



Simultaneous Inference for High-Dimensional Linear Models

Xianyang Zhang & Guang Cheng

To cite this article: Xianyang Zhang & Guang Cheng (2017): Simultaneous Inference for High-Dimensional Linear Models, Journal of the American Statistical Association, DOI: [10.1080/01621459.2016.1166114](https://doi.org/10.1080/01621459.2016.1166114)

To link to this article: <http://dx.doi.org/10.1080/01621459.2016.1166114>

 View supplementary material 

 Accepted author version posted online: 06 Apr 2016.
Published online: 12 Apr 2017.

 Submit your article to this journal 

 Article views: 489

 View Crossmark data 

Simultaneous Inference for High-Dimensional Linear Models

Xianyang Zhang^a and Guang Cheng^b

^aDepartment of Statistics, Texas A&M University, College Station, TX; ^bDepartment of Statistics, Purdue University, West Lafayette, IN

ABSTRACT

This article proposes a bootstrap-assisted procedure to conduct simultaneous inference for high-dimensional sparse linear models based on the recent desparsifying Lasso estimator. Our procedure allows the dimension of the parameter vector of interest to be exponentially larger than sample size, and it automatically accounts for the dependence within the desparsifying Lasso estimator. Moreover, our simultaneous testing method can be naturally coupled with the margin screening to enhance its power in sparse testing with a reduced computational cost, or with the step-down method to provide a strong control for the family-wise error rate. In theory, we prove that our simultaneous testing procedure asymptotically achieves the prespecified significance level, and enjoys certain optimality in terms of its power even when the model errors are non-Gaussian. Our general theory is also useful in studying the support recovery problem. To broaden the applicability, we further extend our main results to generalized linear models with convex loss functions. The effectiveness of our methods is demonstrated via simulation studies. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received June 2015
Revised February 2016

KEYWORDS

Desparsifying Lasso;
High-dimensional linear
models; Multiple testing;
Multiplier Bootstrap;
Simultaneous inference;
Support recovery

1. Introduction

High-dimensional statistics has become increasingly popular due to the rapid development of information technologies and their applications in scientific experiments. There is a huge body of work on sparse estimation of high-dimensional models. The statistical properties of these procedures have been extensively studied in the literature; see Bühlmann and van de Geer (2011). The research effort has recently turned to statistical inference such as constructing confidence interval and hypothesis testing for regression coefficients in high-dimensional sparse models. For instance, Wasserman and Roeder (2009), and Meinshausen, Meier, and Bühlmann (2009) proposed significance tests for high-dimensional regression coefficients based on sample splitting. Lockhart et al. (2014) derived a significance test for variables along the Lasso solution path. Lee et al. (2016) proposed an exact postselection inference procedure based on the idea of polyhedral selection. Another research line for conducting inference is to exploit the idea of low-dimensional projection or inverting the Karush-Kuhn-Tucker condition (see, e.g., Belloni, Chernozhukov, and Kato 2014; Javanmard and Montanari 2014; van de Geer et al. 2014; Zhang and Zhang 2014).

In the high-dimensional regime, grouping of variables and exploiting group structure is quite natural (Yuan and Lin 2006; Meier, van de Geer, and Bühlmann 2008). Leading examples include multifactor analysis-of-variance and additive modeling. From a practical viewpoint, when research interest concerns not only a single variable but rather a group of variables, it seems indispensable to go beyond an approach of inferring individual regression coefficients. The problem of conducting simultaneous inference or inference for groups of variables in high-dimensional models has gained some recent

attention. Based on the basis pursuit solution, Meinshausen (2015) proposed an interesting procedure to construct confidence intervals for groups of variables without restrictive assumptions on the design. Mandozzi and Bühlmann (2014) extended the hierarchical testing method in Meinshausen (2008) to the high-dimensional setting, which is able to detect the smallest groups of variables and asymptotically control the family-wise error rate (FWER).

In this article, we propose a bootstrap-assisted procedure to conduct simultaneous inference in sparse linear models with (possibly) non-Gaussian errors

$$Y = \mathbf{X}\beta^0 + \epsilon,$$

where Y is a response, \mathbf{X} is a design matrix, and $\beta^0 = (\beta_1^0, \dots, \beta_p^0)^T$ is a vector of unknown regression coefficients. Specifically, we consider the following simultaneous testing

$$H_{0,G} : \beta_j^0 = \tilde{\beta}_j \quad \text{for all } j \in G \subseteq \{1, 2, \dots, p\}$$

versus the alternative $H_{a,G} : \beta_j^0 \neq \tilde{\beta}_j$ for some $j \in G$, where $\tilde{\beta}_j$ with $j \in G$ is a vector of prespecified values (e.g., by domain experts). We point out that two extreme cases, where $G = \{1, 2, \dots, p\}$ or $|G|$ is small, have been considered in the recent literature; see, for example, Arias-Castro, Candès, and Plan (2011), Zhong and Chen (2011), Feng et al. (2013), Zhang and Zhang (2014), and van de Geer et al. (2014). However, relatively few attention has been paid to the case where G lies between these two extremes. Our method allows G to be an arbitrary subset of $\{1, 2, \dots, p\}$, and thus can be applied to a wider range of real problems such as testing the significance of a growing set of genes (associated with certain clinical outcome) conditional on another set of genes, whose size is allowed to grow

as well. In comparison with Meinshausen (2015) and Mandozzi and Bühlmann (2014), our method is asymptotically exact and relies on suitable assumptions (Assumptions 2.1 and 2.2) on the design which are different from the highly correlated design considered in Mandozzi and Bühlmann (2014).

Our general framework is built upon the desparsifying Lasso estimator, denoted as $\check{\beta} = (\check{\beta}_1, \dots, \check{\beta}_p)^T$, recently developed in van de Geer et al. (2014) and Zhang and Zhang (2014), whose review is given in Section 2.1. In Section 2.2, a test statistic is proposed as $T_{n,G} := \max_{j \in G} \sqrt{n} |\check{\beta}_j - \tilde{\beta}_j|$, whose critical values are obtained via a simple multiplier bootstrap method. Based on the asymptotic linear expansion of $\check{\beta}$, we show that the proposed multiplier bootstrap method consistently approximates the null limiting distribution of $T_{n,G}$, and thus the testing procedure achieves the prespecified significance level asymptotically. It is worth mentioning that the proposed bootstrap-assisted procedure is adaptive to the dimension of the component of interest, and it automatically accounts for the dependence within the desparsifying Lasso estimators. In theory, we also prove that our testing procedure enjoys certain minimax optimality (Verzelen 2012) in terms of its power even when the model errors are non-Gaussian and the cardinality of G is exponentially larger than sample size.

Moreover, our new methodology is readily applicable to some other important statistical problems in the high-dimensional setting, such as support recovery, testing for sparse signals, and multiple testing. The support recovery procedure is proposed in Section 3.1 as an important by-product of our general theory. It has been shown through simulations that the proposed procedure can be more accurate in recovering signals than Lasso, the stability selection (Meinshausen and Bühlmann 2010), and the screen and clean procedure (Wasserman and Roeder 2009). The above bootstrap-assisted test method can also be coupled with the margin screening in Fan and Lv (2008) to enhance the power performance in sparse testing with a reduced computational cost, which is very attractive in the ultra-high-dimensional setting. Hence, in Section 3.2 we propose a three-step procedure that first randomly splits the sample into two subsamples, screens out the irrelevant variables based on the first subsample, and finally performs the above maximum-type testing on the reduced model based on the second subsample. Another application is a multiple testing problem: for each $j \in G$

$$H_{0,j} : \beta_j^0 \leq \tilde{\beta}_j \quad \text{versus} \quad H_{a,j} : \beta_j^0 > \tilde{\beta}_j.$$

To obtain a strong control of the FWER, we incorporate the above bootstrap idea into the step-down method (Romano and Wolf 2005) in Section 3.3. As noted in Chernozhukov, Chetverikov, and Kato (2013), this hybrid method is asymptotically nonconservative as compared to the Bonferroni-Holm procedure since the correlation amongst the test statistics has been taken into account.

To broaden the applicability, we further extend our main results to generalized linear models with convex loss functions in Section 4. The usefulness of the above simultaneous inference methods is illustrated via simulation studies in Section 5. The technical details and additional numerical results are included

in a supplement file. Some theoretical derivations in this article rely on an impressive Gaussian approximation (GAR) theory recently developed in Chernozhukov, Chetverikov, and Kato (2013). The application of GAR theory is nontrivial as one needs to verify suitable moment conditions on the leading term of the desparsifying Lasso estimator, and quantify the estimation effect as well as the impact of the remainder term in (7) below. We also want to point out that the GAR theory is applied without conditioning on the design as random design is considered throughout the article. Our results complement Belloni, Chernozhukov, and Kato (2014) who establish the validity of uniform confidence band in the high-dimensional least absolute deviation regression.

Finally, we introduce some notation. For a $p \times p$ matrix $B = (b_{ij})_{i,j=1}^p$, let $\|B\|_\infty = \max_{1 \leq i, j \leq p} |b_{ij}|$ and $\|B\|_1 = \max_{1 \leq j \leq p} \sum_{i=1}^p |b_{ij}|$. Denote by $\|a\|_q = (\sum_{i=1}^p |a_i|^q)^{1/q}$ and $\|a\|_\infty = \max_{1 \leq j \leq p} |a_j|$ for $a = (a_1, \dots, a_p)^T \in \mathbb{R}^p$ and $q > 0$. For a set \mathcal{A} , denote its cardinality by $|\mathcal{A}|$. Denote by $\lfloor a \rfloor$ the integer part of a positive real number a . For two sequences $\{a_n\}$ and $\{b_n\}$, write $a_n \asymp b_n$ if there exist positive constants c and C such that $c \leq \liminf_n (a_n/b_n) \leq \limsup_n (a_n/b_n) \leq C$. Also write $a_n \lesssim b_n$ if $a_n \leq C b_n$ for some constant $C > 0$ independent of n (and p). The symbol $N_p(\mu, \Sigma)$ is reserved for a p -dimensional multivariate normal distribution with mean μ and covariance matrix Σ .

2. Main Theory

2.1. Desparsifying Lasso Estimator

In this section, we review the desparsifying (debiased) estimator proposed in van de Geer et al. (2014), which is essentially the same as the estimator proposed in Zhang and Zhang (2014) but motivated from a different viewpoint, that is, inverting the Karush-Kuhn-Tucker (KKT) condition of Lasso. Consider a high-dimensional sparse linear model

$$Y = \mathbf{X}\beta^0 + \epsilon, \quad (1)$$

with a response $Y = (Y_1, \dots, Y_n)^T$, an $n \times p$ design matrix $\mathbf{X} := [X_1, \dots, X_p]$, an error $\epsilon = (\epsilon_1, \dots, \epsilon_n)^T$ independent of \mathbf{X} , and an unknown $p \times 1$ regression vector $\beta^0 = (\beta_1^0, \dots, \beta_p^0)^T$. The parameter dimension p can be much larger than sample size n . Suppose \mathbf{X} has iid rows having mean zero and covariance matrix $\Sigma = (\sigma_{ij})_{i,j=1}^p$ with $\Sigma^{-1} := \Theta = (\theta_{ij})_{i,j=1}^p$. We denote the active set of variables by $\mathcal{S}_0 = \{1 \leq j \leq p : \beta_j^0 \neq 0\}$ and its cardinality by $s_0 = |\mathcal{S}_0|$. The Lasso estimator (Tibshirani 1996) is defined as

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

for some tuning parameter $\lambda > 0$.

The desparsifying estimator is obtained by inverting the KKT condition,

$$\check{\beta} = \hat{\beta} + \hat{\Theta}\mathbf{X}^T(Y - \mathbf{X}\hat{\beta})/n, \quad (3)$$

where $\hat{\Theta}$ is a suitable approximation for the inverse of the Gram matrix $\hat{\Sigma} := \mathbf{X}^T\mathbf{X}/n$. In what follows, we consider the approximate inverse $\hat{\Theta}$ given by Lasso for the nodewise regression on the design matrix \mathbf{X} ; see Meinshausen and Bühlmann (2006).

