\boldsymbol{\beta}_{0S_0}^\top [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0} \geq \frac{1}{4}$. For the second term just as before, conditionally on $\mathbb{X}_{,S_0}$ we have

$$n^{-1} \boldsymbol{\beta}_{0S_0}^\top \mathbb{X}_{,S_0}^{*\top} \mathbb{X}_{,S_0} [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0} \sim \mathcal{N}(0, n^{-1} \boldsymbol{\beta}_{0S_0}^\top [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0}).$$

Then (given that $\|n(\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0})^{-1}\|_{2,2} \leq \frac{1}{(1-2\sqrt{\frac{s}{n}})^2}$) by a standard tail bound we have that the second term is $\leq 4\sqrt{\frac{\log n}{n}}$ with probability at least $1 - \frac{2}{n}$. Putting everything together we have:

$$1 - \boldsymbol{\beta}_{0S_0}^\top \mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0} \geq \frac{1}{4} - 4\sqrt{\frac{\log n}{n}}.$$

The last expression is clearly bigger than $\frac{1}{5}$ for large enough values of n . Hence we conclude that with high probability we have:

$$\begin{aligned} & \| [n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbb{X}_{,S_0} + \boldsymbol{\beta}_{0S_0} \boldsymbol{\beta}_{0S_0}^\top]^{-1} - [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \|_{\infty, \infty} \\ & \leq 5 \| [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \boldsymbol{\beta}_{0S_0} \|_1 \| \boldsymbol{\beta}_{0S_0}^\top \mathbf{M}[n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1} \|_\infty \end{aligned}$$

For the first term, by the definition of matrix $\|\cdot\|_{2,2}$ norm we further have:

$$\begin{aligned} \|[n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}\boldsymbol{\beta}_{0S_0}\|_1 &\leq \sqrt{s}\|[n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}\boldsymbol{\beta}_{0S_0}\|_2 \leq \sqrt{s}\|\boldsymbol{\beta}_{0S_0}\|_2\|[n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}\|_{2,2} \\ &\leq \frac{\sqrt{s}}{(1-2\sqrt{\frac{s}{n}})^2}. \end{aligned}$$

Combining this inequality with our previous bound we get:

$$\|[n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbb{X}_{,S_0} + \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top]^{-1} - [n^{-1}\mathbb{X}_{,S_0}^\top \mathbb{X}_{,S_0}]^{-1}\|_{\infty,\infty} \leq \frac{5\sqrt{s}}{(1-2\sqrt{\frac{s}{n}})^2} \left(C_1 \frac{s}{n} \|\boldsymbol{\beta}_{0S_0}\|_\infty + \tilde{C}_2 \sqrt{\frac{\log p}{n}} \right).$$

Next we show that $[n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbb{X}_{,S_0} + \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top]^{-1}$ is also close to $[n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1}$. Another usage of Woodbury's matrix identity yields:

$$[n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbb{X}_{,S_0} + \boldsymbol{\beta}_{0S_0}\boldsymbol{\beta}_{0S_0}^\top]^{-1} - [n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} = \frac{[n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} \tilde{\mathbf{M}} \boldsymbol{\beta}_{0S_0} \boldsymbol{\beta}_{0S_0}^\top [n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1}}{1 - \boldsymbol{\beta}_{0S_0}^\top [n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} \tilde{\mathbf{M}} \boldsymbol{\beta}_{0S_0}},$$

where $\tilde{\mathbf{M}} = n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0} - \mathbb{I} - n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbb{X}_{,S_0}$. Note that since $\tilde{\mathbb{X}}_{,S_0}^\top \perp \mathbb{X}_{,S_0} \boldsymbol{\beta}_{0S_0}$, the same argument as before goes through. Combining the bounds with a triangle inequality completes the proof, using the fact that $\sqrt{\frac{s}{n}} \leq \frac{1}{4}$. \blacksquare

