

ON ASYMPTOTICALLY OPTIMAL CONFIDENCE REGIONS AND TESTS FOR HIGH-DIMENSIONAL MODELS

BY SARA VAN DE GEER, PETER BÜHLMANN,
YA'ACOV RITOV¹ AND RUBEN DEZEURE

ETH Zürich, ETH Zürich, The Hebrew University of Jerusalem and ETH Zürich

We propose a general method for constructing confidence intervals and statistical tests for single or low-dimensional components of a large parameter vector in a high-dimensional model. It can be easily adjusted for multiplicity taking dependence among tests into account. For linear models, our method is essentially the same as in Zhang and Zhang [*J. R. Stat. Soc. Ser. B Stat. Methodol.* **76** (2014) 217–242]: we analyze its asymptotic properties and establish its asymptotic optimality in terms of semiparametric efficiency. Our method naturally extends to generalized linear models with convex loss functions. We develop the corresponding theory which includes a careful analysis for Gaussian, sub-Gaussian and bounded correlated designs.

1. Introduction. Much progress has been made over the last decade in high-dimensional statistics where the number of unknown parameters greatly exceeds sample size. The vast majority of work has been pursued for point estimation such as consistency for prediction [7, 21], oracle inequalities and estimation of a high-dimensional parameter [6, 11, 12, 24, 33, 34, 47, 51] or variable selection [17, 30, 49, 53]. Other references and exposition to a broad class of models can be found in [18] or [10].

Very little work has been done for constructing confidence intervals, statistical testing and assigning uncertainty in high-dimensional sparse models. A major difficulty of the problem is the fact that sparse estimators such as the lasso do not have a tractable limiting distribution: already in the low-dimensional setting, it depends on the unknown parameter [25] and the convergence to the limit is not uniform. Furthermore, bootstrap and even subsampling techniques are plagued by noncontinuity of limiting distributions. Nevertheless, in the low-dimensional setting, a modified bootstrap scheme has been proposed; [13] and [14] have recently proposed a residual based bootstrap scheme. They provide consistency guarantees for the high-dimensional setting; we consider this method in an empirical analysis in Section 4.

Received March 2013; revised January 2014.

¹Supported by the Forschungsinstitut für Mathematik (FIM) at ETH Zürich and from the Israel Science Foundation (ISF).

MSC2010 subject classifications. Primary 62J07; secondary 62J12, 62F25.

Key words and phrases. Central limit theorem, generalized linear model, lasso, linear model, multiple testing, semiparametric efficiency, sparsity.

Some approaches for quantifying uncertainty include the following. The work in [50] implicitly contains the idea of sample splitting and corresponding construction of p -values and confidence intervals, and the procedure has been improved by using multiple sample splitting and aggregation of dependent p -values from multiple sample splits [32]. Stability selection [31] and its modification [41] provides another route to estimate error measures for false positive selections in general high-dimensional settings. An alternative method for obtaining confidence sets is in the recent work [29]. From another and mainly theoretical perspective, the work in [24] presents necessary and sufficient conditions for recovery with the lasso $\hat{\beta}$ in terms of $\|\hat{\beta} - \beta^0\|_\infty$, where β^0 denotes the true parameter: bounds on the latter, which hold with probability at least say $1 - \alpha$, could be used in principle to construct (very) conservative confidence regions. At a theoretical level, the paper [35] derives confidence intervals in ℓ_2 for the case of two possible sparsity levels. Other recent work is discussed in Section 1.1 below.

We propose here a method which enjoys optimality properties when making assumptions on the sparsity and design matrix of the model. For a linear model, the procedure is as the one in [52] and closely related to the method in [23]. It is based on the lasso and is “inverting” the corresponding KKT conditions. This yields a nonsparse estimator which has a Gaussian (limiting) distribution. We show, within a sparse linear model setting, that the estimator is optimal in the sense that it reaches the semiparametric efficiency bound. The procedure can be used and is analyzed for high-dimensional sparse linear and generalized linear models and for regression problems with general convex (robust) loss functions.

1.1. Related work. Our work is closest to [52] who proposed the semiparametric approach for distributional inference in a high-dimensional linear model. We take here a slightly different view-point, namely by inverting the KKT conditions from the lasso, while relaxed projections are used in [52]. Furthermore, our paper extends the results in [52] by: (i) treating generalized linear models and general convex loss functions; (ii) for linear models, we give conditions under which the procedure achieves the semiparametric efficiency bound and our analysis allows for rather general Gaussian, sub-Gaussian and bounded design. A related approach as in [52] was proposed in [8] based on ridge regression which is clearly suboptimal and inefficient with a detection rate (statistical power) larger than $1/\sqrt{n}$.

