

Debiasing the debiased Lasso with bootstrap

Sai Li

Department of Biostatistics, University of Pennsylvania, Philadelphia, PA19104
e-mail: saili@penmedicine.upenn.edu

Abstract: We consider statistical inference for a single coordinate of regression coefficients in high-dimensional linear models. Recently, the debiased estimators are popularly used for constructing confidence intervals and hypothesis testing in high-dimensional models. However, some representative numerical experiments show that they tend to be biased for large coefficients, especially when the number of large coefficients dominates the number of small coefficients. In this paper, we propose a modified debiased Lasso estimator based on bootstrap. Let us denote the proposed estimator BS-DB for short. We show that, under the irrepresentable condition and other mild technical conditions, the BS-DB has smaller order of bias than the debiased Lasso in existence of a large proportion of strong signals. If the irrepresentable condition does not hold, the BS-DB is guaranteed to perform no worse than the debiased Lasso asymptotically. Confidence intervals based on the BS-DB are proposed and proved to be asymptotically valid under mild conditions. Our study on the inference problems integrates the properties of the Lasso on variable selection and estimation novelly. The superior performance of the BS-DB over the debiased Lasso is demonstrated via extensive numerical studies.

Keywords and phrases: Confidence intervals, high-dimensional models, debiased Lasso.

Received June 2019.

1. Introduction

1.1. Background

High-dimensional linear models have broad applications in many fields, such as biology, genetics, and machine learning. A number of statistical methods have been introduced to solve the problems on prediction, estimation, and variable selection regarding regression coefficients. On the other hand, statistical inference in high-dimensional models has recently caught a lot of research interests and efforts for its importance in providing uncertainty assessment and the nontrivial statistical challenges.

The Lasso estimator [26] has been a popular tool for modeling high-dimensional data. When the number of covariates p is fixed, however, it has been shown to have no closed form for its limiting distribution in the low dimensional setting [19]. Chatterjee and Lahiri [7] showed the inconsistency of bootstrapping the Lasso if at least one coefficient is zero. Thus, there is substantial difficulty in drawing valid inference based on the Lasso estimates directly. Nevertheless,

Chatterjee and Lahiri [8] developed a modified bootstrap estimator based on the Lasso as well as a bootstrap estimator based on Adaptive Lasso [38]. For p increasing with the sample size n , Chatterjee and Lahiri [9] showed the bootstrap approximation consistency for Adaptive Lasso estimators under some technical conditions.

In the $p \gg n$ scenario, Zhang and Zhang [33] proposed asymptotically Gaussian-distributed estimators of low-dimensional parameters in high-dimensional linear regression models. Such estimators are known as the “debiased Lasso” or the “de-sparsifying Lasso”. In the same paper, they proposed an optimization scheme for calculating the “correction score”, whose properties are carefully studied in [17] for both fixed designs and sub-Gaussian designs. Along this line of research, many recent papers have studied relevant generalizations for the debiased approach. Van de Geer et al. [28] considered the debiased Lasso estimator in generalized linear models with convex loss functions. Bühlmann and van de Geer [4] and Jankova and Van De Geer [16] studied statistical inference in misspecified high-dimensional linear models and graphical models, respectively. Fang et al. [15] considered the debiased method in high-dimensional Cox models. From the minimax perspective, Cai and Guo [5] studied the optimal expected lengths of confidence intervals for linear combinations of regression coefficients in sparse high-dimensional linear models. Javanmard and Montanari [18] considered sample size conditions for the debiased Lasso method in high-dimensional linear models with Gaussian design and Gaussian noise. Under some regularity conditions, they show a potentially weaker sample size condition when the true precision matrix of the design is sparse. Related approaches are also actively studied in the context of econometrics and causal inference [1, 3, 10].

The present work is motivated by the connections between statistical inference and variable selection problems. Variable selection has become an active research topic in high-dimensional literature for decades. Many established variable selection methods have been proposed and studied [26, 6, 14, 31, 24, 32]. It is known that if the nonzero coefficients can be consistently selected, least square estimators based on the selected model can lead to asymptotically valid inference procedures. However, the consistency of variable selection always requires the beta-min condition, which assumes the strengths of nonzero coefficients are uniformly larger than certain threshold. This condition is uncheckable and can be hard to fulfill in applications. In a recent paper [36], the post-Lasso least squares is justified for asymptotic valid statistical inference in high-dimensional linear models. Their analysis is based on the conditions guaranteeing the set of variables selected by the Lasso is deterministic with high probability. In the current work, we explore the interaction between variable selection and inference problems and demonstrate the benefits of having a large proportion of strong signals for statistical inference under proper conditions. Specially, we do not require the beta-min condition and the selected set of variables is not necessarily deterministic.

Another philosophy for inference in the high-dimensional setting is based on selective inference, whose focus is on making inference conditional on the selected model [21, 12, 20, 27]. However, it is not considered in the current work.

Our proposed approach is closely related to the bootstrap procedures for inference. Bootstrap has been widely used in high-dimensional models for conducting statistical inference. Mammen [22] considered estimating the distribution of linear contrasts and of F-test statistics when p increases with n . Chernozhukov et al. [11] developed theories for multiplier bootstrap to approximate the maximum of a sum of high-dimensional random vectors. Dezeure et al. [13] proposed residual, paired and wild multiplier bootstrap methodology for the debiased Lasso estimators. Zhang and Cheng [34] proposed a bootstrap-assisted debiased Lasso estimator to conduct simultaneous inference for non-Gaussian errors. The purpose of using bootstrap in aforementioned two papers mainly concern dealing with heteroscedastic errors as well as simultaneous inference. In the present work, we show the bias correction effect of bootstrap in high-dimensional inference.

1.2. The debiased approach

The debiased Lasso [33, 28, 17] for high-dimensional linear models can be described as follows. Consider a linear regression model

$$y_i = x_i^T \beta + \epsilon_i, \quad (1)$$

where $\beta \in \mathbb{R}^p$ is a vector of regression coefficients and $\epsilon_1, \dots, \epsilon_n$ are *i.i.d.* random variables with mean 0 and variance σ^2 . We consider the high-dimensional scenario where p can be larger or much larger than n . We assume β is sparse with support S such that $|S| = s$. Let $X \in \mathbb{R}^{n \times p}$ be the design matrix with the i -th row being x_i^T and $y = (y_1, \dots, y_n)^T$. Let X_j denote the j -th column of X . Let $\Sigma^n = X^T X/n$. Let $\Sigma = \mathbb{E}[\Sigma^n]$ denote the population gram matrix which is positive definite and $\Theta = \Sigma^{-1}$. Let $\Theta_j = \Theta_{\cdot, j}$ denote the j -th column of Θ .

The Lasso [26] estimator of β is defined as

$$\hat{\beta} = \arg \min_{b \in \mathbb{R}^p} \left\{ \frac{1}{2n} \|y - Xb\|_2^2 + \lambda \|b\|_1 \right\} \quad (2)$$

for some tuning parameter $\lambda > 0$. Some of its variations have been proposed and studied [38, 2, 25]. Suppose that we are interested in making inference of β_j for some $j \in \{1, \dots, p\}$. The debiased Lasso estimator of β_j can be written as

$$\hat{\beta}_j^{(DB)} = \hat{\beta}_j + \frac{\langle \hat{z}_j, y - X\hat{\beta} \rangle}{n}, \quad (3)$$

where $\hat{z}_j \in \mathbb{R}^n$ is the so-called correction score and can be computed via another Lasso regression. Specifically, define

$$\hat{\gamma} = \arg \min_{\gamma \in \mathbb{R}^p} \left\{ \frac{1}{2n} \|X\gamma\|_2^2 + \lambda_j \|\gamma_{-j}\|_1 : \gamma_j = 1 \right\} \quad (4)$$

for some tuning parameter $\lambda_j > 0$ and

$$\hat{z}_j = \begin{cases} X\Theta_j & \text{if } \Theta \text{ is known} \\ \frac{nX\hat{\gamma}}{\langle X_j, X\hat{\gamma} \rangle} & \text{if } \Theta \text{ is unknown.} \end{cases} \quad (5)$$

In fact, the correction score can also be realized via a quadratic optimization approach [33, 17] when Θ is unknown. The difference is that the optimization in (4) can utilize the sparsity of Θ_j , if Θ_j is sparse indeed, and can achieve semiparametric efficiency under proper conditions [28].

A $100 \times (1 - \alpha)\%$ two-sided confidence interval for β_j can be constructed as

$$\hat{\beta}_j^{(DB)} \pm q_{\alpha/2} \hat{\sigma} \|\hat{z}_j\|_2/n, \tag{6}$$

where q_τ is τ -th quantile of standard normal distribution and $\hat{\sigma}$ is some consistent estimate of the noise level σ .

When Θ is unknown, it has been proved that the asymptotic normality of $\hat{\beta}_j^{(DB)}$ requires $n \gg (s \log p)^2$ and some other technical conditions, say, in Van de Geer et al. [28]. These conditions guarantee that the remaining bias of the debiased Lasso is asymptotically sufficiently small such that $n^{-1/2}$ -length confidence intervals are achievable. In Cai and Guo [5], it has been shown that the minimax optimal confidence intervals for β_j has lengths of order $n^{-1/2} + s \log p/n$. That is, to achieve a confidence interval with length of order $n^{-1/2}$, the condition $n \gg (s \log p)^2$ is unavoidable in the minimax sense. This reveals the optimality of the debiased Lasso procedure.

Actually, when investigating the worst case scenario considered in Cai and Guo [5] (Theorem 3), one can see that it concerns the case where $\underline{s} = |\{j : |\beta_j| \asymp \sqrt{\log p/n}\}| \asymp s$. That is, loosely speaking, the number of “weak” coefficients dominates the number of “strong” coefficients. Therefore, if there exists a large proportion of strong signals, the debiased Lasso may not be optimal. If not, it is unknown what is a better procedure for statistical inference. This question is of significant practical value in view of the following numerical experiments, which suggest that the debiased Lasso estimators can be severely biased for large signals.

1.2.1. The bias of the debiased Lasso related to signal strengths

We consider a simplified setting where the observations are generated from model (1) with known noise level $\sigma = 1$ and known Σ . We set $n = 100$ and $p = 300$. Each row of X is *i.i.d.* from a multivariate Gaussian distribution with mean zero and covariance matrix Σ . The noise vector ϵ are *i.i.d.* from standard Gaussian distribution. We consider seven levels of sparsity, i.e. $s = 4, 6, \dots, 16$. For each s , we set $S = \{1, \dots, s\}$ and $\beta_S = (4, 2, 4, 2, \dots, 4, 2, 4, 0.2)^T$. Comparing with the inflated noise level $\sigma \sqrt{2 \log p/n} \approx 0.338$, first $s - 1$ coefficients can be viewed as strong signals and the s -th coefficient can be viewed as weak. We consider two formats of Σ . The first one is $\Sigma = I_p$, i.e. the identity matrix. The second one is $\Sigma = \Sigma^\circ$, where Σ° is a block diagonal matrix with $\Sigma_{S^c, S^c}^\circ = I_{p-s}$ and $\Sigma_{S, S}^\circ$ is Toeplitz with the first row equals $(1, -0.1, \mathbf{0}_{s-3}^T, -0.1)$. We compute the debiased Lasso (DB) based on (2) and (3) with tuning parameter $\lambda = 2\sigma \sqrt{\log p/n}$ and the oracle correction score $\hat{z}_j = X \Theta_j$.

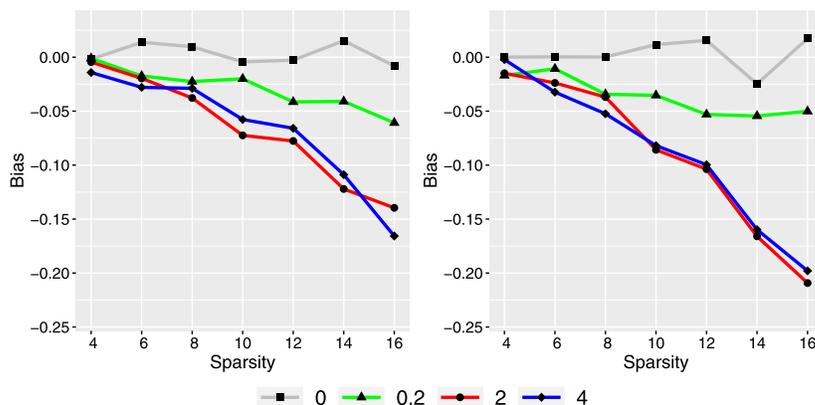


FIG 1. Estimation errors of the debiased Lasso for $\beta_j \in \{4, 2, 0.2, 0\}$ at different sparsity levels with $\Sigma = I_p$ (left panel) and $\Sigma = \Sigma^o$ (right panel). Each point is the median estimate based on 500 independent experiments.

When Σ is identity (left panel of Figure 1), distinct patterns are observed for signals with different strengths. The debiased Lasso for β_j , $j \in S^c$ has bias floating around zero. For β_j , $j \in S$, the debiased estimators are always negatively biased across different sparsity levels. This implies that the debiased Lasso can lead to low chance of discovering true signals under the current set-up. Moreover, the debiased estimators for strong signals are more severely biased as the sparsity level s gets larger. The right panel of Figure 1 displays the bias of the debiased Lasso when $\Sigma = \Sigma^o$. The difference of the left and right panel of Figure 1 precludes the effect of ρ_S defined in (17) where $\rho_S = 1$ with $\Sigma = I_p$ and $\rho_S = 1.25$ with $\Sigma = \Sigma^o$ at all sparsity levels. When ρ_S gets larger, the bias of the debiased Lasso on the true support gets larger.

Motivated by this numerical experiment, we look into the effects of signal strengths on statistical inference in high-dimensional models. Under proper conditions, we provide a new error analysis of the debiased Lasso which shows its distinct behaviors on and off the true support.

More importantly, we introduce a bootstrapped debiased Lasso approach (BS-DB), which can have smaller order of bias than the debiased Lasso when there are a large proportion of strong signals. Figure 2 unveils the bias correction effect of the BS-DB estimator in comparison to the debiased Lasso in the simulation settings considered above. It is not hard to see that the bias for strong signals are significantly reduced with our proposal. The BS-DB estimator is constructed according to Section 2.1 with the same tuning parameter as in debiased Lasso. More numerical results on constructing confidence intervals with unknown covariance matrix and unknown noise level are presented in Section 4. In next subsection, we summarize the major contributions of the current work.

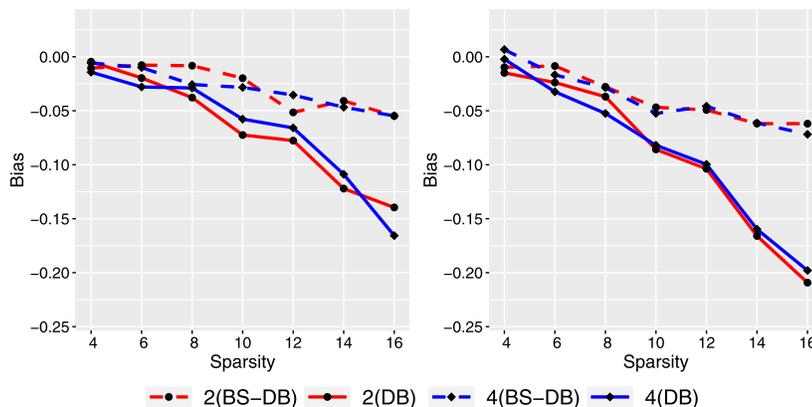


FIG 2. Estimation errors of the debiased Lasso (DB) and of the BS-DB for $\beta_j \in \{4, 2\}$ at different sparsity levels with identify covariance matrix (left panel) and $\Sigma = \Sigma^o$ (right panel). Each point is the median estimate based on 500 independent experiments.

1.3. Summary of our contributions

For the proposed BS-DB estimator defined in (11), we will prove that

$$\hat{\beta}_j^{(BS-DB)} - \beta_j = \widehat{rem}_j + \hat{\eta}_j, \tag{7}$$

where \widehat{rem}_j is the remaining bias of $\hat{\beta}_j^{(BS-DB)}$ and $\sqrt{n}\hat{\eta}_j$ is asymptotically normal. We upper bound the magnitude of \widehat{rem}_j assuming sub-Gaussian noise (Condition 3.2).

(a) Suppose that the design matrix is row-wise independent Gaussian (Condition 3.1) and the irrerepresentable condition (Condition 3.3) holds, if $n \gg s \log p$, then

$$|\widehat{rem}_j| = O_p \begin{cases} \frac{\sqrt{s_- \log p}}{n} + \frac{\rho_S \sqrt{s_- \log p}}{n} \mathbb{1}(j \in S) & \text{if } \Theta \text{ is known} \\ \frac{s_- \log p}{n} & \text{if } \Theta \text{ is unknown,} \end{cases} \tag{8}$$

where $\mathbb{1}(\cdot)$ is an indicator function, $s_- = |\{j \in S : |\beta_j| \leq C \rho_S \lambda\}|$ for a large enough constant C , and ρ_S defined in (17). It implies that, under proper conditions, asymptotically valid confidence intervals can be constructed based on $\hat{\beta}_j^{(BS-DB)}$ if $s \log p \lesssim n$ and $(s_- \log p)^2 \ll n$ for unknown Θ . This sample size condition is weaker than the one for the debiased Lasso if $s_- \ll s$. We will bring more discussions on the demand and potential relaxations of the irrerepresentable condition in Section 3. As a byproduct, we get a new error expansion of the debiased Lasso under current conditions in Theorem 3.3, which shows that its remaining bias can be of different magnitude for $j \in S$ and $j \notin S$.

(b) When the design is sub-Gaussian (Condition 3.4), without assuming the irrerepresentable condition, we prove that the remaining bias of the BS-DB is $O_P(s \log p/n)$ under mild conditions, which is asymptotically “no worse” than the debiased Lasso (Theorem 3.7).