Proof [Proof of Lemma E.0.2] We first perform a singular value decomposition on the $\tilde{\mathbb{X}}_{,S_0} = \mathbf{U}_{n \times s} \mathbf{D}_{s \times s} \mathbf{V}_{s \times s}^\top$ matrix. Note that since multiplying $\tilde{\mathbb{X}}_{,S_0}$ by a unitary $s \times s$ matrix on the right, or by a unitary $n \times n$ matrix on the left doesn't change the distribution of $\tilde{\mathbb{X}}_{,S_0}$ we conclude that the matrices \mathbf{U} , \mathbf{D} and \mathbf{V} are independent. This representation gives us that $(n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0})^{-1} - \mathbb{I} = \mathbf{V}(n\mathbf{D}^{-2} - \mathbb{I})\mathbf{V}^\top$. With this notation we can rewrite:

$$(n[\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} - \mathbb{I})n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbf{Y} = \mathbf{V} \underbrace{(n\mathbf{D}^{-2} - \mathbb{I})n^{-1/2}\mathbf{D}}_{\mathbf{W}} n^{-1/2} \mathbf{U}^\top \mathbf{Y}.$$

We recall that by construction $\tilde{\mathbb{X}}_{,S_0}$ is independent of \mathbf{Y} . The elements of the matrix \mathbf{W} can be bounded in a simple manner. We have $\|\mathbf{W}\|_{2,2} \leq \|(n\mathbf{D}^{-2} - \mathbb{I})\|_{2,2} \|n^{-1/2}\mathbf{D}\|_{2,2}$, and by Lemma A.0.1, as before we have: $\|(n\mathbf{D}^{-2} - \mathbb{I})\|_{2,2} \leq \frac{1}{(1-2\sqrt{\frac{s}{n}})^2} - 1 \leq \frac{4\sqrt{\frac{s}{n}}}{(1-2\sqrt{\frac{s}{n}})^2}$ and $\|n^{-1/2}\mathbf{D}\|_{2,2} \leq 1 + 2\sqrt{\frac{s}{n}}$ with probability at least $1 - 2\exp(-s/2)$. We will condition on the event $\|\mathbf{W}\|_{2,2} \leq \frac{4\sqrt{\frac{s}{n}}}{(1-2\sqrt{\frac{s}{n}})^2} (1 + 2\sqrt{\frac{s}{n}}) < 9\sqrt{\frac{s}{n}}$, with the last inequality holding for $\sqrt{\frac{s}{n}} \leq \frac{1}{8}$. Since every random variable in the above display is independent from \mathbf{W} , the distributions of \mathbf{V} , \mathbf{U} and \mathbf{Y} stay unchanged under this conditioning. Let \mathbf{e}_i be a unit vector with 1 on the i^{th} position. Since we are interested in bounding the $\|\cdot\|_\infty$ we will start with the following:

$$\mathbf{e}_i^\top (n[\tilde{\mathbb{X}}_{,S_0}^\top \tilde{\mathbb{X}}_{,S_0}]^{-1} - \mathbb{I})n^{-1}\tilde{\mathbb{X}}_{,S_0}^\top \mathbf{Y} = \mathbf{v}_i^\top \mathbf{W} [n^{-1/2} \mathbf{U}^\top \mathbf{Y}],$$

where \mathbf{v}_i^\top is the i^{th} row of the matrix \mathbf{V} . Condition on the vector $n^{-1/2} \mathbf{U}^\top \mathbf{Y}$. Since \mathbf{v}_i is independent of $n^{-1/2} \mathbf{U}^\top \mathbf{Y}$ it follows that the distribution of \mathbf{v}_i is uniform on the unit

sphere in \mathbb{R}^s . We next show that the function $F(\mathbf{v}_i) = \mathbf{v}_i^\top \mathbf{W} [n^{-1/2} \mathbf{U}^\top \mathbf{Y}]$ is Lipschitz. We have:

$$\begin{aligned} \|\nabla F\|_2 &\leq \|\mathbf{W}\|_{2,2} \|n^{-1/2} \mathbf{U}^\top \mathbf{Y}\|_2 \leq 9 \sqrt{\frac{s}{n}} n^{-1/2} \sqrt{\sum_{i=1}^s (\mathbf{u}_i^\top \mathbf{Y})^2} \\ &\leq 9 \sqrt{\frac{s}{n}} n^{-1/2} \|\mathbf{Y}\|_2, \end{aligned}$$

where the last inequality follows from the fact that the vectors \mathbf{u}_i are orthonormal and hence $\sum_{i=1}^s (\mathbf{u}_i^\top \mathbf{Y})^2 \leq \|\mathbf{Y}\|_2^2$. Since Y_i are assumed to have finite second moment, by Chebyshev's inequality we have that:

$$\mathbb{P}(|n^{-1} \|\mathbf{Y}\|_2^2 - \sigma^2| \geq t) \leq \frac{\eta}{nt^2}.$$