Let \mathbf{X}_{-j} be the design matrix without the j th column. For $j = 1, 2, \dots, p$, consider

$$\hat{\gamma}_j := \operatorname{argmin}_{\gamma \in \mathbb{R}^{p-1}} (\|\mathbf{X}_j - \mathbf{X}_{-j}\gamma\|_2^2/n + 2\lambda_j \|\gamma\|_1) \quad (4)$$

with $\lambda_j > 0$, where we denote $\hat{\gamma}_j = \{\hat{\gamma}_{j,k} : 1 \leq k \leq p, k \neq j\}$. Let $\hat{\mathbf{C}} = (\hat{c}_{i,j})_{i,j=1}^p$ be a $p \times p$ matrix with $\hat{c}_{i,i} = 1$ and $\hat{c}_{i,j} = -\hat{\gamma}_{i,j}$ for $i \neq j$. Let $\hat{\tau}_j^2 = \|\mathbf{X}_j - \mathbf{X}_{-j}\hat{\gamma}_j\|_2^2/n + \lambda_j \|\hat{\gamma}_j\|_1$ and write $\hat{\mathbf{T}}^2 = \operatorname{diag}(\hat{\tau}_1^2, \dots, \hat{\tau}_p^2)$ as a diagonal matrix. Finally, the nodewise Lasso estimator for Θ is constructed as $\hat{\Theta} = \hat{\mathbf{T}}^{-2}\hat{\mathbf{C}}$.

Denote by $\gamma_j = \operatorname{argmin}_{\gamma \in \mathbb{R}^{p-1}} \mathbb{E}\|\mathbf{X}_j - \mathbf{X}_{-j}\gamma\|_2^2$, and define $\eta_j = \mathbf{X}_j - \mathbf{X}_{-j}\gamma_j = (\eta_{1,j}, \dots, \eta_{n,j})^T$. Define $\tau_j^2 = \mathbb{E}\|\eta_j\|_2^2/n = 1/\theta_{j,j}$ for $j = 1, 2, \dots, p$. Let $s_j = |\{1 \leq k \leq p : k \neq j, \theta_{jk} \neq 0\}|$. Denote the j th row of \mathbf{X} and $\hat{\Theta}$ by $\tilde{X}_j = (X_{j1}, \dots, X_{jp})^T$ and $\hat{\Theta}_j$, respectively.

Assumption 2.1. The design matrix \mathbf{X} has either iid sub-Gaussian rows (i.e., $\sup_{\|a\|_2 \leq 1} \mathbb{E} \exp\{|\sum_{j=1}^p a_j X_{ij}|^2/C\} \leq 1$ for some large enough positive constant C) or iid rows satisfying for some $K_n \geq 1$, $\max_{1 \leq i \leq n, 1 \leq j \leq p} |X_{ij}| \leq K_n$ (strongly bounded case), where K_n is allowed to grow with n . In the strongly bounded case, we assume in addition that $\max_j \|\mathbf{X}_{-j}\gamma_j\|_\infty \leq K_n$ and $\max_j \mathbb{E}\eta_{1,j}^4 \leq K_n^4$.

Assumption 2.2. The smallest eigenvalue Λ_{\min}^2 of Σ satisfies that $c < \Lambda_{\min}^2$, and $\max_j \sum_{i,j} s_j \leq C$, where c, C are some positive constants.

Theorem 2.1 (**Theorem 2.4** of van de Geer et al. 2014). Suppose **Assumptions 2.1–2.2** hold. Assume in the sub-Gaussian case, it holds that $\max_{1 \leq j \leq p} \sqrt{s_j \log(p)/n} = o(1)$ and in the strongly bounded case, $\max_{1 \leq j \leq p} K_n^2 s_j \sqrt{\log(p)/n} = o(1)$. Then with suitably chosen $\lambda_j \asymp K_0 \sqrt{\log(p)/n}$ uniformly for $j = 1, 2, \dots, p$, where $K_0 = 1$ in the sub-Gaussian case and $K_0 = K_n$ in the strongly bounded case, we have

$$\|\hat{\Theta}_j - \Theta_j\|_1 = O_p\left(K_0 s_j \sqrt{\log(p)/n}\right),$$

$$\|\hat{\Theta}_j - \Theta_j\|_2 = O_p\left(K_0 \sqrt{s_j \log(p)/n}\right), \quad (5)$$

$$|\hat{\tau}_j^2 - \tau_j^2| = O_p\left(K_0 \sqrt{s_j \log(p)/n}\right), \quad (6)$$

uniformly for $1 \leq j \leq p$. Furthermore, suppose $\{\epsilon_i\}$ are iid with $c' < \sigma_\epsilon^2 < c$ and in the sub-Gaussian case for \mathbf{X} , $\mathbb{E} \exp\{|\epsilon_i|/C\} \leq 1$ for some positive constants $c, c', C > 0$. Assume that λ is suitably chosen such that $\lambda \asymp K_0 \sqrt{\log(p)/n}$, and $K_0 s_0 \log(p)/\sqrt{n} = o(1)$ and $\max_{1 \leq j \leq p} K_0 s_j \sqrt{\log(p)/n} = o(1)$. Then

$$\sqrt{n}(\hat{\beta} - \beta^0) = \hat{\Theta} \mathbf{X}^T \epsilon / \sqrt{n} + \Delta, \quad \|\Delta\|_\infty = o_p(1), \quad (7)$$

where $\Delta = (\Delta_1, \dots, \Delta_p)^T = -\sqrt{n}(\hat{\Theta}\hat{\Sigma} - I)(\hat{\beta} - \beta^0)$.

Theorem 2.1 provides an explicit expansion for the desparsifying Lasso estimator and states that the remainder term Δ can be well controlled in the sense that $\|\Delta\|_\infty = o_p(1)$, which is very useful in the subsequent derivation.

2.2. Simultaneous Inference Procedures

In the high-dimensional regime, it is natural to test the hypothesis

$$H_{0,G} : \beta_j^0 = \tilde{\beta}_j \quad \text{for all } j \in G \subseteq \{1, 2, \dots, p\}$$

versus the alternative $H_{a,G} : \beta_j^0 \neq \tilde{\beta}_j$ for some $j \in G$. For example, we consider the sparse testing, that is, $\tilde{\beta}_j = 0$, in **Section 3.2**. Throughout the article, we allow $|G|$ to grow as fast as p , which can be of an exponential order w.r.t. n . Hence, our results go beyond the existing ones in van de Geer et al. (2014), Zhang and Zhang (2014), and Javanmard and Montanari (2014), where $|G|$ is fixed. Simultaneous inference has been considered earlier via Bonferroni adjustment, which leads to very conservative procedures. In contrast, our method is asymptotically nonconservative.

In this section, we propose the test statistic

$$\max_{j \in G} \sqrt{n} |\check{\beta}_j - \tilde{\beta}_j|$$

with $\check{\beta}_j$ being the desparsifying estimator. As will be seen in **Sections 3.2** and **3.3**, this test statistic can be naturally coupled with the margin screening (Fan and Lv 2008) to enhance its power in sparse testing and also reduce the computational cost in nodewise Lasso, or with the step-down method (Romano and Wolf 2005) to provide a strong control for the FWER. We next describe a simple multiplier bootstrap method to obtain an accurate critical value. The asymptotic linear expansion in (7) can be rewritten as

$$(\hat{\Theta} \mathbf{X}^T \epsilon / \sqrt{n})_j = \sum_{i=1}^n \hat{\Theta}_j^T \tilde{X}_i \epsilon_i / \sqrt{n} = \sum_{i=1}^n \hat{\xi}_{ij} / \sqrt{n},$$

$$\hat{\xi}_{ij} = \hat{\Theta}_j^T \tilde{X}_i \epsilon_i,$$

where $(a)_j = a_j$ for $a = (a_1, \dots, a_p)^T$. Generate a sequence of random variables $\{e_i\}_{i=1}^n \stackrel{\text{iid}}{\sim} N(0, 1)$ and define the multiplier bootstrap statistic,

$$W_G = \max_{j \in G} \sum_{i=1}^n \hat{\Theta}_j^T \tilde{X}_i \hat{\sigma}_\epsilon e_i / \sqrt{n},$$

where $\hat{\sigma}_\epsilon^2$ is a consistent estimator of the error variance σ_ϵ^2 , e.g., the variance estimator from the scaled Lasso (Sun and Zhang 2012). The bootstrap critical value is given by $c_G(\alpha) = \inf\{t \in \mathbb{R} : P(W_G \leq t | (Y, \mathbf{X})) \geq 1 - \alpha\}$.

The validity of the above bootstrap method requires the following two assumptions.

Assumption 2.3. (i) If \mathbf{X} has iid sub-Gaussian rows, assume that $(\log(pn))^7/n \leq C_1 n^{-c_1}$ for some constants $c_1, C_1 > 0$. In this case, suppose $\{\epsilon_i\}$ are iid sub-Gaussian with $c' < \sigma_\epsilon^2 < c$ for $c, c' > 0$. (ii) If $\max_{1 \leq i \leq n, 1 \leq j \leq p} |X_{ij}| \leq K_n$, assume that $\max_{1 \leq j \leq p} s_j K_n^2 (\log(pn))^7/n \leq C_2 n^{-c_2}$ for some constants $c_2, C_2 > 0$. In this case, suppose $\{\epsilon_i\}$ are iid subexponential, that is, $\mathbb{E} \exp\{|\epsilon_i|/C'\} \leq 1$ and $c' < \sigma_\epsilon^2 < c$ for some constants $c, c', C' > 0$.

Assumption 2.4. There exists a sequence of positive numbers $\alpha_n \rightarrow +\infty$ such that $\alpha_n/p = o(1)$, $\alpha_n (\log p)^2 \max_j \lambda_j \sqrt{s_j} = o(1)$ and $P(\alpha_n (\log p)^2 |\hat{\sigma}_\epsilon^2 - \sigma_\epsilon^2| > 1) \rightarrow 0$.

Assumption 2.3 requires that (i) the regressors and errors are both sub-Gaussian or (ii) the regressors are strongly bounded while the errors are subexponential. It also imposes suitable restrictions on the growth rate of p . We point out that the existence of a factor $(\log(p))^2$ in **Assumption 2.4** is due to an application of **Theorem 2.1** in Chernozhukov, Chetverikov, and Kato (2014) regarding the comparison for the maxima of two Gaussian random vectors. **Assumption 2.4** is a very mild technical condition. For example if $|\widehat{\sigma}_\epsilon^2 - \sigma_\epsilon^2| = O_p(1/\sqrt{n})$ and $\lambda_j \asymp K_0\sqrt{\log(p)/n}$ uniformly for all j , then **Assumption 2.4** holds provided that $K_0(\log p)^{5/2} \max_j \sqrt{s_j/n} = o(1)$. It is worth noting that the \sqrt{n} convergence rate for $\widehat{\sigma}_\epsilon^2$ can be achieved in the high-dimensional setting (e.g., by the scaled Lasso in Sun and Zhang 2012). Recall that $K_0 = 1$ in the sub-Gaussian case and $K_0 = K_n$ in the strongly bounded case. **Theorem 2.2** below establishes the validity of the bootstrap procedure for one-sided test.