1.4. Notations

For a set $A \subseteq \{1, \dots, p\}$, $X_A \in \mathbb{R}^{n \times |A|}$ is the submatrix formed by X_j , $j \in A$. Let $A^c = \{1, \dots, p\} \setminus A$. For $j = 1, \dots, p$, let $\{-j\} = \{1, \dots, p\} \setminus \{j\}$. For a vector $v \in \mathbb{R}^k$, let $\|v\|_p$ denote the standard ℓ_p -norm of v . For another vector $u \in \mathbb{R}^k$, let $\langle v, u \rangle = \sum_{i=1}^k v_i u_i$. For a symmetric matrix $D \in \mathbb{R}^{k_1 \times k_1}$, let $\Lambda_{\max}(D)$ and $\Lambda_{\min}(D)$ be the largest and smallest eigenvalues of D . Let $\text{Tr}(D)$ denote the trace of D . Let I_p be the $p \times p$ identity matrix. We use e_j to refer to the j -th standard basis element, i.e. $e_1 = (1, 0, \dots, 0)^T$. We use $e_{S_1} \in \mathbb{R}^{p \times |S_1|}$ to refer to the submatrix of an identity matrix, which contains j -th column of I_p for $\forall j \in S_1$. For a random variable v , let $\|v\|_{\psi_2} = \sup_{t \geq 1} t^{-1/2} (\mathbb{E}[|v|^t])^{1/t}$ denote its sub-Gaussian norm. For a random vector $V \in \mathbb{R}^k$, let $V_{\psi_2} = \sup_{\|b\|_2=1} \|\langle V, b \rangle\|_{\psi_2}$. Let $f(n)$ and $g(n)$ be two functions. We use $f(n) \ll g(n)$ to refer to “ $f(n) = o(g(n))$ ”. The notation $f(n) \gg g(n)$ can be defined analogously. We use $f(n) \lesssim g(n)$ to refer to $f(n) = O(g(n))$ and $f(n) \gtrsim g(n)$ can be defined analogously. We use c, c_0, c_1, \dots and C_0, C_1, \dots to denote generic constants that can vary from one position to the other.

1.5. Organization of the rest of the paper

The rest of the paper is organized as follows. In Section 2, we introduce the proposed approach for constructing confidence intervals in high-dimensional linear models. In Section 3, we prove the theoretical properties of the proposed approach with and without the irrepresentable condition. In Section 4, we demonstrate the empirical performance of the BS-DB for statistical inference in comparison to the debiased Lasso in various settings. In Section 5, we bring more discussions to the topics related to this paper. The proofs are provided in Section 6 and in the Appendix.

2. Bootstrapping the debiased Lasso

In this section, we introduce the proposed procedure and bring some intuitions to its merits.

2.1. The procedure for constructing confidence intervals

- (i) (Fitting the modified debiased Lasso) Let $\hat{\beta}$ be the initial Lasso estimator computed via (2) with tuning parameter λ . Let \hat{w}_j be the correction score defined in (15) (with known or unknown precision matrix). We compute a modified debiased Lasso estimator based on $\hat{\beta}$ and \hat{w}_j :

$$\hat{\beta}_j^{(mDB)} = \hat{\beta}_j + \frac{\langle \hat{w}_j, y - X\hat{\beta} \rangle}{n}. \quad (9)$$

- (ii) (Bootstrapping the Lasso) Let $\hat{\beta}^*$ be the Lasso estimate with input $(X, X\hat{\beta})$ and tuning parameter λ . We estimate the bias of $\hat{\beta}_j^{(mDB)}$ with

$$\hat{b}_j^* = \hat{\beta}_j^* + \frac{\langle \hat{w}_j, X(\hat{\beta} - \hat{\beta}^*) \rangle}{n} - \hat{\beta}_j. \tag{10}$$

Subtracting \hat{b}_j^* from $\hat{\beta}_j^{(mDB)}$, we arrive at a bootstrapped debiased Lasso estimator

$$\hat{\beta}_j^{(BS-DB)} = \hat{\beta}_j^{(mDB)} - \hat{b}_j^*. \tag{11}$$

- (iii) An $100 \times (1 - \alpha)\%$ two-sided confidence interval for β_j can be constructed as

$$\hat{\beta}_j^{(BS-DB)} \pm q_{\alpha/2} \hat{\sigma} \|\hat{w}_j\|_2 / n, \tag{12}$$

where $\hat{\sigma}$ can be any consistent estimator of σ .

The modified debiased Lasso computed in Step (i) is based on a different correction score in comparison to the original debiased Lasso. The specific expression and rationale are illustrated in detail in Section 2.2. In Step (ii), the estimator $\hat{\beta}^*$ is a noiseless Lasso based on the empirical estimates. In the usual parametric bootstrap, the response vector is constructed as

$$y_i^* = x_i^T \hat{\beta} + \hat{\sigma} \xi_i^*, \quad i = 1, \dots, n, \tag{13}$$

where ξ_i^* are synthetic *i.i.d.* standard Gaussian random variables and $\hat{\sigma}$ is a consistent estimate of noise level. Hence, the estimator $\hat{\beta}^*$ employed in Step (ii) can be viewed as a noiseless Lasso based on parametric bootstrap. In fact, our $\hat{\beta}^*$ can be replaced with an average of the usual noisy parametric bootstrap estimators. Although the noisy parametric bootstrap estimates can be used to generate empirical confidence intervals, our main purpose of bootstrap is for bias correction rather than uncertainty quantification. Hence, we focus on the proposed noiseless version which can simplify the computation and does not introduce extra randomness. (I would like to thank an anonymous referee for pointing this out.)

We also comment that the current framework and the proposed approach are different from the results of bootstrapping the Adaptive Lasso considered in Chatterjee and Lahiri [9]. In fact, in the current work, the beta-min condition is not required and selection consistency is not needed. As a consequence, some direct methods, such as least squares after selection, cannot lead to valid inference procedures under current conditions in general.

2.2. The proposed correction score

Before articulating the format of the proposed correction score \hat{w}_j , let us bring some intuitions to its construction. Recall that the success of the debiased Lasso relies on choosing a correction score $v \in \mathbb{R}^n$ such that [33]

$$e_j^T - v^T X/n \approx 0. \tag{14}$$

One possible realization of v is \hat{z}_j defined in (5). In fact, it is ideal but not amendable to have an “exactly equal” relationship in (14). It can be seen from later discussion (22) that if v is such that $v^T X_A/n = 0$ for a set $A \subseteq \{-j\}$, then the debiased Lasso based on the correction score v is free from the bias of $\hat{\beta}_A$. However, an “exactly equal” relationship is not achievable in general since $\Theta^n = (\Sigma^n)^{-1}$ is not well-defined when $p \geq n$. As a compromise, we would like to obtain a correction score v such that

$$(e_j^T - v^T X/n)_S = 0 \quad \text{and} \quad (e_j^T - v^T X/n)_{S^c} \approx 0.$$

That is, on a “small” set which has large bias, S , we require the exact equality to hold in (14) and on S^c , we allow for “approximately equal”. As the true support S is unknown, we replace it by an estimate based on the Lasso, $\hat{S} = \{j : \hat{\beta}_j \neq 0\}$, which is not necessarily consistent but satisfies certain desirable properties. Let $\hat{A}_j = \hat{S} \cap \{-j\}$. Let $P_{\hat{A}_j}^\perp = I_n - X_{\hat{A}_j} \left(X_{\hat{A}_j}^T X_{\hat{A}_j} \right)^{-1} X_{\hat{A}_j}^T$ denote the projection operator, where $X_{\hat{A}_j}^T X_{\hat{A}_j}/n$ will be shown to be invertible with high probability. Formally, our proposed correction score \hat{w}_j can be defined as

$$\hat{w}_j = \begin{cases} \frac{n P_{\hat{A}_j}^\perp X_{\Theta_j}}{\langle X_j, P_{\hat{A}_j}^\perp X_{\Theta_j} \rangle} & \text{if } \Theta \text{ is known} \\ \frac{n X_{\hat{\kappa}}}{\langle X_j, X_{\hat{\kappa}} \rangle} & \text{if } \Theta \text{ is unknown,} \end{cases} \quad (15)$$

where

$$\hat{\kappa} = \arg \min_{\kappa \in \mathbb{R}^p} \left\{ \frac{1}{2n} \|X\kappa\|_2^2 + \lambda_j \|\kappa_{\hat{S}^c \setminus \{j\}}\|_1 : \kappa_j = 1 \right\}. \quad (16)$$

The optimization in (16) is different from (4) as it does not penalize the coefficients in \hat{A}_j . It is easy to see from (15) and the KKT condition of (16) that $(e_j^T - \hat{w}_j^T X/n)_{\hat{S}} = 0$ no matter Θ is known or not. We demonstrate the theoretical advantages of the proposed BS-DB estimator in the next section.

3. Theoretical properties

In this section, we study the theoretical properties of the proposed BS-DB estimator. We will first show some preliminary results under the irrerepresentable condition and then justify the asymptotic validness of the proposed confidence interval (12). We will further study the robustness of our proposal to the violation of the irrerepresentable condition. For the main results, we assume the following conditions.

Condition 3.1 (Gaussian designs). *Each row of X is i.i.d. Gaussian distributed with mean 0 and covariance matrix Σ , where Σ satisfies that $0 < C_{\min} \leq \Lambda_{\min}(\Sigma) \leq \Lambda_{\max}(\Sigma) \leq C_{\max} < \infty$ and $\max_{1 \leq j \leq p} \Sigma_{j,j} \leq K_1 < \infty$.*

The smallest eigenvalue of Σ is assumed to be bounded away from zero, which is used in upper bounding the remaining bias of the debiased Lasso. The largest

eigenvalue of Σ is assumed to be finite, which is used in justify the asymptotic normality. The Gaussian distribution is needed for technical convenience and is also required in closely related works [30, 18].

Condition 3.2 (Sub-Gaussian errors). ϵ_i ($i = 1, \dots, n$) are *i.i.d. sub-Gaussian* with mean 0 and variance $\sigma^2 > 0$ such that $\max_{1 \leq i \leq n} \|\epsilon_i\|_{\psi_2} \leq K_2$.

Condition 3.3 (Irrepresentable condition).

$$\phi = \max_{k \in S^c} \|\Sigma_{k,S}(\Sigma_{S,S})^{-1}\|_1 < 1.$$

As previously mentioned, Condition 3.3 is crucial in our main analysis. It has been used in analyzing the variable selection properties of the Lasso [35, 30] and it is equivalent to the “neighborhood stability” condition [23]. In fact, the purpose of irrepresentable condition is to guarantee that the support estimated by the Lasso is a subset of the true support with high probability (Lemma 3.1 and Lemma 3.2). Some concave penalized methods, say Zhang [32], require weaker versions of Condition 3.3 for guaranteeing such properties and they can also be suitable for debiasing. For conciseness and computational convenience, we focus on the Lasso.

It is worth mentioning that Condition 3.3 does not rule out the most difficult scenarios which yield the minimax optimal sample size condition $n \gg (s \log p)^2$. Indeed, the lower bound results for the confidence intervals in high-dimensional linear models [5, 18] are based on some worst cases satisfying $\phi = 0$. Therefore, our results are comparable to the existing lower bound results achieved by the debiased Lasso.

3.1. Preliminary lemmas

We first prove some preliminary results as consequences of the irrepresentable condition. For $b \in \mathbb{R}^p$, let $sgn(b)$ be an element of the sub-differential of the ℓ_1 -norm of b . That is, $(sgn(b))_j = sgn(b_j) = b_j/|b_j|$ if $b_j \neq 0$ and $(sgn(b))_j \in [-1, 1]$ if $b_j = 0$. Let

$$\rho_S = \max_{j \in S} \|e_j^T(\Sigma_{S,S})^{-1}\|_1. \tag{17}$$

Notice that $\rho_S \in [c_1, c_2\sqrt{s}]$ for some positive constants c_1 and c_2 under Condition 3.1.

Lemma 3.1 (Sparsity pattern recovery for the Lasso). *Assume that Condition 3.1–Condition 3.3 are satisfied. If the sequence (n, p, s, λ) satisfies that*

$$n \geq \frac{cK_1s \log p}{C_{\min}(1 - \phi)^2} \quad \text{and} \quad \lambda \geq \frac{c_0\sigma}{1 - \phi} \sqrt{\frac{K_1 \log p}{n}} \tag{18}$$

for large enough constants c and c_0 , then

$$\Omega_0 := \{\hat{S} \subseteq S \text{ and } \|\hat{\beta}_S - \beta_S\|_\infty \leq C\rho_S\lambda\} \tag{19}$$

holds with probability at least $1 - c_1/p - \exp(-c_2n) - c_3/n$ for some large enough constants C, c_1, c_2, c_3 , and ρ_S defined in (17).

Lemma 3.1 is more general than Theorem 3 in Wainwright [30] in the sense that the ℓ_∞ -bound of the estimation error in (19) does not require the beta-min condition. This is a nontrivial generalization as $\text{sgn}(\hat{\beta}_S) \neq \text{sgn}(\beta_S)$ in general and dependence between $\text{sgn}(\hat{\beta}_S)$ and X_S makes it challenging to get the desired bound on $\hat{\beta}_S - \beta_S$. Our idea is to decouple $\hat{\beta}_S$ with X_k for each $k \in S$ with a “leave-one-out” argument.

The sparsity pattern recovery can be similarly derived for the bootstrap version $\hat{\beta}^*$. We summarize an important consequence of Lemma 3.1 and its bootstrap counterpart in the next lemma. Let $\hat{S}^* = \{j : \hat{\beta}_j^* \neq 0\}$.

Lemma 3.2. *Assume that Condition 3.1–Condition 3.3 are satisfied. If (18) holds, then*

$$\Omega_0^* := \{\hat{S}^* \subseteq \hat{S} \text{ and } \|\hat{\beta}_{\hat{S}^*}^* - \hat{\beta}_S\|_\infty \leq C^* \rho_S \lambda\}. \quad (20)$$

holds with probability at least $1 - c_1/p - c_2/n - \exp(-c_3n)$ for some large enough constants c_1, c_2, c_3 , and C^* . In the event $\Omega_0 \cap \Omega_0^*$, $S_+ \subseteq \hat{S} \subseteq S$, $S_+ \subseteq \hat{S}^* \subseteq S$, and $\text{sgn}(\beta_j) = \text{sgn}(\hat{\beta}_j) = \text{sgn}(\hat{\beta}_j^*)$ for $j \in S_+$, where

$$S_+ = \{j : |\beta_j| > (C + C^*)\rho_S \lambda, j = 1, \dots, p\} \quad (21)$$

for C in (19) and C^* in (20).

Lemma 3.2 implies that the estimated supports based on $\hat{\beta}$ and $\hat{\beta}^*$ can be upper and lower bounded by two unknown but deterministic sets, namely S and S_+ , with high probability, where S_+ can be viewed as the set of strong signals. In addition, the sign vectors of $\hat{\beta}_{S_+}$ and $\hat{\beta}_{S_+}^*$ are asymptotically deterministic. These observations are crucial for the success of our proposed bias correction as will be seen from the next subsection.

3.2. Illustrations of the proposed bias correction

In this subsection, we layout some key steps for analyzing the proposed approach based on the results proved in Section 3.1. Let us focus on the case where Θ is unknown. We first review some key steps in the typical analysis of the debiased Lasso. From (3) and (19), we have

$$\hat{\beta}_j^{(DB)} - \beta_j = \underbrace{(e_j^T - \hat{z}_j^T X/n)}_{\widehat{rem}_j^o} (\hat{\beta} - \beta) + \frac{\langle \hat{z}_j, \epsilon \rangle}{n},$$

where $\widehat{rem}_j^o = O_P(s \log p/n)$ by an $\ell_\infty - \ell_1$ splitting using the KKT condition of (4) and ℓ_1 -bound on $\hat{\beta} - \beta$. For the asymptotic normality of $\sqrt{n}(\hat{\beta}_j^{(DB)} - \beta_j)$ to be true, one needs $|\widehat{rem}_j^o| = o_P(n^{-1/2})$, which gives the typical sample size condition $(s \log p)^2 \ll n$, and the asymptotic normality of $\langle \hat{z}_j, \epsilon \rangle / \sqrt{n}$, which holds under mild conditions.

Next, we illustrate that, with the proposed BS-DB, the bias coming from strong signals is removed in two steps. Let $W^n = X^T \epsilon/n$. The event Ω_0 (19) and the KKT condition of $\hat{\beta}$ together imply that

$$\begin{aligned} \hat{\beta}_j^{(mDB)} - \beta_j &= (e_j^T - \hat{w}_j^T X/n)_S (\hat{\beta}_S - \beta_S) + \frac{\langle \hat{w}_j, \epsilon \rangle}{n} \\ &= -\lambda \underbrace{(e_j^T - \hat{w}_j^T X/n)_S (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S)}_{\hat{b}_j} \\ &\quad + \underbrace{(e_j^T - \hat{w}_j^T X/n)_S (\Sigma_{S,S}^n)^{-1} W_S^n}_{\hat{r}_j} + \frac{\langle \hat{w}_j, \epsilon \rangle}{n}, \end{aligned} \tag{22}$$

where \hat{b}_j is the dominant bias term and \hat{r}_j and the last term are linear combinations of ϵ . For S_+ defined in Lemma 3.2, let $S_- = S \setminus S_+$ denote the set of weak signals. Let $s_+ = |S_+|$ and $s_- = |S_-|$. To see the bias contributed by strong signals, we rewrite \hat{b}_j as

$$\hat{b}_j = -\lambda (e_j^T - \hat{w}_j^T X/n)_S \underbrace{\begin{pmatrix} (\Sigma_{S,S}^n)^{-1}_{S_+,S_+} \text{sgn}(\hat{\beta}_{S_+}) & (\Sigma_{S,S}^n)^{-1}_{S_+,S_-} \text{sgn}(\hat{\beta}_{S_-}) \\ (\Sigma_{S,S}^n)^{-1}_{S_-,S_+} \text{sgn}(\hat{\beta}_{S_+}) & (\Sigma_{S,S}^n)^{-1}_{S_-,S_-} \text{sgn}(\hat{\beta}_{S_-}) \end{pmatrix}}_{G^n},$$

where $\text{sgn}(\hat{\beta}_{S_+}) = \text{sgn}(\beta_{S_+})$ with high probability under the conditions of Lemma 3.2. We will denote G^n as a 2×2 block matrix such that $G^n_{1,1} = (G^n)_{S_+,S_+}$. Loosely speaking, the purpose of bootstrap is to remove $G^n_{1,1}$ and $G^n_{2,1}$ and the purpose of proposed correction score \hat{w}_j is to remove $G^n_{2,1}$. Specifically, in event Ω_0^* (20), we can similarly show that for \hat{b}_j^* defined in (10),

$$\hat{b}_j^* = -\lambda (e_j^T - \hat{w}_j^T X/n)_S \underbrace{\begin{pmatrix} (\Sigma_{S,S}^n)^{-1}_{S_+,S_+} \text{sgn}(\hat{\beta}_{S_+}^*) & (\Sigma_{S,S}^n)^{-1}_{S_+,S_-} \text{sgn}(\hat{\beta}_{S_-}^*) \\ (\Sigma_{S,S}^n)^{-1}_{S_-,S_+} \text{sgn}(\hat{\beta}_{S_+}^*) & (\Sigma_{S,S}^n)^{-1}_{S_-,S_-} \text{sgn}(\hat{\beta}_{S_-}^*) \end{pmatrix}}_{G^*}.$$

Since $\text{sgn}(\hat{\beta}_{S_+}^*) = \text{sgn}(\hat{\beta}_{S_+})$ by Lemma 3.2, we have $G^n_{.,1} = G^*_{.,1}$. Moreover, the second column of G^n and that of G^* come from the incorrect sign estimation of weak signals. The term $G^n_{1,2}$ is caused by the possible correlation between X_{S_+} and X_{S_-} and it can be large if s_+ is large. To get rid of $G^n_{1,2}$ and $G^*_{1,2}$, we invoke that the proposed correction score \hat{w}_j satisfies $(e_j^T - \hat{w}_j^T X/n)_{\hat{S}} = 0$. If $S_+ \subseteq \hat{S}$, then the effects of $G^n_{1,2}$ and $G^*_{1,2}$ are removed by the proposed correction score. To summarize, (22) and above analysis together imply that in $\Omega_0 \cap \Omega_0^*$,

$$\hat{\beta}_j^{(mDB)} - \beta_j - \hat{b}_j^* = -\lambda (e_j^T - \hat{w}_j^T X/n)_{S_-} (G^n_{2,2} - G^*_{2,2}) + \hat{r}_j + \frac{\langle \hat{w}_j, \epsilon \rangle}{n}.$$

We will show that the remaining bias, i.e. the first two terms on the right hand side on the above expression, are $O_P(s_- \log p/n)$. Finally, the validness of the confidence interval considered in (12) also relies on the asymptotic normality of $\langle \hat{w}_j, \epsilon \rangle/\sqrt{n}$. We will employ the central limit theorem to show desirable results under proper conditions.