Selecting $t = \sigma^2$ is sufficient to keep the above probability going to 0, and furthermore for n large enough guarantees that $n^{-1} \|\mathbf{Y}\|_2^2 \leq 2\sigma^2$ and hence $n^{-1/2} \|\mathbf{Y}\|_2 \leq \sqrt{2}\sigma$. Thus conditional on this event the function F is Lipschitz with a constant equal to $\sqrt{2}9\sigma\sqrt{\frac{s}{n}}$. Since the expectation of the function F is 0, by concentration of measure for Lipschitz functions on the sphere (Ledoux, 2005; Ledoux and Talagrand, 2013), for any $t > 0$ we have:

$$\mathbb{P}(|F(\mathbf{v}_i)| \geq t\sigma) \leq 2 \exp\left(-\tilde{c}s \frac{t^2}{162 \frac{s}{n}}\right),$$

for some absolute constant $\tilde{c} > 0$. Taking a union bound the above becomes:

$$\mathbb{P}(\max_{i \in [s]} |F(\mathbf{v}_i)| \geq t\sigma) \leq 2s \exp\left(-\tilde{c} \frac{t^2 n}{162}\right).$$

Selecting $t = 18\sqrt{\frac{\log p}{cn}}$, keeps the probability vanishing at a rate faster than $2s/p^2 \leq 2/p$ and completes the proof. ■

Proof [Proof of Corollary 2.3.9] Tracing the proof of Theorem 2.3.4 we realize that it suffices to show the following two quantities remain well controlled under the usage of \hat{g} :

- i. $|n^{-1} \sum_{i=1}^n \mathbf{X}_i^\top \beta_0 \hat{g}(Y_i) - c_0| \leq O(\sqrt{\log p/n})$,
- ii. $n^{-1} \sum_{i=1}^n \hat{g}^2(Y_i) = O(1)$,

with probability at least $1 - O(p^{-1})$ and $1 - O(n^{-1})$ correspondingly. To deal with i. observe that:

$$\begin{aligned} \left| n^{-1} \sum_{i=1}^n \mathbf{X}_i^\top \boldsymbol{\beta}_0 \widehat{g}(Y_i) - c_0 \right| &\leq \left| n^{-1} \sum_{i=1}^n \mathbf{X}_i^\top \boldsymbol{\beta}_0 g(Y_i) - c_0 \right| + \left| n^{-1} \sum_{i=1}^n \mathbf{X}_i^\top \boldsymbol{\beta}_0 (g(Y_i) - \widehat{g}(Y_i)) \right| \\ &\leq \underbrace{\left| n^{-1} \sum_{i=1}^n \mathbf{X}_i^\top \boldsymbol{\beta}_0 g(Y_i) - c_0 \right|}_{I_1} \\ &\quad + \underbrace{\left\{ n^{-1} \sum_{i=1}^n (\mathbf{X}_i^\top \boldsymbol{\beta}_0)^2 \right\}^{1/2} \left\{ n^{-1} \sum_{i=1}^n (\widehat{g}(Y_i) - g(Y_i))^2 \right\}^{1/2}}_{I_2}. \end{aligned}$$

The term I_1 remains controlled by the proof of Theorem 2.3.4, while for the term I_2 we have:

$$I_2 \leq O(\sqrt{\log p/n}),$$

with probability at least $1 - O(p^{-1})$, where we used the assumption on \widehat{g} and the fact that the random variables $(\mathbf{X}_i^\top \boldsymbol{\beta}_0)^2 \sim \chi_1^2$ and hence concentrate exponentially about their mean — 1, by a standard tail bound (Boucheron et al., 2013).

