

Simultaneous Support Recovery in High Dimensions: Benefits and Perils of Block ℓ_1/ℓ_∞ -Regularization

Sahand N. Negahban and Martin J. Wainwright, *Senior Member, IEEE*

Abstract—Given a collection of $r \geq 2$ linear regression problems in p dimensions, suppose that the regression coefficients share partially common supports of size at most s . This set-up suggests the use of ℓ_1/ℓ_∞ -regularized regression for joint estimation of the $p \times r$ matrix of regression coefficients. We analyze the high-dimensional scaling of ℓ_1/ℓ_∞ -regularized quadratic programming, considering both consistency rates in ℓ_∞ -norm, and how the minimal sample size n required for consistent variable selection scales with model dimension, sparsity, and overlap between the supports. We first establish bounds on the ℓ_∞ -error as well sufficient conditions for exact variable selection for fixed design matrices, as well as for designs drawn randomly from general Gaussian distributions. Specializing to the case $r = 2$ linear regression problems with standard Gaussian designs whose supports overlap in a fraction $\alpha \in [0, 1]$ of their entries, we prove that ℓ_1/ℓ_∞ -regularized method undergoes a phase transition characterized by the rescaled sample size $\theta_{1,\infty}(n, p, s, \alpha) = n/\{(4 - 3\alpha)s \log(p - (2 - \alpha)s)\}$. An implication is that the use of ℓ_1/ℓ_∞ -regularization yields improved statistical efficiency if the overlap parameter is large enough ($\alpha > 2/3$), but has worse statistical efficiency than a naive Lasso-based approach for moderate to small overlap ($\alpha < 2/3$). Empirical simulations illustrate the close agreement between theory and actual behavior in practice. These results show that caution must be exercised in applying ℓ_1/ℓ_∞ block regularization: if the data does not match its structure very closely, it can impair statistical performance relative to computationally less expensive schemes.

Index Terms— ℓ_1 -constraints, compressed sensing, convex relaxation, group Lasso, high-dimensional inference, model selection, phase transitions, sparse approximation, subset selection.

I. INTRODUCTION

THE past decade has witnessed a flurry of research on the use of ℓ_1 -based and related convex relaxations for solving problems of high-dimensional statistical inference. These methods are suitable when the problem exhibits some

type of sparsity, and such structure can lead to substantial gains in statistical efficiency. For instance, while it is impossible to consistently estimate a general p -dimensional vector unless the number of samples n is larger than the dimension, the additional side-information of sparsity allows for consistent estimation even in the high-dimensional setting $p \gg n$.

This paper deals with high-dimensional scaling in the context of solving multiple regression problems, where the regression vectors are assumed to have shared sparse structure. More specifically, suppose that we are given a collection of r different linear regression models in p dimensions, with regression vectors $\bar{\beta}^i \in \mathbb{R}^p$, for $i = 1, \dots, r$. Let $S(\bar{\beta}^i) = \{j \mid \bar{\beta}_j^i \neq 0\}$ denote the support set of $\bar{\beta}^i$, and suppose that each such support has cardinality at most s . In many applications—among them sparse approximation, graphical model selection, and image reconstruction—it is natural to impose a sparsity constraint, corresponding to restricting the cardinality $|S(\bar{\beta}^i)|$ of each support set. Moreover, one might expect some amount of overlap between the sets $S(\bar{\beta}^i)$ and $S(\bar{\beta}^j)$ for indices $i \neq j$ since they correspond to the sets of active regression coefficients in each problem. Let us consider some examples to illustrate.

- Image denoising or compression is often performed using a wavelet transform or some other type of multiresolution basis [18], and natural images tend to have sparse representations in such bases [30]. Moreover, similar images—say the same scene taken from multiple cameras—would be expected to share a similar subset of active features in the reconstruction. Consequently, one might expect that using a block-regularizer that enforces such joint sparsity could lead to improved image denoising or compression.
- Consider the problem of identifying the structure of a Markov network [11] based on a collection of samples (e.g., such as observations of a social network). For networks with a single parameter per edge—including Gaussian Markov random fields models and binary graphical models—a line of recent work [20], [28] has shown that ℓ_1 -based methods can be successful in recovering the network structure. However, many graphical models have multiple parameters per edge (e.g., for discrete models with nonbinary state spaces), and it is natural that the subset of parameters associated with a given edge are zero (or nonzero) in a grouped manner. Thus, any method for recovering the graph structure should impose a block-structured regularization that groups together the subset of parameters associated with a single edge.
- Finally, a standard problem in genetic analysis is to use a set of gene expression arrays, where each array corresponds to a different patient but the same underlying tissue

Manuscript received May 04, 2009; revised May 21, 2010; accepted September 16, 2010. Date of current version May 25, 2011. This work was supported in part by the National Science Foundation under Grants DMS-0605165 and CCF-0545862, and by the Air Force Office of Scientific Research under Grant AFOSR-09NL184. The work of S. N. Negahban was also supported in part by a Vodafone Fellowship.

This work was presented in part at the NIPS 2008 Conference, Vancouver, BC, Canada, December 2008. This work was posted in May 2009 as Technical Report 774, Department of Statistics, University of California, Berkeley, and on arxiv.org as 0905.0642.

S. N. Negahban is with the Department of Electrical Engineering and Computer Sciences, Wireless Foundations Center, University of California, Berkeley, CA 94720-1770 USA (e-mail: sahand_n@eecs.berkeley.edu).

M. J. Wainwright is with the Department of Statistics and the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA (e-mail: wainwrig@eecs.berkeley.edu).

Communicated by A. Krzyzak, Associate Editor for Pattern Recognition, Statistical Learning, and Inference.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2011.2144150

type (e.g., tumor), in order to discover the subset of features relevant for tumorous growths. This problem can be expressed as a joint regression problem, again with a shared sparsity constraint coupling together the different patients. In this context, the recent work of Liu *et al.* [15] shows that imposing additional structural constraints can be beneficial (e.g., they are able to greatly reduce the number of expressed genes while maintaining the same prediction performance).

Given these structural conditions of shared sparsity in these and other applications, it is reasonable to consider how this common structure can be exploited so as to increase the statistical efficiency of estimation procedures.

There is now a substantial and relatively mature body of work on ℓ_1 -regularization for sparse regression models, dating back to the introduction of the Lasso and basis pursuit [7], [31]. The literature contains several lines of work, distinguished by the type of loss function considered. Some authors seek to obtain a regression estimate $\hat{\beta}$ that has low prediction error, whereas others seek to obtain a consistent estimate of the true regression vector β^* in ℓ_2 -norm. For these problems, it is known that it is both necessary and sufficient to impose certain lower bounds on the singular values of the design matrix; we refer the reader to the papers [5], [34] and references therein for discussion of the weakest known conditions, and to the paper [26] for necessary and sufficient conditions in the minimax setting. In contrast to these lines of work, of most relevance to this paper is the criterion of *variable selection*—that is, returning an estimate that correctly identifies the support of the unknown regression vector [20], [36], [39].

A more recent line of work (e.g., [24], [32], [33], [37], [38]), motivated by applications in which block or hierarchical structure arises, has proposed the use of block $\ell_{a,b}$ norms for various $a, b \in [1, \infty]$. Other researchers (e.g., [1], [3], [23]) have studied the use of the nuclear or trace norm as a regularizer for multitask regression, suitable for problems with low-rank structure. Of particular relevance to this paper is the block ℓ_1/ℓ_∞ norm, proposed initially by Turlach *et al.* [33] and Tropp *et al.* [32]. This form of block regularization is a special case of the more general family of composite or hierarchical penalties, as studied by Zhao *et al.* [38]. Various authors have empirically demonstrated that block regularization schemes can yield better performance for different data sets [15], [24], [38]. Some recent work by Bach [2] has provided consistency results for ℓ_1/ℓ_2 block-regularization schemes under classical scaling, meaning that $n \rightarrow +\infty$ with p fixed. Meier *et al.* [19] has established high-dimensional consistency for the predictive risk of ℓ_1/ℓ_2 block-regularized logistic regression. The papers [16], [22], [27] have provided high-dimensional consistency results for ℓ_1/ℓ_q block regularization for support recovery using fixed design matrices, but the rates do not provide sharp differences between the case $q = 1$ and $q > 1$.

At the initial submission of this paper, there was relatively limited amount of theoretical work characterizing if and when the use of block regularization schemes actually leads to gains in statistical efficiency. Some concurrent and subsequent work has addressed this question in the affirmative for the case of ℓ_1/ℓ_2 regularization. For the problem of variable selection consid-

ered here, concurrent work by Obozinski *et al.* [25] (involving a subset of the current authors) has studied whether ℓ_1/ℓ_2 regularization yields improvements for variable selection in multivariate regression. Under incoherence conditions on the design matrices, they showed that it can yield statistical gains up to a factor of r , corresponding to the number of tasks in the multivariate regression. Subsequent work by two separate groups [10], [17] has also studied the case of ℓ_1/ℓ_2 regularization, focusing primarily on estimation of the regression matrix in Frobenius norm. This work again shows gains as a function of the number of tasks, yielding significant gains when r is much larger than $\log p$.

This paper considers the effect of ℓ_1/ℓ_∞ regularization on variable selection performance. Our main contribution is to obtain some precise—and arguably surprising—insights into the benefits and dangers of using block ℓ_1/ℓ_∞ regularization, as compared to simpler ℓ_1 -regularization (separate Lasso for each regression problem). We begin by providing a general set of sufficient conditions for consistent variable selection for both fixed design matrices, and random Gaussian design matrices. In addition to these basic consistency results, we then seek to characterize rates, for the particular case of standard Gaussian designs, in a manner precise enough to address the following questions:

- a) First, under what structural assumptions on the data does the use of ℓ_1/ℓ_∞ block-regularization provide a quantifiable reduction in the scaling of the sample size n , as a function of the problem dimension p and other structural parameters, required for achieving variable selection?
- b) Second, are there any settings in which ℓ_1/ℓ_∞ block-regularization can be harmful relative to computationally less expensive procedures?

Answers to these questions yield useful insight into the tradeoff between computational and statistical efficiency in high-dimensional inference. Indeed, the convex programs that arise from using block-regularization typically require a greater computational cost to solve. Accordingly, it is important to understand under what conditions this increased computational cost guarantees that fewer samples are required for achieving a fixed level of statistical accuracy.

The analysis of this paper gives conditions on the designs and regression matrix \bar{B} for which ℓ_1/ℓ_∞ yields improvements (question (a)), and also shows that if there is sufficient mismatch between the regression matrix \bar{B} and the ℓ_1/ℓ_∞ norm, then use of this regularizer actually impairs statistical efficiency relative to a naive ℓ_1 -approach. As a representative instance of our theory, consider the special case of standard Gaussian design matrices and two regression problems ($r = 2$), with the supports $S(\bar{\beta}^1)$ and $S(\bar{\beta}^2)$ each of size s and overlapping in a fraction $\alpha \in [0, 1]$ of their entries. For this problem, we prove that block ℓ_1/ℓ_∞ regularization undergoes a phase transition—meaning a sharp threshold between success and recovery—that is specified by the rescaled sample size

$$\theta_{1,\infty}(n, p, s, \alpha) := \frac{n}{(4 - 3\alpha)s \log(p - (2 - \alpha)s)}. \quad (1)$$

In words, for any $\delta > 0$ and for scalings of the quadruple (n, p, s, α) such that $\theta_{1,\infty} \geq 1 + \delta$, the probability of suc-




Fig. 1. Probability of success in recovering the joint signed supports plotted against the rescaled sample size $\theta_{\text{Las}} := n/[2s \log(p - (2 - \alpha)s)]$ for linear sparsity $s = 0.1p$. Each stack of graphs corresponds to a fixed overlap α , as labeled on the figure. The three curves within each stack correspond to problem sizes $p \in \{128, 256, 512\}$; note how they all align with each other and exhibit step-like behavior, consistent with Theorem 3. The vertical lines correspond to the thresholds $\theta_{1,\infty}(\alpha)$ predicted by Theorem 3; note the close agreement between theory and simulation.

cessfully recovering both $S(\bar{\beta}^1)$ and $S(\bar{\beta}^2)$ converges to one, whereas for scalings such that $\theta_{1,\infty} \leq 1 - \delta$, the probability of success converges to zero.

Fig. 1 illustrates how the theoretical threshold (1) agrees with the behavior observed in practice. This figure plots the probability of successful recovery using the block ℓ_1/ℓ_∞ approach versus the rescaled sample size $n/\{2s \log[p - (2 - \alpha)s]\}$; the results shown here are for $r = 2$ regression parameters. The plots show 12 curves, corresponding to three different problem sizes $p \in \{128, 256, 512\}$ and four different values of the overlap parameter $\alpha \in \{0.1, 0.3, 0.7, 1\}$. First, let us focus on the set of curves labeled with $\alpha = 1$, corresponding to case of complete overlap between the regression vectors. Notice how the curves for all three problem sizes p , when plotted versus the rescaled sample size, line up with one another; this “stacking effect” shows that the rescaled sample size captures the phase transition behavior. Similarly, for other choices of the overlap, the sets of three curves (over problem size p) exhibit the same stacking behavior. Secondly, note that the results are consistent with the theoretical prediction (1): the stacks of curves shift to the right as the overlap parameter α decreases from 1 towards 0, showing that problems with less overlap require a larger rescaled sample size. More interesting is the sharpness of agreement in quantitative terms: the vertical lines in the center of each stack show the point at which our theory (1) predicts that the method should transition from failure to success.

By comparison to previous theory on the behavior of the Lasso (ordinary ℓ_1 -regularized quadratic programming), the scaling (1) has two interesting implications. For the s -sparse

regression problem with standard Gaussian designs, the Lasso has been shown [36] to transition from success to failure as a function of the rescaled sample size

$$\theta_{\text{Las}}(n, p, s) := \frac{n}{2s \log(p - s)}. \tag{2}$$

In particular, under the conditions imposed here, solving two separate Lasso problems, one for each regression problem, would recover both supports for problem sequences (n, p, s) such that $\theta_{\text{Las}} > 1$. Thus, one consequence of our analysis is to characterize the relative statistical efficiency of ℓ_1/ℓ_∞ regularization versus ordinary ℓ_1 -regularization, as described by the ratio $R := \frac{\theta_{1,\infty}}{\theta_{\text{Las}}}$.

Our theory predicts that (disregarding some $o(1)$ factors) the relative efficiency scales as $R(\alpha) \sim \frac{4-3\alpha}{2}$, which (as we show later) shows excellent agreement with empirical behavior in simulation. Our characterization of $R(\alpha)$ confirms that if the regression matrix \bar{B} is well-aligned with the block ℓ_1/ℓ_∞ regularizer—more specifically for overlaps $\alpha \in [\frac{2}{3}, 1]$ —then block-regularization increases statistical efficiency. On the other hand, our analysis also conveys a *cautionary message*: if the overlap is too small—more precisely, if $\alpha < 2/3$ —then block $\ell_{1,\infty}$ is actually impaired relative to the naive Lasso-based approach. This fact illustrates that some care is required in the application of block regularization schemes.

In terms of proof techniques, the analysis of this paper is considerably more delicate than the analogous arguments required to show support consistency for the Lasso [20], [36], [39]. The major difference—and one that presents substantial technical

challenges—is that the subdifferential¹ of the block ℓ_1/ℓ_∞ norm is a much more subtle object than the subdifferential of the ordinary ℓ_1 -norm. In particular, the ℓ_1 -norm has an ordinary derivative whenever the coefficient vector is nonzero. In contrast, even for nonzero rows of the regression matrix, the block ℓ_1/ℓ_∞ norm may be nondifferentiable, and these nondifferentiable points play a key role in our analysis. (See Section IV-A for more detail on the subdifferential of this block norm.) As we show, it is the Frobenius norm of the subdifferential on the regression matrix support that controls high-dimensional scaling. For the ordinary ℓ_1 -norm, this Frobenius norm is always equal to s , whereas for matrices with $r = 2$ columns and α fraction overlap, this Frobenius norm can be as small as $\frac{(4-3\alpha)s}{2}$. As our analysis reveals, it is precisely the differing structures of these subdifferentials that change the high-dimensional scaling for ℓ_1 versus $\ell_{1,\infty}$ regularization.

The remainder of this paper is organized as follows. In Section II, we provide a precise description of the problem. Section III is devoted to the statement of our main results, some discussion of their consequences, and illustration by comparison to empirical simulations. In Section IV, we provide an outline of the proof, with the technical details of many intermediate lemmas deferred to the appendices.

A. Notational Conventions

For the convenience of the reader, we summarize here some notation to be used throughout the paper. We reserve the index $i \in \{1, \dots, r\}$ as a superscript in indexing the different regression problems, or equivalently the columns of the matrix $\bar{B} \in \mathbb{R}^{p \times r}$. Given a design matrix $X \in \mathbb{R}^{n \times p}$ and a subset $S \subseteq \{1, \dots, p\}$, we use X_S to denote the $n \times |S|$ submatrix obtained by extracting those columns indexed by S . For a pair of matrices $A \in \mathbb{R}^{m \times \ell}$ and $B \in \mathbb{R}^{m \times n}$, we use the notation $\langle A, B \rangle := A^T B$ for the resulting $\ell \times n$ matrix. For a symmetric matrix $A \in \mathbb{R}^{m \times m}$, we denote $\rho_u(A) = \max_{k \in \{1, \dots, m\}} A_{kk}$.

We use the following standard asymptotic notation: for functions f, g , the notation $f(n) = \mathcal{O}(g(n))$ means that there exists a fixed constant $0 < C < +\infty$ such that $f(n) \leq Cg(n)$; the notation $f(n) = \Omega(g(n))$ means that $f(n) \geq Cg(n)$, and $f(n) = \Theta(g(n))$ means that $f(n) = \mathcal{O}(g(n))$ and $f(n) = \Omega(g(n))$.

II. PROBLEM SETUP

We begin by setting up the problem to be studied in this paper, including multivariate regression and family of block-regularized programs for estimating sparse vectors.

A. Multivariate Regression and Block Regularization Schemes

In this paper, we consider the following form of multivariate regression. For each $i = 1, \dots, r$, let $\bar{\beta}^i \in \mathbb{R}^p$ be a regression vector, and consider the r -variate linear regression problem

$$y^i = X^i \bar{\beta}^i + w^i, \quad i = 1, 2, \dots, r. \quad (3)$$

¹As we describe in more detail in Section IV-A, the subdifferential is the appropriate generalization of gradient to convex functions that are allowed to have “corners,” like the ℓ_1 and ℓ_1/ℓ_∞ norms; the standard books [9], [29] contain more background on subdifferentials and their properties.

Here, each $X^i \in \mathbb{R}^{n \times p}$ is a design matrix, possibly different for each vector $\bar{\beta}^i$, and $w^i \in \mathbb{R}^n$ is a noise vector. We assume that the noise vectors w^i and w^j are independent for different regression problems $i \neq j$. In this paper, we assume that each w^i has a multivariate Gaussian $N(0, \sigma^2 I_{n \times n})$ distribution. However, we note that qualitatively similar results will hold for any noise distribution with sub-Gaussian tails (see the book [6] for more background on sub-Gaussian variates).

For compactness in notation, we frequently use \bar{B} to denote the $p \times r$ matrix with $\bar{\beta}^i \in \mathbb{R}^p$ as the i^{th} column. Given a parameter $q \in [1, \infty]$, we define the ℓ_1/ℓ_q block-norm as follows:

$$\|\bar{B}\|_{\ell_1/\ell_q} := \sum_{k=1}^p \|(\bar{\beta}_k^1, \bar{\beta}_k^2, \dots, \bar{\beta}_k^r)\|_q \quad (4)$$

corresponding to applying the ℓ_q norm to each row of \bar{B} , and the ℓ_1 -norm across all of these blocks. We note that all of these block norms are special cases of the CAP family of penalties [38].