Theorem 2.2. Suppose that **Assumptions 2.1–2.4** hold. Assume that

$$\max_{1 \leq j \leq p} K_0^2 s_j^2 (\log(pn))^3 (\log(n))^2 / n = o(1), \quad (8)$$

$$K_0^4 s_0^2 (\log(p))^3 / n = o(1), \quad (9)$$

and λ and λ_j are suitably chosen such that

$$\lambda \asymp K_0\sqrt{\log(p)/n}, \quad \lambda_j \asymp K_0\sqrt{\log(p)/n} \text{ uniformly for } j. \quad (10)$$

Then we have for any $G \subseteq \{1, 2, \dots, p\}$,

$$\sup_{\alpha \in (0,1)} \left| P \left(\max_{j \in G} \sqrt{n} |\check{\beta}_j - \beta_j^0| > c_G(\alpha) \right) - \alpha \right| = o(1). \quad (11)$$

The scaling condition (8) is imposed to control the estimation effect caused by replacing Θ with its nodewise Lasso estimator $\widehat{\Theta}$, while condition (9) is needed to bound the remainder term Δ . One crucial feature of our bootstrap-assisted testing procedure is that it explicitly accounts for the effect of $|G|$ in the sense that the bootstrap critical value $c_G(\alpha)$ depends on G . This is in sharp contrast with the extreme value approach in Cai, Liu, and Xia (2014); see **Remark 2.4**. Hence, our approach is more robust to the change in $|G|$. Since $\max_{j \in G} \sqrt{n} |\beta_j - \check{\beta}_j| = \sqrt{n} \max_{j \in G} \max\{\check{\beta}_j - \beta_j, \beta_j - \check{\beta}_j\}$, similar arguments imply that under $H_{0,G}$,

$$\sup_{\alpha \in (0,1)} \left| P \left(\max_{j \in G} \sqrt{n} |\check{\beta}_j - \beta_j| > c_G^*(\alpha) \right) - \alpha \right| = o(1),$$

where $c_G^*(\alpha) = \inf\{t \in \mathbb{R} : P(W_G^* \leq t | (Y, \mathbf{X})) \geq 1 - \alpha\}$ with $W_G^* = \max_{j \in G} |\sum_{i=1}^n \widehat{\Theta}_j^T \widetilde{X}_i \widehat{\sigma}_\epsilon e_i / \sqrt{n}|$. This result is readily applicable to construct simultaneous confidence intervals for β_j^0 with $j \in G$.

Remark 2.1. Alternatively, we can employ Efron's empirical bootstrap to obtain the critical value. For simplicity, we take $G = \{1, 2, \dots, p\}$. Let $\widehat{h}_i = (\widehat{h}_{i1}, \dots, \widehat{h}_{ip})^T$ with $\widehat{h}_{ij} = \widehat{\Theta}_j^T \widetilde{X}_i \widehat{\sigma}_\epsilon$. Let $\widehat{h}_1^*, \dots, \widehat{h}_n^*$ be a sample from the empirical distribution based on $\{\widehat{h}_i\}_{i=1}^n$. Define the empirical bootstrap statistic as $W_{\text{EB}}^* = \max_{1 \leq j \leq p} |\sum_{i=1}^n (\widehat{h}_{ij}^* - \sum_{i=1}^n \widehat{h}_{ij}^*/n) / \sqrt{n}|$. The empirical bootstrap critical value is then given by $c_{\text{EB}}^*(\alpha) = \inf\{t \in \mathbb{R} : P(W_{\text{EB}}^* \leq t | (Y, \mathbf{X})) \geq 1 - \alpha\}$. Following the arguments in Appendix K of Chernozhukov, Chetverikov, and Kato (2013) and the proof of **Theorem 2.2**, we can establish the asymptotic

equivalence between the empirical bootstrap and the multiplier bootstrap, which theoretically justifies the use of the former. We omit the technical details here to conserve space.

Let $\widehat{\Xi} = (\widehat{\omega}_{ij})_{i,j=1}^p = \widehat{\Theta} \widehat{\Sigma} \widehat{\Theta}^T \widehat{\sigma}_\epsilon^2$. We next consider the studentized statistic $\max_{j \in G} \sqrt{n} (\check{\beta}_j - \widetilde{\beta}_j) / \sqrt{\widehat{\omega}_{jj}}$ for one-sided test. In this case, the bootstrap critical value can be obtained via $\bar{c}_G(\alpha) = \inf\{t \in \mathbb{R} : P(\bar{W}_G \leq t | (Y, \mathbf{X})) \geq 1 - \alpha\}$, where $\bar{W}_G = \max_{j \in G} \sum_{i=1}^n \widehat{\Theta}_j^T \widetilde{X}_i \widehat{\sigma}_\epsilon e_i / \sqrt{n \widehat{\omega}_{jj}}$. The following theorem justifies the validity of the bootstrap procedure for the studentized statistic.

Theorem 2.3. Under the assumptions in **Theorem 2.2**, we have for any $G \subseteq \{1, 2, \dots, p\}$,

$$\sup_{\alpha \in (0,1)} \left| P \left(\max_{j \in G} \sqrt{n} (\check{\beta}_j - \beta_j^0) / \sqrt{\widehat{\omega}_{jj}} > \bar{c}_G(\alpha) \right) - \alpha \right| = o(1). \quad (12)$$

Following the arguments in the proof of **Theorem 2.3**, it is straightforward to show that under $H_{0,G}$,

$$\sup_{\alpha \in (0,1)} \left| P \left(\max_{j \in G} \sqrt{n} |\check{\beta}_j - \widetilde{\beta}_j| / \sqrt{\widehat{\omega}_{jj}} > \bar{c}_G^*(\alpha) \right) - \alpha \right| = o(1),$$

where $\bar{c}_G^*(\alpha) = \inf\{t \in \mathbb{R} : P(\bar{W}_G^* \leq t | (Y, \mathbf{X})) \geq 1 - \alpha\}$ with $\bar{W}_G^* = \max_{j \in G} |\sum_{i=1}^n \widehat{\Theta}_j^T \widetilde{X}_i \widehat{\sigma}_\epsilon e_i / \sqrt{n \widehat{\omega}_{jj}}|$. Compared to **Theorem 2** in Zhang and Zhang (2014), our two-sided testing procedure above is straightforward to implement and asymptotically exact. In particular, the unknown quantile parameter in (29) of Zhang and Zhang (2014) seems not directly obtainable. Our procedure is also asymptotically nonconservative as compared to the method in Meinshausen (2015).

Remark 2.2. Our inferential procedures allow the precision matrix to be sparse. In fact, they work for any estimation method for precision matrix as long as the estimation effect such as $\|\widehat{\Theta}^T - \Theta^T\|_1$ and $\|\widehat{\Theta} \widehat{\Sigma} - I\|_\infty$ can be well controlled under suitable assumptions (see e.g., Cai, Liu, and Luo 2011; Javanmard and Montanari 2014).

Remark 2.3. An alternative way to conduct simultaneous inference is based on the sum of squares type statistics, for example, Chen and Qin (2010), which is expected to have good power against nonsparse alternatives. However, such a type of test statistics may not work well in the current setting due to the accumulation of estimation errors especially when $|G|$ is large, that is, the error term $\sum_{j \in G} |\Delta_j|$ might be out of control.

We next turn to the (asymptotic) power analysis of the above procedure. Note that when $|G|$ is fixed, our procedure is known to be \sqrt{n} -consistent (implicitly implied by **Theorem 2.1**). In fact, even when $|G| \rightarrow \infty$, our test still enjoys certain optimality in the sense that the separation rate $(\sqrt{2} + \varepsilon_0)\sqrt{\log(|G|)/n}$ for any $\varepsilon_0 > 0$ derived in **Theorem 2.4** below is minimax optimal according to **Section 3.2** of Verzelen (2012) under suitable assumptions.

Below we focus on the case where $|G| \rightarrow \infty$ as $n \rightarrow \infty$. Define the separation set

$$\mathcal{U}_G(c_0) = \{\beta = (\beta_1, \dots, \beta_p)^T : \max_{j \in G} |\beta_j - \widetilde{\beta}_j| / \sqrt{\widehat{\omega}_{jj}} > c_0 \sqrt{\log(|G|)/n}\}, \quad (13)$$

where $\omega_{jj} = \sigma_\epsilon^2 \theta_{jj}$. Recall that $\Sigma^{-1} = \Theta = (\theta_{ij})_{i,j=1}^p$. Let $\tilde{\Theta} = (\tilde{\theta}_{ij})$ with $\tilde{\theta}_{ij} = \theta_{ij} / \sqrt{\theta_{ii} \theta_{jj}}$.

Assumption 2.5. Assume that $\max_{1 \leq i \neq j \leq p} |\tilde{\theta}_{ij}| \leq c < 1$ for some constant c .

Theorem 2.4. Under the assumptions in [Theorem 2.3](#) and [Assumption 2.5](#), we have for any $\varepsilon_0 > 0$,

$$\inf_{\beta^0 \in \mathcal{U}_G(\sqrt{2} + \varepsilon_0)} P \left(\max_{j \in G} \sqrt{n} |\check{\beta}_j - \tilde{\beta}_j| / \sqrt{\hat{\omega}_{jj}} > \tilde{c}_G^*(\alpha) \right) \rightarrow 1. \tag{14}$$

[Theorem 2.4](#) says that the correct rejection of our bootstrap-assisted test can still be triggered even when there exists only one entry of $\beta^0 - \tilde{\beta}$ with a magnitude being larger than $(\sqrt{2} + \varepsilon_0) \sqrt{\log(|G|)/n}$ in G . Hence, our procedure is very sensitive in detecting sparse alternatives. As pointed out by one referee, when Σ is an identity matrix, the constant $\sqrt{2}$ turns out to be asymptotically optimal in the minimax sense (see [Arias-Castro, Candès, and Plan 2011](#); [Ingster, Tsybakov, and Verzelen 2010](#)). We also note that our procedure is more powerful in detecting significant variables when $|G|$ gets smaller in view of the lower bound in (13). This observation partly motivates us to consider the screening procedure in [Section 3.2](#).

Remark 2.4. An important byproduct of the proof of [Theorem 2.4](#) is that the distribution of $\max_{1 \leq j \leq p} \sqrt{n} |\check{\beta}_j - \beta_j^0| / \sqrt{\hat{\omega}_{jj}}$ can be well approximated by $\max_{1 \leq j \leq p} |Z_j|$ with $Z = (Z_1, \dots, Z_p) \sim^d N(0, \tilde{\Theta})$. Therefore, we have for any $x \in \mathbb{R}$ and as $p \rightarrow +\infty$,

$$P \left(\max_{1 \leq j \leq p} n |\check{\beta}_j - \beta_j^0|^2 / \hat{\omega}_{jj} - 2 \log(p) + \log \log(p) \leq x \right) \rightarrow \exp \left\{ -\frac{1}{\sqrt{\pi}} \exp \left(-\frac{x}{2} \right) \right\}. \tag{15}$$

In contrast with our method, the above alternative testing procedure has to require p to diverge. In addition, the critical value obtained from the above type I extreme value distribution may not work well in practice since this weak convergence is typically slow. Instead, it is suggested to employ an “intermediate” approximation to improve the rate of convergence in the literature, for example, [Liu, Lin, and Shao \(2008\)](#).

3. Applications

This section is devoted to three concrete applications of the general theoretical results developed in [Section 2](#). Specifically, we consider (i) support recovery; (ii) testing for sparse signals; (iii) multiple testing using the step-down method.

3.1. Application I: Support Recovery

The major goal of this section is to identify signal locations in a prespecified set \tilde{G} , that is, support recovery. It turns out that this support recovery problem is closely related to the reparsifying procedure¹ applied to the desparsified estimator considered in

[van de Geer \(2014\)](#) (see [Lemma 2.3](#) therein). In comparison with Lasso, stability selection ([Meinshausen and Bühlmann 2010](#)), and the screen and clean procedure ([Wasserman and Roeder 2009](#)), simulation results in [Section 5.2](#) illustrate that our procedure below can be more accurate in recovering signals.