Next, for ii., by the triangle inequality we have:

$$\sqrt{n^{-1} \sum_{i=1}^n \widehat{g}^2(Y_i)} \leq \sqrt{n^{-1} \sum_{i=1}^n g^2(Y_i)} + \sqrt{n^{-1} \sum_{i=1}^n (\widehat{g}(Y_i) - g(Y_i))^2}.$$

The first term is well controlled as before and is $O(1)$ with probability at least $1 - O(n^{-1})$ and the second term is at most $O(\sqrt{\log p/n})$ with probability at least $1 - O(p^{-1})$ by assumption which concludes the proof. ■

Proof [Proof of Proposition 2.3.8] First let g be such that $\mathbb{E}\{g(Y)\mathbf{X}^\top \boldsymbol{\beta}_0\} \neq 0$. Recall that $\mathbb{E}(\mathbf{X}^\top \boldsymbol{\beta}_0) = 0$. Hence by Cauchy-Schwartz we have:

$$0 < [\mathbb{E}\{g(Y)\mathbf{X}^\top \boldsymbol{\beta}_0\}]^2 = (\mathbb{E}[g(Y)\mathbb{E}\{\mathbf{X}^\top \boldsymbol{\beta}_0|Y\}])^2 \leq \text{Var}\{g(Y)\} \text{Var}\{\mathbb{E}\{\mathbf{X}^\top \boldsymbol{\beta}_0|Y\}\},$$

and therefore $\text{Var}\{\mathbb{E}\{\mathbf{X}^\top \boldsymbol{\beta}_0|Y\}\} > 0$. In the reverse case put $g(Y) = \mathbb{E}\{\mathbf{X}^\top \boldsymbol{\beta}_0|Y\}$ and apply conditional expectation to obtain $\mathbb{E}\{g(Y)\mathbf{X}^\top \boldsymbol{\beta}_0\} = \text{Var}\{\mathbb{E}\{\mathbf{X}^\top \boldsymbol{\beta}_0|Y\}\} > 0$. ■