Recently, and developed independently, the work in [23] provides a detailed analysis for linear models by considering a very similar procedure as in [52] and in our paper. They show that the detection limit is indeed in the $1/\sqrt{n}$ -range and they provide a minimax test result; furthermore, they present extensive simulation results indicating that the ridge-based method in [8] is overly conservative, which is in line with the theoretical results. Their optimality results are interesting and are complementary to the semiparametric optimality established here. Our results cover a substantially broader range of non-Gaussian designs in linear models, and we provide a rigorous analysis for correlated designs with covariance matrix $\Sigma \neq I$: the SDL-test in [23] assumes that Σ is known while we

carefully deal with the issue when Σ^{-1} has to be estimated (and arguing why, e.g., GLasso introduced in [19] is not good for our purpose). Another way and method to achieve distributional inference for high-dimensional models is given in [1] (claiming semiparametric efficiency). They use a two-stage procedure with a so-called post-double-selection as first and least squares estimation as second stage: as such, their methodology is radically different from ours. At the time of writing of this paper, [22] developed another modification which directly computes an approximate inverse of the Gram matrix. Moreover, [4] extended their approach to logistic regression and [2] to LAD estimation using an instrumental variable approach.

1.2. Organization of the paper. In Section 2, we consider the linear model and the lasso. We describe the desparsifying step in Section 2.1 where we need to use an approximately inverting matrix. A way to obtain this matrix is by applying the lasso with nodewise regression, as given in Section 2.1.1. Assuming Gaussian errors, we represent in Section 2.2 the de-sparsified lasso as sum of a normally distributed term and a remainder term. Section 2.3 considers the case of random design with i.i.d. covariables. We first prove for the case of Gaussian design and Gaussian errors that the remainder term is negligible. We then show in Section 2.3.1 that the results lead to honest asymptotic confidence intervals. Section 2.3.2 discusses the assumptions and Section 2.3.3 asymptotic efficiency. The case of non-Gaussian design and non-Gaussian errors is treated in Section 2.3.4.

In Section 3, we consider the extension to generalized linear models. We start out in Section 3.1 with the procedure, which is again desparsifying the ℓ_1 -penalized estimator. We again use the lasso with nodewise regression to obtain an approximate inverse of the matrix of second order derivatives. The computation of this approximate inverse is briefly described in Section 3.1.1. Section 3.2 presents asymptotic normality under high-level conditions. In Section 3.3, we investigate the consistency of the lasso with nodewise regression as estimator of the inverse of the matrix of second-order derivatives of the theoretical risk evaluated at the true unknown parameter β^0 . We also examine here the consistent estimation of the asymptotic variance. Section 3.3.1 gathers the results, leading to Theorem 3.3 for generalized linear models. Section 4 presents some empirical results. The proofs and theoretical material needed are given in Section 5, while the technical proofs of Section 2.3.3 (asymptotic efficiency) and Section 3.3 (nodewise regression for certain random matrices) are presented in the supplemental article [45].

2. High-dimensional linear models. Consider a high-dimensional linear model

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

with $n \times p$ design matrix $\mathbf{X} =: [X_1, \dots, X_p]$ ($n \times 1$ vectors X_j), $\varepsilon \sim \mathcal{N}_n(0, \sigma_\varepsilon^2 I)$ independent of \mathbf{X} and unknown regression $p \times 1$ vector β^0 . We note that non-Gaussian errors are not a principal difficulty, as discussed in Section 2.3.4.

Throughout the paper, we assume that $p > n$ and in the asymptotic results we require $\log(p)/n = o(1)$. We denote by $S_0 := \{j; \beta_j^0 \neq 0\}$ the active set of variables and its cardinality by $s_0 := |S_0|$.

Our main goal is a pointwise statistical inference for the components of the parameter vector β_j^0 ($j = 1, \dots, p$) but we also discuss simultaneous inference for parameters $\beta_G^0 := \{\beta_j^0; j \in G\}$ where $G \subseteq \{1, \dots, p\}$ is any group. To exemplify, we might want to test statistical hypotheses of the form $H_{0,j} : \beta_j^0 = 0$ or $H_{0,G} : \beta_j^0 = 0$ for all $j \in G$, and when pursuing many tests, we aim for an efficient multiple testing adjustment taking dependence into account and being less conservative than say the Bonferroni–Holm procedure.