Our support recovery procedure is concerned with setting a proper threshold τ in the following set

$$\hat{S}_0(\tau) = \{j \in \tilde{G} : |\check{\beta}_j| > \lambda_j^*(\tau)\},$$

where $\lambda_j^*(\tau) = \sqrt{\tau \hat{\omega}_{jj} \log(p)/n}$ and $\hat{\omega}_{jj} = \hat{\sigma}_\epsilon^2 \hat{\Theta}_j^T \hat{\Sigma} \hat{\Theta}_j$. We consider the most challenging scenario where $\tilde{G} = [p]$ with $[p] := \{1, 2, \dots, p\}$. In [proposition 3.1](#), we show that the above support recovery procedure is consistent if the threshold value is set as $\tau^* = 2$, and further justify the optimality of τ^* .

Proposition 3.1. Under the assumptions in [Theorem 2.3](#) and [Assumption 2.5](#), we have

$$\inf_{\beta^0 \in \Psi(2\sqrt{2})} P(\hat{S}_0(2) = S_0) \rightarrow 1, \tag{16}$$

where $\Psi(c_0) = \{\beta = (\beta_1, \dots, \beta_p)^T : \min_{j \in S_0} |\beta_j| / \sqrt{\hat{\omega}_{jj}} > c_0 \sqrt{\log(p)/n}\}$. Moreover, we have for any $0 < \tau < 2$,

$$\sup_{\beta^0 \in \mathcal{S}^*(s_0)} P(\hat{S}_0(\tau) = S_0) \rightarrow 0, \tag{17}$$

where $\mathcal{S}^*(s_0) = \{\beta = (\beta_1, \dots, \beta_p)^T \in \mathbb{R}^p : \sum_{j=1}^p \mathbf{I}\{\beta_j \neq 0\} = s_0\}$.

A key step in the proof of [Proposition 3.1](#) is (15) in [Remark 2.4](#).

3.2. Application II: Testing for Sparse Signals

In this section, we focus on the testing problem, $H_0, \tilde{G} : \beta_j^0 = 0$ for any $j \in \tilde{G} \subseteq \{1, 2, \dots, p\}$. To improve the efficiency of the testing procedure and reduce the computational cost in the nodewise Lasso, we propose a three-step procedure that first randomly splits the sample into two subsamples, screens out the irrelevant variables (leading to a reduced model) based on the first subsample, and then performs simultaneous testing in [Section 2.2](#) on the reduced model based on the second subsample.

Suppose the predictors are properly centered and studentized with sample mean zero and standard deviation one. The three-step procedure is formally described as follows:

1. *Random sample splitting:* Randomly split the sample into two subsamples $\{(\tilde{X}_i, Y_i)\}_{i \in \mathcal{D}_1}$ and $\{(\tilde{X}_i, Y_i)\}_{i \in \mathcal{D}_2}$, where $\mathcal{D}_1 \cup \mathcal{D}_2 = \{1, 2, \dots, n\}$, $|\mathcal{D}_1| = \lfloor c_0 n \rfloor$ and $|\mathcal{D}_2| = n - \lfloor c_0 n \rfloor$ for some $0 < c_0 < 1$.
2. *Marginal screening based on \mathcal{D}_1 :* Let $\mathbf{X}_{\mathcal{D}_1}$ be the submatrix of \mathbf{X} that contains the rows in \mathcal{D}_1 and let $Y_{\mathcal{D}_1} = (Y_i)_{i \in \mathcal{D}_1}$. Compute the correlation $W = (w_1, \dots, w_p)^T = \mathbf{X}_{\mathcal{D}_1}^T Y_{\mathcal{D}_1}$ and consider the submodel $\mathcal{S}_\gamma = \{1 \leq j \leq p : |w_j| > \gamma\}$ with γ being a positive number such that $|\mathcal{S}_\gamma| = |\mathcal{D}_2| - 1$ or $|\mathcal{S}_\gamma| = \lfloor |\mathcal{D}_2| / \log(|\mathcal{D}_2|) \rfloor$.
3. *Testing after screening based on \mathcal{D}_2 :* Under the reduced model \mathcal{S}_γ , compute the desparsifying Lasso estimator $\{\check{\beta}_j\}_{j \in \mathcal{S}_\gamma}$ and the variance estimator $\hat{\omega}_{jj}$ based on

¹This reparsifying procedure has the merits that it can improve the l_∞ -bounds of Lasso and has l_q -bounds similar to Lasso (under sparsity conditions).

$\{(\tilde{X}_i, Y_i)\}_{i \in \mathcal{D}_2}$. Define $\tilde{G}_\gamma = \tilde{G} \cap \mathcal{S}_\gamma$. Denote by $T_{\text{nst}, \gamma} = \max_{j \in \tilde{G}_\gamma} \sqrt{n} |\check{\beta}_j|$ and $T_{\text{st}, \gamma} = \max_{j \in \tilde{G}_\gamma} \sqrt{n} |\check{\beta}_j| / \sqrt{\check{\omega}_{jj}}$ the nonstudentized and studentized test statistics, respectively (if $\tilde{G}_\gamma = \emptyset$, we simply set $T_{\text{nst}, \gamma} = T_{\text{st}, \gamma} = 0$ and do not reject $H_{0, \tilde{G}}$). Let $c_{\tilde{G}_\gamma}^*(\alpha)$ and $\tilde{c}_{\tilde{G}_\gamma}^*(\alpha)$ be the corresponding bootstrap critical values at level α . Reject the null hypothesis if $T_{\text{nst}, \gamma} > c_{\tilde{G}_\gamma}^*(\alpha)$ ($T_{\text{st}, \gamma} > \tilde{c}_{\tilde{G}_\gamma}^*(\alpha)$).

In the above three-step procedure, we have employed the data-splitting strategy as suggested in Wasserman and Roeder (2009) to reduce the Type I error rate due to the selection effect (the so-called selective Type I error rate). Under suitable assumptions that rule out unfaithfulness (small partial correlations), see, for example, Fan and Lv (2008), we have $\mathcal{S}_0 \subseteq \mathcal{S}_\gamma$ with an overwhelming probability, which justifies the validity of the second step. On the event that $\mathcal{S}_0 \subseteq \mathcal{S}_\gamma$, the validity and optimality of the inference procedure based on $T_{\text{nst}, \gamma}$ and $T_{\text{st}, \gamma}$ have been established in Section 2.2.

From a practical viewpoint, the three-step procedure enjoys two major advantages over the single-step procedure: (1) the nodewise Lasso involves the computation of p Lasso problems, which can be computationally intensive especially when p is very large (e.g., p could be tens of thousands in genomic studies). The three-step procedure lessens this computation burden as the nodewise Lasso is now performed under the reduced model \mathcal{S}_γ with a much smaller size; (2) due to the screening step, a reduced model is created which could lead to more efficient inference in some cases, see Section 5.3. This can also be seen from Theorem 2.4 that our testing procedure is more powerful when $|G|$ gets smaller.

3.3. Application III: Multiple Testing With Strong FWER Control

We are interested in the following multiple testing problem:

$$H_{0,j} : \beta_j^0 \leq \tilde{\beta}_j \text{ versus } H_{a,j} : \beta_j^0 > \tilde{\beta}_j \text{ for all } j \in G.$$

For simplicity, we set $G = [p]$. To obtain a strong control of the FWER, we couple the bootstrap-assisted testing procedure with the step-down method proposed in Romano and Wolf (2005). Our method is a special case of a general methodology presented in Section 5 of Chernozhukov, Chetverikov, and Kato (2013) by setting $\hat{\beta}$ therein as the desparsifying estimator $\check{\beta}$. As pointed out in Chernozhukov, Chetverikov, and Kato (2013), our method has two important features: (i) it applies to models with an increasing dimension; (ii) it is asymptotically nonconservative as compared to the Bonferroni-Holm procedure since the correlation among the test statistics is taken into account. In fact, we will compare the finite sample performance of our method with that of the Bonferroni-Holm procedure in Section 5.4. We also want to point out that any procedure controlling the FWER will also control the false discovery rate (Benjamin and Hochberg 1995) when there exist some true discoveries.

Denote by Ω the space for all data-generating processes, and ω_0 be the true process. Each null hypothesis $H_{0,j}$ is equivalent to $\omega_0 \in \Omega_j$ for some $\Omega_j \subseteq \Omega$. For any $\eta \subseteq [p]$, denote by $\Omega^\eta = (\cap_{j \in \eta} \Omega_j) \cap (\cap_{j \notin \eta} \Omega_j^c)$ with $\Omega_j^c = \Omega \setminus \Omega_j$. The strong control of

the FWER means that,

$$\sup_{\eta \subseteq [p]} \sup_{\omega_0 \in \Omega^\eta} P_{\omega_0}(\text{reject at least one hypothesis } H_{0,j}, j \in \eta) \leq \alpha + o(1). \quad (18)$$

Let $T_j = \sqrt{n}(\check{\beta}_j - \tilde{\beta}_j)$ and denote by $c_\eta(\alpha)$ the bootstrapped estimate for the $1 - \alpha$ quantile of $\max_{j \in \eta} T_j$. The step-down method in Romano and Wolf (2005) in controlling the FWER is described as follows. Let $\eta(1) = [p]$ at the first step. Reject all hypotheses $H_{0,j}$ such that $T_j > c_{\eta(1)}(\alpha)$. If no hypothesis is rejected, then stop. If some hypotheses are rejected, let $\eta(2)$ be the set of indices for those hypotheses not being rejected at the first step. At step l , let $\eta(l) \subseteq [p]$ be the subset of hypotheses that were not rejected at step $l - 1$. Reject all hypotheses $H_{0,j}, j \in \eta(l)$ satisfying that $T_j > c_{\eta(l)}(\alpha)$. If no hypothesis is rejected, then stop. Proceed in this way until the algorithm stops. As shown in Romano and Wolf (2005), the strong control of the family-wise error holds provided that

$$c_\eta(\alpha) \leq c_{\eta'}(\alpha), \text{ for } \eta \subseteq \eta', \quad (19)$$

$$\sup_{\eta \subseteq [p]} \sup_{\omega_0 \in \Omega^\eta} P_{\omega_0} \left(\max_{j \in \eta} T_j > c_\eta(\alpha) \right) \leq \alpha + o(1). \quad (20)$$

Therefore, we can show that the step-down method together with the multiplier bootstrap provide strong control of the FWER by verifying (19) and (20). The arguments are similar to those in the proof of Theorem 2.2; also see Theorem 5.1 of Chernozhukov, Chetverikov, and Kato (2013).

Proposition 3.2. Under the assumptions in Theorem 2.2, the step-down procedure with the bootstrap critical value $c_\eta(\alpha)$ satisfies (18).

4. Generalization

In this section, our results are extended beyond the linear models to a general framework with a convex loss function and a penalty function. For $y \in \mathcal{Y} \subseteq \mathbb{R}$ and $x \in \mathcal{X} \subseteq \mathbb{R}^p$, consider a loss function $L_\beta(y, x) = L(y, x^T \beta)$ which is strictly convex in $\beta \in \mathbb{R}^p$. The regularized estimator based on the penalty function $\rho_\lambda(\cdot)$ is defined as

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \mathbb{E}_n L_\beta + \sum_{j=1}^p \rho_\lambda(|\beta_j|) \right\}, \quad (21)$$

where $\mathbb{E}_n g = \sum_{i=1}^n g(y_i, x_i) / n$ and $\{(y_i, x_i)\}_{i=1}^n$ is a sequence of iid observations. Note that our formulation (21) slightly generalizes the framework in Section 3 of van de Geer et al. (2014) by considering a general penalty function. Again, we want to test the hypothesis $H_{0,G} : \beta_j^0 = \tilde{\beta}_j$ for all $j \in G \subseteq [p]$ versus $H_{0,G} : \beta_j^0 \neq \tilde{\beta}_j$ for some $j \in G$. Our test statistic is given by $\max_{j \in G} \sqrt{n} |\check{\beta}_j - \tilde{\beta}_j|$ with $\check{\beta}_j$ being the desparsifying Lasso estimator defined below. We use analogous notation as in Section 2 but with some modifications for the current context. For example, denote β^0 as the unique minimizer of $\beta \mapsto \mathbb{E} L_\beta$.