This family of block-regularizers (4) suggests a natural family of M -estimators for estimating \bar{B} , based on solving the block- ℓ_1/ℓ_q -regularized quadratic program

$$\hat{B} \in \arg \min_{B \in \mathbb{R}^{p \times r}} \left\{ \frac{1}{2n} \sum_{i=1}^r \|y^i - X^i \beta^i\|_2^2 + \lambda_n \|B\|_{\ell_1/\ell_q} \right\} \quad (5)$$

where $\lambda_n > 0$ is a user-defined regularization parameter. Note that the data term is separable across the different regression problems $i = 1, \dots, r$, due to our assumption of independence on the noise vectors. Any coupling between the different regression problems is induced by the block-norm regularization.

In the special case of univariate regression ($r = 1$), the parameter q plays no role, and the block-regularized scheme (6) reduces to the Lasso [7], [31]. If $q = 1$ and $r \geq 2$, the block-regularization function (like the data term) is separable across the different regression problems $i = 1, \dots, r$, and so the scheme (6) reduces to solving r separate Lasso problems. For $r \geq 2$ and $q = 2$, the program (6) is frequently referred to as the group Lasso [24], [37]. Another important case [32], [33] and the focus of this paper is the setting $q = \infty$ and $r \geq 2$, which we refer to as block ℓ_1/ℓ_∞ regularization.

The motivation for using block ℓ_1/ℓ_∞ regularization is to encourage *shared sparsity* among the columns of the regression matrix B . Geometrically, like the ℓ_1 norm that underlies the ordinary Lasso, the ℓ_1/ℓ_∞ block norm has a polyhedral unit ball. However, the block norm captures potential interactions between the columns β^i in the matrix B . Intuitively, taking the maximum encourages the elements $(\beta_k^1, \beta_k^2, \dots, \beta_k^r)$ in any given row $k = 1, \dots, p$ to be zero simultaneously, or to be nonzero simultaneously. Indeed, if $\beta_k^i \neq 0$ for at least one $i \in \{1, \dots, r\}$, then there is no additional penalty to have $\beta_k^j \neq 0$ as well, as long as $|\beta_k^j| \leq |\beta_k^i|$.

B. Estimation in ℓ_∞ Norm and Support Recovery

For a given $\lambda_n > 0$, suppose that we solve the block ℓ_1/ℓ_∞ program, thereby obtaining an estimate

$$\hat{B} \in \arg \min_{B \in \mathbb{R}^{p \times r}} \left\{ \frac{1}{2n} \sum_{i=1}^r \|y^i - X^i \beta^i\|_2^2 + \lambda_n \|B\|_{\ell_1/\ell_\infty} \right\}. \quad (6)$$

We note that under high-dimensional scaling ($p \gg n$), this convex program (6) is not necessarily strictly convex, since the quadratic term is rank deficient and the block ℓ_1/ℓ_∞ norm is polyhedral, which implies that the program is not strictly convex. However, a consequence of our analysis is that under appropriate conditions, the optimal solution \hat{B} is in fact unique.

In this paper, we study the accuracy of the estimate \hat{B} , as a function of the sample size n , regression dimensions p and r , and the sparsity index $s = \max_{i=1, \dots, r} |S(\bar{\beta}^i)|$. There are various metrics with which to assess the ‘‘closeness’’ of the estimate \hat{B} to the truth \bar{B} , including predictive risk, various types of norm-based bounds on the difference $\hat{B} - \bar{B}$, and variable selection consistency. In this paper, we prove results bounding the ℓ_∞/ℓ_∞ difference

$$\|\hat{B} - \bar{B}\|_{\ell_\infty/\ell_\infty} := \max_{k=1, \dots, p} \max_{i=1, \dots, r} |\hat{B}_k^i - \bar{B}_k^i|.$$

In addition, we prove results on support recovery criteria. Recall that for each vector $\bar{\beta}^i \in \mathbb{R}^p$, we use $S(\bar{\beta}^i) = \{k \mid \bar{\beta}_k^i \neq 0\}$ to denote its support set. The problem of row support recovery corresponds to recovering the set

$$U := \bigcup_{i=1}^r S(\bar{\beta}^i) \tag{7}$$

corresponding to the subset $U \subseteq \{1, \dots, p\}$ of indices that are active in at least one regression problem. Note that the cardinality of $|U|$ is upper bounded by rs , but can be substantially smaller (as small as s) if there is overlap among the different supports.

As discussed at more length in Appendix A, given an estimate of the row support of \bar{B} , it is possible to either use additional structure of the solution \hat{B} or perform some additional computation to recover individual signed supports of the columns of \bar{B} . To be precise, define the sign function

$$\text{sign}(t) = \begin{cases} +1, & \text{if } t > 0 \\ 0, & \text{if } t = 0 \\ -1, & \text{if } t < 0. \end{cases} \tag{8}$$

Then the recovery of individual signed supports means estimating the signed vectors with entries $\text{sign}(\hat{\beta}_k^i)$, for each $i = 1, 2, \dots, r$ and for all $k = 1, 2, \dots, p$. Interestingly, when using block ℓ_1/ℓ_∞ regularization, there are multiple ways in which the support (or signed support) can be estimated, depending on whether we use primal or dual information from an optimal solution.

The dual recovery method involves the following steps. First, solve the block-regularized program (6), thereby obtaining primal solution $\hat{B} \in \mathbb{R}^{p \times r}$. For each row $k = 1, \dots, p$, compute the set $\mathbb{M}_k := \arg \max_{i=1, \dots, r} |\hat{\beta}_k^i|$. Estimate the support union via $\hat{U} = \bigcup_{i=1, \dots, r} S(\hat{\beta}^i)$, and estimate the signed support vectors

$$[\mathbb{S}_{\text{dual}}(\hat{\beta}_k^i)] = \begin{cases} \text{sign}(\hat{\beta}_k^i), & \text{if } i \in \mathbb{M}_k \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

As our development will clarify, this procedure (9) corresponds to estimating the signed support on the basis of a dual optimal solution associated with the optimal primal solution. We discuss

the primal-based recovery method and its differences with the dual-based method at more length in Appendix A.

III. MAIN RESULTS AND THEIR CONSEQUENCES

In this section, we provide precise statements of the main results of this paper. Our first main result (Theorem 1) provides sufficient conditions for deterministic design matrices X^1, \dots, X^r , whereas our second main result (Theorem 2) provides sufficient conditions for design matrices drawn randomly from Gaussian ensembles. Both of these results allow for an arbitrary number r of regression problems, and the random design case allows for random Gaussian designs X^k with i.i.d. rows and covariance matrix $\Sigma^k \in \mathbb{R}^{p \times p}$, $k = 1, \dots, r$. Not surprisingly, these results show that the high-dimensional scaling of block ℓ_1/ℓ_∞ is *qualitatively similar* to that of ordinary ℓ_1 -regularization: for instance, in the case of random Gaussian designs and bounded r , our sufficient conditions ensure that $n = \Omega(s \log p)$ samples are sufficient to recover the union of supports correctly with high probability, which matches known results on the Lasso [36], as well as known information-theoretic results on the problem of support recovery [35].

As discussed in the introduction, we are also interested in the more refined question: can we provide necessary and sufficient conditions that are sharp enough to reveal *quantitative differences* between ordinary ℓ_1 -regularization and block regularization? Addressing this question requires analysis that is sufficiently precise to control the constants in front of the rescaled sample size $n/s \log(p - s)$ that controls the performance of both ℓ_1 and block ℓ_1/ℓ_∞ methods. Accordingly, in order to provide precise answers to this question, our final two results concern the special case of $r = 2$ regression problems, both with supports of size s that overlap in a fraction α of their entries, and with design matrices drawn randomly from the standard Gaussian ensemble. In this setting, our final result (Theorem 3) shows that block ℓ_1/ℓ_∞ regularization undergoes a phase transition—that is, a rapid change from failure to success—specified by the rescaled sample size $\theta_{1,\infty}(n, p, s, \alpha)$ previously defined (1). We then discuss some consequences of these results, and illustrate their sharpness with some simulation results.

A. Sufficient Conditions for General Deterministic and Random Designs

In addition to the sample size n , problem dimensions p and r , sparsity index s and overlap parameter α , our results involve certain quantities associated with the design matrices X^i . To begin, in the deterministic case, we assume that the columns of each design matrix X^i , $i = 1, \dots, r$ are normalized so that

$$\frac{1}{n} \|X_k^i\|_2^2 \leq 2 \quad \text{for all } k = 1, 2, \dots, p. \tag{10}$$

The choice of the factor 2 in this bound is for later technical convenience. More significantly, we require that the *incoherence parameter*

$$\gamma(\Sigma) := 1 - \max_{\ell=1, \dots, |U^c|} \sum_{i=1}^r \|\langle X_\ell^i, X_U^i (\langle X_U^i, X_U^i \rangle^{-1}) \rangle\|_1 \tag{11}$$

be strictly positive. For the case of the ordinary Lasso, conditions of this type are known [20], [36], [39] to be both necessary and sufficient for successful support recovery.²

In addition, the statement of our results involve certain quantities associated with the $|U| \times |U|$ matrices $\frac{1}{n}\langle X_U^i, X_U^i \rangle$; in particular, we define a lower bound on the minimum eigenvalue

$$C_{\min}(X) \leq \min_{i=1, \dots, r} \lambda_{\min}\left(\frac{1}{n}\langle X_U^i, X_U^i \rangle\right) \quad (12)$$

as well as an upper bound maximum $\ell_{\infty, \infty}$ -operator norm of the inverses

$$D_{\max}(X) \geq \max_{i=1, \dots, r} \left\| \left(\frac{1}{n}\langle X_U^i, X_U^i \rangle\right)^{-1} \right\|_{\infty}. \quad (13)$$

Remembering that our analysis applies to sequences $\{X_{n,p}\}$ of design matrices, in the simplest scenario, both of the bounding quantities C_{\min} and D_{\max} do not scale with (n, p, s) . To keep notation compact, we write C_{\min} and D_{\max} in the analysis to follow.

We also define the support minimum value

$$\bar{B}_{\min} = \min_{k \in U} \max_{i=1, \dots, r} |\bar{\beta}_k^i| \quad (14)$$

corresponding to the minimum value of the ℓ_{∞} norm of any row $k \in U$.

Theorem 1 (Sufficient Conditions for Deterministic Designs): Consider the observation model (3) with design matrices X^i satisfying the column bound (10) and incoherence condition (11). Suppose that we solve the block-regularized ℓ_1/ℓ_{∞} convex program (6) with regularization parameter $\lambda_n^2 \geq \frac{4\xi\sigma^2}{\gamma^2} \frac{r^2 + r \log(p)}{n}$ for some $\xi > 1$. Then with probability greater than

$$\phi_1(\xi, p, s) := 1 - 2 \exp(-(\xi - 1)[r + \log p]) - 2 \exp(-(\xi^2 - 1) \log(rs)) \quad (15)$$

we are guaranteed that

- The block-regularized program has a unique solution \hat{B} such that $\bigcup_{i=1}^r S(\hat{\beta}^i) \subseteq U$.
- Moreover, the solution satisfies the elementwise ℓ_{∞} -bound

$$\|\hat{B} - \bar{B}\|_{\ell_{\infty}/\ell_{\infty}} \leq \underbrace{\xi \sqrt{\frac{4\sigma^2}{C_{\min}} \frac{\log |U|}{n}} + D_{\max} \lambda_n}_{b_1(\xi, \lambda_n, n, s)}. \quad (16)$$

Consequently, as long as $\bar{B}_{\min} \geq b_1(\xi, \lambda_n, n, s)$, then $\bigcup_{i=1}^r S(\hat{\beta}^i) = U$, so that the solution \hat{B} correctly specifies the union of supports U .

We now state an analogous result for random design matrices; in particular, consider the observation model (3) with design matrices X^i chosen with i.i.d. rows from covariance matrices

²Some work [21] has shown that multistage methods can allow some relaxation of this incoherence condition; however, as our main interest is in understanding the sample complexity of ordinary ℓ_1 versus ℓ_1/ℓ_{∞} relaxations, we do not pursue such extensions here.

Σ^i . In analogy to definitions (12) and (13) in the deterministic case, we define the lower bound

$$C_{\min}(\Sigma) \leq \min_{i=1, \dots, r} \lambda_{\min}(\Sigma_{UU}^i), \quad (17)$$

as well as an analogous upper bound on ℓ_{∞} -operator norm of the inverses

$$\max_{i=1, \dots, r} \left\| (\Sigma_{UU}^i)^{-1} \right\|_{\infty} \leq D_{\max}. \quad (18)$$

Note that unlike the case of deterministic designs, these quantities are *not* functions of the design matrix X , which is now a random variable. Finally, our results involve an analogous incoherence parameter of the covariance matrices $\Sigma = \{\Sigma^i, i = 1, \dots, r\}$, defined as

$$\gamma(\Sigma) := 1 - \max_{k=1, \dots, |U^c|} \sum_{i=1}^r \left\| \Sigma_{kU}^i (\Sigma_{UU}^i)^{-1} \right\|_1 > 0. \quad (19)$$

With this notation, the following result provides an analog of Theorem 1 for random design matrices:

Theorem 2 (Sufficient Conditions for Random Gaussian Designs): Suppose that we are given n i.i.d. observations from the model (3) with

$$n > \frac{8\kappa |U| \rho_u(\Sigma_{U^c|U})}{C_{\min} \gamma^2} (r + \log p) \quad (20)$$

for some $\kappa > 1$. If we solve the convex program (6) with regularization parameter satisfying $\lambda_n \geq \sqrt{\frac{4\xi\sigma^2 \rho_u(\Sigma^i)}{\gamma^2} \left[\frac{r^2 + r \log(p)}{n} \right]}$ for some $\xi > 1$, then with probability greater than

$$\phi_2(\kappa, \xi, n, p, s) := 1 - 2 \exp\{-2(\xi^2 - 1) \log(rs)\} - 2 \exp\{-\kappa(r + \log p)\} \rightarrow 1 \quad (21)$$

we are guaranteed that

- The block-regularized program (6) has a unique solution \hat{B} such that $\bigcup_{i=1}^r S(\hat{\beta}^i) \subseteq U$.
- We have that the quantity $\|\hat{B} - \bar{B}\|_{\ell_{\infty}/\ell_{\infty}}$ is upper bounded by

$$\underbrace{\xi \sqrt{\frac{100\sigma^2}{C_{\min}} \frac{\log |U|}{n}} + \lambda_n \left[\frac{4s}{\sqrt{n}} + D_{\max} \right]}_{b_2(\xi, \lambda_n, n, s)}. \quad (22)$$

Consequently, if $B_{\min}^* \geq b_2(\xi, \lambda_n, n, s)$, then $\bigcup_{i=1}^r S(\hat{\beta}^i) = U$, so that the solution \hat{B} correctly specifies the union of supports U .

To clarify the interpretation of Theorems 1 and 2, part (a) of each claim guarantees that the estimator has no false inclusions, in that the row support of the estimate \hat{B} is contained within the row support of the true matrix \bar{B} . One consequence of part (b) is that as long as the minimum signal parameter B_{\min}^* decays slowly enough, then the estimators have no false exclusions, so that the true row support is correctly recovered.

In terms of consistency rates in block $\ell_{\infty}/\ell_{\infty}$ norm, assuming that the design-related quantities C_{\min} , D_{\max} and γ do not scale

with p , Theorem 1(a) guarantees consistency in elementwise ℓ_∞ -norm at the rate

$$\|\widehat{B} - \overline{B}\|_{\ell_\infty/\ell_\infty} = \mathcal{O}\left(\sqrt{\frac{\sigma^2 r^2 + r \log p}{n}}\right).$$

Here we have used the fact that

$$\log |U| \leq \log(rs) = o(r \log p).$$

Similarly, Theorem 2(b) guarantees consistency in elementwise ℓ_∞ -norm at a somewhat slower rate, since there are additional terms that arise in controlling random design matrices. At a high level then, our results thus far show that for a fixed number r of regression problems, the ℓ_1/ℓ_∞ method guarantees exact support recovery with $n = \Omega(s \log p)$ samples, and guarantees consistency in an elementwise sense at rate $\mathcal{O}(\sqrt{\frac{\log p}{n}})$. In qualitative terms, these results match the known scaling [36] for the Lasso (ℓ_1 -regularized QP), which is obtained as the special case for univariate regression ($r = 1$). It should be noted that this scaling is known to be optimal in an information-theoretic sense: no algorithm can recover support correctly if the rescaled sample size $\theta_{\text{Las}} = \frac{n}{2s \log(p-s)}$ is below a critical threshold [35].

As raised by one of the referees, an interesting question is determining whether the method fails with high probability if the sample size is too small. Indeed, via an adaptation of argument used previously for the Lasso [36], it can be shown that the ℓ_1/ℓ_∞ method will fail with high probability if the sample size satisfies

$$n < \frac{c |U|}{2C_{\min}(2 - \gamma)^2 r} (r + \log p)$$

for a numerical constant $c > 0$. For $r = \Theta(1)$, this bound matches the sufficient condition (20) from Theorem 2 up to constant factors. If r is allowed to scale, then the upper and lower bounds differ by a factor of r . The challenge in obtaining matching upper and lower bounds in general lies in the complicated structure of the ℓ_1/ℓ_∞ subdifferential. For the case of $r = 2$, we can provide very precise control on this subdifferential, and thereby show that the ℓ_1/ℓ_∞ method undergoes a sharp transition from failure to success.

B. A Phase Transition for Standard Gaussian Ensembles

In order to provide keener insight into the advantages and/or disadvantages associated with using ℓ_1/ℓ_∞ block regularization, we need to obtain even sharper results, ones that are capable of distinguishing constants in front of the rescaled sample size θ_{Las} . With this aim in mind, the following results are specialized to the case of $r = 2$ regression problems, where the corresponding design matrices $X^i, i = 1, 2$ are sampled from the standard Gaussian ensemble—i.e., with i.i.d. rows $N(0, I_{p \times p})$. By studying this simpler class of problems, we can make *quantitative comparisons* to the sample complexity of the Lasso, which provide insight into the benefits and dangers of block ℓ_1/ℓ_∞ regularization.

The main result of this section asserts that there is a phase transition in the performance of ℓ_1/ℓ_∞ quadratic programming for support recovery—by which we mean a sharp transition from failure to success—and provide the exact location of this

transition point as a function of (n, p, s) and the overlap parameter $\alpha \in (0, 1)$. The phase transition involves the *support gap*

$$\overline{B}_{\text{gap}} = \max_{i \in S(\overline{\beta}^1) \cap S(\overline{\beta}^2)} \left| |\overline{\beta}_i^1| - |\overline{\beta}_i^2| \right|. \quad (23)$$

This quantity measures how close the two regression vectors are in absolute value on their *shared support*. Our main theorem treats the case in which this gap vanishes (i.e., $\overline{B}_{\text{gap}} = o(1)$); note that block ℓ_1/ℓ_∞ regularization is best-suited to this type of structure. A subsequent corollary provides more general but technical conditions for the cases of nonvanishing support gaps. Our main result specifies a phase transition in terms of the rescaled sample size

$$\theta_{1,\infty}(n, p, s, \alpha) := \frac{n}{(4 - 3\alpha)s \log(p - (2 - \alpha)s)}, \quad (24)$$

as stated in Theorem 3.

Theorem 3 (Phase Transition): Consider sequences of problems, indexed by (n, p, s, α) drawn from the observation model (3) with random design X drawn with i.i.d. standard Gaussian entries and with $C_{\min} = 1 = D_{\max}$.