3.3. Main theorems

In this subsection, we formally establish the main theorems for statistical inference with the BS-DB approach. As a benchmark, we first present an error analysis of the original debiased Lasso demonstrating its different magnitude of bias on and off the true support.

Theorem 3.3 (The remaining bias of the debiased Lasso). *Suppose that Condition 3.1–Condition 3.3 hold true and $s \log p \ll n$. Let $\hat{\beta}_j^{(DB)}$ be computed via (3) with $\lambda = c_1 \sqrt{\log p/n}$ and $\lambda_j = c_2 \sqrt{\log p/n}$ for some sufficiently large constants c_1 and c_2 such that (18) holds. Then*

$$\hat{\beta}_j^{(DB)} - \beta_j = \frac{\langle \hat{z}_j, \epsilon \rangle}{n} + \widehat{rem}_j^o,$$

where

$$|\widehat{rem}_j^o| =_{O_p} \begin{cases} \frac{\sqrt{s \log p}}{n} + \frac{\rho_S \sqrt{s \log p}}{n} \mathbb{1}(j \in S) & \text{if } \Theta \text{ is known} \\ \frac{s \log p}{n} & \text{if } \Theta \text{ is unknown.} \end{cases} \quad (23)$$

Theorem 3.3 implies that when Θ is known, the remaining bias of $\hat{\beta}_j^{(DB)}$ can be larger for $j \in S$ than for $j \in S^c$ as $1 \lesssim \rho_S \lesssim \sqrt{s}$ under Condition 3.1. Moreover, as ρ_S gets larger, the bias of $\hat{\beta}_j^{(DB)}$ for $j \in S$ gets larger. These results coincide with our observations in Figure 1. We mention that one can get rid of the term which involves ρ_S by estimating $\hat{\beta}$ and \hat{z}_j with two independent subsets of data (Theorem 7 in Cai and Guo [5] and Proposition H.1. in Javanmard and Montanari [18]). In practice, especially when the sample size is relatively small, sample splitting can produce unstable results. Hence, we focus on the version without sample splitting. When Θ is unknown, the results of Theorem 3.3 agree with the existing analysis of the debiased Lasso [33, 28, 17]. With Theorem 3.3 serving as a benchmark, we study the the remaining bias of the proposed BS-DB in the next lemma.

Lemma 3.4 (The remaining bias of BS-DB). *Suppose that Condition 3.1–Condition 3.3 hold true and $s \log p \ll n$. Let $\hat{\beta}_j^{(BS-DB)}$ be computed via (11) with $\lambda = c_1 \sqrt{\log p/n}$ and $\lambda_j = c_2 \sqrt{\log p/n}$ for some sufficiently large constants c_1 and c_2 such that (18) holds. Then the expressions in (7) and (8) hold true.*

Lemma 3.4 shows that the magnitude of the remaining bias of $\hat{\beta}_j^{(BS-DB)}$ is determined by the number of weak signals. Comparing with Theorem 3.3, we see that the remaining bias of $\hat{\beta}_j^{(BS-DB)}$ is much smaller than the remaining bias of the debiased Lasso when $s_- \ll s$. This demonstrates the improvement of our proposal and convinces our observations in Figure 2. When Θ is known, the magnitude of remainder terms of $\hat{\beta}_j^{(BS-DB)}$ can be different for $j \in S$ and $j \notin S$, which is analogous to the results for debiased Lasso.

Next, we move on to establish the limiting distribution of the BS-DB estimator. We first prove the convergence rate of a variance estimator, which can also benefit from the irrepresentable condition. Define

$$\hat{\sigma}^2 = \frac{1}{n - |\hat{S}|} \|P_{\hat{S}}^\perp y\|_2^2, \tag{24}$$

where $P_{\hat{S}}^\perp = I_n - X_{\hat{S}} \left(X_{\hat{S}}^T X_{\hat{S}} \right)^{-1} X_{\hat{S}}^T$.

Lemma 3.5 (Convergence rate of the variance estimator). *Suppose that Condition 3.1–Condition 3.3 hold true and $s \log p \ll n$. Let $\hat{\beta}$ be computed via (2) with $\lambda = c_1 \sqrt{\log p/n}$ for a sufficiently large constant c_1 satisfying (18). For $\hat{\sigma}^2$ defined in (24),*

$$|\hat{\sigma}^2 - \sigma^2| = O_P \left(n^{-1/2} + \frac{\min\{\rho_S^2 s_-, s\} \log p}{n} \right). \tag{25}$$

On the right hand side of (25), the first term $n^{-1/2}$ comes from the randomness of ϵ and the second term comes from the estimation error of weak signals. It is known that a widely used variance estimator, the mean of squared residuals based on $\hat{\beta}$, has convergence rate $O_P(n^{-1/2} + s \log p/n)$. We see that $\hat{\sigma}^2$ is no worse than and can have faster rate of convergence than the mean of squared residuals based on the Lasso if $\rho_S^2 s_- \ll s$. That is, $\hat{\sigma}^2$ can be more accurate when the number of strong signals is dominant and the correlation among X_S is weak. The empirical performance of $\hat{\sigma}^2$ is evaluated in various numerical experiments in Section 4.

In the next theorem, we collect all the preliminary results and prove the asymptotic normality of $\hat{\beta}_j^{(BS-DB)}$ under proper sample size conditions. Let $\Phi(c) = \mathbb{P}(x \leq c)$ for a standard normal variable x .

Theorem 3.6 (Asymptotic normality of BS-DB). *Suppose that Condition 3.1–Condition 3.3 hold true and $s \log p \ll n$. Let $\hat{\beta}_j^{(BS-DB)}$ be computed via (11) with $\lambda = c_1 \sqrt{\log p/n}$ and $\lambda_j = c_2 \sqrt{\log p/n}$ for some sufficiently large constants c_1 and c_2 such that (18) holds. (i) If Θ_j is known, then*

$$\sup_{c \in \mathbb{R}} \left| \mathbb{P} \left(\hat{\beta}_j^{(BS-DB)} - \beta_j \leq c \hat{\sigma} \|\hat{w}_j\|_2/n \right) - \Phi(c) \right| \rightarrow 0. \tag{26}$$

(ii) *When Θ_j is unknown, (26) holds if $\max \{s_-, \|\Theta_{(S_+)^c, j}\|_0\} \log p \ll \sqrt{n}$.*

Theorem 3.6 implies that the proposed confidence interval (12) has nominal coverage probability asymptotically. The condition $s_- \log p \ll \sqrt{n}$ when Θ is unknown guarantees that the remaining bias of BS-DB is $o_P(n^{-1/2})$. The condition $\|\Theta_{(S_+)^c, j}\|_0 \log p \ll \sqrt{n}$ guarantees that $\hat{\kappa}$ converges to its probabilistic limit at a sufficiently fast rate such that $\hat{\eta}_j$ in (7) is asymptotically normal. It is especially needed here as our constructed \hat{w}_j is dependent with ϵ . This condition can be relaxed if \hat{S} used in (15) is independent of X or is asymptotically deterministic. Hence, one can perform a sample splitting and compute \hat{S} with one

fold of the data and compute $\hat{\kappa}$ and other estimates with the other fold of the data. In this way, the sparsity requirement on $\Theta_{(S_+)^c, j}$ can be relaxed. On the other hand, some mild conditions can lead to asymptotically deterministic \widehat{S} . One sufficient condition is that $|\{j : |\beta_j| \asymp \rho_S \lambda\}| = 0$ which guarantees $\widehat{S} = S_+$ asymptotically. The **optimality** of the proposed confidence interval (12) can be partially understood from the established lower bound results. Specifically, the parameter space considered in Theorem 3 of Cai and Guo [5] is

$$\Xi_0(s) = \{(\beta, \sigma, \Sigma) : \|\beta\|_0 \leq s, \sigma^2 \geq c_1, 1/c_2 \leq \Lambda_{\min}(\Sigma) \leq \Lambda_{\max}(\Sigma) \leq c_2\}$$

for some constants $c_1 > 0$ and $c_2 > 1$. Let $L_\alpha^*(\Theta, \beta_j)$ denote the minimax expected length of confidence intervals for β_j ((2.3) of Cai and Guo [5]) over Θ at confidence level α . Recall that $\underline{s} = |\{j : |\beta_j| \asymp \sqrt{\log p/n}\}|$. When Θ is unknown, it has been shown that

$$L_\alpha^*(\Xi_0(s), \beta_j) \gtrsim \frac{1}{\sqrt{n}} + \frac{\underline{s} \log p}{n}.$$

Let us consider a more detailed parameter space

$$\begin{aligned} \Xi(s_+, s_-) = \{(\beta, \sigma, \Sigma) : |S_+| \leq s_+, |S_-| \leq s_-, \sigma^2 \geq c_1, \\ 1/c_2 \leq \Lambda_{\min}(\Sigma) \leq \Lambda_{\max}(\Sigma) \leq c_2, \rho_S \leq c_3\}, \end{aligned}$$

for some constants $c_1 > 1$, $c_2 > 1$, and $c_3 < \infty$. As $s_- = \underline{s}$ for constant ρ_S , it is not hard to see that

$$L_\alpha^*(\Xi(s_+, s_-), \beta_j) \gtrsim \frac{1}{\sqrt{n}} + \frac{s_- \log p}{n}.$$

This shows the optimality of the proposed BS-DB in $\Xi(s_+, s_-)$ when Θ is unknown.

3.4. When irrepresentable condition does not hold

The irrepresentable condition is always hard to check in reality. Therefore, it is important to understand whether the proposed BS-DB is still valid for inference when such a condition is not true. In the next theorem, we justify the theoretical properties of BS-DB without the irrepresentable condition. We also relax the Gaussian assumption in Condition 3.1 to sub-Gaussian designs. For practical concerns, we only prove for the case where Θ is unknown.

Condition 3.4 (Sub-Gaussian designs). *Each row of X is i.i.d from a sub-Gaussian distribution with mean zero and covariance matrix Σ . The matrix Σ satisfies that $0 < C_{\min} \leq \Lambda_{\min}(\Sigma) \leq \Lambda_{\max}(\Sigma) \leq C_{\max} < \infty$. There exists a positive constant K_1 such that $\max_{1 \leq i \leq n} \|x_i \Sigma^{-1/2}\|_{\psi_2} \leq K_1 < \infty$.*

Theorem 3.7 (The remaining bias of BS-DB without irrerepresentable condition). *Assume that Condition 3.2 and Condition 3.4 are true and $s \log p \ll n$. Let $\hat{\beta}_j^{(BS-DB)}$ be computed via (11) with $\lambda = c_1 \sqrt{\log p/n}$ and $\lambda_j = c_2 \sqrt{\log p/n}$ for some sufficiently large constants c_1 and c_2 . When Θ is unknown, it holds that*

$$\hat{\beta}_j^{(BS-DB)} - \beta_j = \frac{\langle \hat{w}_j, \epsilon \rangle}{n} + O_P\left(\frac{s \log p}{n}\right).$$

Theorem 3.7 implies that if $n \gg (s \log p)^2$, then the remaining bias of $\hat{\beta}_j^{(BS-DB)}$ is $o_P(n^{-1/2})$ even if the irrerepresentable condition does not hold. Hence, the remaining bias of BS-DB estimator is no larger than that of the debiased Lasso. The asymptotic normality of $\langle \hat{w}_j, \epsilon \rangle / \sqrt{n}$ can be established based on the proof of Theorem 3.6 under proper conditions. We can summarize that when Θ is unknown, there is no loss in applying the BS-DB asymptotically regardless of the irrerepresentable condition and BS-DB can achieve more accurate confidence intervals in existence of a large proportion of strong signals.

4. Numerical experiments

In this section, we demonstrate the empirical performance of BS-DB in comparison to the debiased Lasso in more practical settings with Σ and σ^2 are both unknown.

We set sample size $n = 100$, the number of covariates $p = 300$, and the noise level $\sigma = 1$. We consider $\Sigma = I_p$, $\Sigma = \Sigma^\circ$ defined in Section 1.2, and another Σ where the irrerepresentable condition does not hold. Each row of X is *i.i.d.* generated from $N(0, \Sigma)$ and $\epsilon_i \sim N(0, \sigma^2)$ for $i = 1, \dots, n$. We consider three levels of sparsity. Specifically, $s = 4k$ for $k \in \{1, 2, 3\}$. Let $s_- = |\{j : \beta_j = 0.2\}|$. For any $k \in \{1, 2, 3\}$, we consider the following two cases

- (i) $(s, s_-) = (4k, 1)$, $\beta_{1:s} = (4, 2, 4, 2, \dots, 4, 0.2)^T$, and $\beta_{s+1:p} = \mathbf{0}_{p-s}$.
- (ii) $(s, s_-) = (4k, 2k)$, $\beta_{1:2k} = (4, 2, \dots, 4, 2)^T$, $\beta_{2k+1:s} = (0.2, \dots, 0.2)^T$, and $\beta_{s+1:p} = \mathbf{0}_{p-s}$.

In case (i), a large proportion of signals are strong and in case (ii), half of the signals are strong. The purpose is to demonstrate the effect of overall sparsity and effect of the proportion of weak signals separately. We construct two-sided 95% confidence intervals for $\beta_1 = 4$, $\beta_2 = 2$, $\beta_s = 0.2$, and $\beta_{s+1} = 0$ in each setting. We report summarized statistics based on 500 independent realizations for each setting.

We will compare the coverage probabilities and lengths of confidence intervals given by the debiased Lasso (DB) procedure and the proposed BS-DB. In both methods, the tuning parameter for the Lasso is $\lambda = 2\hat{\sigma}^{(init)} \sqrt{\log p/n}$, where $\hat{\sigma}^{(init)}$ is computed via the scaled Lasso. The tuning parameter in (4) and (16) are set to be $\hat{\sigma}_x^{(init)} \sqrt{2 \log p/n}$, where $\hat{\sigma}_x^{(init)}$ is computed via the scaled Lasso with response X_j and covariates X_{-j} . The estimated noise level $\hat{\sigma}$ is computed according to (24).

4.1. Identify covariance matrix

We first present the numerical results for the case where $\Sigma = I_p$ (Table 1). One can see that the coverage probabilities of debiased Lasso are less accurate than those of BS-DB in all the scenarios, especially when the number of strong signals dominant. In comparison, BS-DB has coverage probabilities close to the nominal level at different sparsity levels and is especially robust to the existence of a large proportion of strong signals. This agrees with our analysis in Theorem 3.3 and Theorem 3.6. The lengths of confidence intervals produced by the debiased Lasso and BS-DB are comparable. Hence, the difference in coverage probabilities is mainly due to the remaining bias term.

The confidence intervals based on debiased Lasso have relatively poor coverage in our numerical studies comparing with in some previous studies. The main reason is that previous studies always use mean of squared residuals (MSR) to estimate the noise level which tends to inflate the true value and results in wider confidence intervals, which compensate for the bias of debiased Lasso. The performance of the proposed variance estimator (24), the scaled Lasso estimator [25], and the MSR are reported in Table 4. One can see that the proposed estimator of noise level is most accurate in different settings. The MSR can have large bias when the number of strong signals dominants. The scaled-Lasso is more reliable than MSR but not as accurate as the proposed estimator when the sparsity is large.