Define

$$\begin{aligned}\dot{L}_\beta(y, x) &= \frac{\partial}{\partial \beta} L_\beta(y, x) = x \dot{L}(y, x^T \beta), \\ \ddot{L}_\beta(y, x) &= \frac{\partial^2}{\partial \beta \partial \beta^T} L_\beta(y, x) = x x^T \ddot{L}(y, x^T \beta),\end{aligned}$$

where $\dot{L} = \partial L(y, a) / \partial a$ and $\ddot{L} = \partial^2 L(y, a) / \partial^2 a$. Define $\widehat{\Sigma} = \mathbb{E}_n \ddot{L}_{\widehat{\beta}}$ and let $\widehat{\Theta} := \widehat{\Theta}(\widehat{\beta})$ be a suitable approximation for the inverse of $\widehat{\Sigma}$ (see more details in Section 3.1.1 of van de Geer et al. 2014). When the penalized loss function in (21) is convex in β , the KKT condition gives

$$\mathbb{E}_n \dot{L}_{\widehat{\beta}} + \widehat{\kappa}_\lambda = 0, \quad (22)$$

where $\widehat{\kappa}_\lambda = (\widehat{\kappa}_1, \dots, \widehat{\kappa}_p)^T$ with $\widehat{\kappa}_j = \text{sign}(\widehat{\beta}_j) \dot{\rho}_\lambda(|\widehat{\beta}_j|)$ if $\widehat{\beta}_j \neq 0$ and some $|\widehat{\kappa}_j| \leq |\dot{\rho}_\lambda(0+)|$ if $\widehat{\beta}_j = 0$. By Taylor expansion, we have

$$\mathbb{E}_n \dot{L}_{\widehat{\beta}} = \mathbb{E}_n \dot{L}_{\beta^0} + \mathbb{E}_n \ddot{L}_{\widehat{\beta}}(\widehat{\beta} - \beta^0) + \mathcal{R},$$

where \mathcal{R} is the remainder term. Plugging back to the KKT condition, we obtain,

$$\mathbb{E}_n \dot{L}_{\beta^0} + \widehat{\Sigma}(\widehat{\beta} - \beta^0) + \mathcal{R} = -\widehat{\kappa}_\lambda,$$

implying that $\widehat{\beta} - \beta^0 + \Delta / \sqrt{n} + \widehat{\Theta} \widehat{\kappa}_\lambda = -\widehat{\Theta} \mathbb{E}_n \dot{L}_{\beta^0} - \widehat{\Theta} \mathcal{R}$, where $\Delta = \sqrt{n}(\widehat{\Theta} \widehat{\Sigma} - I)(\widehat{\beta} - \beta^0)$.

Following van de Geer et al. (2014), we define the desparsifying/debiased estimator as

$$\check{\beta} = \widehat{\beta} + \widehat{\Theta} \widehat{\kappa}_\lambda = \widehat{\beta} - \widehat{\Theta} \mathbb{E}_n \dot{L}_{\widehat{\beta}}.$$

Note $\check{\beta} - \beta^0 = -\widehat{\Theta} \mathbb{E}_n \dot{L}_{\beta^0} - \widehat{\Theta} \mathcal{R} - \Delta / \sqrt{n}$. With some abuse of notation, define $\xi_{ij} = -\Theta_j^T \dot{L}_{\beta^0}(y_i, x_i) = \Theta_j^T x_i \epsilon_i$ with $\epsilon_i = -\dot{L}(y_i, x_i^T \beta^0)$, and $\widehat{\xi}_{ij} = -\widehat{\Theta}_j^T \dot{L}_{\beta^0}(y_i, x_i)$. To conduct inference, we employ the multiplier bootstrap in the following way. Generate a sequence of iid standard normal random variables $\{e_i\}$ and define the bootstrap statistic,

$$\widetilde{W}_G^* = \max_{j \in G} \left| \sum_{i=1}^n \widehat{\Theta}_j^T x_i \dot{L}(y_i, x_i^T \check{\beta}) e_i / \sqrt{n} \right|,$$

where $\check{\beta}$ is a suitable estimator for β^0 . The bootstrap critical value is given by $\widetilde{c}_G^*(\alpha) = \inf\{t \in \mathbb{R} : P(\widetilde{W}_G^* \leq t | \{(y_i, x_i)\}_{i=1}^n)\} \geq 1 - \alpha\}$.

We first impose Assumptions 4.1–4.4 that are similar to Assumptions (C1)–(C5) in van de Geer et al. (2014). Denote by \mathbf{X} the design matrix with the i th row equal to x_i^T .

Assumption 4.1. The derivatives $\dot{L}(y, a)$ and $\ddot{L}(y, a)$ exist for all y and a . For some δ -neighborhood,

$$\max_{a_0 \in \{x_i^T \beta^0 : x_i \in \mathcal{X}\}} \sup_{|a - a_0| \vee |a' - a_0| \leq \delta} \sup_{y \in \mathcal{Y}} \frac{|\ddot{L}(y, a) - \ddot{L}(y, a')|}{|a - a'|} \leq 1.$$

Assumption 4.2. Assume $\max_{i,j} |x_{ij}| \leq K_n$.

Assumption 4.3. Assume that $\|\widehat{\beta} - \beta^0\|_1 = O_P(s_0 \lambda)$ and $\|\mathbf{X}(\widehat{\beta} - \beta^0)\|_2^2 / n = O_P(s_0 \lambda^2)$.

Assumption 4.4. Assume that $\|\mathbb{E}_n \ddot{L}_{\widehat{\beta}} \widehat{\Theta}_j - 1_j\|_\infty = O_P(\lambda_*)$ for some $\lambda_* > 0$, and $\|\mathbf{X} \widehat{\Theta}_j\|_\infty = O_P(K_n)$ uniformly for j . Here, 1_j denotes the vector with the j -th element one and others zero.

Moreover, we make the following additional assumptions. In particular, Assumptions 4.5–4.6 are parallel to Assumptions 2.3–2.4, while Assumption 4.7 is a technical one motivated by the results in Theorem 3.2 of van de Geer et al. (2014).

Assumption 4.5. Suppose $\max_{1 \leq j \leq p} s_j K_n^2 (\log(pn))^7 / n \leq C_2 n^{-c_2}$ for some constants $c_2, C_2 > 0$. Suppose $\{\epsilon_i\}$ are iid subexponential that is $\mathbb{E} \exp(|\epsilon_i|/C') \leq 1$ for some large enough constants $C' > 0$, and $c' < \sigma_\epsilon^2 < c$ for $c, c' > 0$.

Let $\widetilde{\Gamma} = \max_{1 \leq j, k \leq p} |\widehat{\Theta}_j^T \widehat{\Sigma}_{\widehat{\beta}} \widehat{\Theta}_k - \Theta_j^T \Sigma_{\beta^0} \Theta_k|$, where $\widehat{\Sigma}_{\widehat{\beta}} = \mathbf{X}^T \mathbf{W}_{\widehat{\beta}} \mathbf{X} / n$ with $\mathbf{W}_{\widehat{\beta}} = \text{diag}(\dot{L}^2(y_1, x_1^T \widehat{\beta}), \dots, \dot{L}^2(y_n, x_n^T \widehat{\beta}))$, and $\Sigma_{\beta^0} = \text{cov}(x_i \dot{L}(y_i, x_i^T \beta^0))$.

Assumption 4.6. There exists a sequence of positive numbers $\alpha_n \rightarrow +\infty$ such that $\alpha_n / p = o(1)$ and $P(\alpha_n (\log p)^2 \widetilde{\Gamma} > 1) \rightarrow 0$.

Assumption 4.7. Assume that uniformly for j , it holds that

$$\begin{aligned}\|\widehat{\Theta}_j(\widehat{\beta}) - \Theta_j(\beta^0)\|_1 &= O_P(K_n s_j \sqrt{\log(p)/n}) \\ &+ O_P\left(K_n^2 s_0 ((\lambda^2 / \sqrt{\log(p)/n}) \vee \lambda)\right).\end{aligned}$$

We are now in position to present the main result in this section which justifies the use of the bootstrap procedure.

Theorem 4.1. Suppose Assumptions 4.1–4.7 hold. Suppose

$$\begin{aligned}\sqrt{n \log(p)} K_n s_0 \lambda^2 &= o(1), \quad \sqrt{n \log(p)} \lambda \lambda_* s_0 = o(1), \\ \max_j s_j K_n (\log(p))^{3/2} / \sqrt{n} &= o(1), \\ \sqrt{\log(p)} K_n^2 s_0 (\lambda^2 \sqrt{n} \vee \lambda \sqrt{\log(p)}) &= o(1), \\ K_n^2 \log(p) (\log(n))^2 / n &= o(1).\end{aligned}$$

Then we have for any $G \subseteq \{1, 2, \dots, p\}$,

$$\sup_{\alpha \in (0, 1)} \left| P\left(\max_{j \in G} \sqrt{n} |\check{\beta}_j - \beta_j^0| > \widetilde{c}_G^*(\alpha)\right) - \alpha \right| = o(1). \quad (23)$$

Remark 4.1. Let $\widetilde{W}_G^{**} = \max_{j \in G} |\sum_{i=1}^n \widehat{\Theta}_j^T x_i \dot{L}(y_i, x_i^T \check{\beta}) e_i / \sqrt{n \widehat{\omega}_{jj}}|$, where $\omega_{jj} = \widehat{\Theta}_j^T \widehat{\Sigma}_{\widehat{\beta}} \widehat{\Theta}_j$ and $\{e_i\}$ is a sequence of iid standard normal random variables and $\widehat{\omega}_{jj} = \widehat{\Theta}_j^T \widehat{\Sigma}_{\widehat{\beta}} \widehat{\Theta}_j$. Define the bootstrap critical value $\widetilde{c}_G^{**}(\alpha) = \inf\{t \in \mathbb{R} : P(\widetilde{W}_G^{**} \leq t | \{(y_i, x_i)\}_{i=1}^n)\} \geq 1 - \alpha\}$. Following the arguments in the proofs of Theorem 2.3 and Theorem 4.1, we expect a similar result for the studentized test statistic $\max_{j \in G} \sqrt{n} |(\check{\beta}_j - \beta_j^0) / \sqrt{\widehat{\omega}_{jj}}|$. The technical details are omitted here.

5. Numerical Results

In this section, we conduct some simulation studies to evaluate the finite sample performance of the methods proposed in Section 3. All the results are obtained based on sample size $n = 100$, and 1000 Monte Carlo replications.

To obtain the main Lasso estimator, we implemented the scaled Lasso with the tuning parameter $\lambda_0 = \sqrt{2} \widetilde{L}_n(k_0/p)$ with $\widetilde{L}_n(t) = n^{-1/2} \Phi^{-1}(1-t)$, where Φ is the cumulative distribution function for $N(0, 1)$, and k_0 is the solution to $k = \widetilde{L}_1^4(k/p) + 2 \widetilde{L}_1^2(k/p)$ (see Sun and Zhang 2013). We estimate the

noise level σ^2 using the modified variance estimator defined in (24) below. The tuning parameters λ_j s in the nodewise Lasso are chosen via 10-fold cross-validation among all nodewise regressions throughout the simulation studies.

5.1. Simultaneous Confidence Intervals

We consider the linear models where the rows of \mathbf{X} are fixed iid realizations from $N_p(0, \Sigma)$ with $\Sigma = (\Sigma_{i,j})_{i,j=1}^p$ under two scenarios: (i) Toeplitz: $\Sigma_{i,j} = 0.9^{|i-j|}$; (ii) Exchangeable/Compound symmetric: $\Sigma_{i,i} = 1$ and $\Sigma_{i,j} = 0.8$ for $i \neq j$. The active set is $\{1, 2, \dots, s_0\}$, where $s_0 = |\mathcal{S}_0| = 3$ or 15. The coefficients of the linear regression models are generated according to $\text{Unif}[0, 2]$ (uniform distribution on $[0, 2]$), and the errors are generated from the studentized $t(4)$ distribution, that is, $t(4)/\sqrt{2}$, or centralized and the studentized Gamma(4,1) distribution that is, $(\text{Gamma}(4, 1) - 4)/2$.