2.1. *The method: Desparsifying the lasso.* The main idea is to invert the Karush–Kuhn–Tucker characterization of the lasso.

The lasso [43] is defined as

$$(2) \quad \hat{\beta} = \hat{\beta}(\lambda) := \arg \min_{\beta \in \mathbb{R}^p} (\|Y - \mathbf{X}\beta\|_2^2/n + 2\lambda\|\beta\|_1).$$

It is well known that the estimator in (2) fulfills the Karush–Kuhn–Tucker (KKT) conditions:

$$\begin{aligned} -\mathbf{X}^T(Y - \mathbf{X}\hat{\beta})/n + \lambda\hat{\kappa} &= 0, \\ \|\hat{\kappa}\|_\infty \leq 1 \quad \text{and} \quad \hat{\kappa}_j &= \text{sign}(\hat{\beta}_j) \quad \text{if } \hat{\beta}_j \neq 0. \end{aligned}$$

The vector $\hat{\kappa}$ is arising from the subdifferential of $\|\beta\|_1$: using the first equation we can always represent it as

$$(3) \quad \lambda\hat{\kappa} = \mathbf{X}^T(Y - \mathbf{X}\hat{\beta})/n.$$

The KKT conditions can be rewritten with the notation $\hat{\Sigma} = \mathbf{X}^T\mathbf{X}/n$:

$$\hat{\Sigma}(\hat{\beta} - \beta^0) + \lambda\hat{\kappa} = \mathbf{X}^T\varepsilon/n.$$

The idea is now to use a “relaxed form” of an inverse of $\hat{\Sigma}$. Suppose that $\hat{\Theta}$ is a reasonable approximation for such an inverse, then

$$(4) \quad \hat{\beta} - \beta^0 + \hat{\Theta}\lambda\hat{\kappa} = \hat{\Theta}\mathbf{X}^T\varepsilon/n - \Delta/\sqrt{n},$$

where

$$\Delta := \sqrt{n}(\hat{\Theta}\hat{\Sigma} - I)(\hat{\beta} - \beta^0).$$

We will show in Theorem 2.2 that Δ is asymptotically negligible under certain sparsity assumptions. This suggests the following estimator:

$$(5) \quad \hat{b} = \hat{\beta} + \hat{\Theta}\lambda\hat{\kappa} = \hat{\beta} + \hat{\Theta}\mathbf{X}^T(Y - \mathbf{X}\hat{\beta})/n,$$

using (3) in the second equation. This is essentially the same estimator as in [52] and it is of the same form as the SDL-procedure in [23], when plugging in the

estimate $\hat{\Theta}$ for the population quantity $\Theta := \Sigma^{-1}$ where Σ is the population inner product matrix. With (4), we immediately obtain an asymptotic pivot when Δ is negligible, as is justified in Theorem 2.2 below:

$$(6) \quad \sqrt{n}(\hat{b} - \beta^0) = W + o_{\mathbb{P}}(1), \quad W|\mathbf{X} \sim \mathcal{N}_p(0, \sigma_\varepsilon^2 \hat{\Theta} \hat{\Sigma} \hat{\Theta}^T).$$

An asymptotic pointwise confidence interval for β_j^0 is then given by

$$[\hat{b}_j - c(\alpha, n, \sigma_\varepsilon), \hat{b}_j + c(\alpha, n, \sigma_\varepsilon)],$$

$$c(\alpha, n, \sigma_\varepsilon) := \Phi^{-1}(1 - \alpha/2) \sigma_\varepsilon \sqrt{(\hat{\Theta} \hat{\Sigma} \hat{\Theta}^T)_{j,j}/n},$$

where $\Phi(\cdot)$ denotes the c.d.f. of $\mathcal{N}(0, 1)$. If σ_ε is unknown, we replace it by a consistent estimator.