Appendix F. Additional Simulation Results

F.1 $\Sigma = \mathbb{I}$

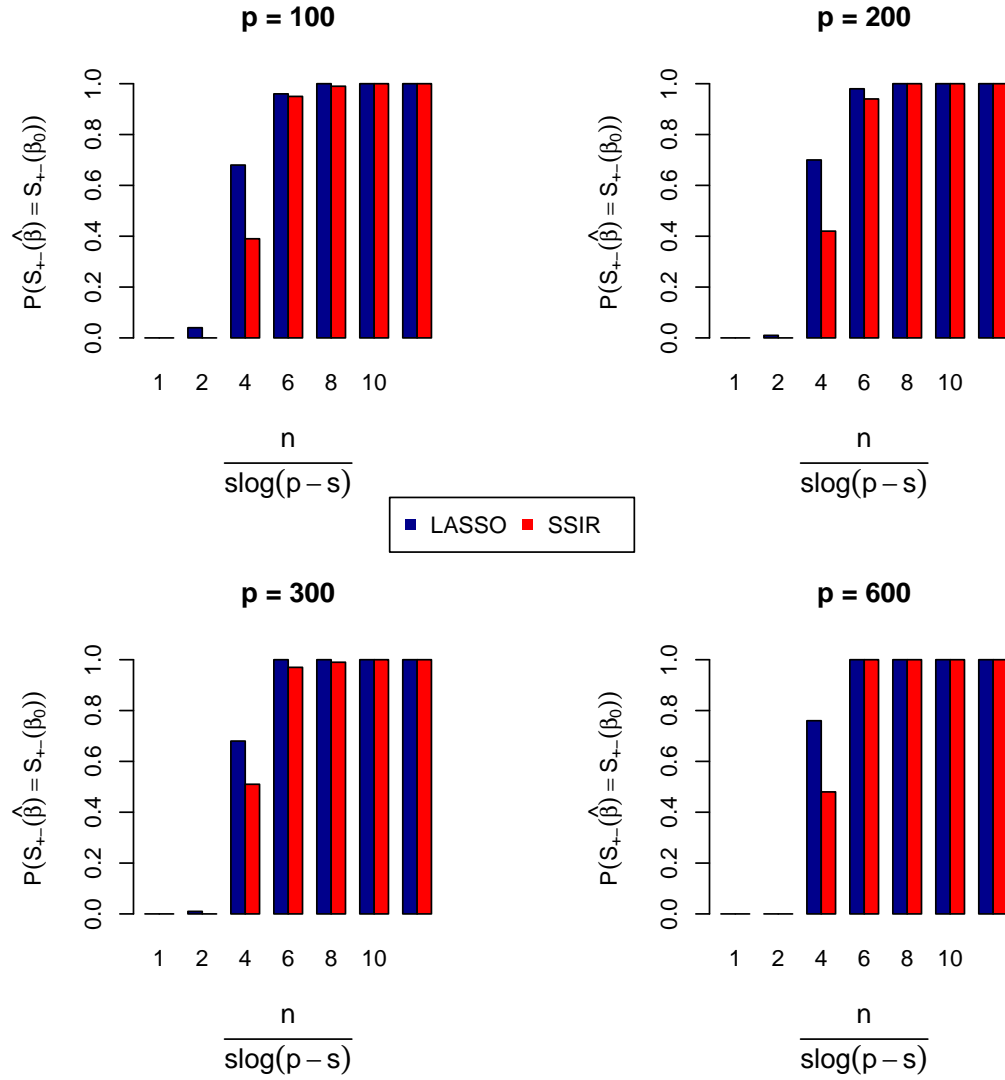


Figure 3: Model (14)

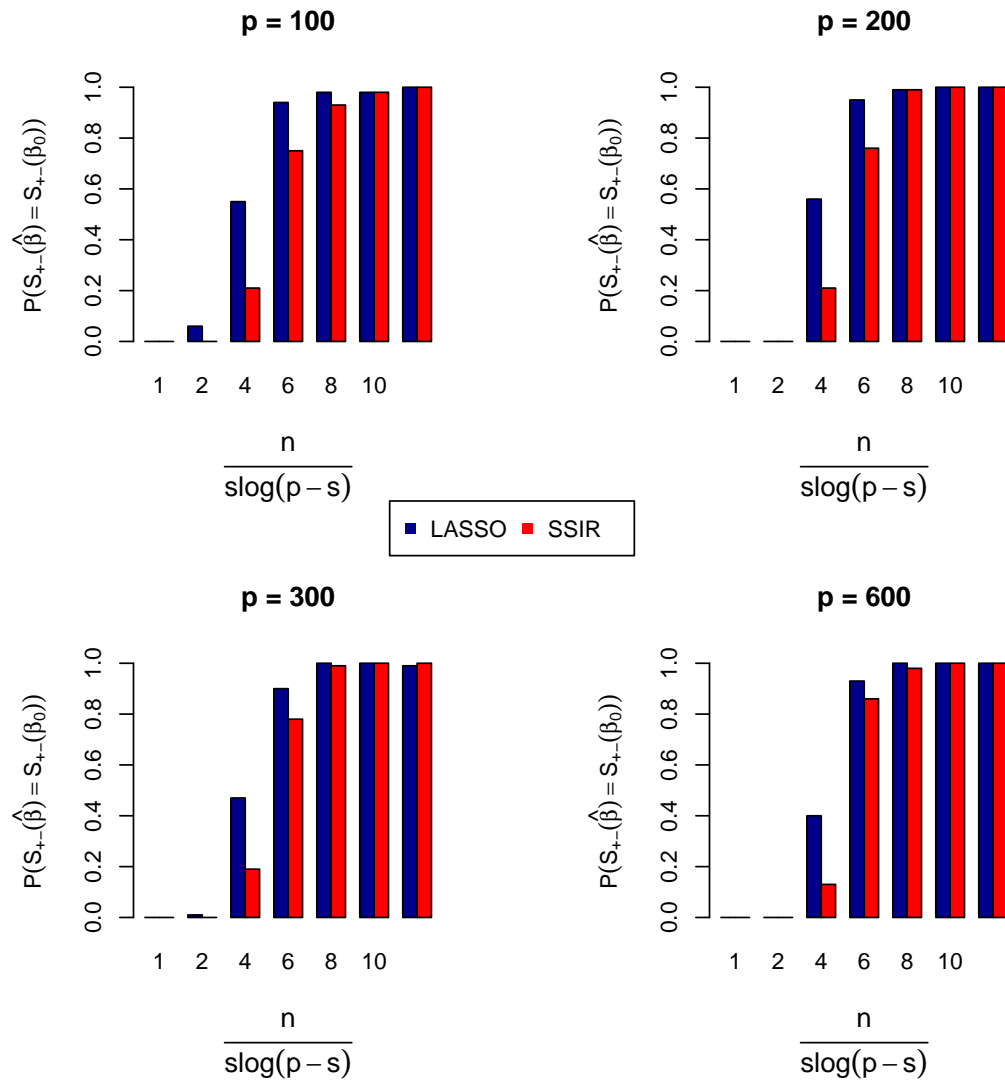


Figure 4: Model (17)