In our simulations, we found that the scaled Lasso tends to underestimate the noise level which could lead to undercoverage. To overcome this problem, we suggest the following modified variance estimator (see a similar estimator in Reid, Tibshirani, and Friedman 2014),

$$\hat{\sigma}^2 = \frac{1}{n - \|\widehat{\beta}_{\text{sc}}\|_0} \|\mathbf{Y} - \mathbf{X}\widehat{\beta}_{\text{sc}}\|_2^2, \quad (24)$$

where $\widehat{\beta}_{\text{sc}}$ denotes the scaled Lasso estimator with the tuning parameter λ_0 . Figure S.1 in the supplement provides boxplots of $\hat{\sigma}/\sigma$ for the variance estimator delivered by the scaled Lasso (denoted by ‘‘SLasso’’) and for the modified variance estimator in (24) (denoted by ‘‘SLasso*’’). Clearly, the modified variance estimator corrects the noise underestimation issue, and thus is more preferable.

In Tables 1–2, we present the coverage probabilities and interval widths for the simultaneous confidence intervals in three different cases: $G = \mathcal{S}_0, \mathcal{S}_0^c$ or $[p]$ (a similar setup was considered in van de Geer et al. 2014). For each simulation run, we record whether the simultaneous confidence interval contains β_j^0 for $j \in G$ and the corresponding interval width. The coverage probabilities and interval widths are then calculated by averaging over 1,000 simulation runs. It is not surprising that the coverage probability is affected by the dimension p , the tuning parameters λ_j (in the nodewise Lasso), the cardinality s_0 , and the covariance matrix Σ . Overall, the nonstudentized method provides satisfactory coverage probability. However, the method based on the extreme value distribution approximation is invalid when the dimension of the components of interest is low; see Remark 2.4. To avoid sending a misleading message, we choose not to provide the results based on the extreme value distribution approximation when $G = \mathcal{S}_0$.

More specifically, we observe from Tables 1–2 that: (i) the coverage is in general more accurate for \mathcal{S}_0^c (as compared to \mathcal{S}_0), even though $|\mathcal{S}_0^c| \gg |\mathcal{S}_0| = s_0$. Similar finding is found in van de Geer et al. (2014) where the coverage for a single coefficient in \mathcal{S}_0^c is more accurate; (ii) the nonstudentized test statistic tends to provide better coverage (but with larger width) as compared to its studentized version when s_0 is large; (iii) when $s_0 = 15$, it becomes difficult to provide accurate coverage for all the active variables among the candidates. In this case, the coverage for

the active set can be significantly lower than the nominal level. Additional numerical results in the supplement (see Figure S.2) indicate that the undercoverage in this case is closely related with the remainder term whose maximum norm generally increases with s_0 ; (iv) compared to the confidence intervals for individual coefficients in van de Geer et al. (2014), the interval widths for simultaneous intervals are wider, which reflects the price we pay for multiplicity adjustment.

5.2. Support Recovery

Below we compare the finite sample performance of the support recovery procedure described in Section 3.1 with those of Lasso, stability selection (Meinshausen and Bühlmann 2010), and the screen and clean procedure in Wasserman and Roeder (2009). Consider the simulation setup in Section 5.1, where the coefficients are now generated from $\text{Unif}[2, 4]$, and the support set $\mathcal{S}_0 = \{u_1, \dots, u_{s_0}\}$ with u_1, \dots, u_{s_0} being a realization of s_0 iid draws without replacement from $\{1, \dots, p\}$. To assess the performance, we consider the following similarity measure

$$d(\widehat{\mathcal{S}}_0, \mathcal{S}_0) = \frac{|\widehat{\mathcal{S}}_0 \cap \mathcal{S}_0|}{\sqrt{|\widehat{\mathcal{S}}_0| \cdot |\mathcal{S}_0|}}.$$

In the implementation of stability selection, we choose the threshold for selection frequency to be 0.6 and the upper bound for the expected number of false positives to be 2.5; see Meinshausen and Bühlmann (2010). For the screen and clean procedure, we randomly split the data into two groups, conduct screening on the first half of the data and cleaning on the second half. This two splits procedure was advocated in the simulations of Wasserman and Roeder (2009). Table 3 summarizes the mean and standard deviation of $d(\widehat{\mathcal{S}}_0, \mathcal{S}_0)$ as well as the numbers of false positives (FP) and false negatives (FN) based on 1000 simulation runs. When $s_0 = 3$, the proposed support recovery procedure clearly outperforms Lasso, and it is comparable to stability selection and the screen and clean procedure. When $s_0 = 15$, the recovery procedure in general outperforms all other three competitors. We note that when $s_0 = 15$ and Σ is exchangeable, the stability selection gives a high number of FN, and thus results in a low value of $d(\widehat{\mathcal{S}}_0, \mathcal{S}_0)$. This is not surprising given that stability selection is mainly designed for conservatively controlling the expected number of FP. The upper bound 2.5 set above might be too small for $s_0 = 15$. The screen and clean procedure generally performs well for $p = 120$, but its performance deteriorates for $p = 500$ and $s_0 = 15$. Overall, the proposed method performs quite well as compared to some existing alternatives.

5.3. Testing for Sparse Signals

This section is devoted to empirically examine the three-step testing procedure proposed in Section 3.2. We consider the following two scenarios: (i) Σ is Toeplitz with $\Sigma_{i,j} = 0.9^{|i-j|}$, and $\beta_j^0 = \sqrt{10 \log(p)/n}$ for $1 \leq j \leq s_0$ and zero otherwise; (ii) Σ is exchangeable, that is, $\Sigma_{i,i} = 1$ and $\Sigma_{i,j} = 0.8$ for $i \neq j$, and $\beta_j^0 = 10\sqrt{\log(p)/n}$ for $1 \leq j \leq s_0$ and zero otherwise.

To implement the three-step testing procedure, we choose the splitting proportion $c_0 = 1/5$ in case (i) and $c_0 = 1/3$ in case

Table 1. Coverage probabilities and interval widths for the simultaneous confidence intervals based on the nonstudentized (“NST”) and studentized (“ST”) test statistics, where $S_0 = \{1, 2, 3\}$ and $p = 120, 500$.

			$S_{0,(i)}$		$S_{0,(i)}^c$		$[p]_{,(i)}$		$S_{0,(ii)}$		$S_{0,(ii)}^c$		$[p]_{,(ii)}$		
			95%	99%	95%	99%	95%	99%	95%	99%	95%	99%	95%	99%	
$p = 120$ $t(4)/\sqrt{2}$	NST _{cv}	Cov	0.82	0.94	0.97	0.99	0.95	0.99	0.91	0.97	0.93	0.98	0.93	0.98	
		Len	0.99	1.22	1.49	1.67	1.50	1.67	0.97	1.18	1.49	1.67	1.49	1.67	
	ST _{cv}	Cov	0.82	0.93	0.97	0.99	0.96	0.99	0.92	0.97	0.92	0.98	0.92	0.98	
		Len	0.97	1.19	1.48	1.64	1.48	1.64	0.96	1.18	1.46	1.62	1.46	1.63	
	EX _{cv}	Cov	NA	NA	0.98	1.00	0.97	0.99	NA	NA	0.94	0.98	0.93	0.98	
		Len	NA	NA	1.51	1.69	1.51	1.69	NA	NA	1.49	1.66	1.49	1.67	
	Gamma	NST _{cv}	Cov	0.84	0.93	0.96	0.99	0.94	0.98	0.91	0.97	0.93	0.98	0.93	0.98
			Len	0.99	1.22	1.50	1.67	1.50	1.67	0.97	1.18	1.50	1.68	1.50	1.68
		ST _{cv}	Cov	0.82	0.92	0.97	0.99	0.95	0.98	0.90	0.97	0.92	0.98	0.92	0.98
			Len	0.97	1.19	1.48	1.65	1.48	1.65	0.97	1.18	1.46	1.63	1.46	1.63
		EX _{cv}	Cov	NA	NA	0.97	0.99	0.96	0.99	NA	NA	0.93	0.99	0.93	0.99
			Len	NA	NA	1.51	1.69	1.51	1.69	NA	NA	1.49	1.67	1.49	1.67
$p = 500$ $t(4)/\sqrt{2}$	NST _{cv}	Cov	0.76	0.90	0.96	0.99	0.94	0.98	0.92	0.98	0.92	0.97	0.92	0.97	
		Len	0.89	1.09	1.47	1.62	1.47	1.62	0.97	1.19	1.65	1.82	1.65	1.82	
	ST _{cv}	Cov	0.77	0.90	0.97	0.99	0.95	0.98	0.92	0.97	0.92	0.97	0.91	0.97	
		Len	0.88	1.08	1.46	1.60	1.46	1.60	0.97	1.18	1.62	1.77	1.62	1.77	
	EX _{cv}	Cov	NA	NA	0.98	0.99	0.96	0.98	NA	NA	0.92	0.97	0.92	0.97	
		Len	NA	NA	1.48	1.63	1.48	1.63	NA	NA	1.64	1.81	1.64	1.81	
	Gamma	NST _{cv}	Cov	0.77	0.90	0.98	0.99	0.96	0.98	0.91	0.97	0.95	0.98	0.94	0.98
			Len	0.90	1.10	1.49	1.63	1.49	1.64	0.98	1.19	1.66	1.83	1.66	1.83
		ST _{cv}	Cov	0.77	0.90	0.98	0.99	0.96	0.98	0.90	0.97	0.93	0.97	0.93	0.97
			Len	0.89	1.09	1.47	1.61	1.47	1.61	0.97	1.19	1.63	1.78	1.63	1.78
		EX _{cv}	Cov	NA	NA	0.98	1.00	0.96	0.99	NA	NA	0.94	0.98	0.94	0.98
			Len	NA	NA	1.49	1.64	1.49	1.64	NA	NA	1.65	1.82	1.65	1.82

NOTE: The tuning parameters λ, s in the nodewise Lasso are chosen to be the same via 10-fold cross-validation among all nodewise regressions for NST_{cv}, ST_{cv}, and EX_{cv}. $t(4)/\sqrt{2}$ and Gamma denote the studentized $t(4)$ distribution and the centralized and studentized Gamma(4,1) distribution, respectively. The coverage probabilities and interval widths are computed based on 1000 simulation runs. For the studentized test, we report the average interval widths over different components. The row “EX” corresponds to the coverage based on the studentized test statistic with the Type I extreme distribution approximation. “Cov” and “Len” denote the coverage probability and interval width respectively. Cases (i) ((ii)) correspond to Toeplitz matrix (exchangeable matrix).

Table 2. Coverage probabilities and interval widths for the simultaneous confidence intervals based on the nonstudentized (“NST”) and studentized (“ST”) test statistics, where $S_0 = \{1, 2, 3, \dots, 15\}$ and $p = 120, 500$.