- a) *Success:* Suppose that the problem sequence (n, p, s, α) satisfies

$$\theta_{1,\infty}(n, p, s, \alpha) > 1 + \delta \quad \text{for some } \delta > 0. \quad (25)$$

If we solve the block-regularized program (6) with $\lambda_n \geq \sqrt{\frac{\xi \sigma^2 \log p}{n}}$ for some $\xi > 2$ and $\overline{B}_{\text{gap}} = o(\lambda_n)$, then with probability greater than $1 - c_1 \exp(-c_2 \delta \log(p - (2 - \alpha)s))$, the block $\ell_{1,\infty}$ -program (6) has a unique solution \widehat{B} such that $S(\widehat{B}) \subseteq U$, and moreover it satisfies the elementwise bound (22) with $C_{\min} = 1 = D_{\max}$. In addition, if $B_{\min}^* > b_2(\xi, \lambda_n, n, s)$, then the unique solution recovers the correct signed support.

- b) *Failure:* For problem sequences (n, p, s, α) such that

$$\theta_{1,\infty}(n, p, s, \alpha) < 1 - \delta \quad \text{for some } \delta > 0 \quad (26)$$

and for any nonincreasing regularization sequence $\lambda_n > 0$, no solution $\widehat{B} = (\widehat{\beta}^1, \widehat{\beta}^2)$ to the block-regularized program (6) has the correct signed support.

In a nutshell, Theorem 3 states that block ℓ_1/ℓ_∞ regularization recovers the correct support with high probability for sequences (n, p, s, α) such that $\theta_{1,\infty}(n, p, s, \alpha) > 1$, and otherwise fails with high probability.

We now consider the case in which the support gap does not vanish, and show that it only further degrades the performance of block ℓ_1/ℓ_∞ regularization. To make the degree of this degradation precise, we define the λ_n -truncated gap vector $T_{\lambda_n}(\overline{B}) \in \mathbb{R}^p$, with elements

$$[T_{\lambda_n}(\overline{B})]_i := \begin{cases} \min \{ \lambda_n, \left| |\overline{\beta}_i^1| - |\overline{\beta}_i^2| \right| \}, & \text{if } i \in S(\overline{\beta}^1) \cap S(\overline{\beta}^2) \\ 0, & \text{otherwise.} \end{cases}$$

Recall that the support overlap $S(\overline{\beta}^1) \cap S(\overline{\beta}^2)$ has cardinality αs by assumption. Therefore, $T_{\lambda_n}(\overline{B})$ has at most αs nonzero




Fig. 2. (a) Plots of the probability $\mathbb{P}[\hat{U} = U]$ of successful joint support recovery versus the sample size n . Each curve corresponds to a different problem size p ; notice how the curves shift to the right as p increases, reflecting the difficulty of solving larger problems. (b) Plots of the same data versus the rescaled sample size $n/[2s \log(p - s)]$; note how all three curves now align with one another, showing that this order parameter is the correct scaling for assessing the method.

entries, and moreover $\|T_{\lambda_n}(\bar{B})\|_2^2 \leq \lambda_n^2 \alpha s$. We then define the rescaled gap limit

$$\Delta(\bar{B}, \lambda_n) := \limsup_{(n,p,s)} \frac{\|T_{\lambda_n}(\bar{B})\|_2^2}{\lambda_n^2 s}. \quad (27)$$

Note that $\Delta(\bar{B}, \lambda_n) \in [0, \alpha]$ by construction. With these definitions, we have the following.

Corollary 1 (Poorer Performance With Nonvanishing Gap): If for any $\delta > 0$, the sample size n is upper bounded as

$$n < (1 - \delta) [(4 - 3\alpha) + \Delta(\bar{B}, \lambda_n)] s \log[p - (2 - \alpha)s] \quad (28)$$

then the dual recovery method (9) fails to recover the individual signed supports.

To understand the implications of this result, suppose that all αs of the gaps $|\beta_i^1| - |\beta_i^2|$ were above the regularization level λ_n . Then by definition, we have $\Delta(\bar{B}, \lambda_n) = \alpha$, so that condition (28) implies that the method fails for all $n < (1 - \delta) [4 - 2\alpha] s \log[p - (2 - \alpha)s]$. Since the factor $(4 - 2\alpha)$ is strictly greater than 2 for all $\alpha < 1$, this scaling is *always worse*³ than the Lasso scaling given by $n \asymp 2s \log(p - s)$ (see (2)), unless there is perfect overlap ($\alpha = 1$), in which case it yields no improvements. Consequently, Corollary 1 shows that the performance ℓ_1/ℓ_∞ regularization is also very sensitive to the numerical amplitudes of the signal vectors.

C. Illustrative Simulations and Some Consequences

In this section, we provide some simulation results to illustrate the phase transition predicted by Theorem 3. Interestingly, these results show that the theory provides an accurate description of practice even for relatively small problem sizes (e.g., $p = 128$). As specified in Theorem 3, we simulate multivariate regression problems with $r = 2$ columns, with the design matrices X^i drawn from the standard Gaussian ensemble. In all

³Here we are assuming that $s/p = o(1)$, so that $\log(p - s) \asymp \log[p - (2 - \alpha)s]$.

cases, we initially solved the ℓ_1/ℓ_∞ program using an implementation of the *quadprog* function in MATLAB, and then verified that the behavior of the solution agreed with the primal-dual optimality conditions specified by our theory. In subsequent simulations, we solved directly for the dual variables, and then checked whether or not the dual feasibility conditions are met.

We first illustrate the difference between unscaled and rescaled plots of the empirical performance, which demonstrate that the rescaled sample size $n/[s \log(p - s)]$ specifies the high-dimensional scaling of block ℓ_1/ℓ_∞ regularization. Fig. 2(a) shows the empirical behavior of the block ℓ_1/ℓ_∞ method for joint support recovery. For these simulations, we applied the method to $r = 2$ regression problems with overlap $\alpha = 1$, and to three different problem sizes $p \in \{128, 256, 512\}$, in all cases with the sparsity index $s = [0.1p]$. Each curve in panel (a) shows the probability of correct support recovery $\mathbb{P}[\hat{U} = U]$ versus the raw sample size n . As would be expected, all the curves initially start at $\mathbb{P}[\hat{U} = U] = 0$, but then transition to 1 as n increases, with the transition taking place at larger and larger sample sizes as p is increased. The purpose of the rescaling is to determine exactly how this transition point depends on the problem size p and other structural parameters (s and α). Fig. 2(b) shows the same simulation results, now plotted versus the rescaled sample size $\theta := n/[2s \log(p - s)]$, which is the appropriate rescaling predicted by our theory. Notice how all three curves now lie on top of another, and moreover transition from failure to success at $\theta \approx 1$, consistent with our theoretical predictions.

We now seek to explore the dependence of the sample size on the overlap fraction $\alpha \in [0, 1]$ of the two regression vectors. For this purpose, we plot the probability of successful recovery versus the rescaled sample size

$$\theta_{1,\infty}(n, p, s, \alpha) = \frac{n}{(4 - 3\alpha)s \log(p - (2 - \alpha)s)}.$$

As shown by Fig. 2(b), when plotted with this rescaling, there is no shift with size p . Moreover, if we choose the sparsity index s

to grow in a fixed way with p (i.e., $s = f(p)$ for some fixed function f), then the only remaining free variable is the overlap parameter α . Note that the theory predicts that the required sample size should decrease as α increases towards 1.

As shown earlier in Section I, Fig. 1 plots the probability of successful recovery of the joint supports versus the rescaled samples size $\theta_{1,\infty}(n, p, s, \alpha)$. Notice that the plot shows four sets of “stacked” curves, where each stack corresponds to a different choice of the overlap parameter, ranging from $\alpha = 1$ (left-most stack) to $\alpha = 0.1$ (right-most stack). Each stack contains three curves, corresponding to the problem sizes $p \in \{128, 256, 512\}$. In all cases, we fixed the support size $s = \lceil 0.1p \rceil$. As with Fig. 2(b), the “stacking” behavior of these curves demonstrates that Theorem 3 isolates the correct dependence on p . Moreover, their step-like behavior is consistent with the theoretical prediction of a phase transition. Notice how the curves shift towards the left as the overlap parameter α increases towards one, reflecting that the problems become easier as the amount of shared sparsity increases. To assess this shift in a qualitative manner, for each choice of overlap $\alpha \in \{0.1, 0.3, 0.7, 1\}$, we plot a vertical line within each group, which is obtained as the threshold value of $\theta_{1,\infty}$ predicted by our theory. Observe how the theoretical value shows excellent agreement with the empirical behavior.

As noted previously in Section I, Theorem 3 has some interesting consequences, particularly in comparison to the behavior of the “naive” Lasso-based individual decoding of signed supports—that is, the method that simply applies the Lasso (ordinary ℓ_1 -regularization) to each column $i = 1, 2$ separately. By known results [36] on the Lasso, the performance of this naive approach is governed by the order parameter $\theta_{\text{Las}}(n, p, s) = \frac{n}{2s \log(p-s)}$, meaning that for any $\delta > 0$, it succeeds for sequences such that $\theta_{\text{Las}} > 1 + \delta$, and conversely fails for sequences such that $\theta_{\text{Las}} < 1 - \delta$. To compare the two methods, we define the relative efficiency coefficient $R(\theta_{1,\infty}, \theta_{\text{Las}}) = \theta_{\text{Las}}(n, p, s) / \theta_{1,\infty}(n, p, s, \alpha)$. A value of $R < 1$ implies that the block method is more efficient, while $R > 1$ implies that the naive method is more efficient. With this notation, we have the following.

Corollary 2: The relative efficiency of the block $\ell_{1,\infty}$ program (6) compared to the Lasso is given by $R(\theta_{1,\infty}, \theta_{\text{Las}}) = \frac{4-3\alpha \log(p-(2-\alpha)s)}{2 \log(p-s)}$. Thus, for sublinear sparsity $s/p \rightarrow 0$, the block scheme has greater statistical efficiency for all overlaps $\alpha \in (2/3, 1]$, but lower statistical efficiency for overlaps $\alpha \in [0, 2/3)$.

Fig. 3 provides an alternative perspective on the data, where we have plotted how the sample size required by block regression changes as a function of the overlap parameter $\alpha \in [0, 1]$. Each set of data points plots a scaled form of the sample size required to hit 50% success, for a range of overlaps, and the straight line $(4 - 3\alpha)/2$ that is predicted by Theorem 3. Note the excellent agreement between the experimental results, for all three problem sizes for $p \in \{128, 256, 512\}$, and the full range of overlaps. The line $(4 - 3\alpha)/2$ also characterizes the relative efficiency R of block regularization versus the naive Lasso-based method, as described in Corollary 2. For overlaps $\alpha > 2/3$, this parameter R drops below 1. On the other hand, for




Fig. 3. Plots of the relative statistical efficiency $R(\alpha)$ of a method based on block- ℓ_1/ℓ_∞ regularization versus the Lasso (ordinary ℓ_1 -regularization). For each value of the parameter $\alpha \in [0, 1]$ that measures overlap between the regression problems, the quantity $R(\alpha)$ is the ratio of sample size required by an estimator based on block ℓ_1/ℓ_∞ -regularization relative to the sample size required by the Lasso (ordinary ℓ_1 -regularization). The error criterion here is recovery of the correct subset of active variables in the regression. Over a range of overlaps, the empirical thresholds of the ℓ_1/ℓ_∞ block regularization method closely align with the theoretical prediction of $(4 - 3\alpha)/2$. The block-based method begins to give benefits versus the “naive” Lasso-based method at the critical overlap $\alpha^* \approx 2/3$, at which point the relative efficiency $R(\alpha)$ first drops below 1. For overlaps $\alpha \in [0, 2/3)$, the joint method actually requires more samples than the naive method.

overlaps $\alpha < 1$, we have $R > 1$, so that applying the joint optimization problem actually decreases statistical efficiency. Intuitively, although there is still some fraction of overlap, the regularization is misleading, in that it tries to enforce a higher degree of shared sparsity than is actually present in the data.

IV. PROOFS

This section contains the proofs of our three theorems. Our proofs are constructive in nature, based on a procedure that constructs pair of matrices $\tilde{B} = (\tilde{\beta}^1, \dots, \tilde{\beta}^r) \in \mathbb{R}^{p \times r}$ and $\tilde{Z} = (\tilde{z}^1, \dots, \tilde{z}^r) \in \mathbb{R}^{p \times r}$. The goal of the construction is to show that matrix \tilde{B} is an optimal primal solution to the convex program (6), and that the matrix \tilde{Z} is a corresponding dual-optimal solution, meaning that it belongs to the subdifferential of the $\ell_{1,\infty}$ -norm (see Lemma 1), evaluated at \tilde{B} . If the construction succeeds, then the pair (\tilde{B}, \tilde{Z}) acts as a witness for the success of the convex program (6) in recovering the correct signed support—in particular, success of the primal-dual witness procedure implies that \tilde{B} is the unique optimal solution of the convex program (6), with its row support contained within U . To be clear, the procedure for constructing this candidate primal-dual solution is *not* a practical algorithm (as it exploits knowledge of the true support sets), but rather a proof technique for certifying the correctness of the block-regularized program.

We begin by providing some background on the subdifferential of the ℓ_1/ℓ_∞ norm; we refer the reader to the books [9], [29] for more background on convex analysis.

A. Structure of ℓ_1/ℓ_∞ -Norm Subdifferential

The subdifferential of a convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ at a point $x \in \mathbb{R}^d$ is the set of all vectors $y \in \mathbb{R}^d$ such that $f(x') \geq$

Primal-dual witness construction:

(A) First, we solve the restricted program

$$\tilde{B} = \arg \min_{B \in \mathbb{R}^{p \times r}, B_{U^c} = 0} \left\{ \frac{1}{2n} \sum_{i=1}^r \|y^i - X^i \beta^i\|_2^2 + \lambda_n \|B\|_{\ell_1/\ell_\infty} \right\}. \quad (29)$$

Given our assumption that the $|U| \times |U|$ sub-matrices $\langle X_U^i, X_U^i \rangle$ are invertible, the solution to this convex program is unique. Moreover, note that $\tilde{B}_{U^c} = 0$ by construction.

(B) We choose $\tilde{Z}_U \in \mathbb{R}^{|U| \times r}$ as an element of the subdifferential $\partial \|\tilde{B}_U\|_{\ell_1/\ell_\infty}$.(C) Using the optimality conditions associated with the original convex program (6), we then solve for the matrix \tilde{Z}_{U^c} , and verify that its rows satisfy the *strict dual feasibility* condition

$$\sum_{i=1}^r |\tilde{z}_k^i| < 1 \quad \text{for all } k \in U^c. \quad (30)$$

(D) A final (optional) step is to verify that \tilde{B}_U satisfies the *sign consistency* conditions $\text{sign}(\tilde{B}_U) = \text{sign}(\tilde{B}_U)$.

Fig. 4. Steps in the primal-dual witness construction. Steps (A) and (B) are straightforward; the main difficulties lie in verifying the strict dual feasibility and sign consistency conditions stated in step (C) and (D).

$f(x) + \langle y, x' - x \rangle$ for all $x' \in \mathbb{R}^d$. See the standard books [9], [29] for background on subdifferentials and their properties.

We state for future reference a characterization of the subdifferential of the ℓ_1/ℓ_∞ block norm.

Lemma 1: The matrix $\tilde{Z} \in \mathbb{R}^{p \times r}$ belongs to the subdifferential $\partial \|\tilde{B}\|_{\ell_1/\ell_\infty}$ if and only if the following conditions hold for each $k = 1, \dots, p$.

i) If $\tilde{\beta}_k^i \neq 0$ for at least one index $i \in \{1, \dots, r\}$, then

$$\tilde{z}_k^i = \begin{cases} t_i \text{sign}(\tilde{\beta}_k^i), & \text{if } i \in \mathbb{M}_k \\ 0, & \text{otherwise.} \end{cases}$$

where $\mathbb{M}_k := \arg \max_{i=1, \dots, r} |\tilde{\beta}_k^i|$, for a set of non-negative scalars $\{t_i, i \in \mathbb{M}_k\}$ such that $\sum_{i \in \mathbb{M}_k} t_i = 1$.

ii) If $\tilde{\beta}_k^i = 0$ for all $i = 1, \dots, r$, then we require $\sum_{i=1}^r |\tilde{z}_k^i| \leq 1$.

B. Primal-Dual Construction

We now describe our method for constructing the matrix pair (\tilde{B}, \tilde{Z}) . Recalling that $U = \bigcup_{i=1}^r S(\tilde{\beta}^i)$ denotes the union of supports of the true regression vectors, let U^c denote the set $\{1, \dots, p\} \setminus U$. With this notation, Fig. 4 provides the four steps of the primal-dual witness construction.

The following lemma summarizes the utility of the primal-dual witness method.

Lemma 2: Suppose that for each $i = 1, \dots, r$, the $|U| \times |U|$ submatrix $\langle X_U^i, X_U^i \rangle$ is invertible. Then for any $\lambda_n > 0$, we have the following correspondences:

i) If steps (A) through (C) of the primal-dual construction succeed, then $(\tilde{B}_U, 0) \in \mathbb{R}^{p \times r}$ is the unique optimal solution of the original convex program (6).

ii) Conversely, suppose that there is a solution $\hat{B} \in \mathbb{R}^{p \times r}$ to the convex program (6) with support contained within U . Then steps (A) through (C) of the primal-dual witness construction succeed.

We provide the proof of Lemma 2 in Appendix D.2. It is convex-analytic in nature, based on exploiting the subgradient optimality conditions associated with both the restricted convex program (29) and the original program (6), and performing some algebra to characterize when the convex program recovers the

correct signed support. Lemma 2 lies at the heart of all three of our theorems. In particular, the positive results of Theorem 1, Theorem 2, and Theorem 3(a) are based on claim (i), which shows that it is sufficient to verify that the primal-dual witness construction succeeds with high probability. The negative result of Theorem 3(b), in contrast, is based on part (ii), which can be restated as asserting that if the primal-dual witness construction fails, then no solution has support contained within U .

Before proceeding to the proofs themselves, we introduce some additional notation and develop some auxiliary results concerning the primal-dual witness procedure, to be used in subsequent development. With reference to steps (A) and (B), we show in Appendix D.2 that the unique solution \tilde{B}_U has the form

$$\tilde{B}_U = \bar{B}_U + \Delta_U \quad (31)$$

where the matrix $\Delta_U \in \mathbb{R}^{|U| \times r}$ has columns

$$\Delta^i := \left(\frac{1}{n} \langle X_U^i, X_U^i \rangle \right)^{-1} \left[\frac{1}{n} \langle X_U^i, w^i \rangle - \lambda_n \tilde{z}_U^i \right] \quad (32)$$

for $i = 1, \dots, r$, and \tilde{z}_U^i is the i^{th} column of the subgradient matrix \tilde{Z}_U .

With reference to step (C), we obtain the candidate dual solution $\tilde{Z}_{U^c} \in \mathbb{R}^{|U^c| \times r}$ as follows. For each $i = 1, \dots, r$, let $\Pi_{X_U^i}$ denote the orthogonal projection onto the range of X_U^i . Using the submatrix $\tilde{Z}_U \in \mathbb{R}^{|U| \times r}$ obtained from step (B), we define column i of the matrix \tilde{Z}_{U^c} as follows:

$$\tilde{z}_{U^c}^i := \frac{1}{\lambda_n n} \langle X_{U^c}^i, (I - \Pi_{X_U^i}) w^i \rangle + \frac{1}{n} \langle X_{U^c}^i, X_U^i \left(\frac{1}{n} \langle X_U^i, X_U^i \rangle \right)^{-1} \tilde{z}_U^i \rangle \quad (33)$$

for $i = 1, \dots, r$. See the end of Appendix D.2 for derivation of this condition.