TABLE 1
Coverage probabilities of BS-DB (cov.bsdb), coverage probabilities of debiased Lasso (cov.db), standard errors of BS-DB (se.bsdb), and standard errors of debiased Lasso (se.db) when $\Sigma = I_p$.

s	β	$s_- = 1$				$s_- = s/2$			
		cov.bsdb	cov.db	se.bsdb	se.db	cov.bsdb	cov.db	se.bsdb	se.db
4	4	0.942	0.888	0.10	0.10	0.936	0.890	0.10	0.10
4	2	0.938	0.882	0.10	0.10	0.932	0.902	0.10	0.10
4	0.2	0.918	0.866	0.10	0.10	0.948	0.926	0.10	0.10
4	0	0.958	0.828	0.10	0.10	0.948	0.886	0.10	0.10
8	4	0.932	0.666	0.11	0.11	0.900	0.838	0.11	0.11
8	2	0.950	0.652	0.11	0.11	0.934	0.800	0.11	0.11
8	0.2	0.940	0.696	0.11	0.11	0.962	0.840	0.11	0.11
8	0	0.942	0.666	0.11	0.11	0.942	0.818	0.11	0.11
12	4	0.936	0.498	0.11	0.11	0.946	0.742	0.11	0.11
12	2	0.922	0.488	0.11	0.11	0.932	0.716	0.11	0.11
12	0.2	0.948	0.544	0.11	0.11	0.956	0.764	0.11	0.11
12	0	0.956	0.552	0.11	0.11	0.962	0.788	0.11	0.11

4.2. Mild correlation on the support

In this subsection, we consider $\Sigma = \Sigma^\circ$ specified in Section 1.2. This is a harder scenario than $\Sigma = I_p$ since ρ_S is increased. Other parameters are set to be the same as in Section 4.1. We see from Table 2 that the debiased Lasso has coverage probabilities much lower than the nominal level at all the sparsity levels while

the BS-DB remains to be robust at most sparsity levels. The most difficult case is $(s, s_-) = (12, 1)$, i.e. the overall sparsity is large and most of them are strong signals, where the confidence intervals given by BS-DB provide lower coverage probabilities for strong signals. One reason is that the estimated noise level given by the scaled-Lasso has large errors in this case. Hence, the choice of λ can be improper and cause large finite sample bias for the initial Lasso estimator. One way to get around this issue is to select λ via cross validation. We can see from Table 4 that the proposed variance estimator has the most reliable performance when $\Sigma = \Sigma^\circ$ in comparison to other two estimators.

TABLE 2
 Coverage probabilities of BS-DB (*cov.bsdb*), coverage probabilities of debiased Lasso (*cov.db*), standard errors of BS-DB (*se.bsdb*), and standard errors of debiased Lasso (*se.db*) when $\Sigma = \Sigma^\circ$.

<i>s</i>	β	$s_- = 1$				$s_- = s/2$			
		<i>cov.bsdb</i>	<i>cov.db</i>	<i>se.bsdb</i>	<i>se.db</i>	<i>cov.bsdb</i>	<i>cov.db</i>	<i>se.bsdb</i>	<i>se.db</i>
4	4	0.936	0.762	0.10	0.10	0.948	0.818	0.10	0.10
4	2	0.950	0.674	0.10	0.10	0.928	0.808	0.10	0.10
4	0.2	0.944	0.662	0.10	0.10	0.924	0.786	0.10	0.10
4	0	0.946	0.846	0.10	0.10	0.940	0.876	0.10	0.10
8	4	0.928	0.530	0.11	0.10	0.926	0.688	0.11	0.11
8	2	0.930	0.464	0.11	0.10	0.946	0.706	0.11	0.11
8	0.2	0.958	0.492	0.11	0.10	0.926	0.702	0.11	0.11
8	0	0.952	0.646	0.11	0.10	0.948	0.810	0.11	0.11
12	4	0.840	0.512	0.23	0.21	0.924	0.622	0.11	0.11
12	2	0.904	0.378	0.23	0.21	0.936	0.610	0.11	0.11
12	0.2	0.948	0.514	0.23	0.21	0.956	0.658	0.11	0.11
12	0	0.940	0.638	0.23	0.21	0.950	0.704	0.11	0.11

4.3. When the irrepresentable condition does not hold

In this subsection, we consider a scenario where the irrepresentable condition does not hold. Specifically, we set $\Sigma_{j,j} = 1, 1 \leq j \leq p$. The upper diagonal of Σ is such that $\Sigma_{j,k} = 0.15$ if $s + 1 \leq k \leq 16$ and $j \leq 16$ and $\Sigma_{j,k} = 0$ otherwise. The true Σ remains to be unknown. It is easy to check that for ϕ defined in Condition 3.3, it holds that $\phi = 0.6$ for $s = 4$, $\phi = 1.2$ for $s = 8$ and $\phi = 1.8$ for $s = 12$. Hence, the irrepresentable condition does not hold with the current Σ when $s \in \{8, 12\}$. We examine the performance of the BS-DB and debiased Lasso in this case. Other parameters are set to be the same as in Section 4.1.

In the current setting, the results in Table 3 show that the BS-DB still provides more accurate coverage than the debiased Lasso but both methods are severely biased when $(s, s_-) = (12, 1)$. One reason is that, as in the case of Section 4.2, the estimated noise levels are largely biased and hence the choices of tuning parameters are improper. The confidence intervals are also wider when $(s, s_-) = (12, 1)$ and one can see that all three variance estimators are severely biased in this case.

TABLE 3
 Coverage probabilities of BS-DB (cov.bsdb), coverage probabilities of debiased Lasso (cov.db), standard errors of BS-DB (se.bsdb), and standard errors of debiased Lasso (se.db) when the irrepresentable condition does not hold.

s	β	$s_- = 1$				$s_- = s/2$			
		cov.bsdb	cov.db	se.bsdb	se.db	cov.bsdb	cov.db	se.bsdb	se.db
4	4	0.934	0.886	0.10	0.10	0.916	0.908	0.10	0.10
4	2	0.908	0.850	0.10	0.10	0.932	0.910	0.10	0.10
4	0.2	0.938	0.874	0.10	0.10	0.950	0.906	0.10	0.10
4	0	0.922	0.356	0.11	0.10	0.940	0.496	0.11	0.10
8	4	0.874	0.396	0.12	0.10	0.850	0.808	0.11	0.10
8	2	0.874	0.390	0.12	0.10	0.856	0.758	0.11	0.10
8	0.2	0.944	0.538	0.12	0.10	0.910	0.820	0.11	0.10
8	0	0.926	0.006	0.12	0.10	0.826	0.076	0.11	0.10
12	4	0.270	0.202	0.32	0.29	0.820	0.620	0.11	0.10
12	2	0.326	0.162	0.33	0.29	0.814	0.608	0.11	0.10
12	0.2	0.278	0.272	0.33	0.29	0.848	0.676	0.11	0.10
12	0	0.084	0.000	0.33	0.29	0.630	0.002	0.11	0.10

TABLE 4
 Median of absolute errors for estimated noise level with identity covariance matrix, $\Sigma = \Sigma^\circ$, and when the irrepresentable condition (IRP) does not hold.

(s, s_-)	$\Sigma = I_p$			$\Sigma = \Sigma^\circ$			IRP does not hold		
	sc-Las	MSR	Prop	sc-Las	MSR	Prop	sc-Las	MSR	Prop
(4,1)	0.06	0.33	0.05	0.07	0.38	0.05	0.06	0.32	0.05
(4,2)	0.06	0.24	0.06	0.06	0.27	0.06	0.05	0.24	0.05
(8,1)	0.23	0.87	0.05	0.35	0.16	0.07	0.22	0.88	0.05
(8,4)	0.12	0.50	0.07	0.14	0.56	0.06	0.09	0.47	0.05
(12,1)	0.69	1.96	0.06	1.42	3.44	0.10	1.99	3.66	1.90
(12,6)	0.26	0.86	0.10	0.31	0.98	0.09	0.20	0.76	0.06

5. Discussion

In this paper, we propose the BS-DB procedure for constructing confidence intervals for regression coefficients in high-dimensional linear models. Our analysis shows that, under the irrepresentable condition and other mild technical conditions, the BS-DB estimator has smaller order of bias in existence of a large proportion of strong signals in comparison to the debiased Lasso considered in the existing literature. If the irrepresentable condition does not hold, then BS-DB is guaranteed to perform no worse than the debiased Lasso asymptotically. Hence, the BS-DB is a robust and competitive alternative of the original debiased Lasso estimator. From our numerical studies, we see that an accurate estimate of noise level is a key ingredient for inference as it can affect the choice of tuning parameters and the uncertainty quantification. It would also be of interest to develop some robust confidence intervals to compensate the bias of the debiased estimators.

6. Proofs of main lemmas and theorems

We first declare some notations. Let $\hat{u} = \hat{\beta} - \beta$ and $\hat{u}^* = \hat{\beta}^* - \hat{\beta}$. For any $k \in S$, let

$$X_k^\perp = X_k - \Sigma_{k,A_k} \Sigma_{A_k,A_k}^{-1} X_{A_k}.$$

Let $A_k = S \setminus \{k\}$ if $j \in S$ and $A_k = S$ if $j \notin S$. Let $\hat{\beta}_{A_k}^{(k)}$ denote the “leave-one-out” Lasso estimate of β_{A_k} given the oracle S :

$$\hat{\beta}^{(k)} = \arg \min_{b \in \mathbb{R}^p} \left\{ \frac{1}{2n} \|y - X_S b_S\|_2^2 + \lambda \|b_{A_j}\|_1 : b_k = \beta_k, b_{S^c} = 0 \right\}. \tag{27}$$

We mention that throughout our proof (except for Theorem 3.7), we repeatedly using the fact that $X\Theta_j$ is independent of X_{-j} when X is Gaussian distributed. This is because the covariance between X_{-j} and $X\Theta_j$ is zero.

6.1. Proof of Theorem 3.3

Proof. First consider the case with **known** Θ_j . (i) If $j \notin S$,

$$\hat{\beta}_j^{(DB)} - \beta_j = (e_j^T - \Theta_j^T \Sigma^n) \hat{u} + \frac{\langle \hat{z}_j, \epsilon \rangle}{n}.$$

In the event that $\{\widehat{S} \subseteq S\}$, $\hat{u}_{S^c} = 0$ and (46) holds. As a result,

$$\hat{\beta}_j^{(DB)} - \beta_j = \underbrace{(-\Theta_j^T \Sigma^n)_S (\Sigma_{S,S}^n)^{-1} W_S^n}_{r_{j,1}} + \underbrace{\lambda (\Theta_j^T \Sigma^n)_S (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S)}_{r_{j,2}} + \frac{\langle \hat{z}_j, \epsilon \rangle}{n}.$$

Conditioning on X , using the sub-Gaussian property of ϵ , we have with probability at least $1 - \exp(-c_1 n)$,

$$\begin{aligned} |r_{j,1}| &\leq C_1 K_2 \sqrt{\frac{\Theta_j^T \Sigma_{\cdot,S}^n (\Sigma_{S,S}^n)^{-1} \Sigma_{S,\cdot}^n \Theta_j}{n}} \leq C_1 K_2 \|\Sigma_{S,\cdot}^n \Theta_j\|_2 \|(\Sigma_{S,S}^n)^{-1}\|_2^{1/2} / \sqrt{n} \\ &\leq C_1 K_2 \sqrt{s} \|(\Sigma_{S,S}^n)^{-1}\|_2^{1/2} / n, \end{aligned}$$

where the last step is due to $X\Theta_j$ is independent of X_S . Using Lemma A.1, given that $n \gg s$,

$$|r_{j,1}| = O_P \left(\frac{K_2 \sqrt{K_1 s}}{n \sqrt{C_{\min}}} \right).$$

We bound $r_{j,2}$ as

$$\begin{aligned} |r_{j,2}| &\leq \lambda \left| \Theta_j^T \Sigma_{\cdot,S}^n (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S) \right| \\ &= \lambda \left| \Theta_{j,S} \text{sgn}(\hat{\beta}_S) + \Theta_{j,S^c} \Sigma_{S^c,S}^n (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S) \right|. \end{aligned}$$

Define

$$X_{S^c}^\perp = X_{S^c} - X_S \Sigma_{S,S}^{-1} \Sigma_{S,S^c}. \tag{28}$$

Then

$$\begin{aligned} (\Sigma_{S,S}^n)^{-1} \Sigma_{S,S^c}^n \Theta_{S^c,j} &= \Sigma_{S,S}^{-1} \Sigma_{S,S^c} \Theta_{S^c,j} + \frac{1}{n} (\Sigma_{S,S}^n)^{-1} X_S^T X_{S^c}^\perp \Theta_{S^c,j} \\ &= -\Theta_{S,j} + \frac{1}{n} (\Sigma_{S,S}^n)^{-1} X_S^T X_{S^c}^\perp \Theta_{S^c,j}, \end{aligned}$$

where the last step is due to $\Theta_{S,j} + \Sigma_{S,S}^{-1} \Sigma_{S,S^c} \Theta_{S^c,j} = 0$ for $j \notin S$. Then we have

$$|r_{j,2}| = \lambda \left| \frac{1}{n} \text{sgn}(\hat{\beta}_S)^T (\Sigma_{S,S}^n)^{-1} X_S^T X_{S^c}^\perp \Theta_{S^c,j} \right|.$$

By the Gaussian property of X , $X_{S^c}^\perp$ is Gaussian and is independent of X_S . Conditioning on X_S and ϵ , we have $x_{i,S^c}^\perp \Theta_{S^c,j} \sim N(0, \Theta_{j,S^c} \Theta_{S^c,S^c}^{-1} \Theta_{S^c,j})$. Notice that

$$\Theta_{j,S^c} \Theta_{S^c,S^c}^{-1} \Theta_{S^c,j} \leq \Theta_{j,j} \leq K_1.$$

Since $\hat{\beta}_S$ is a function of X_S and ϵ , $X_{S^c}^\perp \Theta_{S^c,j}$ is independent of X_S and $\hat{\beta}_S$. Hence,

$$|r_{j,2}| = O_P \left(\lambda \sqrt{\frac{\text{sgn}(\hat{\beta}_S)^T (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S) K_1}{n}} \right) = O_P \left(\lambda \sqrt{\frac{s}{n}} \right),$$

where we use the fact that

$$\text{sgn}(\hat{\beta}_S)^T (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S) \leq \|\text{sgn}(\hat{\beta}_S)\|_2^2 \|(\Sigma_{S,S}^n)^{-1}\|_2 = O_P(s/C_{\min}).$$

Together with the bound for $|r_{j,1}|$, for $j \notin S$ and $\lambda \asymp \sqrt{\log p/n}$, we arrive at

$$|\widehat{rem}_j^o| \leq |r_{j,1}| + |r_{j,2}| = O_P \left(\frac{\sqrt{s \log p}}{n} \right).$$

(ii) If $j \in S$, we use an leave-one-out argument. In the event that $\widehat{S} \subseteq S$, for $\hat{\beta}^{(k)}$ defined via (27), it holds that

$$\begin{aligned} \hat{\beta}_j^{(DB)} - \beta_j &= (e_j^T - \Theta_j^T \Sigma^n) \hat{u} + \frac{\langle \hat{z}_j, \epsilon \rangle}{n} \\ &= \underbrace{(e_j^T - \Theta_j^T \Sigma^n)_{A_j} (\hat{\beta}_{A_j}^{(j)} - \beta_{A_j})}_{r_{j,1}} + \underbrace{(e_j^T - \Theta_j^T \Sigma^n)_S (\hat{\beta}_S - \hat{\beta}_S^{(j)})}_{r_{j,2}} + \frac{\langle \hat{z}_j, \epsilon \rangle}{n}, \end{aligned}$$

where the last step is due to $\hat{\beta}_j^{(j)} = \beta_j$. For $r_{j,1}$, we can bound it similarly as in the case $j \notin S$. In fact, the KKT conditions regarding $\hat{\beta}_{A_j}^{(j)}$ gives

$$\Sigma_{A_j, A_j}^n (\hat{\beta}_{A_j}^{(j)} - \beta_{A_j}) - W_{A_j}^n = -\lambda \text{sgn}(\hat{\beta}_{A_j}^{(j)}), \tag{29}$$

where Σ_{A_j, A_j}^n is invertible with high probability. As a result,

$$r_{j,1} = -\Theta_j^T \Sigma_{\cdot, A_j}^n (\Sigma_{A_j, A_j}^n)^{-1} W_{A_j}^n + \Theta_j^T \Sigma_{\cdot, A_j}^n (\Sigma_{A_j, A_j}^n)^{-1} \lambda \text{sgn}(\hat{\beta}_{A_j}^{(j)}),$$

where the last step is due to $\hat{\beta}_{A_j}^{(j)}$ is a function of X_{-j} and ϵ and $X\Theta_j$ is independent of X_{-j} and ϵ . By similar arguments as for the case $j \in S^c$, we have

$$|r_{j,1}| = O_P \left(\frac{\sqrt{s \log p}}{n} \right).$$

For $r_{j,2}$, we have

$$\begin{aligned} \left| (e_j^T - \Theta_j^T \Sigma^n)_S (\hat{\beta}_S - \hat{\beta}_S^{(j)}) \right| &\leq |1 - \Theta_j^T \Sigma_{\cdot, j}^n| \hat{u}_j + \left| (e_j^T - \Theta_j^T \Sigma^n)_{A_j} (\hat{\beta}_{A_j} - \hat{\beta}_{A_j}^{(j)}) \right| \\ &= O_P(\sqrt{K_1/n}) |\hat{u}_j| + \|\Theta_j^T \Sigma_{\cdot, A_j}^n\|_2 \|\hat{\beta}_{A_j} - \hat{\beta}_{A_j}^{(j)}\|_2 \\ &= O_P(\sqrt{K_1/n}) |\hat{u}_j| + O_P(\sqrt{s/n}) \|\hat{\beta}_{A_j} - \hat{\beta}_{A_j}^{(j)}\|_2, \end{aligned}$$

where last step is due to the independence between $X\Theta_j$ and X_{A_j} . By (46),

$$|\hat{u}_j| \leq \|\hat{u}_S\|_2 = O_P \left(\sqrt{\frac{s \log p}{n}} \right).$$

Using Lemma A.2, $\|\hat{\beta}_{A_j} - \hat{\beta}_{A_j}^{(j)}\|_2 = O_P(\rho_S \lambda)$. Together we have for $j \in S$,

$$|\widehat{rem}_j^o| \leq |r_{j,1}| + |r_{j,2}| = O_P \left(\frac{(\rho_S + 1) \sqrt{s \log p}}{n} \right).$$

It is left to deal with the case where **Θ is unknown**. By definition,

$$\hat{\beta}_j^{(DB)} - \beta_j = (e_j^T - \frac{\hat{\gamma}^T \Sigma^n}{\langle X_j, X \hat{\gamma} \rangle}) \hat{u} + \frac{\hat{z}_j^T \epsilon}{n}.$$

Using KKT condition of (4), we have

$$|\widehat{rem}_j^o| = \left| (e_j^T - \frac{\hat{\gamma}^T \Sigma^n}{\langle X_j, X \hat{\gamma} \rangle}) \hat{u} \right| \leq \left\| e_j^T - \frac{\hat{\gamma}^T \Sigma^n}{\langle X_j, X \hat{\gamma} \rangle} \right\|_\infty \|\hat{u}_S\|_1 = O_P(s \lambda \lambda_j). \quad \square$$

6.2. Proof of Lemma 3.4

Proof. When the results of Lemma 3.2 hold, $\hat{A}_j \subseteq A_j$. Let $\hat{B}_j = A_j \cap \hat{A}_j^c$.