F.2 $\Sigma : \Sigma_{kj} = 2^{-|k-j|}$

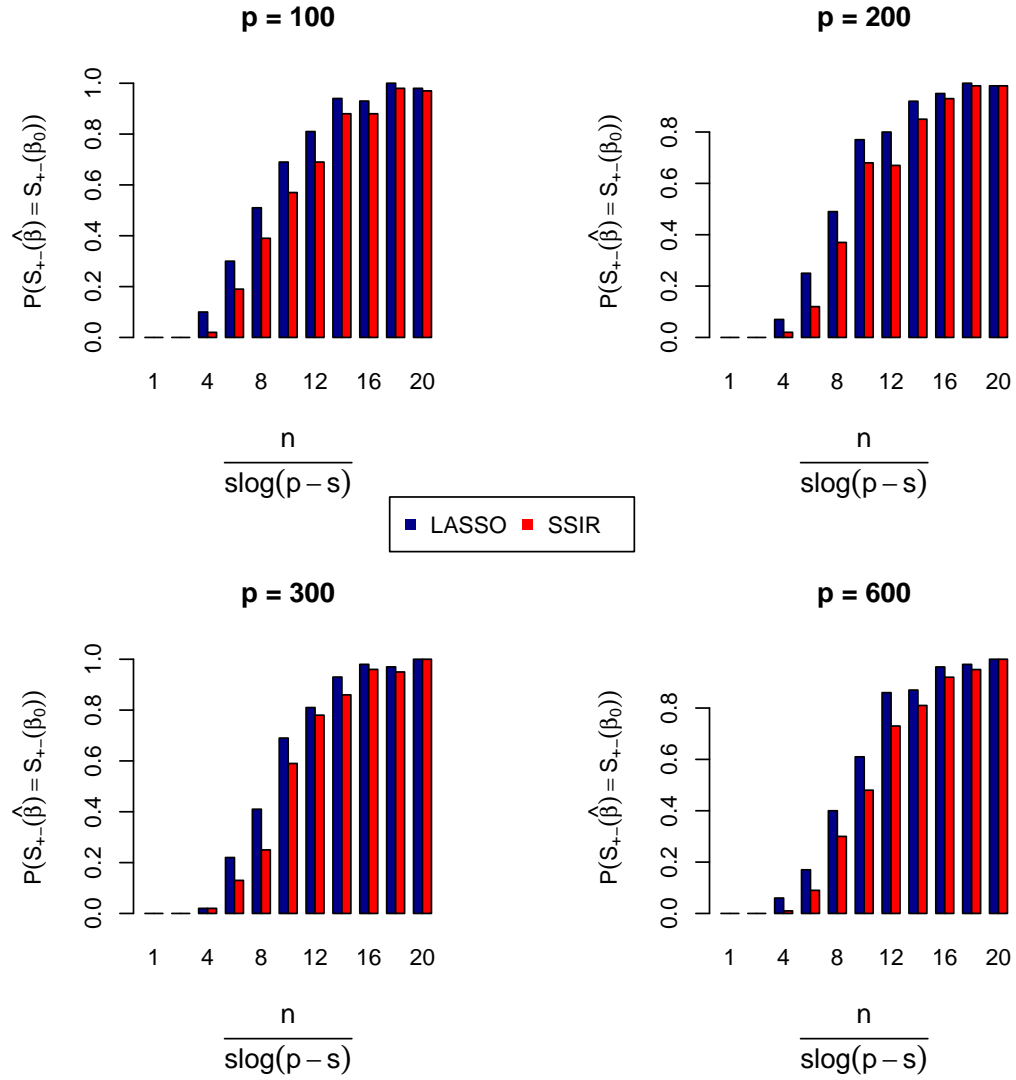


Figure 5: Model (14)