Finally, in order to further simplify notation in our proofs, for each $k \in U^c$, we define the random variable

$$V_k := \sum_{i=1}^r |\tilde{z}_k^i|. \quad (34)$$

With this notation, the strict dual feasibility condition (30) is equivalent to the event $\{\max_{k \in U^c} V_k < 1\}$.

V. PROOF OF THEOREM 1

We begin by establishing a set of sufficient conditions for deterministic design matrices, as stated in Theorem 1.

A. Establishing Strict Dual Feasibility

We begin by obtaining control on the probability of the event $\mathcal{E}(V)$, so as to show that step (C) of the primal-dual witness construction succeeds. Recall that $\Pi_{X_U^i}$ denotes the orthogonal projection onto the range space of X_U^i , and the definition (11) of the incoherence parameter $\gamma \in (0, 1]$. By the mutual incoherence condition (11), we have

$$\max_{k \in U^c} \left\{ \sum_{i=1}^r \left| \frac{1}{n} \langle X_k^i, X_U^i \rangle \left(\left\langle \frac{1}{n} X_U^i, X_U^i \right\rangle \right)^{-1} \tilde{z}_U^i \right| \right\} \leq 1 - \gamma \quad (35)$$

where we have used the fact that $\sum_{i=1}^r |\tilde{z}_j^i| = 1$ for each $j \in U$. Recalling that $V_k = \sum_{i=1}^r |\tilde{z}_k^i|$ and using the definition (33), we have by triangle inequality

$$\mathbb{P}[\max_{k \in U^c} V_k > 1] \leq \mathbb{P}[\mathbf{A}(\gamma)]$$

where we have defined the event

$$\mathbf{A}(\gamma) := \left\{ \max_{k \in U^c} \sum_{i=1}^r \left| \frac{1}{\lambda_n n} \langle X_k^i, (I - \Pi_{X_U^i}) w^i \rangle \right| \geq \gamma \right\}.$$

To analyze this remaining probability, for each index $i = 1, \dots, r$ and $k \in U^c$, define the random variable

$$W_k^i := \frac{1}{\lambda_n n} \langle X_k^i, (I - \Pi_{X_U^i}) w^i \rangle. \quad (36)$$

Since the elements of the n -vector w^i follow a $N(0, \sigma^2)$ distribution, the variable W_k^i is zero-mean Gaussian with variance $\frac{\sigma^2}{\lambda_n^2 n^2} \langle X_k^i, (I - \Pi_{X_U^i}) X_k^i \rangle$. Since $\|X_k^i\|_2^2 \leq 2n$ by assumption and $(I - \Pi_{X_U^i})$ is an orthogonal projection matrix, the variance of each W_k^i is upper bounded by $\frac{2\sigma^2}{\lambda_n^2 n}$.

Consequently, for any choice of sign vector $b \in \{-1, +1\}^r$, the variance of the zero-mean Gaussian $\sum_{i=1}^r b_i W_k^i$ is upper bounded by $\frac{2r\sigma^2}{\lambda_n^2 n}$. Consequently, by taking the union bound over all sign vectors and over indices $k \in U^c$, we have

$$\begin{aligned} \mathbb{P}[\mathbf{A}(\gamma)] &= \mathbb{P} \left[\max_{k \in U^c} \max_{b \in \{-1, +1\}^r} \sum_{i=1}^r b_i W_k^i > \gamma \right] \\ &\leq 2 \exp \left(- \frac{\lambda_n^2 n \gamma^2}{4r\sigma^2} + r + \log p \right). \end{aligned}$$

With the choice $\lambda_n^2 \geq \frac{4\xi\sigma^2}{\gamma^2} \frac{r^2 + r \log(p)}{n}$ for some $\xi > 1$, we conclude that

$$\mathbb{P}[\mathcal{E}(V)] \geq 1 - 2 \exp(-(\xi - 1)[r + \log p]) \rightarrow 1.$$

By Lemma 2(i), this event implies the uniqueness of the solution \hat{B} , and moreover the inclusion of the supports $S(\hat{B}) \subseteq S(\bar{B})$, as claimed.

 B. Establishing ℓ_∞ Bounds

We now turn to establishing the claimed ℓ_∞ -bound (16) on the difference $\hat{B} - \bar{B}$. We have already shown that this difference is exactly zero for rows in U^c ; it remains to analyze the difference $\Delta_U = \hat{B}_U - \bar{B}_U$. It suffices to prove the ℓ_∞ bound for the columns Δ_U^i separately, for each $i = 1, \dots, r$.

We split the analysis of the random variable $\max_{k \in U} |\Delta_k^i|$ into two terms, based on the form of Δ from (32), one involving the dual variables \tilde{z}_U^i , and the other involving the observation noise w^i , as follows:

$$\begin{aligned} \max_{k \in U} |\Delta_k^i| &\leq \underbrace{\left\| \left(\frac{1}{n} \langle X_U^i, X_U^i \rangle \right)^{-1} \frac{1}{n} \langle X_U^i, w^i \rangle \right\|_\infty}_{T_a^i} \\ &\quad + \underbrace{\left\| \left(\left\langle \frac{1}{n} X_U^i, X_U^i \right\rangle \right)^{-1} \lambda_n \tilde{z}_U^i \right\|_\infty}_{T_b^i}. \end{aligned}$$

The second term is easy to control: from the characterization of the subdifferential (Lemma 1), we have $\|\tilde{z}_U^i\|_\infty \leq 1$, so that $T_b^i \leq \lambda_n \left\| \left(\left\langle \frac{1}{n} X_U^i, X_U^i \right\rangle \right)^{-1} \right\|_\infty \leq D_{\max} \lambda_n$.

Turning to the first term T_a^i , we note that since X_U^i is fixed, the $|U|$ -dimensional random vector $Y := \left(\left\langle \frac{1}{n} X_U^i, X_U^i \right\rangle \right)^{-1} \frac{1}{n} \langle X_U^i, w^i \rangle$ is zero-mean Gaussian, with covariance $\frac{1}{n} \left(\left\langle \frac{1}{n} X_U^i, X_U^i \right\rangle \right)^{-1}$. Therefore, the variance of each entry satisfies $\text{var}(Y_k) \leq \frac{1}{C_{\min} n}$, and we can exploit this fact in standard Gaussian tail bounds. By applying the union bound twice, first over $k \in U$, and then over $i \in \{1, 2, \dots, r\}$, we obtain that

$$\begin{aligned} \mathbb{P} \left[\max_{i=1, \dots, r} T_a^i \geq t \right] \\ \leq 2 \exp(-t^2 n C_{\min} / (2) + \log(rs) + \log r), \end{aligned}$$

where we have used the fact that $|U| \leq rs$. Setting $t = \xi \sqrt{\frac{4 \log(rs)}{C_{\min} n}}$ yields that

$$\max_{i=1, \dots, r} \max_{k \in U} |\Delta_k^i| \leq \underbrace{\xi \sqrt{\frac{4}{C_{\min}} \frac{\log rs}{n}}}_{b_1(\xi, \lambda_n, n, s)} + D_{\max} \lambda_n$$

with probability greater than $1 - 2 \exp(-(\xi^2 - 1) \log(rs))$, as claimed.

Finally, to establish support recovery, recall that we proved above that Δ^i is bounded by $b_1(\xi, \lambda_n, n, s)$. Hence, as long as $B_{\min}^* > b_1(\xi, \lambda_n, n, s)$, then we are guaranteed that if $\bar{B}_k^i \neq 0$, then $\hat{B}_k^i \neq 0$.

VI. PROOF OF THEOREM 2

We now turn to the proof of Theorem 2, providing sufficient conditions for general Gaussian ensembles. Recall that for $i = 1, 2, \dots, r$, each $X^i \in \mathbb{R}^{n \times p}$ is a random design matrix, with rows drawn i.i.d. from a zero-mean Gaussian with $p \times p$ covariance matrix Σ^i .

A. Establishing Strict Dual Feasibility

Recalling that $V_k = \sum_{i=1}^r |\tilde{z}_k^i|$ and using the definition (33), we have the decomposition $\max_{k \in U^c} |V_k| \leq M_1 + M_2$, where

$$M_1 := \max_{k \in U^c} \sum_{i=1}^r \left| \frac{1}{\lambda_n n} \langle X_k^i, (I - \Pi_{X_U^i}) w^i \rangle \right|, \quad \text{and}$$

$$M_2 := \max_{k \in U^c} \sum_{i=1}^r \left| \frac{1}{n} \langle X_k^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle \right|.$$

In order to show that $\max_{k \in U^c} |V_k| < 1$ with high probability, we deal with each of these two terms in turn, showing that $M_1 < \gamma/2$, and $M_2 < 1 - \gamma/2$, both with high probability.

In order to bound M_1 , we require the following condition on the columns of the design matrices.

Lemma 3: Let $\sigma_{\max} = \max_i \rho_u(\Sigma^i)$. For $n > 2 \log(rp)$, each column of the design matrices $X^i, i = 1, \dots, r$ has controlled ℓ_2 -norm:

$$\mathbb{P} \left[\max_{i=1, \dots, r} \max_{k=1, \dots, p} \|X_k^i\|_2^2 \leq 2\sigma_{\max} n \right] \leq 2 \exp \left(-\frac{n}{2} + \log(rp) \right).$$

This claim follows immediately by union bound and concentration results for χ^2 -variables; in particular, the bound (69a) in Appendix E.

Under the condition of Lemma 3, each variable $W_k^i := \frac{1}{\lambda_n n} \langle X_k^i, (I - \Pi_{X_U^i}) w^i \rangle$ is zero-mean Gaussian, with variance at most $\frac{2\sigma^2}{\lambda_n^2 n}$. Consequently, for any choice of signs $b \in \{-1, +1\}^r$, the vector $\sum_{i=1}^r b_i W_k^i$ is zero-mean Gaussian, with variance at most $\frac{2\sigma^2 r}{\lambda_n^2 n}$. Therefore, for any $t > 0$, we have

$$\mathbb{P} \left[\max_{k \in U^c} \sum_{i=1}^r |W_k^i| \geq t \right] = \mathbb{P} \left[\max_{k \in U^c} \max_{b \in \{-1, +1\}^r} \sum_{i=1}^r b_i W_k^i \geq t \right] \leq 2 \exp \left(-\frac{\lambda_n^2 n}{4\sigma^2 r} t^2 + r + \log p \right).$$

Setting $t = \gamma/2$ yields that

$$\mathbb{P}[M_1 \geq \gamma/2] \leq 2 \exp \left(-\frac{\lambda_n^2 n}{16\sigma^2 r} \gamma^2 + r + \log p \right).$$

Lemma 4: Suppose that the design covariance matrices $\Sigma^i, i = 1, \dots, r$ satisfy the mutual incoherence condition (11). Then we have $M_2 \leq (1 - \gamma) + M'_2$, where

$$M'_2 := \max_{k \in U^c} \sum_{i=1}^r \left| \frac{1}{n} \langle Y_k^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle \right| \quad (37)$$

and where each random matrix $Y_{U^c}^i \in \mathbb{R}^{n \times (p - |U|)}$ has $N(0, \Sigma_{U^c|U}^i)$ rows independent of w^i and X_U^i .

See Appendix B for the proof of this claim.

It remains to show that the random variable M'_2 defined in (37) is upper bounded by $\gamma/2$ with high probability. Conditioning on X_U^i and w^i , the scalar random variable $\frac{1}{n} \langle Y_k^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle$ is zero-mean Gaussian, with variance upper bounded as

$$\frac{1}{n} \langle \tilde{z}_U^i, (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle \leq \rho_{\max} \frac{\|\tilde{z}_U^i\|_2^2}{C_{\min} n}$$

where $\rho_{\max} = \max_i \rho_u(\Sigma_{U^c|U}^i)$. Recalling that $\sum_{i=1}^r |\tilde{z}_j^i| = 1$, for any choice of signs $b \in \{-1, +1\}^r$, the variable

$$\sum_{i=1}^r b_i \frac{1}{n} \langle Y_k^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle$$

is zero-mean Gaussian, with variance at most $\frac{|U| \rho_{\max}}{C_{\min} n}$. Therefore, the quantity $\mathbb{P}[M'_2 \geq \gamma/2]$ is upper bounded by

$$\mathbb{P} \left[\max_{k \in U^c} \max_{b \in \{-1, +1\}^r} \left| \sum_{i=1}^r b_i \frac{1}{n} \langle Y_k^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle \right| \geq \gamma/2 \right]$$

which is at most

$$2 \exp \left(-\frac{C_{\min} n}{8|U| \rho_{\max}} \gamma^2 + r + \log p \right).$$

This probability vanishes faster than $2 \exp \{ -\kappa(r + \log p) \}$ as long as

$$n > \frac{8\kappa |U| \rho_{\max}}{C_{\min} \gamma^2} (r + \log p)$$

as claimed.

B. Establishing ℓ_∞ Bounds

We now turn to establishing the claimed ℓ_∞ -bound (16) on the difference $\hat{B} - \bar{B}$. As in the analogous portion of the proof of Theorem 1, we use the decomposition $\max_{k \in U} |\Delta_k^i| = T_a^i + T_b^i$, where

$$T_a^i := \left\| \left(\langle \frac{1}{n} X_U^i, X_U^i \rangle \right)^{-1} \frac{1}{n} \langle X_U^i, w^i \rangle \right\|_\infty, \quad \text{and}$$

$$T_b^i := \left\| \left(\langle \frac{1}{n} X_U^i, X_U^i \rangle \right)^{-1} \lambda_n \tilde{z}_U^i \right\|_\infty.$$

In the setting of random design matrices, a bit more work is required to control these terms.

Beginning with the second term T_b^i , by triangle inequality, it is upper bounded by

$$\left\| \left[\left(\langle \frac{1}{n} X_U^i, X_U^i \rangle \right)^{-1} - (\Sigma_{UU}^i)^{-1} \right] \lambda_n \tilde{z}_U^i \right\|_\infty + \left\| (\Sigma_{UU}^i)^{-1} \lambda_n \tilde{z}_U^i \right\|_\infty$$

which is in turn at most

$$\left\| \left[\left(\langle \frac{1}{n} X_U^i, X_U^i \rangle \right)^{-1} - (\Sigma_{UU}^i)^{-1} \right] \right\|_2 \lambda_n \sqrt{s} + D_{\max} \lambda_n$$

where we have used the facts that $\|\tilde{z}_U^i\|_2 \leq \sqrt{s}$, since \tilde{z}_U^i belongs to the subdifferential of the block ℓ_1/ℓ_∞ norm (see Lemma 1) so that $\sum_{i=1}^r |\tilde{z}_j^i| \leq 1$ for all $j \in U$. By, concentration bounds for eigenvalues of Gaussian random matrices (see (72b) in Appendix E), we conclude that

$$T_b^i \leq 4\lambda_n \sqrt{s} \sqrt{\frac{s}{n}} + D_{\max} \lambda_n = \lambda_n \left[\frac{4s}{\sqrt{n}} + D_{\max} \right].$$

Now consider the first term T_a^i : if we condition on X_U^i , then the $|U|$ -dimensional random vector $Y := \left(\langle \frac{1}{n} X_U^i, X_U^i \rangle \right)^{-1} \frac{1}{n} \langle X_U^i, w^i \rangle$ is zero-mean Gaussian,

with covariance $\frac{1}{n}(\langle \frac{1}{n}X_U^i, X_U^i \rangle)^{-1}$. By concentration bounds for eigenvalues of Gaussian random matrices (see (72b) in Appendix E), the quantity $\frac{1}{n}\|(\langle \frac{1}{n}X_U^i, X_U^i \rangle)^{-1}\|_2$ is upper bounded by

$$\frac{1}{n} \left\{ \|(\langle \frac{1}{n}X_U^i, X_U^i \rangle)^{-1} - (\Sigma_{UU}^i)^{-1}\|_2 + \|(\Sigma_{UU}^i)^{-1}\|_2 \right\}$$

which in turn, is at most

$$\frac{4}{C_{\min}n} \sqrt{rs} + \frac{1}{C_{\min}n} \leq \frac{5}{C_{\min}n}$$

since $rs/n \leq 1$. Therefore, we have shown that the variance of each element of Y is upper bounded by $5/(C_{\min}n)$, so that we can apply standard Gaussian tail bounds. By applying the union bound twice, first over $k \in U$, and then over $i \in \{1, 2, \dots, r\}$, we obtain

$$\mathbb{P}[\max_{i=1, \dots, r} T_a^i \geq t] \leq 2 \exp(-t^2 n C_{\min} / (50) + \log |U| + \log r).$$

Setting $t = \xi \sqrt{\frac{100 \log(rs)}{C_{\min}n}}$ yields that $\mathbb{P}[\max_{i=1, \dots, r} T_a^i \geq t]$ is at most

$$2 \exp \left\{ -2\xi^2 \log(rs) + \log(rs) + \log r \right\} \leq 2 \exp \left\{ -2(\xi^2 - 1) \log(rs) \right\}$$

where we have used the fact that $|U| \leq rs$. Combining the pieces, we conclude that

$$\max_{k \in U} |\Delta_k^i| \leq \xi \sqrt{\frac{100 \log(rs)}{C_{\min}n}} + \lambda_n \left[\frac{4s}{\sqrt{n}} + D_{\max} \right]$$

with probability greater than

$$1 - 2 \exp \left\{ -2(\xi^2 - 1) \log(rs) \right\} - c_1 \exp(-c_2 n)$$

as claimed.

VII. PROOF OF THEOREM 3

We now turn to the proof of the phase transition predicted by Theorem 3, which applies to random design matrices X^1 and X^2 drawn from the standard Gaussian ensemble. This proof requires significantly more technical work than the preceding two proofs, since we need to control all the constants exactly, and to establish both necessary and sufficient conditions on the sample size.

A. Proof of Theorem 3(a)

We begin with the achievability result. Our proof parallels that of Theorems 1 and 2, in that we first establish strict dual feasibility, and then turn to proving ℓ_∞ bounds and exact support recovery.

1) *Establishing Strict Dual Feasibility:* Recalling that $V_k = \sum_{i=1}^2 |\tilde{z}_k^i|$, we have $\max_{k \in U^c} |V_k| \leq M_1 + M_2$, where the random variables M_1 and M_2 were defined at the start of Section VI-A. In order to prove that $\max_{k \in U^c} |V_k| < 1$ with high probability for the values of n, s , and p , we will first establish that $M_1 < \epsilon/2$ and $M_2 < 1 - \epsilon$ for an appropriately chosen value of ϵ .