2.1.1. *The lasso for nodewise regression.* A prime example to construct the approximate inverse $\hat{\Theta}$ is given by the lasso for the nodewise regression on the design \mathbf{X} : we use the lasso p times for each regression problem X_j versus \mathbf{X}_{-j} , where the latter is the design submatrix without the j th column. This method was introduced by [30]. We provide here a formulation suitable for our purposes. For each $j = 1, \dots, p$,

$$(7) \quad \hat{\gamma}_j := \arg \min_{\gamma \in \mathbb{R}^{p-1}} (\|X_j - \mathbf{X}_{-j}\gamma\|_2^2/n + 2\lambda_j \|\gamma\|_1),$$

with components of $\hat{\gamma}_j = \{\hat{\gamma}_{j,k}; k = 1, \dots, p, k \neq j\}$. Denote by

$$\hat{C} := \begin{pmatrix} 1 & -\hat{\gamma}_{1,2} & \cdots & -\hat{\gamma}_{1,p} \\ -\hat{\gamma}_{2,1} & 1 & \cdots & -\hat{\gamma}_{2,p} \\ \vdots & \vdots & \ddots & \vdots \\ -\hat{\gamma}_{p,1} & -\hat{\gamma}_{p,2} & \cdots & 1 \end{pmatrix}$$

and write

$$\hat{T}^2 := \text{diag}(\hat{\tau}_1^2, \dots, \hat{\tau}_p^2),$$

where for $j = 1, \dots, p$

$$\hat{\tau}_j^2 := \|X_j - \mathbf{X}_{-j}\hat{\gamma}_j\|_2^2/n + \lambda_j \|\hat{\gamma}_j\|_1.$$

Then define

$$(8) \quad \hat{\Theta}_{\text{Lasso}} := \hat{T}^{-2} \hat{C}.$$

Note that although $\hat{\Sigma}$ is self-adjoint, its relaxed inverse $\hat{\Theta}_{\text{Lasso}}$ is not. In the sequel, we denote by

$$(9) \quad \hat{b}_{\text{Lasso}} = \text{the estimator in (5) with } \hat{\Theta} \text{ the nodewise lasso from (8).}$$

The estimator \hat{b}_{Lasso} corresponds to the proposal in [52].

Let the j th row of $\hat{\Theta}$ be denoted by $\hat{\Theta}_j$ (as a $1 \times p$ vector) and analogously for \hat{C}_j . Then $\hat{\Theta}_{\text{Lasso},j} = \hat{C}_j/\hat{\tau}_j^2$.

The KKT conditions for the nodewise lasso (7) imply that

$$\hat{\tau}_j^2 = (X_j - \mathbf{X}_{-j}\hat{\gamma}_j)^T X_j/n$$

so that

$$X_j^T \mathbf{X} \hat{\Theta}_{\text{Lasso},j}^T / n = 1.$$

These KKT conditions also imply that

$$\|\mathbf{X}_{-j}^T \mathbf{X} \hat{\Theta}_{\text{Lasso},j}^T\|_\infty / n \leq \lambda_j / \hat{\tau}_j^2.$$

Hence, for the choice $\hat{\Theta}_j = \hat{\Theta}_{\text{Lasso},j}$ we have

$$(10) \quad \|\hat{\Sigma} \hat{\Theta}_j^T - e_j\|_\infty \leq \lambda_j / \hat{\tau}_j^2,$$

where e_j is the j th unit column vector. We call this the extended KKT conditions.

We note that using, for example, the GLasso estimator of [19] for $\hat{\Theta}$ may not be optimal because with this choice a bound for $\|\hat{\Sigma} \hat{\Theta}_j^T - e_j\|_\infty$ is not readily available and this means we cannot directly derive desirable componentwise properties of the estimator \hat{b} in (5) as established in Section 2.3. The same can be said about a ridge type of estimator for $\hat{\Theta}$, a choice analyzed in [8]. We note that in (10) the bound depends on $\hat{\tau}_j^2$ and is in this sense not under control.

In [22], a program is proposed which gives an approximate inverse $\hat{\Theta}$ such that $\|\hat{\Sigma} \hat{\Theta}_j^T - e_j\|_\infty$ is bounded by a prescribed constant. We will show in Remark 2.1 that a bound of the form (10) with λ_j proportional (by a prescribed constant) to $\tilde{\tau}_j := \|X_j - \mathbf{X}_{-j}\hat{\gamma}_j\|_2/\sqrt{n}$ gives the appropriate normalization when considering a Studentized version of the estimator \hat{b}_{Lasso} .