			$S_{0,(i)}$		$S_{0,(i)}^c$		$[p]_{,(i)}$		$S_{0,(ii)}$		$S_{0,(ii)}^c$		$[p]_{,(ii)}$		
			95%	99%	95%	99%	95%	99%	95%	99%	95%	99%	95%	99%	
$p = 120$ $t(4)/\sqrt{2}$	NST _{cv}	Cov	0.68	0.87	0.99	1.00	0.87	0.95	0.73	0.88	0.92	0.98	0.89	0.96	
		Len	1.44	1.73	1.68	1.92	1.72	1.97	1.27	1.48	1.67	1.90	1.67	1.91	
	ST _{cv}	Cov	0.50	0.70	0.99	1.00	0.75	0.85	0.68	0.84	0.83	0.92	0.74	0.88	
		Len	1.30	1.50	1.55	1.73	1.56	1.74	1.23	1.42	1.53	1.70	1.53	1.71	
	EX _{cv}	Cov	NA	NA	0.99	1.00	0.77	0.88	NA	NA	0.85	0.94	0.77	0.92	
		Len	NA	NA	1.58	1.77	1.59	1.78	NA	NA	1.55	1.74	1.56	1.75	
	Gamma	NST _{cv}	Cov	0.69	0.88	0.99	1.00	0.86	0.95	0.77	0.90	0.93	0.98	0.90	0.97
			Len	1.44	1.74	1.69	1.93	1.72	1.97	1.27	1.49	1.68	1.91	1.68	1.92
		ST _{cv}	Cov	0.52	0.70	0.99	1.00	0.74	0.86	0.71	0.87	0.85	0.94	0.78	0.91
			Len	1.30	1.51	1.55	1.73	1.57	1.75	1.24	1.43	1.53	1.71	1.54	1.71
		EX _{cv}	Cov	NA	NA	0.99	1.00	0.77	0.88	NA	NA	0.87	0.96	0.81	0.94
			Len	NA	NA	1.59	1.78	1.60	1.79	NA	NA	1.56	1.75	1.57	1.76
$p = 500$ $t(4)/\sqrt{2}$	NST _{cv}	Cov	0.56	0.82	0.99	1.00	0.86	0.95	0.36	0.57	0.93	0.98	0.81	0.94	
		Len	1.35	1.65	1.74	1.95	1.75	1.97	1.32	1.55	2.01	2.38	2.01	2.38	
	ST _{cv}	Cov	0.34	0.57	0.99	1.00	0.76	0.86	0.35	0.55	0.61	0.77	0.49	0.66	
		Len	1.22	1.40	1.58	1.73	1.59	1.74	1.28	1.48	1.71	1.87	1.71	1.87	
	EX _{cv}	Cov	NA	NA	0.99	1.00	0.78	0.87	NA	NA	0.65	0.81	0.53	0.70	
		Len	NA	NA	1.61	1.77	1.61	1.77	NA	NA	1.73	1.90	1.73	1.91	
	Gamma	NST _{cv}	Cov	0.52	0.80	1.00	1.00	0.86	0.95	0.38	0.57	0.94	0.99	0.83	0.96
			Len	1.36	1.66	1.75	1.96	1.76	1.98	1.34	1.57	2.04	2.42	2.04	2.42
		ST _{cv}	Cov	0.32	0.52	0.99	1.00	0.74	0.85	0.34	0.53	0.60	0.77	0.47	0.66
			Len	1.22	1.41	1.59	1.75	1.60	1.75	1.29	1.50	1.73	1.90	1.73	1.90
		EX _{cv}	Cov	NA	NA	0.99	1.00	0.76	0.87	NA	NA	0.64	0.80	0.50	0.69
			Len	NA	NA	1.61	1.78	1.62	1.78	NA	NA	1.75	1.93	1.76	1.93

NOTE: The tuning parameters λ, s in the nodewise Lasso are chosen via 10-fold cross-validation among all nodewise regressions for NST_{cv}, ST_{cv}, and EX_{cv}. $t(4)/\sqrt{2}$ and Gamma denote the studentized $t(4)$ distribution and the centralized and studentized Gamma(4,1) distribution respectively. The coverage probabilities and interval widths are computed based on 1000 simulation runs. For the studentized test, we report the average interval widths over different components.

Table 3. The mean and standard deviation (SD) of $d(\widehat{\mathcal{S}}_0, \mathcal{S}_0)$, and the numbers of false positives (FP) and false negatives (FN).

		$s_0 = 3$				$s_0 = 15$			
		Mean	SD	FP	FN	Mean	SD	FP	FN
$p = 120, (i)$									
$t(4)/\sqrt{2}$	SupRec	0.98	0.05	0.16	0.00	0.98	0.03	0.68	0.02
	Stability	0.93	0.08	0.52	0.00	0.35	0.07	0.63	12.64
	Lasso _{sc}	0.68	0.10	3.94	0.00	0.72	0.04	13.87	0.00
	S&C	0.99	0.04	0.04	0.01	0.87	0.15	0.18	3.14
Gamma	SupRec	0.97	0.06	0.20	0.00	0.98	0.03	0.71	0.00
	Stability	0.93	0.08	0.58	0.00	0.36	0.07	0.61	12.63
	Lasso _{sc}	0.68	0.11	3.98	0.00	0.72	0.04	13.94	0.00
	S&C	0.99	0.03	0.03	0.01	0.87	0.15	0.14	3.30
$p = 120, (ii)$									
$t(4)/\sqrt{2}$	SupRec	0.97	0.06	0.24	0.00	0.98	0.02	0.55	0.01
	Stability	0.97	0.06	0.26	0.00	0.00	0.00	0.13	15.00
	Lasso _{sc}	0.56	0.09	7.09	0.00	0.65	0.04	20.24	0.00
	S&C	0.99	0.04	0.04	0.01	0.80	0.24	0.16	4.41
Gamma	SupRec	0.97	0.06	0.23	0.00	0.98	0.02	0.60	0.00
	Stability	0.96	0.06	0.27	0.00	0.00	0.00	0.13	15.00
	Lasso _{sc}	0.56	0.09	7.07	0.00	0.65	0.04	20.20	0.00
	S&C	0.99	0.03	0.04	0.00	0.81	0.22	0.24	4.40
$p = 500, (i)$									
$t(4)/\sqrt{2}$	SupRec	0.97	0.06	0.20	0.00	0.53	0.04	0.14	10.66
	Stability	0.84	0.10	1.32	0.00	0.36	0.03	1.74	11.97
	Lasso _{sc}	0.62	0.09	5.24	0.00	0.38	0.03	19.25	7.37
	S&C	0.96	0.13	0.36	0.12	0.12	0.17	18.26	14.04
Gamma	SupRec	0.97	0.06	0.18	0.00	0.53	0.04	0.15	10.68
	Stability	0.86	0.09	1.24	0.00	0.36	0.03	1.75	11.93
	Lasso _{sc}	0.62	0.09	5.38	0.00	0.38	0.03	19.15	7.36
	S&C	0.96	0.14	0.55	0.14	0.56	0.08	18.22	14.00
$p = 500, (ii)$									
$t(4)/\sqrt{2}$	SupRec	0.97	0.06	0.25	0.00	0.96	0.04	1.38	0.10
	Stability	0.96	0.07	0.34	0.00	0.08	0.10	0.86	14.59
	Lasso _{sc}	0.41	0.05	15.16	0.00	0.50	0.02	45.64	0.00
	S&C	0.98	0.08	0.05	0.05	0.04	0.11	28.51	14.61
Gamma	SupRec	0.97	0.07	0.27	0.00	0.96	0.04	1.43	0.04
	Stability	0.95	0.07	0.36	0.00	0.08	0.10	0.86	14.58
	Lasso _{sc}	0.41	0.05	15.18	0.00	0.50	0.02	45.41	0.00
	S&C	0.98	0.06	0.05	0.05	0.05	0.11	27.54	14.62

NOTE: The tuning parameters λ, s in the nodewise Lasso are chosen to be the same via 10-fold cross-validation among all nodewise regressions. The subscript "sc" stands for the scaled Lasso. $t(4)/\sqrt{2}$ and Gamma denote the studentized $t(4)$ distribution and studentized Gamma(4,1) distribution respectively. S&C denotes the screen and clean procedure. The mean, SD, FP, and FN are computed based on 1000 simulation runs. Cases (i) (ii) correspond to Toeplitz matrix (exchangeable matrix).

(ii). For the marginal screening step, we pick γ such that $|\mathcal{S}_\gamma| = |\mathcal{D}_2| - 1$. As pointed out in Fan and Lv (2008), the marginal screening may not perform very well in the case of strong pairwise correlation, that is, when Σ is exchangeable. To overcome this problem, we propose the following remedy inspired by the iterative sure independence screening in Fan and Lv (2008). We first select a subset \mathcal{B}_1 of k_1 variables using Lasso based on the subsample \mathcal{D}_1 . Let $\widehat{r}_i = Y_i - \widetilde{X}_i^T \widehat{\beta}$ be the corresponding residuals with $\widehat{\beta}$ being the Lasso estimator, and $\widetilde{X}_i^c = (X_{ij})_{j \notin \mathcal{B}_1}$. We then treat those residuals as the new responses and apply the marginal screening as described in Step 2 to $\{(\widetilde{X}_i^c, \widehat{r}_i)\}_{i \in \mathcal{D}_1}$ to select a subset \mathcal{B}_2 of $|\mathcal{D}_2| - 1 - k_1$ variables. Finally, we let $\mathcal{B} = \mathcal{B}_1 \cup \mathcal{B}_2$, which contains $|\mathcal{D}_2| - 1$ variables. In case (ii) with $s_0 = 3$, this remedy selects all the three relevant variables with probability 0.98 which is much higher than 0.59 delivered by the marginal screening.

The numerical results are reported in Table 4, which compares the performance between the three-step procedure and the one-step procedure without sample splitting and marginal screening. Some remarks are in order regarding the simulation results: (1) for the Toeplitz covariance structure, the three-step procedure has reasonable size for t errors, and its size is slightly upward distorted for Gamma errors. In contrast, the

one-step procedure has downward size distortion; for exchangeable covariance structure, both procedures show upward size distortions; (2) the three-step procedure generates higher power for the Toeplitz covariance structure. We note that the empirical powers of both procedures are close to 1 in case (ii) due to the higher signal-to-noise ratio. Overall, the three-step procedure has better power property in case (i).

5.4. Testing With FWER Control

In this section, we compare the finite sample performance of the step-down method in Section 3.3 with the Bonferroni-Holm procedure whose finite sample performance has been studied in van de Geer et al. (2014). To this end, we consider multiple two sided testing of hypothesis $H_{0,j} : \beta_j^0 = 0$ among all $j = 1, 2, \dots, p$. The data is again generated from the linear model considered in Section 5.1. For Σ , we focus on the following two cases: (i) Toeplitz: $\Sigma_{i,j} = 0.9^{|i-j|}$; (ii) Block diagonal: $\Sigma_{i,i} = 1$, $\Sigma_{i,j} = 0.9$ for $5(k-1) + 1 \leq i \neq j \leq 5k$ with $k = 1, 2, \dots, \lfloor p/5 \rfloor$, and $\Sigma_{i,j} = 0$ otherwise. We employ both the step-down method based on the studentized/nonstudentized test statistic, and the Bonferroni-Holm procedure (based on the studentized test statistic) to control the FWER. Table 5 reports

Table 4. Empirical sizes (upper panel) and powers (lower panel) based on the one-step (without sample splitting and screening) and three-step procedures with the nonstudentized (“NST”) and the studentized (“ST”) statistics, where $S_0 = \{1, 2, 3\}$ and $S_0 = \{1, 2, 3, \dots, 15\}$, and $p = 500$.