By the results from the previous section, we have $M_1 < \epsilon/2$ with probability

$$\mathbb{P} \left[\max_{k \in U^c} \sum_{i=1}^2 |W_k^i| \geq \epsilon/2 \right] \leq 2 \exp \left(-\frac{\lambda_n^2 n \epsilon^2}{32\sigma^2} + 2 + \log p \right).$$

Recall that

$$M_2 = \max_{k \in U^c} \sum_{i=1}^2 \left| \frac{1}{n} \langle X_k^i, X_U^i (\frac{1}{n} \langle X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle \right|$$

and that $X_{U^c}^i$ is independent of X_U^i and w^i . We will show that $M_2 < 1 - \epsilon$ with high probability by using results on Gaussian extrema. Conditioning on (X_U, w, \tilde{z}_U^1) , the random variable $Y_k^i = \frac{1}{n} \langle X_k^i, X_U^i (\frac{1}{n} \langle X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle$ is zero-mean with variance at most

$$\frac{1}{n} \langle \tilde{z}_U^i, (\frac{1}{n} \langle X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle \leq \|(\frac{1}{n} \langle X_U^i, X_U^i \rangle)^{-1}\|_2 \frac{\|\tilde{z}_U^i\|_2^2}{n}.$$

Under the given conditioning, the random variables Y_k^1 and Y_k^2 are independent and for any sign vector $b \in \{-1, +1\}^2$, the random variable $\sum_{i=1}^2 b_i Y_k^i$ is Gaussian, zero-mean and has variance upper bounded by

$$\sum_{i=1}^2 \|(\frac{1}{n} \langle X_U^i, X_U^i \rangle)^{-1}\|_2 \frac{\|\tilde{z}_U^i\|_2^2}{n}.$$

As long as $s/n = o(1)$, Lemma 13 yields the bound

$$\|(\frac{1}{n} \langle X_U^i, X_U^i \rangle)^{-1}\|_2^2 \leq (1 + \delta)$$

with probability at least $1 - c_1 \exp(-c_2 n)$. Consequently, the sum $\sum_{i=1}^2 b_i Y_k^i$ is a zero mean Gaussian with variance at most $\frac{(1+\delta)}{n} \sum_{i=1}^2 \|\tilde{z}_U^i\|_2^2$.

Recall that Z_U was obtained from Step (B) of the primal-dual witness construction. The next lemma provides control over the sum $\sum_{i=1}^2 \|\tilde{z}_U^i\|_2^2$.

Lemma 5: Under the assumptions of Theorem 3 and Corollary 1, as long as $\lambda_n^2 n \rightarrow +\infty$, then $\|\tilde{z}_U^1\|_2^2$ is concentrated: for all $\delta > 0$, then for sufficiently large s and n , the probability

$$\mathbb{P} \left[\|\tilde{z}_U^1\|_2^2 + \|\tilde{z}_U^2\|_2^2 \leq (1 - \delta) \frac{s}{2} \left\{ (4 - 3\alpha) + \frac{1}{\lambda_n^2 s} \|\overline{B}_{\text{diff}}\|_2^2 \right\} \right] \tag{38}$$

converges to zero, and moreover

$$\mathbb{P} \left[\|\tilde{z}_U^1\|_2^2 + \|\tilde{z}_U^2\|_2^2 \geq (1 + \delta) \frac{s}{2} \left\{ (4 - 3\alpha) + \frac{1}{\lambda_n^2 s} \|\overline{B}_{\text{diff}}\|_2^2 \right\} \right] \leq c_1 \exp(-c_2 n). \tag{39}$$

See Appendix C for the proof of this claim.

Consequently, by applying the union bound and using Gaussian tail bounds, we obtain that the probability $\mathbb{P}[M_2 \geq 1 - \epsilon]$ is at most

$$4(p - (2 - \alpha)s) \exp \left\{ -\frac{(1 - \epsilon)^2 n}{(1 + \delta)s \left\{ (4 - 3\alpha) + \frac{\|\overline{B}_{\text{diff}}\|_2^2}{\lambda_n^2 s} \right\}} \right\} + c_1 \exp(-c_2 n)$$

which vanishes with $n \rightarrow \infty$ as long as

$$n > \frac{(1 + \delta)s\{(4 - 3\alpha) + \frac{1}{\lambda_n^2 s} \|\overline{B}_{\text{diff}}\|_2^2\}}{(1 - \epsilon)^2 \log(p - (2 - \alpha)s)}.$$

B. Proof of Theorem 3(b)

We now turn to the proof of the converse claim in Theorem 3. We establish the claim by contradiction. We show that if a solution \widehat{B} exists such that $\widehat{B}_{U^c} = 0$, then under the stated upper bound on the sample size n , there exists some $\epsilon > 0$ such that $\mathbb{P}[\max_{k \in U^c} (|\widehat{z}_k^1| + |\widehat{z}_k^2|) > 1 + \epsilon]$ converges to one. From the definition (33), we see that conditioned on $(X_U, w, \widehat{z}_U^1)$, the variables $\{\widehat{z}_k^1, k \in U^c\}$ are i.i.d. zero-mean Gaussians, with variance given by

$$\text{var}(\widehat{z}_k^1) := \left\| \frac{1}{\lambda_n n} \Pi_{U^\perp} w - \frac{1}{n} X_U \left(\frac{1}{n} X_U^T X_U \right)^{-1} \widehat{z}_U^1 \right\|_2^2.$$

By orthogonality, we have that

$$\text{var}(\widehat{z}_k^1) = \left\| \frac{1}{\lambda_n n} \Pi_{U^\perp} w \right\|_2^2 + \left\| \frac{1}{n} X_U \left(\frac{1}{n} X_U^T X_U \right)^{-1} \widehat{z}_U^1 \right\|_2^2$$

so that the idempotency of projection operators implies that $\text{var}(\widehat{z}_k^1) \geq \sigma^2$, where

$$\sigma^2 := \max \left\{ \frac{1}{\lambda_n^2 n} \frac{\|\Pi_{U^\perp} w\|_2^2}{n}, \lambda_{\min} \left(\left(\frac{1}{n} X_U^T X_U \right)^{-1} \right) \frac{\|\widehat{z}_U^1\|_2^2}{n} \right\}. \quad (40)$$

Note that $\sigma^2 = \sigma^2(X_U, w, \widehat{z}_U^1)$ is a scalar random variable, but fixed under the conditioning. Turning to the variables $\{\widehat{z}_k^2, k \in U^c\}$, a similar argument shows that we have $\text{var}(\widehat{z}_k^2) \geq \widetilde{\sigma}^2$, where $\widetilde{\sigma}^2 = \widetilde{\sigma}^2(\overline{X}_U, \overline{w}, \widehat{z}_U^2)$ is the analogous random variable.

For $k \in U^c$, let $\widehat{z}_k^1 \sim N(0, \sigma^2)$ and $\widehat{z}_k^2 \sim N(0, \widetilde{\sigma}^2)$. We then have

$$\begin{aligned} & \mathbb{P} \left[\frac{\max_{k \in U^c} (|\widehat{z}_k^1| + |\widehat{z}_k^2|)}{1 + \epsilon} > 1 \right] \\ & \stackrel{(a)}{\geq} \mathbb{P}[\max_{k \in U^c} |\widehat{z}_k^1| + |\widehat{z}_k^2| > 1 + \epsilon] \\ & \geq \mathbb{P}[\max_{k \in U^c} (\widehat{z}_k^1 + \widehat{z}_k^2) > 1 + \epsilon] \\ & \stackrel{(b)}{=} \mathbb{P}[\max_{k \in U^c} Z_k > 1 + \epsilon] \end{aligned}$$

where $Z_k \sim N(0, \sigma^2 + \widetilde{\sigma}^2)$. Here inequality (a) follows because σ^2 and $\widetilde{\sigma}^2$ are lower bounds on the variances of $\{\widehat{z}_k^1, k \in U^c\}$ and $\{\widehat{z}_k^2, k \in U^c\}$ respectively, and equality (b) follows since \widehat{z}_k^1 and \widehat{z}_k^2 are independent zero-mean Gaussians with variances σ^2 and $\widetilde{\sigma}^2$, respectively.

To simplify notation, let $N = |U^c| = p - (2 - \alpha)s$. By standard results for Gaussian maxima [14], for any $\delta > 0$, there exists an integer $N(\delta)$ such that for all $N \geq N(\delta)$,

$$\mathbb{E}[\max_{j \in U^c} Z_j] \geq (1 - \delta) \sqrt{2(\sigma^2 + \widetilde{\sigma}^2) \log N}.$$

Moreover, the maximum function is Lipschitz, so that by Gaussian concentration for Lipschitz functions [13], [14], for any $\eta > 0$, we have

$$\mathbb{P} \left[\max_{j \in U^c} Z_j \leq \mathbb{E}[\max_{j \in U^c} Z_j] - \eta \right] \leq \exp \left(- \frac{\eta^2}{2(\sigma^2 + \widetilde{\sigma}^2)} \right).$$

Combining these two statements yields that for all $N \geq N(\delta)$, we have

$$\begin{aligned} \mathbb{P} \left[\max_{j \in U^c} Z_j \leq (1 - \delta) \sqrt{2(\sigma^2 + \widetilde{\sigma}^2) \log N} - \eta \right] \\ \leq \exp \left(- \frac{\eta^2}{2(\sigma^2 + \widetilde{\sigma}^2)} \right). \quad (41) \end{aligned}$$

It remains to show that there exists some $\epsilon > 0$ such that $\mathbb{P}[\max_{k \in U^c} Z_k \leq 1 + \epsilon]$ converges to zero.

Case 1: First suppose that $\lambda_n^2 n = \mathcal{O}(1)$. In this case, we have $\sigma^2 = \Omega\left(\frac{\|\Pi_{U^\perp} w\|_2^2}{n}\right)$. With probability greater than $1 - c_1 \exp(-c_2 n)$, this quantity is lower bounded by a constant, using concentration for χ^2 -variates. In this case, $\sqrt{2(\sigma^2 + \widetilde{\sigma}^2) \log N} - \eta \rightarrow +\infty$ w.h.p., so that the result follows trivially.

Case 2: Otherwise, we must have $\lambda_n^2 n \rightarrow +\infty$. Under this condition, we now establish a lower bound on σ^2 that holds with high probability; it will be seen that a similar lower bound holds for $\widetilde{\sigma}^2$. We begin by noting the lower bound $\sigma^2 \geq \frac{\|\widehat{z}_U^1\|_2^2}{n} \lambda_{\min} \left(\left(\frac{1}{n} X_U^T X_U \right)^{-1} \right)$. To control the minimum eigenvalue, define the event

$$\begin{aligned} T(X_U) := \\ \{X_U \mid \lambda_{\min} \left(\left(\frac{1}{n} X_U^T X_U \right)^{-1} \right) \geq (1 + \sqrt{s/n})^{-2}\}. \quad (42) \end{aligned}$$

By standard random matrix concentration arguments (see Appendix E), for a universal constant $c > 0$, we are guaranteed that $\mathbb{P}[T^c(X_U)] \leq 2 \exp(-cn)$. Consequently, conditioned on $T(X_U)$, we have

$$\sigma^2 + \widetilde{\sigma}^2 \geq \frac{\|\widehat{z}_U^1\|_2^2 + \|\widehat{z}_U^2\|_2^2}{n} (1 + \sqrt{s/n})^{-2}. \quad (43)$$

From Lemma 5, we note that if $s/n = o(1)$, then for any $\delta > 0$, we have the lower bound

$$\begin{aligned} \sigma^2 + \widetilde{\sigma}^2 \geq \\ (1 - \delta) \frac{s}{2n} \{(4 - 3\alpha) + (\Delta(\overline{B}, \lambda_n))^2\} (1 - o(1)). \quad (44) \end{aligned}$$

The following result is the final step in the proof of Theorem 3(b).

Lemma 6: Suppose that $\lambda_n^2 n \rightarrow +\infty$.

- a) If $\frac{s}{n} = \Omega(1)$, then $\mathbb{P}[\max_{k \in U^c} Z_k \leq 2] \rightarrow 0$.
- b) If $\frac{s}{n} \rightarrow 0$, then there exists some $\epsilon > 0$ such that $\mathbb{P}[\max_{k \in U^c} Z_k \leq 1 + \epsilon] \rightarrow 0$.

Proof:

- a) If $\frac{s}{n}$ is bounded below by some constant $c > 0$, then we have

$$\sigma^2 \geq (1 - \alpha) \frac{s}{n} \geq (1 - \alpha)c,$$

which implies that $(\sigma^2 + \tilde{\sigma}^2) \log N \rightarrow +\infty$. Thus, setting $\delta = 1/4$ and $\eta = \frac{1}{2} \sqrt{2(\sigma^2 + \tilde{\sigma}^2) \log N}$ in (41), we obtain

$$\begin{aligned} \mathbb{P} \left[\max_{k \in U^c} Z_k \leq \left(\frac{1}{2} - \delta \right) \sqrt{2(\sigma^2 + \tilde{\sigma}^2) \log N} \right] \\ = \mathbb{P} \left[\max_{k \in U^c} Z_k \leq \frac{1}{4} \sqrt{2(\sigma^2 + \tilde{\sigma}^2) \log N} \right] \end{aligned}$$

which is at most $\exp \left(- \frac{\log N}{4} \right)$. Since $\frac{1}{4} \sqrt{2(\sigma^2 + \tilde{\sigma}^2) \log N} \geq 2$ for N large enough, the claim follows.

- b) In this case, we may apply the lower bound (44), so that, for any $\delta > 0$, we have

$$\begin{aligned} \sigma^2 + \tilde{\sigma}^2 \\ \geq (1 - \delta) \frac{s}{2n} \{ (4 - 3\alpha) + (\Delta(\bar{B}, \lambda_n)) \} (1 - o(1)) \end{aligned}$$

with high probability. Since $n < (1 - \nu)[(4 - 3\alpha) + (\Delta(\bar{B}, \lambda_n))]s \log N$ by assumption, we obtain that $\sqrt{2(\sigma^2 + \tilde{\sigma}^2) \log N}$ is at most

$$(1 - o(1)) \sqrt{(1 - \delta) \frac{s}{n} \{ (4 - 3\alpha) + (\Delta(\bar{B}, \lambda_n)) \} \log N}$$

which is in turn at most $(1 - o(1)) \sqrt{\frac{1 - \delta}{1 - \nu}}$. Consequently, from (41), for any $\eta > 0$ and $\delta > 0$, we have for all $N \geq N(\delta)$,

$$\begin{aligned} \mathbb{P} \left[\max_{k \in U^c} Z_k \leq (1 - \delta) (1 - o(1)) \frac{1}{\sqrt{1 - \nu}} - \eta \right] \\ \exp \left(- \frac{\eta^2}{2(\sigma^2 + \tilde{\sigma}^2)} \right). \quad (45) \end{aligned}$$

Since $\nu > 0$, we may choose $\eta, \delta > 0$ sufficiently small so that for sufficiently large choices of (s, n) , we have

$$(1 - \delta) (1 - o(1)) \frac{1}{\sqrt{1 - \nu}} - \eta \geq 1 + \epsilon$$

for some $\epsilon > 0$. Since from Lemma 5, the condition $s/n = o(1)$ implies that $\sigma^2 + \tilde{\sigma}^2 = o(1)$ w.h.p, we thus conclude that, using these choices of η and δ , we have

$$\mathbb{P} \left[\max_{k \in U^c} Z_k \leq 1 + \epsilon \right] \leq o(1) + \exp \left(- \frac{\eta^2}{2(\sigma^2 + \tilde{\sigma}^2)} \right)$$

which tends to zero as claimed. ■

VIII. DISCUSSION

In this paper, we provided a number of theoretical results that provide a sharp characterization of when, and if so by how much the use of block ℓ_1/ℓ_∞ regularization actually leads to improvements in statistical efficiency in the problem of multivariate regression. As suggested in a body of past work, the

use of block ℓ_1/ℓ_∞ regularization is well-motivated in many application contexts. However, since it involves greater computational cost than more naive approaches, the question of whether this greater computational price yields statistical gains is an important one.

This paper assessed statistical efficiency in terms of the number of samples required to recover the support exactly; however, one could imagine studying the same issue for related loss functions (e.g., ℓ_2 -loss or prediction loss), and it would be interesting to see if the results were qualitatively similar or not. Our results demonstrate that some care needs to be exercised in the application of ℓ_1/ℓ_∞ regularization. Indeed, it can yield improved statistical efficiency when the regression matrix exhibits structured sparsity, with high overlaps among the sets of active coefficients within each column. However, our analysis shows that these improvements are quite sensitive to the exact structure of the regression matrix, and how well it aligns with the regularizing norm. When this alignment is not high enough, then the use of ℓ_1/ℓ_∞ can actually impair performance relative to more naive (and less computationally intensive) schemes based on ℓ_1 -regularization, such as the Lasso. Moreover, whether or not regularization ℓ_1/ℓ_∞ yields statistical improvements is very sensitive to the actual magnitudes of the different regression problems. In comparison to related results obtained by Obozinski *et al.* [25] on block ℓ_1/ℓ_2 regularization, the block ℓ_1/ℓ_∞ norm exhibits some fragility, in that the conditions under which it actually improves statistical efficiency are delicate and easily violated. An interesting open direction is study whether or not it is possible to develop computationally efficient methods that are *fully adaptive* to the sparsity overlap—namely, methods that behave like ordinary ℓ_1 -regularization when there is no or little shared sparsity, and behave like block regularization schemes in the presence of shared sparsity.

APPENDIX A

RECOVERING INDIVIDUAL SIGNED SUPPORTS

In this appendix, we discuss some issues associated with recovering individual signed supports. We begin by observing that once the support union U has been recovered, one can restrict the regression problem to this subset U , and then apply Lasso to each problem separately (with substantially lower cost, since each problem is now low-dimensional) in order to recover the individual signed supports. If one is not willing to perform some extra computation in this way, then the interpretation of Theorems 1 and 2—in terms of recovering the individual signed supports—requires a more delicate treatment, which we discuss in this appendix.

Interestingly, the structure of the block ℓ_1/ℓ_∞ norm permits two ways in which to recover the individual signed supports.

A) ℓ_1/ℓ_∞ *Primal Recovery*: Solve the block-regularized program (6), thereby obtaining a (primal) optimal solution $\hat{B} \in \mathbb{R}^{p \times r}$. Estimate the support union via $\hat{U} := \bigcup_{i=1, \dots, r} S(\hat{\beta}^i)$, and estimate the signed support vectors via

$$[\mathbb{S}_{\text{pri}}(\hat{\beta}^i)]_k := \text{sign}(\hat{\beta}_k^i). \quad (46)$$

B) ℓ_1/ℓ_∞ *Dual Recovery*: Solve the block-regularized program (6), thereby obtaining a primal solution

$\hat{B} \in \mathbb{R}^{p \times r}$. For each row $k = 1, \dots, p$, compute the set $\mathbb{M}_k := \arg \max_{i=1, \dots, r} |\hat{\beta}_k^i|$. Estimate the support union via $\hat{U} = \bigcup_{i=1, \dots, r} S(\hat{\beta}^i)$, and estimate the signed support vectors

$$[\mathbb{S}_{\text{dua}}(\hat{\beta}_k^i)] = \begin{cases} \text{sign}(\hat{\beta}_k^i), & \text{if } i \in \mathbb{M}_k \\ 0, & \text{otherwise.} \end{cases} \quad (47)$$

The procedure (47) corresponds to estimating the signed support on the basis of a dual optimal solution associated with the optimal primal solution.

The dual signed support recovery method (47) is more conservative in estimating the individual support sets. In particular, for any given $i \in \{1, \dots, r\}$, it only allows an index k to enter the signed support estimate $\mathbb{S}_{\text{dua}}(\hat{\beta}^i)$ when $|\hat{\beta}_k^i|$ achieves the maximum magnitude (possibly nonunique) across all indices $i = 1, \dots, r$. Consequently, unlike the primal estimator (47), a corollary of Theorem 1 guarantees that the dual signed support method (47) never suffers from false inclusions in the signed support set. On the other hand, unlike the primal estimator, it may incorrectly exclude indices of some supports—that is, it may exhibit false exclusions.