(i) We first study the case where Θ is known. By (9) and (46), in the event Ω_0 (19) we have

$$\begin{aligned} \hat{\beta}_j^{(mDB)} - \beta_j &= (e_j^T - \hat{w}_j^T X/n)\hat{u} + \frac{\hat{w}_j^T \epsilon}{n} = -\hat{w}_j^T X_{A_j} \hat{u}_{A_j}/n + \frac{\hat{w}_j^T \epsilon}{n} \\ &= \hat{b}_j + \underbrace{\frac{1}{n} \hat{w}_j^T X_{A_j} \hat{u}_{A_j}}_{\hat{r}_j} - \hat{b}_j + \frac{\hat{w}_j^T \epsilon}{n}, \end{aligned} \tag{30}$$

where

$$\hat{b}_j = \frac{1}{n} \hat{w}_j^T X_{A_j} (\Sigma_{A_j, A_j}^n)^{-1} \lambda \text{sgn}(\hat{\beta}_{A_j}^{(j)}) \tag{31}$$

for $\hat{\beta}^{(j)}$ defined in (27). For \hat{r}_j , using the definition of \hat{w}_j in (15), we have

$$|\hat{r}_j| \leq \underbrace{\left| \frac{\Theta_j^T X^T P_{A_j}^\perp X_{A_j} (\hat{\beta}_{A_j} - \hat{\beta}_{A_j}^{(j)})}{\langle X_j, P_{A_j}^\perp X \Theta_j \rangle} \right|}_{\hat{r}_{j,1}} + \underbrace{\left| \frac{\Theta_j^T X^T P_{A_j}^\perp X_{A_j} (\hat{\beta}_{A_j}^{(j)} - \beta_{A_j})}{\langle X_j, P_{A_j}^\perp X \Theta_j \rangle} - \hat{b}_j \right|}_{\hat{r}_{j,2}}$$

For the denominators in $\hat{r}_{j,1}$ and $\hat{r}_{j,2}$, noticing that

$$\begin{aligned} \frac{1}{n} |\langle X_j, P_{A_j}^\perp X \Theta_j \rangle| &\geq \frac{1}{n \Theta_{j,j}} |\langle X \Theta_j, P_{A_j}^\perp X \Theta_j \rangle| - \frac{1}{n} |\langle P_{A_j}^\perp X_{-j} \Sigma_{-j, -j}^{-1} \Sigma_{-j, j}, X \Theta_j \rangle| \\ &= \frac{1}{n \Theta_{j,j}} \|P_{A_j}^\perp X \Theta_j\|_2^2 - \frac{1}{n \Theta_{j,j}} |\langle P_{A_j}^\perp X_{-j} \Theta_{j, -j}, X \Theta_j \rangle| \\ &= 1 - O_P(|\hat{A}_j|/n + n^{-1/2}), \end{aligned} \tag{32}$$

where the last step is due to $P_{A_j}^\perp X_{-j}$ is a function of X_{-j} and is independent of $X \Theta_j$. For the numerator of $\hat{r}_{j,1}$,

$$\begin{aligned} |\Theta_j^T X^T P_{A_j}^\perp X_{A_j} (\hat{\beta}_{A_j} - \hat{\beta}_{A_j}^{(j)})| &\leq \|\Theta_j^T X^T P_{A_j}^\perp X_{\hat{B}_j}\|_2 \|\hat{\beta}_{\hat{B}_j} - \hat{\beta}_{\hat{B}_j}^{(j)}\|_2 \\ &\leq \|\Theta_j^T X^T P_{A_j}^\perp X_{\hat{B}_j}\|_2 O_P(\rho_S \lambda \mathbb{1}(j \in S)), \end{aligned}$$

where the first step is due to the projection and the second step is due to Lemma A.2 if $j \in S$ and $\hat{\beta}_S = \hat{\beta}_S^{(j)}$ if $j \notin S$. Since $P_{A_j}^\perp X_{\hat{B}_j}$ is a function of X_{-j} , which is independent of $X \Theta_j$, we have

$$\|\Theta_j^T X^T P_{A_j}^\perp X_{\hat{B}_j}\|_2^2 = O_P(\Theta_{j,j} \text{Tr}(X_{\hat{B}_j}^T P_{A_j}^\perp X_{\hat{B}_j})) = O_P(|\hat{B}_j|n).$$

Since $|\hat{B}_j| \leq s_-$ with high probability, we have

$$|\hat{r}_{j,1}| = O_P(\rho_S \sqrt{s_-} \lambda \mathbb{1}(j \in S) / \sqrt{n}).$$

For $\hat{r}_{j,2}$, by the definition of $\hat{\beta}_{A_j}^{(j)}$ in (27), it holds that

$$\hat{\beta}_{A_j}^{(j)} - \beta_{A_j} = (\Sigma_{A_j, A_j}^n)^{-1} X_{A_j}^T \epsilon/n - \lambda (\Sigma_{A_j, A_j}^n)^{-1} \text{sgn}(\hat{\beta}_{A_j}^{(j)})$$

and hence

$$|\hat{r}_{j,2}| = \left| \frac{\Theta_j^T X^T P_{\hat{A}_j}^\perp X_{A_j} (\Sigma_{A_j, A_j}^n)^{-1} W_{A_j}^n}{\langle X_j, P_{\hat{A}_j}^\perp X \Theta_j \rangle} \right|.$$

Using the sub-Gaussian property of ϵ conditioning on X , we have

$$|\hat{r}_{j,2}| = O_P(\sigma \|P_{A_j} P_{\hat{A}_j}^\perp X \Theta_j\|_2/n) = O_P(\sqrt{s_-}/n).$$

where the last step is due to the independence between $X \Theta_j$ and $P_{\hat{A}_j}^\perp X_{A_j}$ and

$$\|P_{A_j} P_{\hat{A}_j}^\perp X \Theta_j\|_2^2 = O_P(\text{Tr}(P_{A_j} P_{\hat{A}_j}^\perp)) \leq s_-.$$

Together we have,

$$|\hat{r}_j| = O_P(\rho_S \sqrt{s_-} \lambda \mathbb{1}(j \in S)/\sqrt{n}) + o_P(\sqrt{s_-}/n).$$

For the bootstrapped estimate of bias, \hat{b}_j^* defined in (10), for \hat{b}_j defined in (31)

$$\hat{b}_j^* = -\hat{w}_j^T X_{A_j} \hat{u}_{A_j}^*/n = \hat{b}_j - \underbrace{\left(\frac{1}{n} \hat{w}_j^T X_{A_j} \hat{u}_{A_j}^* + \hat{b}_j \right)}_{\hat{r}_j^*}. \tag{33}$$

Notice that

$$\begin{aligned} |\hat{r}_j^*| &\leq \underbrace{\frac{1}{n} |\hat{w}_j^T X_{A_j} (\hat{\beta}_{A_j}^* - \hat{\beta}_{A_j}^{*(j)})|}_{\hat{r}_{j,1}^*} + \underbrace{\frac{1}{n} |\hat{w}_j^T X_{A_j} (\hat{\beta}_{A_j}^{*(j)} - \hat{\beta}_{A_j}^{(j)}) - \hat{b}_j|}_{\hat{r}_{j,2}^*} \\ &\quad + \underbrace{\frac{1}{n} |\hat{w}_j^T X_{A_j} (\hat{\beta}_{A_j}^{(j)} - \hat{\beta}_{A_j})|}_{\hat{r}_{j,3}^*}. \end{aligned}$$

where $\hat{r}_{j,3}^* = \hat{r}_{j,1}$. By Lemma A.3, $\hat{r}_{j,1}^*$ can be similarly bounded as for $\hat{r}_{j,1}$. Therefore,

$$|\hat{r}_{j,1}^*| + |\hat{r}_{j,3}^*| = O_P(\sqrt{s_-} \rho_S \lambda \mathbb{1}(j \in S)/\sqrt{n}).$$

For $\hat{r}_{j,2}^*$, we have

$$\begin{aligned} |\hat{r}_{j,2}^*| &= \left| \frac{1}{n} \hat{w}_j^T X_{A_j} (\hat{\beta}_{A_j}^{*(j)} - \hat{\beta}_{A_j}^{(j)}) + \frac{1}{n} \lambda \hat{w}_j^T X_{A_j} (\Sigma_{A_j, A_j}^n)^{-1} \text{sgn}(\hat{\beta}_{A_j}^{(j)}) \right| \\ &\leq \underbrace{\left| \frac{\lambda}{n \lambda} \hat{w}_j^T X_{A_j} (\Sigma_{A_j, A_j}^n)^{-1} W_{A_j}^* \right|}_{\hat{r}_{j,2}^*} + \underbrace{\left| \frac{\lambda}{n} \hat{w}_j^T X_{A_j} (\Sigma_{A_j, A_j}^n)^{-1} [\text{sgn}(\hat{\beta}_{A_j}^{(j)}) - \text{sgn}(\hat{\beta}_{A_j}^{*(j)})] \right|}_{\hat{r}_{j,2}^*}, \end{aligned}$$

where the first term can be similarly bounded as $\hat{r}_{j,2}$. The second term can be rewritten using the definition of \hat{w}_j :

$$\tilde{r}_{j,2}^* = \frac{\lambda}{n} \frac{|\Theta_j^T X^T P_{\hat{A}_j}^\perp X_{A_j} (\Sigma_{A_j, A_j}^n)^{-1} [\text{sgn}(\hat{\beta}_{A_j}^{(j)}) - \text{sgn}(\hat{\beta}_{A_j}^{*(j)})]|}{|X_j^T P_{\hat{A}_j}^\perp X \Theta_j|}$$

Noticing that $\hat{\beta}_{A_j}^{*(j)}$, and $\hat{\beta}_{A_j}^{(j)}$ are functions of X_{-j} and ϵ . Hence, $X\Theta_j$ is independent of $P_{\hat{A}_j}^\perp X_{A_j}$, $\hat{\beta}_{A_j}^{*(j)}$, and $\hat{\beta}_{A_j}^{(j)}$ no matter $j \in S$ or not. Hence,

$$\begin{aligned} \tilde{r}_{j,2}^* &= O_P\left(\frac{\lambda}{n}\|P_{\hat{A}_j}^\perp X_{A_j}(\Sigma_{A_j,A_j}^n)^{-1}[sgn(\hat{\beta}_{A_j}^{(j)}) - sgn(\hat{\beta}_{A_j}^{*(j)})]\|_2\right) \\ &= O_P\left(\frac{\lambda}{\sqrt{n}}\|(\Sigma_{A_j,A_j}^n)^{-1}\|_2^{1/2}\|sgn(\hat{\beta}_{A_j}^{(j)}) - sgn(\hat{\beta}_{A_j}^{*(j)})\|_2\right). \end{aligned}$$

If $j \notin S$, $\hat{\beta}_S^{(j)} = \check{\beta}_S$ defined in (43) and $\hat{\beta}_S^{*(j)} = \check{\beta}_S^*$ defined in (48). We have

$$\begin{aligned} &\|sgn(\hat{\beta}_{A_j}^{(j)}) - sgn(\hat{\beta}_{A_j}^{*(j)})\|_2 \\ &\leq \|sgn(\hat{b}_{A_j}^*) - sgn(\beta_{A_j})\|_2 + \|sgn(\hat{\beta}_{A_j}^*) - sgn(\beta_{A_j})\|_2 \leq \sqrt{s_-}. \end{aligned}$$

If $j \in S$, we have for a large enough constant C , it holds with high probability that

$$\|\hat{\beta}_{A_j}^{(j)} - \hat{\beta}_{A_j}^{*(j)}\|_\infty \leq \|\hat{\beta}_{A_j}^{(j)} - \hat{\beta}_{A_j}\|_\infty + \|\hat{\beta}_{A_j}^* - \hat{\beta}_{A_j}^{*(j)}\|_\infty + \|\hat{\beta}_{A_j}^* - \hat{\beta}_{A_j}\|_\infty \leq C\rho_S\lambda,$$

where the last step is due to Lemma A.2, Lemma A.3 and Lemma 3.2. Therefore, with high probability,

$$\begin{aligned} \|sgn(\hat{\beta}_{A_j}^{(j)}) - sgn(\hat{\beta}_{A_j}^{*(j)})\|_2^2 &\leq |\{k \in A_j : |\hat{\beta}_k^{(j)}| \leq C^*\rho_S\lambda\}| \\ &\leq |\{k \in A_j : |\beta_k| \leq (C + C^*)\rho_S\lambda\}| \leq |S_-|. \end{aligned}$$

Hence, we have proved

$$|\hat{r}_j^*| = O_P(\sqrt{s_-}\lambda/\sqrt{n} + \sqrt{s_-}\rho_S\lambda/\sqrt{n}\mathbb{1}(j \in S)).$$

Invoking (30) and (33),

$$\hat{\beta}_j^{(BS-DB)} - \beta_j = \frac{\hat{w}_j^T \epsilon}{n} + \hat{b}_j - \hat{r}_j - (\hat{b}_j - \hat{r}_j^*) = \frac{\hat{w}_j^T \epsilon}{n} + |\hat{r}_j| + |\hat{r}_j^*|.$$

The proof for known Θ is complete.

(ii) Next we consider **unknown Θ** . By (9) and Lemma 3.1 we have

$$\begin{aligned} \hat{\beta}_j^{(mDB)} - \beta_j &= (e_j^T - \hat{w}_j^T X/n)\hat{u} + \frac{\hat{w}_j^T \epsilon}{n} = -\hat{w}_j^T X_{A_j} \hat{u}_{A_j}/n + \frac{\hat{w}_j^T \epsilon}{n} \\ &= \frac{1}{n}\hat{w}_j^T X_{A_j} e_{A_j}^T (\Sigma_{S,S}^n)^{-1}[\lambda sgn(\hat{\beta}_S) - W_S^n] + \frac{\hat{w}_j^T \epsilon}{n} \\ &= \hat{b}_j - \underbrace{\frac{1}{n}\hat{w}_j^T X_{A_j} e_{A_j}^T (\Sigma_{S,S}^n)^{-1} W_S^n}_{\hat{r}_j} + \frac{\hat{w}_j^T \epsilon}{n}, \end{aligned}$$

where

$$\hat{b}_j = -\frac{1}{n}\hat{w}_j^T X_{A_j} e_{A_j}^T (\Sigma_{S,S}^n)^{-1} \lambda sgn(\hat{\beta}_S). \tag{34}$$

By the KKT condition of (16),

$$\Sigma_{\hat{A}_j}^n \hat{\kappa} = 0 \text{ and } \Sigma_{\hat{S}^c \setminus \{j\}}^n \hat{\kappa} = -\lambda_j \text{sgn}(\hat{\kappa}_{\hat{S}^c \setminus \{j\}}). \quad (35)$$

As a result,

$$\begin{aligned} |\hat{r}_j| &= \left| \frac{1}{n} \hat{w}_j^T X_{\hat{B}_j} e_{\hat{B}_j}^T (\Sigma_{S,S}^n)^{-1} W_S^n \right| \leq |\hat{B}_j| \|\hat{w}_j^T X_{\hat{B}_j} / n\|_\infty \|e_{\hat{B}_j}^T (\Sigma_{S,S}^n)^{-1} W_S^n\|_\infty \\ &\leq |\hat{B}_j| \|\hat{w}_j^T X_{\hat{B}_j} / n\|_\infty \max_{k \in S_-} \|e_k^T (\Sigma_{S,S}^n)^{-1} W_S^n\|_\infty = O_P(s_- \lambda_j \lambda), \end{aligned}$$

where the first and last step are due to the (35), the third step is due to $\hat{B}_j \subseteq S_-$ with high probability.

For the bootstrapped estimate of bias,

$$\begin{aligned} \hat{b}_j^* &= -\hat{w}_j^T X_{A_j} \hat{u}_{A_j}^* / n = \frac{\lambda}{n} \hat{w}_j^T X_{A_j} e_{A_j}^T (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S^*) \\ &= \hat{b}_j + \underbrace{\frac{1}{n} \hat{w}_j^T X_{A_j} e_{A_j}^T (\Sigma_{S,S}^n)^{-1} [\lambda \text{sgn}(\hat{\beta}_S^*) - \lambda \text{sgn}(\hat{\beta}_S)]}_{\hat{r}_j^*}. \end{aligned}$$

Now we bound \hat{r}_j^* . Using (35) again, we have

$$\begin{aligned} |\hat{r}_j^*| &= \lambda \left| \frac{1}{n} \hat{w}_j^T X_{A_j} e_{A_j}^T (\Sigma_{S,S}^n)^{-1} [\text{sgn}(\hat{\beta}_S^*) - \text{sgn}(\hat{\beta}_S)] \right| \\ &= \frac{n \lambda \lambda_j}{\langle X_j, P_{A_j}^\perp X \hat{\kappa} \rangle} \left| \text{sgn}(\hat{\kappa}_{\hat{B}_j})^T e_{\hat{B}_j}^T (\Sigma_{S,S}^n)^{-1} e_{S_-} (\text{sgn}(\hat{\beta}_{S_-}^*) - \text{sgn}(\hat{\beta}_{S_-})) \right| \\ &\leq \lambda \lambda_j (1 + o_P(1)) \|\text{sgn}(\hat{\kappa}_{\hat{B}_j})\|_2 \|(\Sigma_{S,S}^n)^{-1}\|_2 \|\text{sgn}(\hat{\beta}_{S_-}^*) - \text{sgn}(\hat{\beta}_{S_-})\|_2 \\ &\leq O_P(s_- \lambda \lambda_j / C_{\min}), \end{aligned}$$

where the first step is by (35), the second step is due to (32) and the last step is due to Lemma 3.2. Therefore, $\hat{r}_j^* = O_P(s_- \lambda \lambda_j)$. \square

Proof of Theorem 3.6. By Lemma 3.4, we are left to establish the asymptotic normality of $\langle \hat{w}_j, \epsilon \rangle / \sqrt{n}$.