2.2. *Theoretical result for fixed design.* We provide here a first result for fixed design \mathbf{X} . A crucial identifiability assumption on the design is the so-called compatibility condition [44]. To describe this condition, we introduce the following notation. For a $p \times 1$ vector β and a subset $S \subseteq \{1, \dots, p\}$, define β_S by

$$\beta_{S,j} := \beta_j 1\{j \in S\}, \quad j = 1, \dots, p.$$

Thus, β_S has zeroes for the components outside the set S . The compatibility condition for $\hat{\Sigma}$ requires a positive constant $\phi_0 > 0$ such that for all β satisfying $\|\beta_{S_0^c}\|_1 \leq 3\|\beta_{S_0}\|_1$ (the constant 3 is relatively arbitrary, it depends on the choice of the tuning parameter λ)

$$\|\beta_{S_0}\|_1^2 \leq s_0 \beta^T \hat{\Sigma} \beta / \phi_0^2.$$

The value ϕ_0^2 is called the compatibility constant.

We make the following assumption:

(A1) The compatibility condition holds for $\hat{\Sigma}$ with compatibility constant $\phi_0^2 > 0$. Furthermore, $\max_j \hat{\Sigma}_{j,j} \leq M^2$ for some $0 < M < \infty$.

The assumption (A1) is briefly discussed in Section 2.3.2. We then obtain the following result where we use the notation $\|A\|_\infty := \max_{j,k} |A_{j,k}|$ for the element-wise sup-norm for a matrix A .

THEOREM 2.1. *Consider the linear model in (1) with Gaussian error $\varepsilon \sim \mathcal{N}_n(0, \sigma_\varepsilon^2 I)$, and assume (A1). Let $t > 0$ be arbitrary. When using the lasso in (2) with $\lambda \geq 2M\sigma_\varepsilon\sqrt{2(t^2 + \log(p))}/n$ and the lasso for nodewise regression in (8) we have:*

$$\begin{aligned} \sqrt{n}(\hat{b}_{\text{Lasso}} - \beta^0) &= W + \Delta, \\ W &= \hat{\Theta}_{\text{Lasso}} \mathbf{X}^T \varepsilon / \sqrt{n} \sim \mathcal{N}_n(0, \sigma_\varepsilon^2 \hat{\Omega}), \quad \hat{\Omega} := \hat{\Theta} \hat{\Sigma} \hat{\Theta}^T, \\ \mathbb{P} \left[\|\Delta\|_\infty \geq 8\sqrt{n} \left(\max_j \frac{\lambda_j}{\hat{\tau}_j^2} \right) \frac{\lambda s_0}{\phi_0^2} \right] &\leq 2 \exp[-t^2]. \end{aligned}$$

A proof is given in Section 5.2.

REMARK 2.1. In practice, one will use a Studentized version of \hat{b}_{Lasso} . Let us consider the j th component. One may verify that $\hat{\Omega}_{j,j} = \tilde{\tau}_j^2 / \hat{\tau}_j^4$, where $\tilde{\tau}_j^2$ is the residual sum of squares $\tilde{\tau}_j^2 := \|X_j - \mathbf{X}_{-j} \hat{\gamma}\|_2^2 / n$. Under the conditions of Theorem 2.1,

$$\begin{aligned} \frac{\sqrt{n}(\hat{b}_{\text{Lasso},j} - \beta_j^0)}{\hat{\Omega}_{j,j}^{1/2} \sigma_\varepsilon} &= V_j + \tilde{\Delta}_j, \\ V_j &\sim \mathcal{N}(0, 1), \\ \mathbb{P} \left[|\tilde{\Delta}_j| \geq 8\sqrt{n} \left(\frac{\lambda_j}{\tilde{\tau}_j} \right) \left(\frac{\lambda}{\sigma_\varepsilon} \right) \frac{s_0}{\phi_0^2} \right] &\leq 2 \exp[-t^2]. \end{aligned}$$

A Studentized version has the unknown variance σ_ε^2 replaced by a consistent estimator, $\tilde{\sigma}_\varepsilon^2$ say. Thus, the bound for $\tilde{\Delta}_j$ depends on the normalized tuning parameters $\lambda_j / \tilde{\tau}_j$ and $\lambda / \tilde{\sigma}_\varepsilon$. In other words, the standardized estimator is standard normal with a standardized remainder term. The appropriate choice for λ makes $\lambda / \tilde{\sigma}_\varepsilon$ scale independent. Scale independence for $\lambda_j / \tilde{\tau}_j$ can be shown under certain conditions, as we will do in the next subsection. Scale independent regularization can also be achieved numerically by using the square-root lasso introduced in [3], giving an approximate inverse, $\hat{\Theta}_{\sqrt{\text{Lasso}}}$ say, as alternative for $\hat{\Theta}_{\text{Lasso}}$. Most of the theory that we develop in the coming subsections goes through with the choice $\hat{\Theta}_{\sqrt{\text{Lasso}}}$ as well. To avoid digressions, we do not elaborate on this.