To provide a concrete illustration of this distinction, suppose that $p = 4$ and $r = 3$, and that the true matrix \bar{B} and estimate take the following form:

$$\bar{B} = \begin{bmatrix} 2 & 0 & -3 \\ 2 & 4 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \text{ and } \hat{B} = \begin{bmatrix} 1.9 & 0.1 & -2.9 \\ 1.7 & 3.9 & -0.1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Consistent with the claims of Theorem 1, the estimate \hat{B} correctly recovers the support union—viz. $S(\hat{B}) = \hat{U} = \{1, 2\} = S(\bar{B})$. The primal (46) and dual (47) methods return the following estimates of the individual signed supports:

$$\mathbb{S}_{\text{pri}}(\hat{B}) = \begin{bmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \text{ and } \mathbb{S}_{\text{dua}}(\hat{B}) = \begin{bmatrix} 0 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Consequently, the primal estimate includes false nonzeros in positions (1, 2) and (2, 3), whereas the dual estimate includes false zeros in positions (1, 1) and (2, 1).

We note that it is possible to ensure that under some conditions that the dual support method (47) will correctly recover each of the individual signed supports, without any incorrect exclusions. However, as illustrated by Theorem 3 and Corollary 1, doing so requires additional assumptions on the size of the gap $|\bar{\beta}_k^i| - |\bar{\beta}_k^j|$ for indices $k \in B := S(\bar{\beta}^i) \cap S(\bar{\beta}^j)$.

APPENDIX B PROOF OF LEMMA 4

Note that conditioned on X_U , the rows of the random matrix $X_{U^c}^i$ are i.i.d. Gaussian random vectors with mean $\Sigma_{U^c U}^i (\Sigma_{UU}^i)^{-1} (X_U^i)^T$ and covariance

$$\Sigma_{U^c|U}^i = \Sigma_{U^c U^c}^i - \Sigma_{U^c U}^i (\Sigma_{UU}^i)^{-1} \Sigma_{UU^c}^i.$$

Consequently, we can write

$$X_{U^c}^i = X_U^i (\Sigma_{UU}^i)^{-1} \Sigma_{UU^c}^i + Y_{U^c}^i$$

where the rows of $Y_{U^c}^i$ are distributed $N(0, \Sigma_{U^c|U}^i)$. Putting together the pieces, we can rewrite $\frac{1}{n} \langle X_{U^c}^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \rangle$ as

$$\Sigma_{U^c U}^i (\Sigma_{UU}^i)^{-1} + \frac{1}{n} \langle Y_{U^c}^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} \rangle.$$

Using these expressions and triangle inequality, we obtain that M is upper bounded by

$$\max_{k \in U^c} \left\{ \sum_{i=1}^r \|e_k^T \Sigma_{U^c U}^i (\Sigma_{UU}^i)^{-1}\|_1 \right\} + \max_{k \in U^c} \sum_{i=1}^r \left| \frac{1}{n} \langle Y_k^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} z_U^i \rangle \right|.$$

Applying the mutual incoherence assumption (19) yields

$$M \leq (1 - \gamma) + \max_{k \in U^c} \sum_{i=1}^r \left| \frac{1}{n} \langle Y_k^i, X_U^i (\langle \frac{1}{n} X_U^i, X_U^i \rangle)^{-1} z_U^i \rangle \right| \quad (48)$$

as claimed.

APPENDIX C PROOF OF LEMMA 5

Recall that $\tilde{z}_B^1 = (\tilde{z}_{B^c}^1, \tilde{z}_B^1)$, $\|\tilde{z}_{B^c}^1\|_2^2 = (1 - \alpha)s$, and that B is the set where $|\hat{\beta}_B^1| = |\hat{\beta}_B^2|$. Thus, the claim is equivalent to showing that $\|\tilde{z}_B^1\|_2^2$ is concentrated. If $\alpha = 0$, then the claim is trivial, so that we may assume that $\alpha > 0$. Recall moreover that

$$S^1 \tilde{z}_B^1 = \frac{1}{\lambda_n} \{ M^2 [M^1 + M^2]^{-1} f^1 - M^1 [M^2 + M^1]^{-1} f^2 \} + M^1 [M^1 + M^2]^{-1} \bar{I} - \frac{1}{\lambda_n} [(M^1)^{-1} + (M^2)^{-1}]^{-1} \bar{B}_{\text{diff}}. \quad (49)$$

Using $\|\cdot\|_2$ to denote the spectral norm, we first claim that as long as $s/n \rightarrow 0$, then the following events hold with probability greater than $1 - c_1 \exp(-c_2 n)$:

$$\|M^1 - I\|_2 = \mathcal{O}(\sqrt{s/n}), \quad (50a)$$

$$\|[(M^1)^{-1} + (M^2)^{-1}]^{-1} - I/2\|_2 = \mathcal{O}(\sqrt{s/n}), \quad (50b)$$

and

$$\|M^1 [M^1 + M^2]^{-1} - I/2\|_2 = \mathcal{O}(\sqrt{s/n}) \quad (50c)$$

as well as the analogous events with M^1 and M^2 interchanged.

To verify the bound (50a), we first diagonalize the projection matrix. All of its eigenvalues are 0 or 1, and it has rank $(n - s)$ w.p. one, so that we may write $\Pi_{B^c \perp} = U^T D U$ for some orthogonal matrix U , and the diagonal matrix $D = \text{diag}\{1_{n-s}, 0_s\}$:

$$M = n^{-1} X_B^T U^T D U X_B.$$

But the projection $\Pi_{B^c \perp}$ is independent of X_B , which implies that the random rotation matrix U is independent of X_B , and hence $X_B \stackrel{d}{=} U X_B$. Since D is diagonal with $(n - s)$ ones and s zeros, $M \stackrel{d}{=} n^{-1} W^T W$, where $W \in \mathbb{R}^{(n-s) \times |B|}$ is a standard Gaussian random matrix. Consequently, the quantity $\|M - I\|_2$

is equal in distribution to the quantity $\|n^{-1}W^TW - I\|_2$, and moreover, we have

$$\|n^{-1}W^TW - I\|_2 \leq \left| \frac{n-s}{n} - 1 \right| \left\| \frac{1}{n-s} W^TW \right\|_2 + \left\| \frac{1}{n-s} W^TW - I \right\|_2$$

a quantity which scales as $\mathcal{O}(\sqrt{s/n})$, since $\|W^TW/(n-s)\|_2 = \mathcal{O}(1)$, and

$$\left\| \frac{1}{n-s} W^TW - I \right\|_2 = \mathcal{O}\left(\sqrt{\frac{s}{n-s}}\right) = \mathcal{O}(\sqrt{s/n})$$

using concentration arguments for random matrices (see Lemma 13 in Appendix E).

Turning to the claim (50b), we first note that

$$\begin{aligned} \|[M^{-1} + \widetilde{M}^{-1}]^{-1} - I/2\|_2 &= \underbrace{\|[M^{-1} + \widetilde{M}^{-1}]^{-1}(I - [M^{-1} + \widetilde{M}^{-1}]/2)\|_2}_A. \end{aligned}$$

By triangle inequality and the submultiplicativity of the norm, we have the equation shown at the bottom of the page. Finally, since $\|[M^{-1} + \widetilde{M}^{-1}]^{-1}\|_2 = \mathcal{O}(1)$, we see that (50b) holds.

In order to establish the bound (50c), note that

$$\|M[M + \widetilde{M}]^{-1} - I/2\|_2 = \|(M/2 - \widetilde{M}/2)[M + \widetilde{M}]^{-1}\|_2$$

which is upper bounded by

$$\begin{aligned} \frac{1}{2} \{ \|M - I\|_2 + \|\widetilde{M} - I\|_2 \} \|[M + \widetilde{M}]^{-1}\|_2 &= \|[M + \widetilde{M}]^{-1}\|_2 \times \mathcal{O}(\sqrt{s/n}). \end{aligned}$$

Since $\|[M + \widetilde{M}] - 2I\|_2 = \mathcal{O}(\sqrt{s/n}) \rightarrow 0$, we have $\|[M + \widetilde{M}]^{-1}\|_2 = \mathcal{O}(1)$, which establishes the claim (50c).

We are now ready to establish the claims of the lemma. From the representation (49), we apply triangle inequality and our bounds on spectral norms, thereby obtaining that $\sqrt{\|\widetilde{z}_B^1\|_2^2 + \|\widetilde{z}_B^2\|_2^2}$ is upper bounded by

$$\begin{aligned} \left[\left\| \frac{\widetilde{I}}{2} - \frac{1}{2\lambda_n} (|\widetilde{\beta}_B^2| - |\widetilde{\beta}_B^1|) \right\|_2^2 + \left\| \frac{\widetilde{I}}{2} + \frac{1}{2\lambda_n} (|\widetilde{\beta}_B^2| - |\widetilde{\beta}_B^1|) \right\|_2^2 \right]^{1/2} + 2\|r\| \end{aligned}$$

which is in turn at most

$$\sqrt{s} \left\{ \sqrt{\frac{\alpha}{2} + \frac{1}{2s\lambda_n^2} \left(\|\widetilde{\beta}_B^1\|_2^2 - |\widetilde{\beta}_B^2| \right)} + \frac{2}{\sqrt{s}} \|r\|_2 \right\}$$

with probability greater than $1 - c_1 \exp(-c_2 n)$, where $r = \widetilde{z}_B^1 - \frac{1}{2}(\widetilde{I} - \frac{1}{\lambda_n} (|\widetilde{\beta}_B^2| - |\widetilde{\beta}_B^1|))$. By the decomposition of \widetilde{z}_B^1 in (49) and applying bounds (50), we obtain that $\|r\|_2$ is at most

$$\begin{aligned} \sqrt{s} \{ \mathcal{O}(\sqrt{s/n}) + \frac{1}{\lambda_n \sqrt{s}} \mathcal{O}(\sqrt{s/n}) \|\widetilde{\beta}_B^1\|_2 - |\widetilde{\beta}_B^2| \} &+ \frac{1}{2\sqrt{s}\lambda_n} (1 + \mathcal{O}(\sqrt{s/n})) (\|f\|_2 + \|\widetilde{f}\|_2). \end{aligned}$$

Since $s/n = o(1)$, in order to establish the upper bound (39) it suffices to show that $\|f\|_2 + \|\widetilde{f}\|_2 = o(\sqrt{s}\lambda_n)$ with high probability. Similarly, in the other direction, we have

$$\begin{aligned} &\sqrt{\|\widetilde{z}_B^1\|_2^2 + \|\widetilde{z}_B^2\|_2^2} \\ &\geq \sqrt{s} \left\{ \sqrt{\frac{\alpha}{2} + \frac{1}{2s\lambda_n^2} \left(\|\widetilde{\beta}_B^1\|_2^2 - |\widetilde{\beta}_B^2| \right)} - \frac{2}{\sqrt{s}} \|r\|_2 \right\}. \end{aligned}$$

Following the same line of reasoning, in order to prove the lower bound (38), it suffices to show that $\|f\|_2 + \|\widetilde{f}\|_2 = o(\sqrt{s}\lambda_n)$ with high probability. Since $\|f\|_2$ and $\|\widetilde{f}\|_2$ behave similarly, it suffices to show that $\|f\|_2 = o(\lambda_n \sqrt{s})$. From the definition (55), we see that conditioned on $(X_{B^c}, w, \widetilde{z}_{B^c}^1)$, the random vector f is zero-mean Gaussian, with i.i.d. elements with variance

$$\sigma^2 := \frac{\lambda_n^2}{n} (\widetilde{z}_{B^c}^1)^T (X_{B^c}^T X_{B^c} / n)^{-1} \widetilde{z}_{B^c}^1 + \frac{1}{n} w^T \Pi_{B^c \perp} w.$$

Recalling that $\|\widetilde{z}_{B^c}^1\|_2^2 = (1 - \alpha)s$, the variance σ^2 is at most

$$\frac{\lambda_n^2 (1 - \alpha)s}{n} \lambda_{\max}((X_{B^c}^T X_{B^c} / n)^{-1}) + \frac{1}{n} \frac{\|\Pi_{B^c \perp}(w)\|_2^2}{n}.$$

By random matrix concentration (see the discussion following Lemma 13 in Appendix E), we have $\lambda_{\max}((X_{B^c}^T X_{B^c} / n)^{-1}) \leq 1 + \mathcal{O}(\sqrt{s/n})$ w.h.p., and by χ^2 tail bounds (see Lemma 12 in Appendix E), we have $\frac{\|\Pi_{B^c \perp}(w)\|_2^2}{n} = \mathcal{O}(1)$ w.h.p. Consequently, with high probability, we have $\sigma^2 = \mathcal{O}(\frac{\lambda_n^2 s}{n} + \frac{1}{n})$. Since the Gaussian random vector f has length $|B| = \Theta(s)$, again by concentration for χ^2 random variables, we have (with

$$\begin{aligned} A &\leq \|[M^{-1} + \widetilde{M}^{-1}]^{-1}\|_2 \|I - [M^{-1} + \widetilde{M}^{-1}]/2\|_2 \\ &\leq \frac{1}{2} \{ \|I/2 - M^{-1}/2\|_2 + \|I/2 - \widetilde{M}^{-1}/2\|_2 \} \|[M^{-1} + \widetilde{M}^{-1}]^{-1}\|_2 \\ &= \|[M^{-1} + \widetilde{M}^{-1}]^{-1}\|_2 \mathcal{O}(\sqrt{s/n}) \end{aligned}$$

probability greater than $1 - c_1 \exp(-c_2 s)$, $\|f\|_2^2 = \mathcal{O}(\sigma^2 s)$. Combining the pieces, we conclude that w.h.p.

$$\begin{aligned} \|f\|_2^2 &= \mathcal{O}(\lambda_n^2 s \frac{s}{n} + \frac{s}{n}) \\ &= \mathcal{O}(\lambda_n^2 s [\frac{s}{n} + \frac{1}{\lambda_n^2 n}]) = o(\lambda_n^2 s) \end{aligned}$$

where the final equality follows since $s/n = o(1)$ and $1/(\lambda_n^2 n) = o(1)$.

APPENDIX D CONVEX-ANALYTIC CHARACTERIZATION OF OPTIMAL SOLUTIONS

This section is devoted to the development of various properties of the optimal solution(s) of the block ℓ_1/ℓ_∞ -regularized problem (6).

A) Basic Optimality Conditions: By standard conditions for optimality in convex programs [29], the zero-vector must belong to the subdifferential of the objective function in the convex program (6), or equivalently, we must have for each $i = 1, 2, \dots, r$

$$\frac{1}{n} \langle X^i, X^i \rangle \hat{\beta}^i - \frac{1}{n} (X^i)^T y^i + \lambda_n \tilde{z}^i = 0 \quad (51)$$

where $\tilde{Z} \in \mathbb{R}^{p \times r}$ must be an element of the subdifferential $\partial \|\hat{B}\|_{\infty,1}$. Substituting the relation $y^i = X^i \bar{\beta}^i + w^i$, we obtain

$$\frac{1}{n} \langle X^i, X^i \rangle (\hat{\beta}^i - \bar{\beta}^i) - \frac{1}{n} (X^i)^T w^i + \lambda_n \tilde{z}^i = 0. \quad (52)$$

B) Proof of Lemma 2: We begin with the proof of part i): suppose that steps (A)–(C) of the primal-witness construction succeed. By definition, it outputs a primal pair, of the form $(\hat{B}_U, 0)$, along with a candidate dual optimal solution $\{(\tilde{Z}_U, \tilde{Z}_{U^c})\}$. Note that the conditions defining the ℓ_1/ℓ_∞ subdifferential apply in an elementwise manner, to each index $i = 1, \dots, p$. Since the subvector \tilde{Z}_U was chosen from the subdifferential of the restricted optimal solution, it is dual feasible. Moreover, since the strict dual feasibility condition (30) holds, the matrix \tilde{Z}_{U^c} constructed in step (C) is dual feasible for the zero-solution in the subblock U^c . Therefore, we conclude that $(\hat{B}_U, 0)$ is a primal optimal solution for the full block-regularized program (6).

It remains to establish uniqueness of this solution. Define the ball

$$\mathbb{K} = \{ \tilde{Z} \in \mathbb{R}^{p \times r} \mid \sum_{i=1}^r |\tilde{z}_k^i| \leq 1 \quad \forall k = 1, \dots, p \}$$

and observe that we have the variational representation

$$\|B\|_{1,\infty} = \sup_{\tilde{Z} \in \mathbb{K}} \langle \tilde{Z}, B \rangle$$

where $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product. With this notation, the block-regularized program (6) is equivalent to the saddle-point problem

$$\inf_{B \in \mathbb{R}^{p \times r}} \sup_{\tilde{Z} \in \mathbb{K}} \left\{ \frac{1}{2n} \sum_{i=1}^r \|y^i - X^i \beta^i\|_2^2 + \lambda_n \langle \tilde{Z}, B \rangle \right\}.$$

Since this saddle-point problem is strictly feasible and convex-concave, it has a value. Moreover, given any dual optimal solution—in particular, \tilde{Z} from the primal-dual construction—any optimal primal solution \hat{B} must satisfy the saddle point condition

$$\|\hat{B}\|_{1,\infty} = \sup_{\tilde{Z} \in \mathbb{K}} \langle \tilde{Z}, \hat{B} \rangle.$$

But this condition can only hold if $\forall i \in \{1, 2, \dots, r\}$, $\beta_k^i = 0$ for any index $k \in \{1, \dots, p\}$ such that $\sum_{i=1}^r |\tilde{z}_k^i| < 1$. Therefore, any optimal primal solution must satisfy $\hat{B}_{U^c} = 0$, so that solving the original program (6) is equivalent to solving the restricted program (29). Lastly, if the matrices $\langle X_U^i, X_U^i \rangle$ are invertible for each $i \in \{1, 2, \dots, r\}$, then the restricted problem (29) is strictly convex, and so has a unique solution, thereby completing the proof of Lemma 2(i).

We now prove part ii) of Lemma 2. Suppose that we are given an estimate \hat{B} of the true parameters \bar{B} by solving the convex program (6) such that $\hat{B}_{U^c} = 0$.

Since \hat{B} is an optimal solution to the convex program (6), the optimality conditions of equation (52), must be satisfied. We may rewrite those conditions as

$$\begin{aligned} \frac{1}{n} \langle X_U^i, X_U^i \rangle (\Delta^i) - \frac{1}{n} (X_U^i)^T w^i + \lambda_n \tilde{z}_U^i &= 0 \\ \frac{1}{n} \langle X_{U^c}^i, X_U^i \rangle (\Delta^i) - \frac{1}{n} (X_{U^c}^i)^T w^i + \lambda_n \tilde{z}_{U^c}^i &= 0 \end{aligned}$$

where $\Delta^i = \hat{\beta}^i - \bar{\beta}^i$. Recalling that $\hat{B}_{U^c} = \bar{B}_{U^c} = 0$, we obtain

$$\frac{1}{n} \langle X_U^i, X_U^i \rangle (\Delta^i_U) - \frac{1}{n} (X_U^i)^T w^i + \lambda_n \tilde{z}_U^i = 0, \text{ and} \quad (53a)$$

$$\frac{1}{n} \langle X_{U^c}^i, X_U^i \rangle (\Delta^i_U) - \frac{1}{n} (X_{U^c}^i)^T w^i + \lambda_n \tilde{z}_{U^c}^i = 0. \quad (53b)$$

Again, by standard conditions for optimality in convex programs [4], [9], the first of these two equations is exactly the condition that must be satisfied by an optimal solution of the restricted program (29). However, we have already shown that the candidate solution \hat{B}_U satisfies this condition, so that it must also be an optimal solution of the convex program (29). Additionally, the value of \tilde{Z}_U that satisfies (53a) for each $i \in \{1, 2, \dots, r\}$ is an element of $\partial \|\hat{B}\|_{\infty,1}$. We have thus shown that steps (B) and (C) of the primal-witness construction succeed. It remains to establish uniqueness in part (A). However, we note that $\langle X_U^i, X_U^i \rangle$ is invertible for each i . Hence, for any solution \hat{B} such that $\hat{B}_{U^c} = 0$:

$$\Delta^i_U = \left(\frac{1}{n} \langle X_U^i, X_U^i \rangle \right)^{-1} \left[\frac{1}{n} (X_U^i)^T w^i - \lambda_n \tilde{z}_U^i \right]$$

is well-defined and unique, noting that $\Delta^i_{U^c} = 0$. Thus, we have established the equality (32) and that \hat{B}_U is unique. Therefore, \hat{B} gives solutions to steps (A) and (B) when solving the restricted convex program over the set U .