First consider the case where Θ_j is known. Since $P_{\hat{A}_j}^\perp \epsilon$ is a function of ϵ and X_{-j} , it is independent of $X \Theta_j$. As $\langle X \Theta_j, P_{\hat{A}_j}^\perp \epsilon \rangle | X_{-j}, \epsilon \sim N(0, \|P_{\hat{A}_j}^\perp \epsilon\|_2^2 \Theta_{j,j})$ and $\langle X_j, P_{\hat{A}_j}^\perp X \Theta_j \rangle / n = \Theta_{j,j} + o_P(1)$, we have for any $t \in \mathbb{R}$,

$$\mathbb{P}(\langle \hat{w}_j, \epsilon \rangle / \sqrt{n} \leq t | X_{-j}, \epsilon) = \Phi(\Theta_{j,j}^{1/2} t / \sigma) + o(1).$$

Together with Lemma 3.4,

$$\mathbb{P}\left(\sqrt{n}(\hat{\beta}_j^{(BS-DB)} - \beta_j) \leq t\right) = \Phi(\Theta_{j,j}^{1/2} t / \sigma) + o(1).$$

The fact that $\|\hat{w}_j\|_2 / \sqrt{n} = \Theta_{j,j}^{1/2} + o_P(1)$ can be similarly shown as above.

Next, we consider the case where Θ is unknown. Let $\kappa^* = \Theta_{j,j}^{-1}\Theta_j$ and κ be such that

$$X_{\hat{A}_j}^T X \kappa = 0 \text{ and } \kappa_{\hat{A}_j^c} = \kappa_{\hat{A}_j^c}^*.$$

Notice that $\kappa_{\hat{A}_j} = -(\Sigma_{\hat{A}_j, \hat{A}_j}^n)^{-1} \Sigma_{\hat{A}_j, \hat{A}_j^c}^n \kappa_{\hat{A}_j^c} = -(\Sigma_{\hat{A}_j, \hat{A}_j}^n)^{-1} \Sigma_{\hat{A}_j, \hat{A}_j^c}^n \kappa_{\hat{A}_j^c}^*$ and $\kappa_j = \kappa_j^* = 1$. Let $\hat{v} = \hat{\kappa} - \kappa$. For some positive constant ϕ_0 , define

$$\mathcal{E}_1 = \left\{ \|\Sigma_{\hat{S}^c \cap \{-j\}}^n \kappa\|_\infty \leq \lambda_j/2, \right. \\ \left. \inf_{|J| \leq \|\kappa_{\hat{S}^c \cap \{-j\}}\|_0} \inf_{\|u_{J^c}\|_1 \leq 3\|u_J\|_1 \neq 0} \frac{\|P_{\hat{A}_j}^\perp X_{\hat{S}^c \cap \{-j\}} u\|_2^2}{n\|u_J\|_2^2} \geq \phi_0 > 0 \right\}. \quad (36)$$

We first prove the desired results in the event \mathcal{E}_1 . At the end of the proof, we verify that \mathcal{E}_1 holds with high probability.

In view of (16), we have

$$\begin{aligned} \frac{1}{2n} \|X \hat{v}\|_2^2 &\leq |\hat{v}^T \Sigma^n \kappa| + \lambda_j \|\kappa_{\hat{S}^c \cap \{-j\}}\|_1 - \lambda_j \|\hat{\kappa}_{\hat{S}^c \cap \{-j\}}\|_1 \\ &= |\hat{v}_{\hat{A}_j^c}^T \Sigma_{\hat{A}_j^c}^n \kappa| + \lambda_j \|\kappa_{\hat{S}^c \cap \{-j\}}\|_1 - \lambda_j \|\hat{\kappa}_{\hat{S}^c \cap \{-j\}}\|_1 \\ &\leq \|\hat{v}_{\hat{S}^c \cap \{-j\}}\|_1 \|\Sigma_{\hat{S}^c \cap \{-j\}}^n \kappa\|_\infty + \lambda_j \|\kappa_{\hat{S}^c \cap \{-j\}}\|_1 - \lambda_j \|\hat{\kappa}_{\hat{S}^c \cap \{-j\}}\|_1, \end{aligned}$$

where the second step follows from the definition of κ and the last step is due to $\kappa_j = \hat{\kappa}_j = 1$. In event \mathcal{E}_1 , the following oracle inequality holds:

$$\frac{1}{2} \hat{v}^T \Sigma^n \hat{v} \leq \lambda_j \|\kappa_{\hat{S}^c \cap \{-j\}}\|_1 - \lambda_j \|\hat{\kappa}_{\hat{S}^c \cap \{-j\}}\|_1 + \frac{\lambda_j}{2} \|\hat{v}_{\hat{S}^c \cap \{-j\}}\|_1, \quad (37)$$

where

$$\hat{v}^T \Sigma^n \hat{v} = \frac{1}{n} \|P_{\hat{A}_j}^\perp X_{\hat{S}^c \cap \{-j\}} \hat{v}_{\hat{S}^c \cap \{-j\}}\|_2^2.$$

Let J denote the support of $\kappa_{\hat{S}^c \cap \{-j\}}$. Standard decomposition of the right hand side of (37) leads to

$$\frac{1}{2n} \|P_{\hat{A}_j}^\perp X_{\hat{S}^c \cap \{-j\}} \hat{v}_{\hat{S}^c \cap \{-j\}}\|_2^2 \leq \frac{3\lambda_j}{2} \sum_{k \in J} |\hat{v}_k| - \frac{\lambda_j}{2} \sum_{k \in \hat{S}^c \cap \{-j\} \setminus J} |\hat{v}_k|.$$

Using the second statement in event \mathcal{E}_1 , we have

$$\frac{\phi_0}{2} \|\hat{v}_J\|_2^2 \leq \frac{3\lambda_j}{2} \sum_{k \in J} |\hat{v}_k| \leq \frac{3\lambda_j}{2} \sqrt{|J|} \|\hat{v}_J\|_2.$$

Hence,

$$\begin{aligned} \|\hat{v}_J\|_2 &\leq \frac{3\lambda_j \sqrt{|J|}}{\phi_0}, \quad \|\hat{v}_J\|_1 \leq \sqrt{|J|} \|\hat{v}_J\|_2 \leq \frac{3\lambda_j |J|}{\phi_0}, \text{ and} \\ \|\hat{v}_{\hat{S}^c \cap \{-j\}}\|_1 &\leq (1+3) \|\hat{v}_J\|_1 \leq \frac{12\lambda_j |J|}{\phi_0}. \end{aligned}$$

According to our choice of λ_j , we arrive at

$$\|\hat{v}_{\hat{A}_j^c}\|_1 = O_P\left(\frac{\|\kappa_{\hat{S}^c \cap \{-j\}}\|_0}{\phi_0} \sqrt{\frac{\log p}{n}}\right).$$

Therefore,

$$\begin{aligned} \frac{\langle \epsilon, X \hat{\kappa} \rangle}{n} &= \frac{\langle \epsilon, X \kappa \rangle}{n} + \frac{\langle \epsilon, X_{\hat{S}^c \cap \{-j\}} \hat{v}_{\hat{S}^c \cap \{-j\}} \rangle}{n} \\ &= \frac{\langle \epsilon, X \kappa \rangle}{n} + \left\| \frac{X^T \epsilon}{n} \right\|_{\infty} \|\hat{v}_{\hat{S}^c \cap \{-j\}}\|_1 = \frac{\langle \epsilon, X \kappa \rangle}{n} + O_P\left(\frac{\|\Theta_{\hat{S}^c \cap \{-j\}, j}\|_0 \log p}{n}\right), \end{aligned}$$

where $\langle \epsilon, X \kappa \rangle/n = \langle \epsilon, P_{\hat{A}_j}^\perp X_{\hat{A}_j^c} \kappa_{\hat{A}_j^c}^* \rangle/n$. The rest of proof follows from the above statement for known Θ .

It is left to verify $\mathbb{P}(\mathcal{E}_1) \rightarrow 1$ for \mathcal{E}_1 defined in (36). We first notice that

$$\begin{aligned} \Sigma_{\hat{S}^c \cap \{-j\}}^n \kappa &= \frac{1}{n} \langle X_{\hat{S}^c \cap \{-j\}}, X_j + X_{\hat{A}_j} \kappa_{\hat{A}_j} + X_{\hat{S}^c \cap \{-j\}} \kappa_{\hat{S}^c \cap \{-j\}} \rangle \\ &= \frac{1}{n} \langle X_{\hat{S}^c \cap \{-j\}}, P_{\hat{A}_j}^\perp (X_j + X_{\hat{S}^c \cap \{-j\}} \kappa_{\hat{S}^c \cap \{-j\}}) \rangle \\ &= \frac{1}{n} \langle X_{\hat{S}^c \cap \{-j\}}, P_{\hat{A}_j}^\perp X_{\hat{A}_j^c} \kappa_{\hat{A}_j^c} \rangle = \frac{1}{n} \langle X_{\hat{S}^c \cap \{-j\}}, P_{\hat{A}_j}^\perp X \kappa^* \rangle, \end{aligned}$$

where the last step is due to $P_{\hat{A}_j}^\perp X_{\hat{A}_j^c} \kappa_{\hat{A}_j^c} = P_{\hat{A}_j}^\perp X_{\hat{A}_j^c} \kappa_{\hat{A}_j^c} + P_{\hat{A}_j}^\perp X_{\hat{A}_j} \kappa_{\hat{A}_j}^*$ and $\kappa_{\hat{A}_j^c} = \kappa_{\hat{A}_j^c}^*$. By the Gaussian property of X , $X \kappa^*$ is independent of X_{-j} . As $P_{\hat{A}_j}^\perp X_{\hat{S}^c \cap \{-j\}}$ is a function of X_{-j} ,

$$\mathbb{P}\left(\|\Sigma_{\hat{S}^c \cap \{-j\}}^n \kappa\|_{\infty} \geq c_1 \sqrt{\log p/n}\right) = o(1)$$

for some positive constant c_1 . Hence, for $\lambda_j = C \sqrt{\log p/n_0}$ with large enough constant C , the first statement in event \mathcal{E}_1 holds with high probability.

For the second statement, in the event that $S_+ \subseteq \tilde{S} \subseteq S$, it holds that

$$\begin{aligned} &\mathbb{P}\left(\inf_{|J| \leq \|\kappa_{\hat{S}^c \cap \{-j\}}\|_0} \inf_{\|u_{J^c}\|_1 \leq 3\|u_J\|_1 \neq 0} \frac{\|P_{\hat{A}_j}^\perp X_{\hat{S}^c \cap \{-j\}} u\|_2^2}{n\|u_J\|_2^2} \leq \phi_0\right) \\ &\leq 2^{s^-} \max_{S_+ \subseteq \tilde{S} \subseteq S} \mathbb{P}\left(\inf_{|J| \leq \|\kappa_{\hat{S}^c \cap \{-j\}}\|_0} \inf_{\|u_{J^c}\|_1 \leq 3\|u_J\|_1 \neq 0} \frac{\|P_{\hat{A}_j}^\perp X_{\hat{S}^c \cap \{-j\}} u\|_2^2}{n\|u_J\|_2^2} \leq \phi_0\right) + o(1), \end{aligned}$$

where $\tilde{A}_j = \tilde{S} \cap \{-j\}$ and inequality is because there are at most 2^{s^-} different \tilde{S} . We further use the fact that

$$\inf_{\|u_{J^c}\|_1 \leq 3\|u_J\|_1 \neq 0} \frac{\|P_{\tilde{A}_j}^\perp X_{\hat{S}^c \cap \{-j\}} u\|_2^2}{n\|u_J\|_2^2} = \inf_{\|u_{J^c}\|_1 \leq 3\|u_J\|_1 \neq 0} \frac{\|P_{\tilde{A}_j}^\perp X_{\tilde{S}^c \cap \{-j\}} u\|_2^2}{n\|u_J\|_2^2},$$

where $X_{\tilde{S}^c \cap \{-j\}}^\perp = X_{\hat{S}^c \cap \{-j\}} - X_{\tilde{A}_j} \Sigma_{\tilde{A}_j, \tilde{A}_j}^{-1} \Sigma_{\tilde{A}_j, \hat{S}^c \cap \{-j\}}$ is independent of $X_{\tilde{A}_j}$. Let $\{g_1, \dots, g_{n-|\tilde{A}_j|}\}$ be an orthonormal basis for the complement of the column

space of $X_{\tilde{A}_j}$. Then we have

$$\|P_{\tilde{A}_j}^\perp X_{\tilde{S}^c \cap \{-j\}}^\perp u\|_2^2 = \sum_{k=1}^{n-|\tilde{A}_j|} \|g_k^T X_{\tilde{S}^c \cap \{-j\}}^\perp u\|_2^2,$$

where $g_k^T X_{\tilde{S}^c \cap \{-j\}}^\perp |X_{\tilde{A}_j}$ is Gaussian with mean zero and variance $g_k^T \Sigma_{\tilde{S}^c \cap \{-j\}, \tilde{S}^c \cap \{-j\}}^\perp g_k$, where

$$\Sigma_{\tilde{S}^c \cap \{-j\}, \tilde{S}^c \cap \{-j\}}^\perp = \Sigma_{\tilde{S}^c \cap \{-j\}, \tilde{S}^c \cap \{-j\}} - \Sigma_{\tilde{S}^c \cap \{-j\}, \tilde{A}_j} \Sigma_{\tilde{A}_j, \tilde{A}_j}^{-1} \Sigma_{\tilde{A}_j, \tilde{S}^c \cap \{-j\}}.$$

Moreover, $v_k^T X_{\tilde{S}^c \cap \{-j\}}^\perp$ is independent of $v_{k'}^T X_{\tilde{S}^c \cap \{-j\}}^\perp$ for $k \neq k'$ conditioning on $X_{\tilde{A}_j}$. We can then use Theorem 1.6 of Zhou [37] to show that

$$\begin{aligned} & 2^{s-} \mathbb{P} \left(\inf_{|J| \leq \|\kappa_{\tilde{S}^c \cap \{-j\}}\|_0} \inf_{\|u_{J^c}\|_1 \leq 3\|u_J\|_1 \neq 0} \sum_{k=1}^{n-|\tilde{A}_j|} \|g_k^T X_{\tilde{S}^c \cap \{-j\}}^\perp u\|_2^2 \leq \phi_0 \right) \\ & \leq 2^{s-} \mathbb{P} \left(\inf_{|J| \leq \|\kappa_{(S_+)^c}\|_0} \inf_{\|u_{J^c}\|_1 \leq 3\|u_J\|_1 \neq 0} \sum_{k=1}^{n-s} \|g_k^T X_{\tilde{S}^c \cap \{-j\}}^\perp u\|_2^2 \leq \phi_0 \right) = o(1) \end{aligned}$$

if $n-s \gg \max\{\max\{\|\kappa_{(S_+)^c}\|_0, 1\} \log p, s_-\}$ and $\Lambda_{\min}(\Sigma_{\tilde{S}^c \cap \{-j\}, \tilde{S}^c \cap \{-j\}}^\perp) \geq 2\phi_0$. In fact,

$$\{\Sigma_{\tilde{S}^c \cap \{-j\}, \tilde{S}^c \cap \{-j\}}^\perp\}^{-1} = \{(\Sigma_{-j, -j})^{-1}\}_{\tilde{S}^c \cap \{-j\}, \tilde{S}^c \cap \{-j\}},$$

which is positive definite for any $S_+ \subseteq \tilde{S} \subseteq S$. \square

Acknowledgements

The author would like to thank Professor Cun-Hui Zhang for insightful discussions on the technical part of the paper as well as the presentation.

References

- [1] Belloni, A., V. Chernozhukov, and C. Hansen (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies* 81(2), 608–650. [MR3207983](#)
- [2] Belloni, A., V. Chernozhukov, and L. Wang (2011). Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika* 98(4), 791–806. [MR2860324](#)
- [3] Belloni, A., V. Chernozhukov, and Y. Wei (2016). Post-selection inference for generalized linear models with many controls. *Journal of Business & Economic Statistics* 34(4), 606–619. [MR3547999](#)
- [4] Bühlmann, P. and S. van de Geer (2015). High-dimensional inference in misspecified linear models. *Electronic Journal of Statistics* 9(1), 1449–1473. [MR3367666](#)

- [5] Cai, T. T. and Z. Guo (2017). Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity. *The Annals of Statistics* 45(2), 615–646. [MR3650395](#)
- [6] Candès, E. and T. Tao (2007). The dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics* 35(6), 2313–2351. [MR2382644](#)
- [7] Chatterjee, A. and S. Lahiri (2010). Asymptotic properties of the residual bootstrap for lasso estimators. *Proceedings of the American Mathematical Society* 138(12), 4497–4509. [MR2680074](#)
- [8] Chatterjee, A. and S. N. Lahiri (2011). Bootstrapping lasso estimators. *Journal of the American Statistical Association* 106(494), 608–625. [MR2847974](#)
- [9] Chatterjee, A. and S. N. Lahiri (2013). Rates of convergence of the adaptive lasso estimators to the oracle distribution and higher order refinements by the bootstrap. *The Annals of Statistics* 41(3), 1232–1259. [MR3113809](#)
- [10] Chernozhukov, V., D. Chetverikov, et al. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics* 21(1), C1–C68. [MR3769544](#)
- [11] Chernozhukov, V., D. Chetverikov, and K. Kato (2013). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics* 41(6), 2786–2819. [MR3161448](#)
- [12] Chernozhukov, V., C. Hansen, and M. Spindler (2015). Valid post-selection and post-regularization inference: An elementary, general approach. *Annu. Rev. Econ.* 7(1), 649–688.
- [13] Dezeure, R., P. Bühlmann, and C.-H. Zhang (2017). High-dimensional simultaneous inference with the bootstrap. *Test* 26(4), 685–719. [MR3713586](#)
- [14] Fan, J. and R. Li (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association* 96(456), 1348–1360. [MR1946581](#)
- [15] Fang, E. X., Y. Ning, and H. Liu (2017). Testing and confidence intervals for high dimensional proportional hazards models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79(5), 1415–1437. [MR3731669](#)
- [16] Jankova, J. and S. Van De Geer (2015). Confidence intervals for high-dimensional inverse covariance estimation. *Electronic Journal of Statistics* 9(1), 1205–1229. [MR3354336](#)
- [17] Javanmard, A. and A. Montanari (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *Journal of Machine Learning Research* 15, 2869–2909. [MR3277152](#)
- [18] Javanmard, A. and A. Montanari (2018). Debiasing the lasso: Optimal sample size for gaussian designs. *The Annals of Statistics* 46(6A), 2593–2622. [MR3851749](#)
- [19] Knight, K. and W. Fu (2000). Asymptotics for lasso-type estimators. *The Annals of Statistics* 28(5), 1356–1378. [MR1805787](#)
- [20] Lee, J. D., D. L. Sun, et al. (2016). Exact post-selection inference, with application to the lasso. *The Annals of Statistics* 44(3), 907–927. [MR3485948](#)