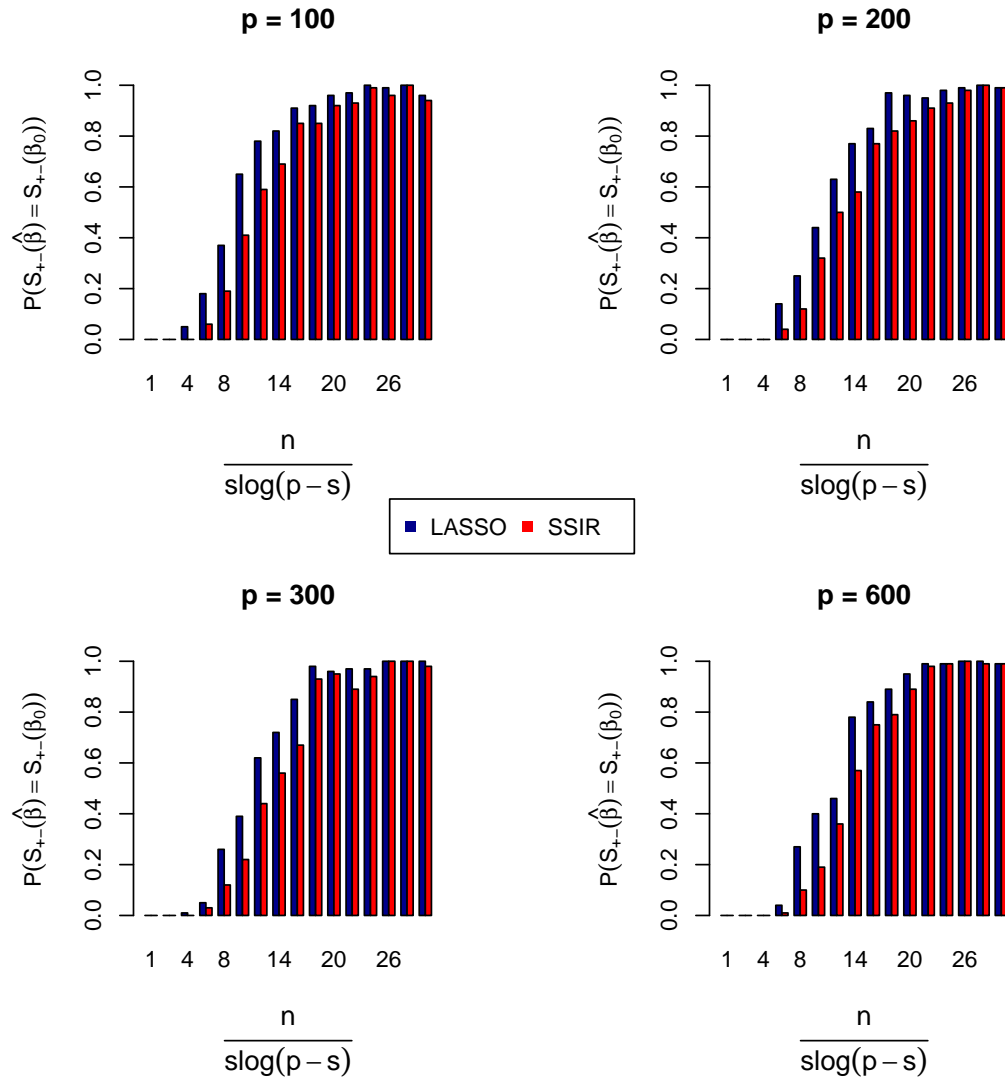


Figure 6: Model (17)