Finally, we derive the form of the dual solution $\tilde{z}_{U^c}^i$, as a function of $\langle X_U^i, X_U^i \rangle$, \tilde{Z}_U , and $\hat{B} - \bar{B}$. Recall that $\langle X_U^i, X_U^i \rangle$ is in-

vertible, \tilde{Z}_U is an element of the subdifferential of $\partial\|\tilde{B}_U\|_{\ell_1/\ell_\infty}$, and $\hat{B}_{U^c} = \bar{B}_{U^c} = 0$. From equation (32), we have

$$\begin{aligned} \tilde{z}_{U^c}^i &= \frac{1}{\lambda_n n} \langle X_{U^c}^i, (I - \Pi_{X_U^i}) w^i \rangle \\ &\quad + \frac{1}{n} \langle X_{U^c}^i, X_U^i (\frac{1}{n} \langle X_U^i, X_U^i \rangle)^{-1} \tilde{z}_U^i \rangle \end{aligned} \quad (54)$$

valid for $i = 1, \dots, r$. The claimed form of the dual solution follows by substituting (32) into (53b).

C) *Subgradients on the Support:* In this section, we focus on the specific form of the dual variables \tilde{z}_U^i . Our approach is to construct a candidate set of dual variables, and then show that they are valid. We begin by defining the sets $B = S(\bar{\beta}^i) \cap S(\bar{\beta}^j)$, corresponding to the intersection of the supports, and the set $B^c = U \setminus B$ corresponding to elements in one (but not both) of the supports. For $i = 1, 2$, we let $S^i \in \mathbb{R}^{\alpha s \times \alpha s}$ is a diagonal matrix whose diagonal entries correspond to $\text{sign}(\bar{\beta}_B^i)$. In addition, we define the vectors $f^i \in \mathbb{R}^{\alpha s}$ via

$$\begin{aligned} f^i := S^i \left[\frac{1}{n} \langle X_B^i, X_{B^c}^i \rangle (\frac{1}{n} \langle X_{B^c}^i, X_{B^c}^i \rangle)^{-1} \lambda_n \tilde{z}_{B^c}^i \right. \\ \left. - \frac{1}{n} \langle X_B^i, I - \Pi_{X_{B^c}^i} \rangle w^i \right] \end{aligned} \quad (55)$$

as well as the matrices $M^i \in \mathbb{R}^{\alpha s \times \alpha s}$ via

$$M^i := \frac{1}{n} S^i \langle X_B^i, (I - \Pi_{(X_{B^c}^i)}) X_B^i \rangle S^i.$$

Given these definitions, we have the following lemma.

Lemma 7: Assume that $r = 2$, and that $|\hat{\beta}_B^1| = |\hat{\beta}_B^2|$. If $\hat{B}_{U^c} = \bar{B}_{U^c} = 0$, then the dual variable \tilde{z}^1 satisfies the relation

$$\begin{aligned} S^1 \tilde{z}_B^1 &= \\ \frac{1}{\lambda_n} \{ &M^2 [M^1 + M^2]^{-1} f^1 - M^1 [M^2 + M^1]^{-1} f^2 \} \\ &+ M^1 [M^1 + M^2]^{-1} \bar{1} - \\ &\frac{1}{\lambda_n} [(M^1)^{-1} + (M^2)^{-1}]^{-1} \bar{B}_{\text{diff}} \end{aligned} \quad (56)$$

and $\tilde{z}_{B^c}^2 = \mathbb{S}_\pm(\bar{\beta}_{B^c}^2)$, with analogous results holding for \tilde{z}^2 .

Given these forms for $S^1 \tilde{z}_B^1$ and $S^2 \tilde{z}_B^2$, it remains to show that the relation $S^1 \tilde{z}_B^1 + S^2 \tilde{z}_B^2 = 1$ holds under the conditions of Theorem 3(a). Intuitively, this condition should hold since under the conditions of theorem 3(a), the matrix M^i is approximately the identity, and the vector f^i is approaching 0. Finally, we expect that $\bar{B}_{\text{diff}} := |\bar{\beta}_B^2| - |\bar{\beta}_B^1|$ is very small, hence the final term is also very small. Therefore, on the set B , both $S^1 \tilde{z}_B^1$ and $S^2 \tilde{z}_B^2$ are approximately equal to $\frac{1}{2}$. We formalize this rough intuition in the following lemma.

Lemma 8: Under the assumptions of Theorem 3(a), and for sufficiently large (n, p, s) , each of the following conditions hold with probability greater than $1 - c_1 \exp(-c_2 n)$:

$$\left\| \frac{1}{\lambda_n} [(M^1)^{-1} + (M^2)^{-1}]^{-1} (\bar{B}_{\text{diff}}) \right\|_\infty \leq \epsilon, \quad (57)$$

$$\begin{aligned} \left\| \frac{1}{\lambda_n} \{ M^2 [M^1 + M^2]^{-1} f^1 \right. \\ \left. - M^1 [M^2 + M^1]^{-1} f^2 \} \right\|_\infty \leq \epsilon \end{aligned} \quad (58)$$

and

$$\left\| M^1 [M^1 + M^2]^{-1} \bar{1} - \frac{1}{2} \right\|_\infty \leq \frac{1}{2} - 3\epsilon. \quad (59)$$

Given Lemmas 7 and 8, we can conclude that the definition for the dual variables on the support is valid. The remaining subsections in this appendix are dedicated to verifying the above results: in particular, we prove Lemma 7 in Appendix D4 and Lemma 8 in Appendix D5.

D) *Proof of Lemma 7:* We now proceed to establish the validity of the closed form expressions for \tilde{z}_U^1 and \tilde{z}_U^2 . From (53a), we have

$$\begin{aligned} \Delta^1_{B^c} &= (\frac{1}{n} \langle X_{B^c}^i, X_{B^c}^i \rangle)^{-1} (X_{B^c}^1)^T w^1 \\ &- (\frac{1}{n} \langle X_{B^c}^i, X_{B^c}^i \rangle)^{-1} \left[\frac{1}{n} \langle X_{B^c}^1, X_B^1 \rangle \Delta^1_B + \lambda_n \tilde{z}_{B^c}^1 \right]. \end{aligned}$$

Substituting back into (53a) yields

$$\begin{aligned} \frac{1}{n} \langle X_B^1, X_B^1 \rangle \Delta^1_B + \frac{1}{n} \langle X_B^1, X_{B^c}^1 \rangle \Delta^1_{B^c} \\ - \frac{1}{n} (X_B^1)^T w^1 + \lambda_n \tilde{z}_B^1 = 0 \end{aligned}$$

so that we obtain

$$M^1 \Delta^1_B = f^1 - \lambda_n \tilde{z}_B^1 \quad \text{and similarly,} \quad (60a)$$

$$M^2 \Delta^2_B = f^2 - \lambda_n \tilde{z}_B^2. \quad (60b)$$

Recall that by assumption, we have

$$S^1 \hat{\beta}_B^1 = |\hat{\beta}_B^1| = |\hat{\beta}_B^2| = S^2 \hat{\beta}_B^2$$

as well as $S \tilde{z}_B^1 + \tilde{S} \tilde{z}_B^2 = 1$. Subtracting $M^1 S^1 \bar{\beta}_B^1$ and $M^2 S^2 \bar{\beta}_B^2$ from (60a) and (60b)

$$M^1 S^1 (\Delta^1_B - \bar{\beta}_B^1) = f^1 - \lambda_n S^1 \tilde{z}_B^1 - M^1 S^1 \bar{\beta}_B^1 \quad (61a)$$

$$M^2 S^2 (\Delta^2_B - \bar{\beta}_B^2) = f^2 - \lambda_n S^2 \tilde{z}_B^2 - M^2 S^2 \bar{\beta}_B^2. \quad (61b)$$

Applying the fact that $S^1 (\Delta^1_B - \bar{\beta}_B^1) = S^2 (\Delta^2_B - \bar{\beta}_B^2)$, we obtain

$$\begin{aligned} (M^1 + M^2) S^1 (\Delta^1_B - \bar{\beta}_B^1) &= (f^1 + f^2) \\ &- \lambda_n \bar{1} - M^1 S^1 \bar{\beta}_B^1 - M^2 S^2 \bar{\beta}_B^2 \end{aligned}$$

where $\bar{1} \in \mathbb{R}^{\alpha s}$. Then solving for $S^1 (\Delta^1_B - \bar{\beta}_B^1)$ letting $S^1 \bar{\beta}_B^1 - S^2 \bar{\beta}_B^2 = \bar{B}_{\text{diff}}$ and substituting back into (61a), we obtain

$$\begin{aligned} \lambda_n S^1 \tilde{z}_B^1 &= M^1 [M^1 + M^2]^{-1} \lambda_n \bar{1} \\ &- [(M^1)^{-1} + (M^2)^{-1}]^{-1} (\bar{B}_{\text{diff}}) \\ &+ M^2 [M^1 + M^2]^{-1} f^1 \\ &- M^1 [M^1 + M^2]^{-1} f^2. \end{aligned} \quad (62)$$

E) *Proof of Lemma 8:* We have the decomposition $\frac{1}{\lambda_n}[(M^1)^{-1} + (M^2)^{-1}]^{-1} \bar{B}_{\text{diff}} = T_1 + T_2$, where

$$T_1 = \frac{1}{\lambda_n}([(M^1)^{-1} + (M^2)^{-1}]^{-1} - I/2) \bar{B}_{\text{diff}}$$

and $T_2 = \frac{\bar{B}_{\text{diff}}}{2\lambda_n}$. Under the assumptions of Theorem 3(a), we have $|\frac{\bar{B}_{\text{diff}}}{2\lambda_n}| \rightarrow 0$, hence, $T_2 \leq \epsilon/4$ for s large enough.

In order to bound T_1 , we note that

$$\|[(M^1)^{-1} + (M^2)^{-1}]^{-1} - I/2\|_2 = \mathcal{O}(\sqrt{s/n})$$

with probability greater than $1 - c_1 \exp(-c_2 n)$. (In particular, see the bound (50b) from Appendix C.) Consequently, we may decompose $[(M^1)^{-1} + (M^2)^{-1}]^{-1} - I/2$ as QDQ^T where Q and D are independent and Q is distributed uniformly over all orthogonal matrices, and $\|D\|_2 = \mathcal{O}(\sqrt{s/n})$. Using this decomposition, the following lemma, proved in Appendix D6, allows us to obtain the necessary control on the quantity $\|T_1\|_\infty$:

Lemma 9: Let $Q \in \mathbb{R}^{s \times s}$ be a matrix chosen uniformly at random from the space of all orthogonal matrices. Consider a second random matrix A , independent of Q . If $s/n = o(1)$, then for any fixed vector $x \in \mathbb{R}^s$ and fixed $\epsilon > 0$, we have

a) If $\|A\|_2 \leq \sqrt{\frac{s}{n}}$, then

$$\mathbb{P}\left[\|Q^T A Q x\|_\infty \geq \frac{\epsilon}{2}\right] \leq c_1 \exp\left(-c_2 \epsilon^2 \frac{n}{s\|x\|_2^2} + \log(s)\right).$$

b) If $\|A\|_2 \leq \frac{s}{n}$, then

$$\mathbb{P}\left[\|Q^T A Q x\|_\infty \geq \frac{\epsilon}{2}\right] \leq c_1 \exp\left(-c_2 \epsilon^2 \frac{n^2}{s^2\|x\|_2^2} + \log(s)\right).$$

With reference to the problem of bounding $\|T_1\|_\infty$, we may apply part (a) of this lemma with $A = D$ and $x = \frac{\bar{B}_{\text{diff}}}{2\lambda_n}$ to conclude that $\|T_1\|_\infty \leq \epsilon/2$ with high probability, thereby establishing the bound (57).

We now turn to proving the bound (58). We begin by decomposing the terms involved in this equation as

$$\begin{aligned} \frac{1}{\lambda_n} M^2 [M^1 + M^2]^{-1} f^1 &= \frac{f^1}{2\lambda_n} \\ &+ \frac{1}{\lambda_n} \left[M^2 [M^1 + M^2]^{-1} - \frac{I}{2} \right] f^1 \end{aligned}$$

and

$$\begin{aligned} \frac{1}{\lambda_n} M^1 [M^1 + M^2]^{-1} f^2 &= \frac{f^2}{2\lambda_n} \\ &+ \frac{1}{\lambda_n} \left[M^1 [M^1 + M^2]^{-1} - \frac{I}{2} \right] f^2. \end{aligned}$$

Now recall the form of $f^i \in \mathbb{R}^{\alpha s}$. Conditioned on $X_{B^c}^i$ and w^i , the variable $f^i/(2\lambda_n)$ is zero mean Gaussian with variance

$$\sigma^2 := \frac{1}{4} \langle \tilde{z}_{B^c}^i, \frac{1}{n} \left(\frac{1}{n} \langle X_{B^c}^i, X_{B^c}^i \rangle \right)^{-1} \tilde{z}_{B^c}^i \rangle I_{\alpha s} + \|w^i\|_2^2 / (n^2 \lambda_n^2) I_{\alpha s}.$$

By Lemmas 12 and 13 from Appendix E, as well as the fact that $\|\tilde{z}_s^i\|_2^2 = (1 - \alpha)s$, we have

$$\sigma^2 \leq \frac{1}{4} \langle \tilde{z}_{B^c}^i, \frac{1}{n} \left(\frac{1}{n} \langle X_{B^c}^i, X_{B^c}^i \rangle \right)^{-1} \tilde{z}_{B^c}^i \rangle + \|w^i\|_2^2 / (n^2 \lambda_n^2) \quad (63)$$

$$\leq \frac{1}{4} (1 - \alpha) \frac{s}{n} (1 + \delta) + \frac{1}{2} \frac{1}{n \lambda_n^2} \quad (64)$$

with probability greater than $1 - c_1 \exp(-c_2 n)$. Hence, by standard Gaussian tail bounds, the inequalities $\|f^1/(2\lambda_n)\|_\infty < \epsilon/4$ and $\|f^2/(2\lambda_n)\|_\infty < \epsilon/4$ both hold with probability greater than $1 - c_1 \exp(-\delta' \log(p - 2s))$.

Now to bound the first term in the decomposition we begin by diagonalizing $M^2 = Q^T D Q$. Note that Q is independent of X^1 and D and by symmetry $X_B^1 \stackrel{d}{=} Q X_B^1$. Following some algebra, we find that

$$\begin{aligned} \frac{1}{\lambda_n} \left[M^2 [M^1 + M^2]^{-1} - \frac{I}{2} \right] f^1 \\ = \frac{1}{\lambda_n} Q^T \left[D [Q M^1 Q^T + D]^{-1} - \frac{I}{2} \right] Q f^1. \end{aligned}$$

The random vector f^1 is independent of Q and $Q f^1$ is independent of Q by symmetry. Hence, the vector

$$v := \frac{1}{2} [2D(D + Q M^1 Q^T)^{-1} - I] Q \frac{1}{\lambda_n} f^1$$

is independent of Q . For a given constant c_3 , let us define the event

$$\mathcal{S} := \{ \|v\|_2^2 \leq c_3^2 \frac{s^2}{n} \left[\frac{s}{n} + \frac{1}{\lambda_n^2 n} \right] \}.$$

We can then write

$$\mathbb{P}[\|Q^T v\|_\infty \geq \epsilon] \leq \mathbb{P}[\|Q^T v\|_\infty \geq \epsilon \mid \mathcal{S}] + \mathbb{P}[\mathcal{S}^c]. \quad (65)$$

By Lemma 13, we have $\|D\|_2 = \mathcal{O}(1)$ with high probability. Moreover, we claim that

$$\| [2D(D + Q M^1 Q^T)^{-1} - I] \|_2 = \mathcal{O}(\sqrt{s/n}).$$

To establish this claim, note that we can write

$$\begin{aligned} [2D(D + Q M^1 Q^T)^{-1} - I] &= (D - I) \\ &+ 2D((D + Q M^1 Q^T)^{-1} - I/2). \end{aligned}$$

Note that $\|D - I\|_2 = \mathcal{O}(\sqrt{s/n})$ and moreover,

$$\|((D + Q M^1 Q^T)^{-1} - I/2)\|_2 = \mathcal{O}(\sqrt{s/n})$$

by (71a). Consequently, the sum of the two random matrices is also $\mathcal{O}(\sqrt{s/n})$, as claimed.

Recall the variance bound (64), and note that each component of f^1 is independent. Applying the concentration result from Lemma 12 for χ -squared random variables yields that

$$\|f^1\|_2^2 \leq \frac{1}{4} (1 + \delta) \frac{s^2}{n} + \frac{1}{2} \frac{s}{n \lambda_n^2}$$

with high probability. Hence, under the above conditions, we have

$$\begin{aligned} & \left\| \frac{1}{2} [2D(D + QM^1Q^T)^{-1} - I] Q \frac{1}{\lambda_n} f^1 \right\|_2^2 \\ & \leq \left\| \frac{1}{2} [2D(D + QM^1Q^T)^{-1} - I] \right\|_2^2 \left\| Q \frac{1}{\lambda_n} f^1 \right\|_2^2 \end{aligned}$$

which is in turn at most

$$c_3^2 \frac{s^2}{n} \left[\frac{s}{n} + \frac{1}{\lambda_n^2 n} \right]$$

with high probability. This bound implies that \mathcal{S} holds with high probability, or more precisely that $\mathbb{P}[\mathcal{S}^c] \leq c_1 \exp(-c_2 s)$.

It remains to control the first term in the decomposition (65). We do so using the following lemma, which is proved in Appendix D7.

Lemma 10: Let $Q \in \mathbb{R}^{m \times m}$ be a matrix chosen uniformly at random from the space of orthogonal matrices. Let $v \in \mathbb{R}^m$ be a random vector independent of Q , such that $\|v\|_2 \leq \xi$ with probability one. Then we have

$$\mathbb{P}[\|Q^T v\|_\infty \geq 2\xi \sqrt{\frac{\log m}{m}}] = o(1).$$

We now apply this lemma to the random vector v with $m = s$, and $\xi = c_3 \frac{s}{\sqrt{n}} \sqrt{\frac{s}{n} + \frac{1}{\lambda_n^2 n}}$. Note that

$$2\xi \sqrt{\frac{\log s}{s}} = 2c_3 \sqrt{\frac{\log s}{n}} \sqrt{\frac{s}{n} + \frac{1}{\lambda_n^2 n}} = o(1)$$

from which the second claim (58) in Lemma 8 follows.