- [21] Lockhart, R., J. Taylor, R. J. Tibshirani, and R. Tibshirani (2014). A significance test for the lasso. *The Annals of statistics* 42(2), 413–468. [MR3210970](#)
- [22] Mammen, E. (1993). Bootstrap and wild bootstrap for high dimensional linear models. *The Annals of Statistics* 21(1), 255–285. [MR1212176](#)
- [23] Meinshausen, N. and P. Bühlmann (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics* 34(3), 1436–1462. [MR2278363](#)
- [24] Meinshausen, N. and P. Bühlmann (2010). Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(4), 417–473. [MR2758523](#)
- [25] Sun, T. and C.-H. Zhang (2012). Scaled sparse linear regression. *Biometrika* 99(4), 879–898. [MR2999166](#)
- [26] Tibshirani, R. (1996). Regression selection and shrinkage via the lasso. *Journal of the Royal Statistical Society B* 58(1), 267–288. [MR1379242](#)
- [27] Tibshirani, R. J., J. Taylor, R. Lockhart, and R. Tibshirani (2016). Exact post-selection inference for sequential regression procedures. *Journal of the American Statistical Association* 111(514), 600–620. [MR3538689](#)
- [28] Van de Geer, S., P. Bühlmann, et al. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics* 42(3), 1166–1202. [MR3224285](#)
- [29] Vershynin, R. (2010). Introduction to the non-asymptotic analysis of random matrices. [arXiv:1011.3027](#). [MR2963170](#)
- [30] Wainwright, M. J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Transactions on Information Theory* 55(5), 2183–2202. [MR2729873](#)
- [31] Wasserman, L. and K. Roeder (2009). High dimensional variable selection. *Annals of statistics* 37(5A), 2178–2201. [MR2543689](#)
- [32] Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of statistics* 38(2), 894–942. [MR2604701](#)
- [33] Zhang, C.-H. and S. Zhang (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 76(1), 217–242. [MR3153940](#)
- [34] Zhang, X. and G. Cheng (2017). Simultaneous inference for high-dimensional linear models. *Journal of the American Statistical Association* 112(518), 757–768. [MR3671768](#)
- [35] Zhao, P. and B. Yu (2006). On model selection consistency of Lasso. *The Journal of Machine Learning Research* 7, 2541–2563. [MR2274449](#)
- [36] Zhao, S., A. Shojaie, and D. Witten (2017). In defense of the indefensible: a very naive approach to high-dimensional inference. [arXiv:1705.05543](#).
- [37] Zhou, S. (2009). Restricted eigenvalue conditions on subgaussian random matrices. [arXiv:0912.4045](#).
- [38] Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association* 101(476), 1418–1429. [MR2279469](#)

Appendix A: Proof of lemmas and theorems in Section 3

A.1. Some technical lemmas

Lemma A.1. *Under Condition 3.1–Condition 3.3, we have the following results. Let $c_1 > 4$, $c_2 > 0$. For $n > c_1 s$, with probability at least $1 - 2 \exp(-c_2 n)$,*

$$C_{\min}/4 \leq \Lambda_{\min}(\Sigma_{S,S}^n) \leq \Lambda_{\max}(\Sigma_{S,S}^n) \leq 4C_{\max}. \tag{38}$$

Proof of Lemma A.1. The proof is standard. See Corollary 5.35 of [29]. \square

Lemma A.2. *Under conditions 3.1–3.3, if $s \log p \leq c_0 n$ for a small enough constant c_0 , then it holds that with probability at least $1 - c_1/p - c_2 s/n - \exp(-c_3 n)$, there exists a large enough constant C such that for $\hat{\beta}_{A_k}^{(k)}$ defined in (27),*

$$\max_{k \in S} \left\| \hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)} \right\|_2 \leq C \rho_S \lambda.$$

Proof of Lemma A.2. We justify Lemma A.2 in the event $\{C_1 \leq \Lambda_{\min}(\Sigma_{S,S}^n) \leq \Lambda_{\max}(\Sigma_{S,S}^n) \leq C_2\}$, which holds with high probability by Lemma A.1. We similarly show the results of Lemma 6.4–Lemma 6.7 in Javanmard and Montanari [18] under current conditions. Specifically, define

$$u_k(b) = \frac{X_k^T (\epsilon + X_{A_k} (\beta_{A_k} - b))}{n}.$$

By (82) of Javanmard and Montanari [18], we have

$$\frac{1}{2n} \left\| P_{X_k}^\perp X_{A_k} (\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}) \right\|_2^2 \leq |u(\hat{\beta}_{A_k}^{(k)})| \left| u(\hat{\beta}_{A_k}) - u(\hat{\beta}_{A_k}^{(k)}) \right|. \tag{39}$$

Noticing that

$$\begin{aligned} \max_{k \in S} |u(\hat{\beta}_{A_k}^{(k)})| &= \max_{k \in S} \left| \frac{X_k^T (\epsilon + X_{A_k} (\beta_{A_k} - \hat{\beta}_{A_k}^{(k)}))}{n} \right| \\ &\leq \max_{k \in S} \left| \frac{X_k^T \epsilon}{n} \right| + \max_{k \in S} \left| \Sigma_{k, A_k} \Sigma_{A_k, A_k}^{-1} \Sigma_{A_k, A_k}^n (\beta_{A_k} - \hat{\beta}_{A_k}^{(k)}) \right| \\ &\quad + \max_{k \in S} \left| \frac{\langle X_k^\perp, X_{A_k} (\beta_{A_k} - \hat{\beta}_{A_k}^{(k)}) \rangle}{n} \right|, \end{aligned}$$

where the first term is no larger than $c_1 \sqrt{\log p/n}$ with probability at least $1 - c_2/p$, the second term can be bounded by

$$\begin{aligned} &\max_{k \in S} \frac{1}{e_k^T \Sigma_{S,S}^{-1} e_k} |e_k^T \Sigma_{S,S}^{-1} e_{A_k} \Sigma_{A_k, A_k}^n (\beta_{A_k} - \hat{\beta}_{A_k}^{(k)})| \\ &\leq \frac{1}{C_{\min}} \rho_S \max_{k \in S} \|\Sigma_{A_k, A_k}^n (\beta_{A_k} - \hat{\beta}_{A_k}^{(k)})\|_\infty \leq \frac{\lambda \rho_S}{C_{\min}}, \end{aligned}$$

where the last step is due to the KKT condition of $\hat{\beta}_{A_k}^{(k)}$. With probability at least $1 - c_3/p$, it holds that

$$\max_{k \in S} \left| \frac{\langle X_k^\perp, X_{A_k}(\beta_{A_k} - \hat{\beta}_{A_k}^{(k)}) \rangle}{n} \right| \leq c_4 K_1 \sqrt{\log p/n} \max_{k \in S} \|X_{A_k}(\beta_{A_k} - \hat{\beta}_{A_k}^{(k)})\|_2.$$

It is left to show that

$$\frac{1}{\sqrt{n}} \max_{k \in S} \|X_{A_k}(\beta_{A_k} - \hat{\beta}_{A_k}^{(k)})\|_2 = o_P(1).$$

Using the KKT condition of (27),

$$\hat{\beta}_{A_k}^{(k)} - \beta_{A_k} = (\Sigma_{A_k, A_k}^n)^{-1} W_{A_k}^n - \lambda (\Sigma_{A_k, A_k}^n)^{-1} \text{sgn}(\hat{\beta}_{A_k}^{(k)})$$

and hence

$$\frac{1}{n} \|X_{A_k}(\hat{\beta}_{A_k}^{(k)} - \beta_{A_k})\|_2^2 \leq \frac{2}{n} \|P_{A_k} \epsilon\|_2^2 + 2\lambda^2 \text{sgn}(\hat{\beta}_{A_k}^{(k)})^T (\Sigma_{A_k, A_k}^n)^{-1} \text{sgn}(\hat{\beta}_{A_k}^{(k)}),$$

where $P_{A_k} = X_{A_k} (X_{A_k}^T X_{A_k})^{-1} X_{A_k}^T$. Using the fact that $A_k \subseteq S$ for any $k \in S$, we further have

$$\begin{aligned} & \max_{k \in S} \frac{1}{n} \|X_{A_k}(\hat{\beta}_{A_k}^{(k)} - \beta_{A_k})\|_2^2 \\ & \leq \max_{k \in S} \frac{2}{n} \|P_{A_k} \epsilon\|_2^2 + 2\lambda^2 \max_{k \in S} \text{sgn}(\hat{\beta}_{A_k}^{(k)})^T (\Sigma_{A_k, A_k}^n)^{-1} \text{sgn}(\hat{\beta}_{A_k}^{(k)}) \\ & \leq \frac{2}{n} \|P_S \epsilon\|_2^2 + 2s\lambda^2 \max_{k \in S} \|(\Sigma_{A_k, A_k}^n)^{-1}\|_2 \leq \frac{c_4 K_2^2 s}{n} + \frac{4s\lambda^2}{C_{\min}}, \end{aligned} \tag{40}$$

with probability at least $1 - c_5 s/n - \exp(-c_6 n)$. The last step is by Cheybeshev's inequality and the fact that

$$(\Sigma_{A_k, A_k}^n)^{-1} \preceq (\Sigma_{S, S}^n)^{-1}_{A_k, A_k} \text{ and } \|(\Sigma_{S, S}^n)^{-1}_{A_k, A_k}\|_2 \leq \|(\Sigma_{S, S}^n)^{-1}\|_2.$$

We arrive at if $s \log p/n \leq c_7$ for small enough constant c_7 , then with probability at least $1 - c_3/p - c_5 s/n - \exp(-c_6 n)$, there exists a large enough constant C such that

$$\max_{k \in S} |u(\hat{\beta}_{A_k}^{(k)})| \leq C \rho_S \lambda. \tag{41}$$

Note that

$$\max_{k \in S} \frac{|u(\hat{\beta}_{A_k}) - u(\hat{\beta}_{A_k}^{(k)})|}{\|\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}\|_2} \leq \max_{k \in S} \left\| \frac{X_k^T X_{A_k}}{n} \right\|_2 \leq C_2.$$

Together with (39) and (41), we have

$$\begin{aligned} \max_{k \in S} \frac{\|P_k^\perp X_{A_k}(\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)})\|_2^2}{2n \|\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}\|_2} & \leq \max_{k \in S} |u(\hat{\beta}_{A_k}^{(k)})| \max_{k \in S} \frac{|u(\hat{\beta}_{A_k}) - u(\hat{\beta}_{A_k}^{(k)})|}{\|\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}\|_2} \\ & \leq C_2 C \rho_S \lambda \end{aligned} \tag{42}$$

with high probability. Define $\mu^{(k)} \in \mathbb{R}^s$ such that

$$\begin{aligned} \mu_k^{(k)} &= - \left\langle \frac{X_k}{\|X_k\|_2^2}, X_{A_k}(\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}) \right\rangle, \quad \mu_{A_k}^{(k)} = \hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}. \\ \max_{k \in S} \frac{\|P_k^\perp X_{A_k}(\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)})\|_2^2}{2n \|\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}\|_2^2} &= \max_{k \in S} \frac{\|X_S \mu^{(k)}\|_2^2}{2n \|\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}\|_2^2} \geq \frac{C_1}{2}. \end{aligned}$$

Together with (42), we have

$$\|\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}\|_2 \leq \frac{2C_2 C \rho_S \lambda}{C_1}.$$

with probability at least $1 - c_3/p - c_5 s/n - \exp(-c_6 n)$. □

A.2. Proof of Lemma 3.1

Proof of Lemma 3.1. (i) First consider a restricted Lasso problem

$$\check{b}_S = \arg \min_{b \in \mathbb{R}^s} \left\{ \frac{1}{2n} \|y - X_S b\|_2^2 + \lambda \|b\|_1 \right\}. \tag{43}$$

Define

$$T_k^o = X_k^T \left(X_S (X_S^T X_S)^{-1} \text{sgn}(\check{b}_S) + \frac{P_S^\perp \epsilon}{n\lambda} \right). \tag{44}$$

By Wainwright [30], if $|T_k^o| < 1$ for $\forall k \in S^c$, then Lasso has a unique solution such that $\hat{S} \subseteq S$.

Since \check{b}_S is only a function of X_S and ϵ , conditioning on X_S and ϵ , T_k^o is a Gaussian random variable with mean $\Sigma_{k,S} \Sigma_{S,S}^{-1} \text{sgn}(\check{b}_S)$ and variance

$$\begin{aligned} \text{Var}(T_k^o | X_S, \epsilon) &\leq \Sigma_{k,k} \left\| X_S (X_S^T X_S)^{-1} \text{sgn}(\check{b}_S) + \frac{P_S^\perp \epsilon}{n\lambda} \right\|_2^2 \\ &\leq \Sigma_{k,k} (\text{sgn}(\check{b}_S)^T (X_S^T X_S)^{-1} \text{sgn}(\check{b}_S) + \|\epsilon\|_2^2 / (n\lambda)^2). \end{aligned}$$

Let

$$\Omega_2 = \{ \|\epsilon\|_2^2 / n \leq 1.1\sigma^2, \Lambda_{\max}((\Sigma_{S,S}^n)^{-1}) \leq 4/C_{\min} \}. \tag{45}$$

In Ω_2 ,

$$\text{sgn}(\check{b}_S)^T (X_S^T X_S)^{-1} \text{sgn}(\check{b}_S) \leq \|\text{sgn}(\check{b}_S)\|_2^2 \|(X_S^T X_S)^{-1}\|_2^2 \leq \frac{4s}{nC_{\min}}.$$

And hence

$$\text{Var}(T_k^o | X_S, \epsilon) \leq \Sigma_{k,k} \left(\frac{4s}{nC_{\min}} + \frac{1.1\sigma^2}{n\lambda^2} \right).$$

Thus,

$$\begin{aligned} & \mathbb{P} \left(\max_{k \in S^c} |T_k^o| \geq \max_{k \in S^c} |\Sigma_{k,S} \Sigma_{S,S}^{-1} \text{sgn}(\check{b}_S)| + t | \Omega_2 \right) \\ & \leq 2(p-s) \exp \left\{ -\frac{t^2}{2K_1 \left(\frac{4s}{nC_{\min}} + \frac{1.1\sigma^2}{n\lambda^2} \right)} \right\}. \end{aligned}$$

By setting $t = \frac{1-\phi}{2}$, we have for

$$n \geq \frac{128K_1 s \log p}{C_{\min}(1-\phi)^2} \quad \text{and} \quad \lambda \geq \frac{8.8\sigma}{1-\phi} \sqrt{\frac{K_1 \log p}{n}},$$

it holds that

$$\mathbb{P} \left(\max_{k \in S^c} |T_k^o| > \frac{1+\phi}{2} | \Omega_2 \right) \leq 2 \exp(-\log(p-s)).$$

By Lemma A.1,

$$P(\Omega_2) \geq 1 - c_1/n - \exp(-c_2n).$$

Therefore,

$$\mathbb{P}(\hat{S} \subseteq S) \geq 1 - c_0/p - c_1/n - \exp(-c_2n).$$

In the event that $\{\hat{S} \subseteq S\}$, the KKT condition of (2) yields

$$\hat{u}_S = (\Sigma_{S,S}^n)^{-1} W_S^n - \lambda (\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S). \tag{46}$$

(ii) Noticing that for any $k \in S$, $\mathbb{P}\{\Sigma_{k,k}^n \geq C_{\min}/2\} \geq 1 - \exp(-c_2n)$ by Lemma A.1 and

$$\begin{aligned} \hat{u}_k &= \frac{1}{\Sigma_{k,k}^n} W_k^n - \frac{\lambda}{\Sigma_{k,k}^n} \text{sgn}(\hat{\beta}_k) - \frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n \hat{u}_{A_k} \\ &= \underbrace{\frac{1}{\Sigma_{k,k}^n} W_k^n - \frac{\lambda}{\Sigma_{k,k}^n} \text{sgn}(\hat{\beta}_k)}_{R_{1,k}} - \underbrace{\frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n (\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)})}_{R_{2,k}} - \underbrace{\frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n (\hat{\beta}_{A_k}^{(k)} - \beta_{A_k})}_{R_{3,k}}, \end{aligned}$$

where $\hat{\beta}_{A_k}^{(k)}$ is the leave-one-out estimator (27). Using the sub-Gaussian property of ϵ , it is easy to show that for some large enough c_1 ,

$$\mathbb{P} \left(\max_{k \in S} |R_{1,k}| > c_1(\lambda + \sqrt{\log s/n})/C_{\min} \right) \leq 2 \exp(-c_3 \log p).$$

For $R_{2,k}$, it holds that for some large enough c_2 ,

$$\max_{k \in S} |R_{2,k}| \leq \max_{k \in S} \left\| \frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n \right\|_2 \|\hat{\beta}_{A_k} - \hat{\beta}_{A_k}^{(k)}\|_2 \leq 4 \sqrt{\frac{C_{\max}}{C_{\min}}} \rho_S \lambda$$

with probability at least $1 - c_1/p - \exp(-c_2n) - c_3/n$ due to Lemma A.1 and A.2. For $R_{3,k}$, it holds that

$$\begin{aligned} & \max_{k \in S} \left| \Sigma_{k, A_k}^n (\hat{\beta}_{A_k}^{(k)} - \beta_{A_k}) \right| \\ & \leq \max_{k \in S} \left| \Sigma_{k, A_k} \Sigma_{A_k, A_k}^{-1} \Sigma_{A_k, A_k}^n (\hat{\beta}_{A_k}^{(k)} - \beta_{A_k}) \right| + \max_{k \in S} \frac{1}{n} \left| \langle X_k^\perp, X_{A_k} \rangle (\hat{\beta}_{A_k}^{(k)} - \beta_{A_k}) \right| \\ & \leq \rho_S \max_{k \in S} \|\Sigma_{A_k, A_k}^n (\hat{\beta}_{A_k}^{(k)} - \beta_{A_k})\|_\infty + O_P(\max_{k \in S} \|X_{A_k} (\hat{\beta}_{A_k}^{(k)} - \beta_{A_k})\|_2/n) \\ & \leq \rho_S \lambda + \frac{c_5 \sqrt{s} \lambda}{\sqrt{n}} \end{aligned}$$

with probability $1 - c_2s/n - \exp(-c_3n)$ if $s \log p \leq c_4n$ for a small enough constant c_4 . The last step is due to the KKT condition of $\hat{\beta}_{A_k}^{(k)}$ and (40).