References

- Pierre Alquier and Gérard Biau. Sparse single-index model. *The Journal of Machine Learning Research*, 14(1):243–280, 2013.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. OUP Oxford, 2013.
- Tony Cai, Weidong Liu, and Xi Luo. A constrained l1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106(494):594–607, 2011.
- Stamatis Cambanis, Steel Huang, and Gordon Simons. On the theory of elliptically contoured distributions. *Journal of Multivariate Analysis*, 11(3):368–385, 1981.
- Emmanuel Candes and Terence Tao. The dantzig selector: statistical estimation when p is much larger than n . *The Annals of Statistics*, pages 2313–2351, 2007.
- Rita M Cantor, Kenneth Lange, and Janet S Sinsheimer. Prioritizing gwas results: a review of statistical methods and recommendations for their application. *The American Journal of Human Genetics*, 86(1):6–22, 2010.
- R Dennis Cook and Liqiang Ni. Sufficient dimension reduction via inverse regression. *Journal of the American Statistical Association*, 100(470), 2005.
- Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- Jianqing Fan and Jinchi Lv. Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5):849–911, 2008.
- Jianqing Fan, Rui Song, et al. Sure independence screening in generalized linear models with np -dimensionality. *The Annals of Statistics*, 38(6):3567–3604, 2010.
- Kai-Tai Fang, Samuel Kotz, and Kai Wang Ng. *Symmetric multivariate and related distributions*. Chapman and Hall, 1990.
- Fang Han and Honglang Wang. Provable smoothing approach in high dimensional generalized regression model. *arXiv preprint arXiv:1509.07158*, 2015.
- Joel L Horowitz. *Semiparametric and nonparametric methods in econometrics*. Springer, 2009.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Michel Ledoux. *The concentration of measure phenomenon*. Number 89 in Mathematical Surveys and Monographs. American Mathematical Society, 2005.

- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer Science & Business Media, 2013.
- Jason D Lee, Yuekai Sun, and Jonathan E Taylor. On model selection consistency of m-estimators with geometrically decomposable penalties. *arXiv preprint arXiv:1305.7477*, 2013.
- Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.
- Ker-Chau Li and Naihua Duan. Regression analysis under link violation. *The Annals of Statistics*, pages 1009–1052, 1989.
- Lexin Li and Christopher J Nachtsheim. Sparse sliced inverse regression. *Technometrics*, 48(4), 2006.
- Peter McCullagh and John A Nelder. *Generalized linear models*, volume 37. CRC press, 1989.
- Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, pages 1436–1462, 2006.
- Matey Neykov, Qian Lin, and Jun S Liu. Signed support recovery for single index models in high-dimensions. *arXiv preprint, arXiv:1511.02270*, 2015.
- Heng Peng and Tao Huang. Penalized least squares for single index models. *Journal of Statistical Planning and Inference*, 141(4):1362–1379, 2011.
- Yaniv Plan and Roman Vershynin. The generalized lasso with non-linear observations. *arXiv preprint, arXiv:1502.04071*, 2015.
- Peter Radchenko. High dimensional single index models. *Journal of Multivariate Analysis*, 139:266–282, 2015.
- Charles M Stein. Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151, 1981.
- Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. The lasso with non-linear measurements is the lasso with non-linear measurements is equivalent to one with linear measurements. *arXiv preprint, arXiv:1506.02181*, 2015.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Martin J Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, 2009.

- J Wang, AT Bansal, M Martin, S Germer, R Benayed, L Essioux, JS Lee, A Begovich, A Hemmings, A Kenwright, et al. Genome-wide association analysis implicates the involvement of eight loci with response to tocilizumab for the treatment of rheumatoid arthritis. *The pharmacogenomics journal*, 13(3):235–241, 2013.
- Tao Wang and LiXing Zhu. A distribution-based lasso for a general single-index model. *Science China Mathematics*, 58(1):109–130, 2015.
- Tao Wang, Pei-Rong Xu, and Li-Xing Zhu. Non-convex penalized estimation in high-dimensional models with single-index structure. *Journal of Multivariate Analysis*, 109: 221–235, 2012.
- Yingcun Xia and WK Li. On single-index coefficient regression models. *Journal of the American Statistical Association*, 94(448):1275–1285, 1999.
- Zhuoran Yang, Zhaoran Wang, Han Liu, Yonina C. Eldar, and Tong Zhang. Sparse nonlinear regression: Parameter sparse nonlinear regression: Parameter estimation and asymptotic inference. *arXiv preprint arXiv: 1511:04514*, 2015.
- Xinyang Yi, Zhaoran Wang, Constantine Caramanis, and Han Liu. Optimal linear estimation under unknown nonlinear transform. *arXiv preprint arXiv:1505.03257*, 2015.
- Jingyuan Zhao, Simone Gupta, Mark Seielstad, Jianjun Liu, and Anbupalam Thalamuthu. Pathway-based analysis using reduced gene subsets in genome-wide association studies. *BMC bioinformatics*, 12(1):1, 2011.
- Peng Zhao and Bin Yu. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.