Finally, we turn to proving the third claim (59) in Lemma 8. Following some algebra, we obtain

$$\begin{aligned} & \left\| M^1 [M^1 + M^2]^{-1} \vec{1} - \frac{1}{2} \vec{1} \right\|_\infty = \frac{1}{4} (M^1 - M^2) \vec{1} \\ & \quad + \frac{1}{2} (M^1 - M^2) (I/2 - (M^1 + M^2)^{-1}) \vec{1}. \end{aligned} \quad (66)$$

We diagonalize the matrix $M^1 = Q^T D Q$, where D is diagonal. Since the random matrix M^1 has a spherically symmetric distribution, the matrix Q has a uniform distribution over the space of orthogonal matrices and is independent of D . Using this decomposition, we can rewrite the second term in (66) as

$$\begin{aligned} & \frac{1}{2} Q^T (D - QM^2Q^T) \left(\frac{I}{2} - (D + QM^2Q^T)^{-1} \right) Q \vec{1} \\ & = Q^T R Q \vec{1} \end{aligned} \quad (67)$$

where

$$R := \frac{1}{4} (D - QM^2Q^T) (I - 2(D + QM^2Q^T)^{-1}).$$

We note that R is independent of Q , because D and M^2 are independent of Q . This independence follows from the spherical symmetry of M^2 and the fact that $M^2 \stackrel{d}{=} QM^2Q^T$.

Defining the event $\mathcal{T} := \{\|R\|_2 \leq 4s/n\}$, we claim that

$$\mathbb{P}[\mathcal{T}^c] \leq c_1 \exp(-c_2 n) \rightarrow 0. \quad (68)$$

In order to establish this claim, we note that submultiplicativity and triangle inequality imply that $\|R\|_2$ is upper bounded by

$$\begin{aligned} & \frac{1}{4} \|D - QM^2Q^T\|_2 \\ & \times \|(D + QM^2Q^T)/2 - I\|_2 \|2(D + QM^2Q^T)^{-1}\|_2 \end{aligned}$$

which is at most

$$2(\|D - I\|_2 + \|I - QM^2Q^T\|_2) \times \|(D + QM^2Q^T)/2 - I\|_2$$

where we have used the fact that

$$\|2(Q^T D Q + QM^2Q^T)^{-1}\|_2 \leq 2$$

with probability greater than $1 - c_1 \exp(-c_2 n)$, which follows from Lemma 13. Similarly, from this same result, we have

$$\begin{aligned} \|D - I\|_2 & = \mathcal{O}(\|I - QM^2Q^T\|_2) \\ & = \mathcal{O}(\|(D + QM^2Q^T)/2 - I\|_2) \\ & \leq 2\sqrt{\frac{s}{n}} \end{aligned}$$

so that the claim (68) follows.

Using the decomposition (67) and the tail bound (68), we have

$$\begin{aligned} \mathbb{P}[\|Q^T R Q \vec{1}\|_\infty \geq \epsilon] & = \mathbb{P}[\|Q^T R Q \vec{1}\|_\infty \geq \epsilon \mid \mathcal{T}] + \mathbb{P}[\mathcal{T}^c] \\ & \leq \mathcal{O}\left(\frac{1}{s}\right) + \mathcal{O}(\exp(-c(\epsilon)n)) \end{aligned}$$

where Lemma 9 (proved in Appendix D6) provides control on the first term in the inequality.

F) Proof of Lemma 9: We provide the proof for part a) of the Lemma and note that part b) is analogous. By union bound, we have

$$\mathbb{P}[\|Q^T A Q x\|_\infty \geq \epsilon] \leq s \max_{i=1, \dots, s} \mathbb{P}[|e_i^T Q^T A Q x| \geq \epsilon].$$

We will derive a bound on the probability $\mathbb{P}[|e_i^T Q^T A Q x| \geq \epsilon]$ that holds for all $e_i, i = 1, \dots, s$. We write $e_i^T Q^T A Q x = x_1 v_1^T A v_1 + v_1^T A v_2$, where v_1 denotes the first column of Q , and $v_2 = \sum_{k=2}^s x_k Q_k$ denotes the weighted sum of the remaining $(k - 1)$ columns of Q . Since Q is orthogonal, the vector v_1 has unit norm $\|v_1\|_2 = 1$, the vector v_2 is orthogonal to v_1 , and moreover $\|v_2\|_2^2 \leq \|x\|_\infty^2 (s - 1)$. Owing to the bound on the spectral norm of A , we have

$$|x_1 v_1^T A v_1| \leq \|x\|_\infty \sqrt{\frac{s}{n}}$$

which is less than $\epsilon/2$ for (s, n) sufficiently large, since $s/n = o(1)$.

We now turn to the second term. Note that conditioned on v_2 , the vector v_1 is uniformly distributed over an $(s-1)$ -dimensional unit sphere, contained within the subspace orthogonal to v_2 . Still conditioning on v_2 , consider the function $f(v_1) = v_1^T A v_2$. For any pair of vectors v_1, v'_1 on the unit sphere, we have

$$\begin{aligned} |f(v_1) - f(v'_1)|^2 &= |(v_1 - v'_1)^T A v_2|^2 \\ &\leq \|A\|_2^2 \|x\|_\infty^2 (s-1) \|v_1 - v'_1\|_2^2 \end{aligned}$$

which is at most $\|A\|_2^2 \|x\|_\infty^2 (s-1) [2(1 - \cos(d(v_1, v'_1)))]$, where $d = \arccos(v_1^T v'_1)$ is the geodesic distance. Using the inequality $\cos(d) \geq 1 - d^2/2$, valid for $d \in [0, \pi]$, and the assumption $\|A\|_2 \leq \sqrt{s/n}$, and taking square roots, we obtain

$$|f(v_1) - f(v'_1)| \leq \sqrt{\frac{s}{n}} \|x\|_\infty \sqrt{(s-1)} d(v_1, v'_1)$$

so that f is a Lipschitz function on the unit sphere (with dimension $s-1$) with constant $L = \|x\|_\infty \sqrt{\frac{s}{n}} \sqrt{(s-1)}$. Consequently, by Levy's theorem [13], for any $\epsilon > 0$, we have

$$\begin{aligned} \mathbb{P}[|f(v_1)| \geq \epsilon] &\leq 2 \exp\left(- (s-2) \frac{n}{\|x\|_\infty^2 s(s-1)} \epsilon^2\right) \\ &\leq 2 \exp\left(- c_1 \frac{n}{\|x\|_\infty^2 s} \epsilon^2\right). \end{aligned}$$

As a final side remark, we note that under the scaling of Theorem 3(b), we have $\frac{n}{s} \epsilon^2 - \log(s) \rightarrow \infty$ as $n \rightarrow \infty$, so that the probability in question vanishes.

G) Proof of Lemma 10: By union bound and symmetry of the distribution Q , for any $t > 0$, we have

$$\begin{aligned} \mathbb{P}[\|Q^T v\|_\infty \geq t] &\leq m \mathbb{P}[e_1^T Q^T v \geq t] \\ &= m \mathbb{P}[q_1^T v \geq t] \end{aligned}$$

where q_1 is the first column of Q . Note that q_1 is a random vector distributed uniformly over the unit sphere S^{m-1} in m dimensions. Viewing the vector $v \in \mathbb{R}^m$ as fixed, consider the function $f(q) = q^T v$ defined over S^{m-1} . As in Lemma 9, some calculation shows the Lipschitz constant of g over S^{m-1} is at most $L = \|v\|_2$. Applying Levy's theorem [13], we conclude that for any $\epsilon > 0$,

$$m \mathbb{P}[|f(q_1)| \geq t] \leq 2 \exp\left(- (m-1) \frac{t^2}{2\|v\|_2^2} + \log m\right).$$

Since $\|v\|_2 \leq \xi$ by assumption, it suffices to set $t = 2\xi \sqrt{\frac{\log m}{m}}$.

APPENDIX E SOME LARGE DEVIATION BOUNDS

In this appendix, we state some known large deviation bounds for the Gaussian variates, χ^2 -variates, as well as the eigenvalues of random matrices. The following Gaussian tail bound is standard.

Lemma 11: For a Gaussian variable $Z \sim N(0, \sigma^2)$, for all $t > 0$:

$$\mathbb{P}[|Z| \geq t] \leq 2 \exp\left(- \frac{t^2}{2\sigma^2}\right).$$

The following tail bounds on chi-squared variates are also useful.

Lemma 12: Let X be a χ -squared random variable with d degrees of freedom. Then for all $t > 0$, we have

$$\mathbb{P}\left[\frac{X}{d} \geq (1+t)^2\right] \leq \exp\left(-\frac{dt^2}{2}\right), \quad \text{and} \quad (69a)$$

$$\mathbb{P}\left[\frac{X}{d} \leq (1-2t)\right] \leq \exp(-dt^2). \quad (69b)$$

Proof: These tail bounds are immediate consequences of results due to Laurent and Massart [12], who prove that for all $x > 0$, we have

$$\mathbb{P}[X \geq x + (\sqrt{x} + \sqrt{d})^2] \leq \exp(-x), \quad \text{and} \quad (70a)$$

$$\mathbb{P}(X - d \leq -2\sqrt{dx}) \leq \exp(-x). \quad (70b)$$

Letting $x = dt^2/2$ in (70a), we have

$$\begin{aligned} \exp\left(-\frac{dt^2}{2}\right) &\geq \mathbb{P}\left[\frac{X}{d} \geq \sqrt{2}t + 1 + t^2\right] \\ &\geq \mathbb{P}\left[\frac{X}{d} \geq (1+t)^2\right] \end{aligned}$$

thereby establishing (69a). With the same choice of x , Equation (70b) implies the bound (69b) immediately. ■

Finally, the following type of large deviations bound on the eigenvalues of Gaussian random matrices is standard (e.g., [8]).

Lemma 13: Let $X \in \mathbb{R}^{n \times s}$ be a random matrix from the standard Gaussian ensemble (i.e., $X_{ij} \sim N(0, 1)$, i.i.d). Then with probability greater than $1 - c_1 \exp(-c_2 n)$, for any $\delta > 0$, its eigenspectrum satisfies the bounds

$$\begin{aligned} (1-\delta) \left[1 - \sqrt{\frac{s}{n}}\right]^2 &\leq \Lambda_{\min}\left(\frac{X^T X}{n}\right), \quad \text{and} \\ \Lambda_{\max}\left(\frac{X^T X}{n}\right) &\leq (1+\delta) \left[1 + \sqrt{\frac{s}{n}}\right]^2. \end{aligned}$$

Note that this lemma implies similar bounds for eigenvalues of the inverse:

$$\begin{aligned} \frac{1}{(1+\delta) \left[1 + \sqrt{\frac{s}{n}}\right]^2} &\leq \Lambda_{\min}\left(\left(\frac{X^T X}{n}\right)^{-1}\right), \quad \text{and} \\ \Lambda_{\max}\left(\left(\frac{X^T X}{n}\right)^{-1}\right) &\leq \frac{1}{(1-\delta) \left[1 - \sqrt{\frac{s}{n}}\right]^2}. \end{aligned}$$

From the above two sets of inequalities, we conclude for $s/n \leq 1$, we have with probability greater than $1 - c_1 \exp(-c_2 n)$

$$\left\| \frac{1}{n} X^T X - I \right\|_2 \leq 4\sqrt{\frac{s}{n}}, \quad \text{and} \quad (71a)$$

$$\left\| \left(\frac{1}{n} X^T X \right)^{-1} - I \right\|_2 \leq 4 \sqrt{\frac{s}{n}}. \quad (71b)$$

For random matrices where each row is distributed $N(0, \Sigma)$ and $\Lambda_{\min}(\Sigma) > C_{\min}$ and $\Lambda_{\max}(\Sigma) \leq C_{\max}$, we have

$$\left\| \frac{1}{n} X^T X - \Sigma \right\|_2 \leq 4 \lambda_{\max}(\Sigma) \sqrt{\frac{s}{n}}, \quad \text{and} \quad (72a)$$

$$\left\| \left(\frac{1}{n} X^T X \right)^{-1} - \Sigma^{-1} \right\|_2 \leq \frac{4}{\lambda_{\min}(\Sigma)} \sqrt{\frac{s}{n}}. \quad (72b)$$

ACKNOWLEDGMENT

The authors thank the anonymous referees for helpful comments, and D. Omidiran and G. Reeves for helpful discussions.

REFERENCES

[1] A. Argyriou, T. Evgeniou, and M. Pontil, "Multi-task feature learning," in *Proc. 20th Neural Information Processing Systems Conf.*, 2006.

[2] F. Bach, "Consistency of the group Lasso and multiple kernel learning," *J. Mach. Learn. Res.*, vol. 9, pp. 1179–1225, 2008.

[3] F. Bach, "Consistency of trace norm minimization," *J. Mach. Learn. Res.*, vol. 9, pp. 1019–1048, Jun. 2008.

[4] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA: Athena Scientific, 1995.

[5] P. Bickel, Y. Ritov, and A. Tsybakov, "Simultaneous analysis of Lasso and Dantzig selector," *Ann. Statist.*, vol. 37, no. 4, pp. 1705–1732, 2009.

[6] V. V. Buldygin and Y. V. Kozachenko, *Metric Characterization of Random Variables and Random Processes*. Providence, RI: AMS, 2000.

[7] S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.

[8] K. R. Davidson and S. J. Szarek, "Local operator theory, random matrices, and Banach spaces," in *Handbook of Banach Spaces*. Amsterdam, NL: Elsevier, 2001, vol. 1, pp. 317–336.

[9] J. Hiriart-Urruty and C. Lemaréchal, *Convex Analysis and Minimization Algorithms*. New York: Springer-Verlag, 1993, vol. 1.

[10] J. Huang and T. Zhang, "The Benefit of Group Sparsity Rutgers Univ. , Piscataway, NJ, 2009, Tech. Rep. arXiv:0901.2962.

[11] M. Jordan, Ed., *Learning in Graphical Models*. Cambridge, MA: MIT Press, 1999.

[12] B. Laurent and P. Massart, "Adaptive estimation of a quadratic functional by model selection," *Ann. Statist.*, vol. 28, no. 5, pp. 1303–1338, 1998.

[13] M. Ledoux, *The Concentration of Measure Phenomenon*. Providence, RI: AMS, 2001, Mathematical Surveys and Monographs.

[14] M. Ledoux and M. Talagrand, *Probability in Banach Spaces: Isoperimetry and Processes*. New York: Springer-Verlag, 1991.

[15] H. Liu, J. Lafferty, and L. Wasserman, "Nonparametric regression and classification with joint sparsity constraints," in *Neural Info. Proc. Systems (NIPS)*, Vancouver, BC, Canada, Dec. 2008, vol. 22.

[16] H. Liu and J. Zhang, "On $\ell_1 - \ell_q$ Regularized Regression Carnegie Mellon Univ., Pittsburgh, PA, 2008, Tech. Rep. arXiv:0802.1517v1.

[17] K. Lounici, M. Pontil, A. B. Tsybakov, and S. van de Geer, "Taking Advantage of Sparsity in Multi-Task Learning ETH, Zurich, 2009, Technical Report arXiv:0903.1468.

[18] S. G. Mallat, *A Wavelet Tour of Signal Processing*. New York: Academic, 1998.

[19] L. Meier, S. van de Geer, and P. Bühlmann, "The group lasso for logistic regression," *J. Royal Statist. Soc. B*, vol. 70, pp. 53–71, 2008.

[20] N. Meinshausen and P. Bühlmann, "High-dimensional graphs and variable selection with the lasso," *Ann. Statist.*, vol. 34, no. 2, pp. 1436–1462, 2006.

[21] N. Meinshausen and B. Yu, "Lasso-type recovery of sparse representations for high-dimensional data," *Ann. Statist.*, 2008, to be published.

[22] Y. Nardi and A. Rinaldo, "On the asymptotic properties of the group lasso estimator for linear models," *Electron. J. Statist.*, vol. 2, pp. 605–633, 2008.

[23] S. N. Negahban and M. J. Wainwright, "Estimation of (near) low-rank matrices with noise and high-dimensional scaling," *Ann. Statist.*, 2010, to be published.

[24] G. Obozinski, B. Taskar, and M. Jordan, "Joint Covariate Selection for Grouped Classification. Technical Report Dept. Statistics, Univ. California, Berkeley, Tech. Rep., 2007.

[25] G. Obozinski, M. J. Wainwright, and M. I. Jordan, "Union support recovery in high-dimensional multivariate regression," *Ann. Statist.*, 2010, to be published.

[26] G. Raskutti, M. J. Wainwright, and B. Yu, "Minimax Rates of Estimation for High-Dimensional Linear Regression over ℓ_q -Balls Dept. Statistics, Univ. California, Berkeley, 2009, Technical Report arXiv:0910.2042.

[27] P. Ravikumar, H. Liu, J. Lafferty, and L. Wasserman, "SpAM: Sparse additive models," *J. Royal Statist. Soc. B*, to be published.

[28] P. Ravikumar, M. J. Wainwright, and J. Lafferty, "High-dimensional Ising model selection using ℓ_1 -regularized logistic regression," *Ann. Statist.*, vol. 38, no. 3, pp. 1287–1319, 2010.

[29] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ: Princeton Univ. Press, 1970.

[30] E. P. Simoncelli, "Bayesian denoising of visual images in the wavelet domain," in *Bayesian Inference in Wavelet based Models*, P. Müller and B. Vidakovic, Eds. New York: Springer-Verlag, 1999, vol. 141, Lecture Notes in Statistics, ch. 18, pp. 291–308.

[31] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Royal Statist. Soc. B*, vol. 58, no. 1, pp. 267–288, 1996.

[32] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation," *Signal Process.*, vol. 86, pp. 572–602, Apr. 2006.

[33] B. Turlach, W. N. Venables, and S. J. Wright, "Simultaneous variable selection," *Technometrics*, vol. 27, pp. 349–363, 2005.

[34] S. van de Geer and P. Bühlmann, "On the conditions used to prove oracle results for the lasso," *Electron. J. Statist.*, vol. 3, pp. 1360–1392, 2009.

[35] M. J. Wainwright, "Information-theoretic bounds on sparsity recovery in the high-dimensional and noisy setting," *IEEE Trans. Inf. Theory*, vol. 55, pp. 5728–5741, Dec. 2009.

[36] M. J. Wainwright, "Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso)," *IEEE Trans. Inf. Theory*, vol. 55, pp. 2183–2202, May 2009.

[37] Y. Kim, J. Kim, and Y. Kim, "Blockwise sparse regression," *Statist. Sinica*, vol. 16, no. 2, 2006.

[38] P. Zhao, G. Rocha, and B. Yu, "Grouped and hierarchical model selection through composite absolute penalties," *Ann. Statist.*, vol. 37, no. 6A, pp. 3468–3497, 2009.

[39] P. Zhao and B. Yu, "On model selection consistency of Lasso," *J. Mach. Learn. Res.*, vol. 7, pp. 2541–2567, 2006.

Sahand N. Negahban received the B.S. degree in electrical engineering and computer sciences from the University of California at Berkeley, where he is currently pursuing the Ph.D. degree in the Department of Electrical Engineering and Computer Sciences (EECS). His research interests include machine learning, mathematical statistics, statistical signal processing, convex optimization, and information theory, with a current emphasis on high-dimensional statistical inference and optimization.

Mr. Negahban has been awarded a Vodafone Fellowship, a Vodafone Scholarship, and a Yahoo! KSC Award.

Martin J. Wainwright (SM'10) received the B.S. degree in mathematics from the University of Waterloo, Waterloo, ON, Canada, and the Ph.D. degree in electrical engineering and computer science (EECS) from the Massachusetts Institute of Technology (MIT), Cambridge.

He is currently a Professor at the University of California, Berkeley, with a joint appointment between the Department of Statistics and the Department of Electrical Engineering and Computer Sciences. His research interests include coding and information theory, machine learning, mathematical statistics, and statistical signal processing.

Dr. Wainwright has been awarded an Alfred P. Sloan Foundation Fellowship, an NSF CAREER Award, the George M. Sprowls Prize for his dissertation research (EECS Department, MIT), a Natural Sciences and Engineering Research Council of Canada 1967 Fellowship, an IEEE Signal Processing Society Best Paper Award in 2008, and several outstanding conference paper awards.