Together we have, for some large enough constant C , we have

$$\max_{k \in S} |u_k| \leq C \rho_S \lambda$$

with probability at least $1 - c_1/p - \exp(-c_2n) - c_3/n$ for some positive constants $c_1 - c_3$. □

A.3. Proof of Lemma 3.2

For $\hat{\beta}^{(k)}$ defined in (27), define a constrained noiseless Lasso estimator

$$\hat{\beta}^{*,(k)} = \arg \min_{b \in \mathbb{R}^p} \left\{ \frac{1}{2n} \|X_S \hat{\beta}_S^{(k)} - X_S b\|_2^2 + \lambda \|b\|_1 : b_k = \beta_k, b_{S^c} = 0 \right\}. \quad (47)$$

We first prove the following technical lemma.

Lemma A.3. *For $\hat{\beta}^{*,(k)}$ defined in (47), it holds that if $n \geq c_0 s \log p$ for large enough c_0 , then with probability at least $1 - c_3/p - \exp(-c_4n) - c_5/n$ there exists a large enough constant c_1 such that*

$$\max_{k \in S} \|\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*\|_2 \leq c_1 \rho_S \lambda.$$

Proof of Lemma A.3. In event Ω_0^* , KKT condition of (47) and that of $\hat{\beta}^*$ give that

$$\begin{aligned} -\lambda \operatorname{sgn}(\hat{\beta}_S^{*,(k)}) + \lambda \operatorname{sgn}(\hat{\beta}_S^*) &= \Sigma_{S,S}^n (\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^{(k)}) - \Sigma_{S,S}^n (\hat{\beta}_S^* - \hat{\beta}_S) \\ &= \Sigma_{S,S}^n (\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*) + \Sigma_{S,S}^n (\hat{\beta}_S - \hat{\beta}_S^{(k)}). \end{aligned}$$

We arrive at

$$\begin{aligned} 0 &\geq -\lambda \langle \hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*, \operatorname{sgn}(\hat{\beta}_S^{*,(k)}) - \operatorname{sgn}(\hat{\beta}_S^*) \rangle \\ &\quad (\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*)^T \Sigma_{S,S}^n (\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*) - (\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*)^T \Sigma_{S,S}^n (\hat{\beta}_S - \hat{\beta}_S^{(k)}), \end{aligned}$$

where the first step is by the definition of sub-differential. Therefore,

$$\begin{aligned} \Lambda_{\min}(\Sigma_{S,S}^n) \|\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*\|_2^2 &\leq \|\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*\|_2 \|\Sigma_{S,S}^n\|_2 \|\hat{\beta}_S - \hat{\beta}_S^{(k)}\|_2 \\ &\leq \frac{\Lambda_{\min}(\Sigma_{S,S}^n)}{2} \|\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*\|_2^2 + \frac{2\|\Sigma_{S,S}^n\|_2^2}{\Lambda_{\min}(\Sigma_{S,S}^n)} \|\hat{\beta}_S - \hat{\beta}_S^{(k)}\|_2^2, \end{aligned}$$

where the last step follows from the Young's inequality. In the event that $\{C_1 \leq \Lambda_{\min}(\Sigma_{S,S}^n) \leq \Lambda_{\max}(\Sigma_{S,S}^n) \leq C_2\}$, we arrive at

$$\max_{k \in S} \|\hat{\beta}_S^{*,(k)} - \hat{\beta}_S^*\|_2 \leq \frac{C_2}{C_1} \max_{k \in S} \|\hat{\beta}_S - \hat{\beta}_S^{(k)}\|_2 \leq c_1 \rho_S \lambda$$

for large enough constant c with probability at least $1 - c_3/p - \exp(-c_4n) - c_5/n$. The last step is due to Lemma A.2. \square

Proof of Lemma 3.2. (i) Define a restricted Lasso problem with observations (X, y^*) .

$$\check{b}_S^* = \arg \min_{b \in \mathbb{R}^s} \left\{ \frac{1}{2n} \|y^* - X_S b\|_2^2 + \lambda \|b_S\|_1 \right\}. \tag{48}$$

One can use same arguments in the proof of Lemma 3.1 to get that

$$\mathbb{P}\{\hat{S}^* \subseteq S\} \geq 1 - c_1/p - c_2/n - \exp(-c_3n).$$

(ii) The KKT condition for the bootstrapped Lasso gives that

$$\hat{u}_S^* = -\lambda(\Sigma_{S,S}^n)^{-1} \text{sgn}(\hat{\beta}_S^*).$$

For any $k \in S$, for $\hat{\beta}^{*,(k)}$ defined in (47), we have

$$\begin{aligned} \hat{u}_k^* &= -\lambda \frac{1}{\Sigma_{k,k}^n} \text{sgn}(\hat{\beta}_k^*) - \frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n \hat{u}_{A_k}^* \\ &= \underbrace{-\lambda \frac{1}{\Sigma_{k,k}^n} \text{sgn}(\hat{\beta}_k^*)}_{R_{1,k}} - \underbrace{\frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n [\hat{\beta}_{A_k}^* - \hat{\beta}_{A_k}^{*,(k)}]}_{R_{2,k}} - \underbrace{\frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n [\hat{\beta}_{A_k}^{*,(k)} - \hat{\beta}_{A_k}^{(k)}]}_{R_{3,k}} \\ &\quad - \underbrace{\frac{1}{\Sigma_{k,k}^n} \Sigma_{k,A_k}^n [\hat{\beta}_{A_k}^{(k)} - \hat{\beta}_{A_k}]}_{R_{4,k}}. \end{aligned}$$

Note that $R_{1,k} = O_P(\lambda)$. For $R_{2,k}$, by Lemma A.3 we have with probability at least $1 - c_4/p - c_5/n - \exp(-c_6n)$, for some large enough constant c_7 ,

$$\max_{k \in S} |R_{2,k}| \leq \|\Sigma_{S,S}^n\|_2 \max_{k \in S} \|\hat{\beta}_{A_k}^* - \hat{\beta}_{A_k}^{*,(k)}\|_2 = O_P(\rho_S \lambda).$$

For $R_{3,k}$, by the KKT condition of $\hat{\beta}^{*,(k)}$, we have

$$\begin{aligned} & \max_{k \in S} |R_{3,k}| \\ & \leq \max_{k \in S} \left| \Sigma_{k, A_k} \Sigma_{A_k, A_k}^{-1} \Sigma_{A_k, A_k}^n [\hat{\beta}_{A_k}^{*,(k)} - \hat{\beta}_{A_k}^{(k)}] \right| + \max_{k \in S} \frac{1}{n} |\langle X_k^\perp, X_{A_k} \rangle [\hat{\beta}_{A_k}^{*,(k)} - \hat{\beta}_{A_k}^{(k)}]| \\ & \leq \rho_S \max_{k \in S} \|\Sigma_{A_k, A_k}^n [\hat{\beta}_{A_k}^{*,(k)} - \hat{\beta}_{A_k}^{(k)}]\|_\infty + O_P(\max_{k \in S} \|X_{A_k} (\hat{\beta}_{A_k}^{*,(k)} - \hat{\beta}_{A_k}^{(k)})\|_2/n) \\ & = \rho_S \lambda + o_P(n^{-1/2}), \end{aligned}$$

where the second step is due to $\hat{\beta}_{A_k}^{*,(k)} - \hat{\beta}_{A_k}^{(k)}$ is independent of X_k and the last step is due to a similar argument for (40). $\max_{k \in S} |R_{4,k}|$ can be similarly bounded as for $R_{2,k}$ using Lemma A.2. Hence, we have with probability at least $1 - c_7/p - c_8/n - \exp(-c_9n)$, there exists a large enough constant C^* such that

$$\max_{k \in S} |\hat{u}_k^*| \leq C^* \rho_S \lambda.$$

Therefore Ω_0^* holds with high probability. In view of Ω_0 (19), for $j \in S_+$,

$$\begin{aligned} \hat{\beta}_j &= \beta_j + \hat{u}_j \geq \beta_j - \max_{1 \leq j \leq p} |u_j| \geq \beta_j - C \rho_S \lambda > C' \rho_S \lambda \text{ for } \beta_j \geq 0. \\ \hat{\beta}_j &= \beta_j + \hat{u}_j \leq \beta_j + \max_{1 \leq j \leq p} |u_j| \leq \beta_j + C \rho_S \lambda < -C' \rho_S \lambda \text{ for } \beta_j < 0. \end{aligned}$$

In view of Ω_0^* (20) and repeating previous arguments, we can get $\text{sgn}(\hat{\beta}_j) = \text{sgn}(\hat{\beta}_j^*)$ for $j \in S_+$. \square

Proof of Theorem 3.7. We first notice that for \hat{w}_j defined in (15) with Θ unknown,

$$\hat{\beta}_j^{(mDB)} - \beta_j = (e_j^T - \hat{w}_j^T X/n) \hat{u} + \frac{\hat{w}_j^T \epsilon}{n}.$$

We have

$$|(e_j^T - \hat{w}_j^T X/n) \hat{u}| \leq \|\hat{w}_j^T X_{-j}/n\|_\infty \|\hat{u}\|_1 \leq \lambda_j \|\hat{u}\|_1 = O_P(s \lambda \lambda_j),$$

where the second step is by the KKT condition of (16) and the last step is by the standard analysis of the Lasso. Similarly,

$$|\hat{b}_j^*| = |(e_j^T - \hat{w}_j^T X/n) \hat{u}^*| \leq \lambda_j \|\hat{u}^*\|_1.$$

To obtain an upper bound on $\|\hat{u}^*\|_1$, consider the oracle inequality of $\hat{\beta}^*$:

$$\begin{aligned} \frac{1}{2n} \|X \hat{u}^*\|_2^2 &\leq \lambda \|\hat{\beta}\|_1 - \lambda \|\hat{\beta}^*\|_1 \leq \lambda \|\hat{u}_S^*\|_1 + \lambda \|\hat{\beta}_{S^c}\|_1 - \lambda \|\hat{\beta}_{S^c}^*\|_1 \\ &\leq \lambda \|\hat{u}_S^*\|_1 + 2\lambda \|\hat{\beta}_{S^c}\|_1 - \lambda \|\hat{u}_{S^c}^*\|_1, \end{aligned}$$

where we use $\|\hat{u}_{S^c}^*\|_1 \leq \|\hat{\beta}_{S^c}^*\|_1 + \|\hat{\beta}_{S^c}\|_1$ in the last step. If $2\|\hat{\beta}_{S^c}\|_1 \geq \|\hat{u}_S^*\|_1$, then the right hand side gives $\|\hat{u}_{S^c}^*\|_1 \leq 4\|\hat{\beta}_{S^c}\|_1$. Hence,

$$\|\hat{u}^*\|_1 \leq \frac{9}{2} \|\hat{\beta}_{S^c}\|_1 = \frac{9}{2} \|\hat{u}_{S^c}\|_1 \leq \frac{9}{2} \|\hat{u}\|_1 = O_P(s \sqrt{\log p/n}).$$

If $2\|\hat{\beta}_{S^c}\|_1 < \|\hat{u}_S^*\|_1$, then we arrive at

$$\frac{1}{2n}\|X\hat{u}^*\|_2^2 \leq \lambda\|\hat{\beta}\|_1 - \lambda\|\hat{\beta}^*\|_1 \leq 2\lambda\|\hat{u}_S^*\|_1 - \lambda\|\hat{u}_{S^c}^*\|_1,$$

which is the usual oracle inequality. Standard analysis gives $\|\hat{u}_S^*\|_1 = O_P(s\sqrt{\log p/n})$. As a result,

$$\hat{\beta}_j^{(BS-DB)} - \beta_j = \hat{\beta}_j^{(mDB)} - \beta_j - \hat{b}_j^* = \frac{\hat{w}_j^T \epsilon}{n} + O_P\left(\frac{s \log p}{n}\right). \quad \square$$

A.4. Proof of Lemma 3.5

Proof of Lemma 3.5. We note that

$$P_{\hat{S}}^\perp y = P_{\hat{S}}^\perp (y - X_{\hat{S}} \hat{\beta}_{\hat{S}}) = P_{\hat{S}}^\perp (y - X \hat{\beta}).$$

When Lemma 3.2 holds true,

$$\frac{1}{n - |\hat{S}|} \|P_{\hat{S}_+}^\perp (y - X \hat{\beta})\|_2^2 \leq \hat{\sigma}^2 = \frac{1}{n - |\hat{S}|} \|P_{\hat{S}}^\perp (y - X \hat{\beta})\|_2^2 \leq \frac{1}{n - |\hat{S}|} \|P_{\hat{S}}^\perp (y - X \hat{\beta})\|_2^2.$$

For the left hand side,

$$\begin{aligned} \frac{1}{n - |\hat{S}|} \|P_{\hat{S}_+}^\perp (y - X \hat{\beta})\|_2^2 &= \frac{1}{n - |\hat{S}|} \|P_{\hat{S}_+}^\perp \epsilon\|_2^2 + \frac{1}{n - |\hat{S}|} \|P_{\hat{S}_+}^\perp X \hat{u}\|_2^2 \\ &\quad - \frac{2}{n - |\hat{S}|} \langle P_{\hat{S}_+}^\perp \epsilon, P_{\hat{S}_+}^\perp X \hat{u} \rangle, \end{aligned}$$

where $\hat{u} = \hat{\beta} - \beta$. We have

$$\frac{1}{n} \|P_{\hat{S}_+}^\perp X \hat{u}\|_2^2 \leq \frac{1}{n} \|P_{\hat{S}_+}^\perp X_{S_-} \hat{u}_{S_-}\|_2^2 \leq \|\Sigma_{S_-, S_-}^n\|_2 \|\hat{u}_{S_-}\|_2^2 = O_P(s_- \rho_S^2 \log p/n).$$

We can obtain another bound by noticing $\|\hat{u}_{S_-}\|_2^2 \leq \|\hat{u}_S\|_2^2 = O_P(s \log p/n)$. Therefore,

$$\frac{1}{n} \|P_{\hat{S}_+}^\perp X \hat{u}\|_2^2 = O_P(\min\{\rho_S^2 s_-, s\} \log p/n).$$

Moreover,

$$\begin{aligned} \left| \frac{1}{n} \langle P_{\hat{S}_+}^\perp \epsilon, P_{\hat{S}_+}^\perp X_{S_-} \hat{u}_{S_-} \rangle \right| &\leq \max_{j \in \hat{S}_-} \left| \frac{1}{n} \langle P_{\hat{S}_+}^\perp \epsilon, P_{\hat{S}_+}^\perp X_j \rangle \right| \|\hat{u}_{S_-}\|_1 \\ &= O_P\left(\sqrt{\frac{\log p}{n}}\right) \|\hat{u}_{S_-}\|_1, \end{aligned}$$

where

$$\|\hat{u}_{S_-}\|_1 \leq \min\{s_- \|\hat{u}_{S_-}\|_\infty, \sqrt{s_-} \|\hat{u}_S\|_2\} = O_P\left(\min\{s_- \rho_S \lambda, \sqrt{s_-} \sqrt{\log p/n}\}\right).$$

Therefore,

$$\left| \frac{1}{n} \langle P_{S_+}^\perp \epsilon, P_{S_+}^\perp X_{S_-} \hat{u}_{S_-} \rangle \right| = O_P \left(\min\{\rho_{S S_-}, s\} \frac{\log p}{n} \right).$$

Finally,

$$\frac{1}{n - |\hat{S}|} \|P_{S_+}^\perp \epsilon\|_2^2 = \sigma^2 \frac{n - s_+}{n - \hat{S}} + O_P((n - s)^{-1/2}) = \sigma^2 + O_P(s_-/n + n^{-1/2})$$

for $n \gg s \geq |\hat{S}|$. Hence,

$$\hat{\sigma}^2 \geq \sigma^2 - O_P \left(n^{-1/2} + \frac{\min\{\rho_{S S_-}^2, s\} \log p}{n} \right).$$

By similar arguments, one can show that

$$\frac{1}{n - |\hat{S}|} \|P_{\hat{S}}^\perp y\|_2^2 = \sigma^2 + O_P \left(n^{-1/2} + \frac{s_-}{n} \right)$$

for $n \gg s$. In view of the first inequality in the proof, the proof is complete. \square

Appendix B: More simulation results

We report some numerical results with a “nonparametric configuration” of true coefficients. Specifically, $\beta_k = k^{-0.5}$ for $k = 1, \dots, s$. We consider the true covariance matrix of X being identity or Σ° for Σ° defined in Section 1.2. We report the average coverage probabilities on and off the true support and the average confidence interval lengths on and off the true support. In Table 5, we see that the proposed BS-DB has average coverage probabilities close to the nominal level when $\Sigma = I_p$ and the debiased Lasso method has coverage lower than the nominal level. For $\Sigma = \Sigma^\circ$, the coverage probabilities given by BS-DB is closer to the nominal level than those given by the debiased Lasso but both methods have the average coverage probabilities for $\beta_j, j \in S$, lower than the nominal level.

TABLE 5

Coverage probabilities of BS-DB (cov.bsdb) and debiased Lasso (cov.db) and standard errors of BS-DB (se.bsdb) and debiased Lasso (se.db). The rows with $j \in S$ are the average results for the coefficients on the true support and the rows with $j \in S^c$ are the average results for the coefficients off the true support.

s	j	$\Sigma = I_p$				$\Sigma = \Sigma^\circ$			
		cov.bsdb	cov.db	se.bsdb	se.db	cov.bsdb	cov.db	se.bsdb	se.db
4	S	0.919	0.891	0.15	0.15	0.893	0.810	0.13	0.13
4	S^c	0.943	0.887	0.15	0.15	0.943	0.908	0.13	0.13
8	S	0.944	0.910	0.15	0.15	0.883	0.844	0.15	0.15
8	S^c	0.950	0.907	0.15	0.15	0.948	0.904	0.15	0.15
12	S	0.939	0.894	0.15	0.15	0.898	0.869	0.16	0.15
12	S^c	0.949	0.894	0.15	0.15	0.944	0.914	0.16	0.15