	α	One-step procedure				Three-step procedure			
		$t(4)/\sqrt{2}$		Gamma		$t(4)/\sqrt{2}$		Gamma	
		NST	ST	NST	ST	NST	ST	NST	ST
(i), S_0^c	5%	0.03	0.03	0.02	0.03	0.06	0.06	0.07	0.07
	1%	0.01	0.00	0.00	0.01	0.01	0.01	0.03	0.02
(i), \tilde{S}_0^c	5%	0.02	0.01	0.02	0.02	0.05	0.04	0.06	0.05
	1%	0.01	0.00	0.01	0.01	0.02	0.02	0.03	0.02
(iii), S_0^c	5%	0.09	0.08	0.08	0.08	0.10	0.11	0.09	0.10
	1%	0.02	0.03	0.03	0.03	0.04	0.04	0.03	0.03
(i), $\{3\} \cup S_0^c$	5%	0.66	0.67	0.61	0.62	0.74	0.72	0.73	0.72
	1%	0.54	0.56	0.49	0.51	0.62	0.60	0.63	0.61
(i), $\{2, 3\} \cup S_0^c$	5%	0.89	0.88	0.88	0.87	0.93	0.92	0.94	0.93
	1%	0.80	0.80	0.79	0.78	0.87	0.85	0.87	0.85
(i), $\{15\} \cup \tilde{S}_0^c$	5%	0.63	0.64	0.58	0.59	0.73	0.72	0.70	0.68
	1%	0.53	0.54	0.48	0.50	0.63	0.61	0.60	0.57
(i), $\{14, 15\} \cup \tilde{S}_0^c$	5%	0.86	0.86	0.82	0.82	0.93	0.92	0.93	0.92
	1%	0.76	0.76	0.71	0.72	0.86	0.85	0.87	0.85
(ii), $\{3\} \cup S_0^c$	5%	1.00	1.00	1.00	1.00	0.99	0.99	1.00	1.00
	1%	1.00	1.00	1.00	1.00	0.99	0.99	1.00	1.00
(ii), $\{2, 3\} \cup S_0^c$	5%	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	1%	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

NOTE: The tuning parameters λ, s in the nodewise Lasso are chosen to be the same via 10-fold cross-validation among all nodewise regressions. $t(4)/\sqrt{2}$ and Gamma denote the studentized $t(4)$ distribution and the centralized and studentized Gamma(4,1) distribution, respectively. The sizes and powers are computed based on 1000 simulation runs. Cases (i) ((iii)) correspond to Toeplitz matrix (exchangeable matrix). The nominal levels are 5% and 1%.

Table 5. FWER and power of multiple testing based on the step-down method with the nonstudentized (“NST”) and studentized (“ST”) test statistics, and based on the Bonferroni-Holm procedure (“BH”), where $p = 500$, and the nominal level is 5%. Case (i) ((iii)) corresponds to Toeplitz matrix (block diagonal matrix).

		$s_0 = 3, (i)$		$s_0 = 15, (i)$		$s_0 = 3, (ii)$		$s_0 = 15, (ii)$	
		FWER	Power	FWER	Power	FWER	Power	FWER	Power
$t(4)/\sqrt{2}$	NST	0.037	0.548	0.008	0.732	0.046	0.594	0.005	0.701
	ST	0.028	0.534	0.009	0.722	0.048	0.560	0.005	0.685
	BH	0.024	0.528	0.006	0.717	0.040	0.555	0.004	0.678
Gamma	NST	0.034	0.535	0.014	0.725	0.039	0.581	0.001	0.701
	ST	0.033	0.513	0.016	0.714	0.046	0.549	0.004	0.685
	BH	0.023	0.506	0.011	0.708	0.038	0.545	0.004	0.680

NOTE: The tuning parameters λ, s in the nodewise Lasso are chosen to be the same via 10-fold cross-validation among all nodewise regressions. $t(4)/\sqrt{2}$ and Gamma denote the studentized $t(4)$ distribution and the centralized and studentized Gamma(4,1) distribution respectively. The FWER and powers are computed based on 1,000 simulation runs.

both the FWER and the average power, which is defined as $\sum_{j \in S_0} \mathbf{I}\{H_{0,j} \text{ is rejected}\} / s_0$ based on 1,000 simulated datasets. As seen from the table, the two procedures provide similar control on the FWER. And the step-down method delivers slightly higher average power across all cases considered here.

6. Conclusion

In this article, we have introduced a bootstrap-assisted procedure to conduct simultaneous inference for high-dimensional components of a large parameter vector in sparse linear models. Our procedure is proved to achieve the prespecified significance level asymptotically and to enjoy certain optimality in terms of its power. Our general theory has been successfully applied to three concrete examples, namely support recovery, testing for sparse signals, and multiple testing using the step-down method. Below we point out a few future research directions. The first direction is the automatic and efficient selection of the tuning parameters (e.g., λ, s in the nodewise Lasso) in the context of hypothesis testing and confidence interval construction (see,

e.g., Appendix A.1 of Dezeure et al. 2014). The second direction is to adapt the proposed method to conduct inference on high-dimensional concentration matrix (see, e.g., Jankova and van de Geer 2014). Finally, it is also interesting to extend our results to more complicated models, for example, Cox model, after some technical modifications.

Supplementary Materials

The supplementary materials contain technical details and additional numerical results.

Acknowledgements

The authors thank the Editor, Associate Editor and reviewers for their constructive comments and helpful suggestions, which substantially improved the article.

Funding

Research Sponsored by NSF CAREER Award DMS-1151692, DMS-1418042, Simons Fellowship in Mathematics, Office of Naval Research

(ONR N00014-15-1-2331) and a grant from Indiana Clinical and Translational Sciences Institute. Guang Cheng was on sabbatical at Princeton while part of this work was carried out; he would like to thank the Princeton ORFE department for its hospitality and support.

References

- Arias-Castro, E., Candès, E. J., and Plan, Y. (2011), “Global Testing Under Sparse Alternatives: ANOVA, Multiple Comparisons and the Higher Criticism,” *Annals of Statistics*, 39, 2533–2556. [1,5]
- Belloni, A., Chernozhukov, V., and Kato, K. (2014), “Uniform Post Selection Inference for LAD Regression and Other Z-Estimation Problems,” *Biometrika*, to appear. [1,2]
- Benjamin, Y., and Hochberg, Y. (1995), “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing,” *Journal of the Royal Statistical Society, Series B*, 57, 289–300. [6]
- Bühlmann, P. m., and van de Geer, S. (2011), *Statistics for High-Dimensional Data: Methods, Theory and Applications*, New York: Springer. [1]
- Cai, T., Liu, W., and Luo, X. (2011), “A Constrained ℓ_1 Minimization Approach to Sparse Precision Matrix Estimation,” *Journal of the American Statistical Association*, 106, 594–607. [4]
- Cai, T. T., Liu, W. D., and Xia, Y. (2014), “Two-Sample Test of High Dimensional Means Under Dependence,” *Journal of the Royal Statistical Society, Series B*, 76, 349–372. [4]
- Chen, S. X., and Qin, Y.-L. (2010), “A Two Sample Test for High Dimensional Data With Applications to Gene-Set Testing,” *Annals of Statistics*, 38, 808–835. [4]
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2013), “Gaussian Approximations and Multiplier Bootstrap for Maxima of Sums of High-Dimensional Random Vectors,” *Annals of Statistics*, 41, 2786–2819. [2,4,6]
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2014), “Comparison and Anti-Concentration Bounds for Maxima of Gaussian Random Vectors,” *Probability Theory and Related Fields*, 162, 47–70. [4]
- Dezeure, R., Bühlmann, P., Meier, L., and Meinshausen, N. (2014), “High-Dimensional Inference: Confidence Intervals, p -Values and R-Software hdi,” arXiv:1408.4026v1. [11]
- Fan, J., and Lv, J. (2008), “Sure Independence Screening for Ultra-High Dimensional Feature Space,” *Journal of the Royal Statistical Society, Series B*, 70, 849–911. [2,3,6,10]
- Feng, L., Zou, C., Wang, Z., and Chen, B. (2013), “Rank-Based Score Tests for High-Dimensional Regression Coefficients,” *Electronic Journal of Statistics*, 7, 2131–2149. [1]
- Ingster, Y. I., Tsybakov, A. B., and Verzelen, N. (2010), “Detection Boundary in Sparse Regression,” *Electronic Journal of Statistics*, 4, 1476–1526. [5]
- Jankova, J., and van de Geer, S. (2014), “Confidence Intervals for High-Dimensional Inverse Covariance Estimation,” arXiv:1403.6752. [11]
- Javanmard, A., and Montanari, A. (2014), “Confidence Intervals and Hypothesis Testing for High-Dimensional Regression,” arXiv:1306.3171. [1,3,4]
- Lee, J. D., Sun, D. L., Sun, Y., and Taylor, J. E. (2016), “Exact Post-Selection Inference With Application to the Lasso,” *Annals of Statistics*, 44, 907–927. [1]
- Liu, W.-D., Lin, Z., and Shao, Q.-M. (2008), “The Asymptotic Distribution and Berry Esseen Bound of a New Test for Independence in High Dimension With an Application to Stochastic Optimization,” *Annals of Applied Probability*, 18, 2337–2366. [5]
- Lockhart, R., Taylor, J., Tibshirani, R., and Tibshirani, R. (2014), “A Significance Test for the Lasso,” *Annals of Statistics*, 42, 413–468. [1]
- Mandozzi, J., and Bühlmann, P. (2014), “Hierarchical Testing in the High-Dimensional Setting With Correlated Variables,” arXiv:1312.5556. [1]
- Meier, L., Van De Geer, S., and Bühlmann, P. (2008), “The Group Lasso for Logistic Regression,” *Journal of the Royal Statistical Society, Series B*, 70, 53–71. [1]
- Meinshausen, N. (2008), “Hierarchical Testing of Variable Importance,” *Biometrika*, 95, 265–278. [1]
- (2015), “Group-Bound: Confidence Intervals for Groups of Variables in Sparse High-Dimensional Regression Without Assumptions on the Design,” *Journal of the Royal Statistical Society, Series B*, to appear. [1,4]
- Meinshausen, N., and Bühlmann, P. (2006), “High-Dimensional Graphs and Variable Selection With the Lasso,” *Annals of Statistics*, 34, 1436–1462. [2]
- (2010), “Stability Selection,” *Journal of the Royal Statistical Society*, 72, 417–473. [2,5,8]
- Meinshausen, N., Meier, L., and Bühlmann, P. (2009), “P-Values for High-Dimensional Regression,” *Journal of the American Statistical Association*, 104, 1671–1681. [1]
- Reid, S., Tibshirani, R., and Friedman, J. (2014), “A Study of Error Variance Estimation in Lasso Regression,” arXiv:1311.5274. [8]
- Romano, J., and Wolf, M. (2005), “Exact and Approximate Step-Down Methods for Multiple Hypothesis Testing,” *Journal of the American Statistical Association*, 100, 94–108. [2,3,6]
- Sun, T., and Zhang, C. H. (2012), “Scaled Sparse Linear Regression,” *Biometrika*, 99, 879–898. [3]
- (2013), “Sparse Matrix Inversion With Scaled Lasso,” *Journal of Machine Learning Research*, 14, 3385–3418. [7]
- Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [2]
- van de Geer, S. (2014), “Worst Possible Sub-Directions in High-Dimensional Models,” arXiv:1403.7023. [5]
- van de Geer, S., Bühlmann, P., Ritov, Y., and Dezeure, R. (2014), “On Asymptotically Optimal Confidence Regions and Tests for High Dimensional Models,” *Annals of Statistics*, 42, 1166–1202. [1,2,3,6,7,8,10]
- Verzelen, N. (2012), “Minimax Risks for Sparse Regressions: Ultra-High Dimensional Phenomenons,” *Electronic Journal of Statistics*, 6, 38–90. [2,4]
- Wasserman, L., and Roeder, K. (2009), “High Dimensional Variable Selection,” *Annals of Statistics*, 37, 2178–2201. [1,2,5,6,8]
- Yuan, M., and Lin, Y. (2006), “Model Selection and Estimation in Regression With Grouped Variables,” *Journal of the Royal Statistical Society, Series B*, 68, 49–67. [1]
- Zhang, C. H., and Zhang, S. S. (2014), “Confidence Intervals for Low-Dimensional Parameters With High-Dimensional Data,” *Journal of the Royal Statistical Society, Series B*, 76, 217–242. [1,2,3,4]
- Zhong, P. S., and Chen, S. X. (2011), “Testing for High-Dimensional Regression Coefficients With Factorial Designs,” *Journal of the American Statistical Association*, 106, 260–274. [1]