Theorem 2.2 presents conditions that ensure that $\hat{\tau}_j$ as well as $1/\hat{\tau}_j^2$ are asymptotically bounded uniformly in j (see Lemma 5.3 in Section 5) and that asymptotically one may choose λ as well as each λ_j of order $\sqrt{\log(p)/n}$. Then, if the sparsity s_0 satisfies $s_0 = o(\sqrt{n}/\log p)$, the correct normalization factor for \hat{b}_{Lasso} is \sqrt{n} (as used in the above theorem) and the error term $\|\Delta\|_\infty = o_{\mathbb{P}}(1)$ is negligible. The details are discussed next.

2.3. *Random design and optimality.* In order to further analyze the error term Δ from Theorem 2.1, we consider an asymptotic framework with random design. It uses a scheme where $p = p_n \geq n \rightarrow \infty$ in model (1), and thus, $Y = Y_n, \mathbf{X} = \mathbf{X}_n, \beta^0 = \beta_n^0$ and $\sigma_\varepsilon^2 = \sigma_{\varepsilon,n}^2$ are all (potentially) depending on n . In the sequel, we usually suppress the index n . We make the following assumption.

(A2) The rows of \mathbf{X} are i.i.d. realizations from a Gaussian distribution whose p -dimensional inner product matrix Σ has strictly positive smallest eigenvalue Λ_{\min}^2 satisfying $1/\Lambda_{\min}^2 = \mathcal{O}(1)$. Furthermore, $\max_j \Sigma_{j,j} = \mathcal{O}(1)$.

The Gaussian assumption is relaxed in Section 2.3.4.

We will assume below sparsity with respect to rows of $\Theta := \Sigma^{-1}$ and define

$$s_j := |\{k \neq j : \Theta_{j,k} \neq 0\}|.$$

Recall the notation $\hat{\Omega} := \hat{\Theta}_{\text{Lasso}} \hat{\Sigma} \hat{\Theta}_{\text{Lasso}}^T$. We then have the following main result.

THEOREM 2.2. *Consider the linear model (1) with Gaussian error $\varepsilon \sim \mathcal{N}_n(0, \sigma_\varepsilon^2 I)$ where $\sigma_\varepsilon^2 = \mathcal{O}(1)$. Assume (A2) and the sparsity assumptions $s_0 = o(\sqrt{n}/\log(p))$ and $\max_j s_j = o(n/\log(p))$. Consider a suitable choice of the regularization parameters $\lambda \asymp \sqrt{\log(p)/n}$ for the lasso in (2) and $\lambda_j \asymp \sqrt{\log(p)/n}$ uniformly in j for the lasso for nodewise regression in (8). Then*

$$\begin{aligned} \sqrt{n}(\hat{b}_{\text{Lasso}} - \beta^0) &= W + \Delta, \\ W|\mathbf{X} &\sim \mathcal{N}_p(0, \sigma_\varepsilon^2 \hat{\Omega}), \\ \|\Delta\|_\infty &= o_{\mathbb{P}}(1). \end{aligned}$$

Furthermore, $\|\hat{\Omega} - \Sigma^{-1}\|_\infty = o_{\mathbb{P}}(1)$.

A proof is given in Section 5.5.

Theorem 2.2 has various implications. For a one-dimensional component β_j^0 (with j fixed), we obtain for all $z \in \mathbb{R}$

$$(11) \quad \mathbb{P} \left[\frac{\sqrt{n}(\hat{b}_{\text{Lasso};j} - \beta_j^0)}{\sigma_\varepsilon \sqrt{\hat{\Omega}_{j,j}}} \leq z \mid \mathbf{X} \right] - \Phi(z) = o_{\mathbb{P}}(1).$$

Furthermore, for any fixed group $G \subseteq \{1, \dots, p\}$ which is potentially large, we have that for all $z \in \mathbb{R}$

$$\mathbb{P} \left[\max_{j \in G} \frac{\sqrt{n} |\hat{b}_{\text{Lasso};j} - \beta_j^0|}{\sigma_\varepsilon \sqrt{\hat{\Omega}_{j,j}}} \leq z \mid \mathbf{X} \right] - \mathbb{P} \left[\max_{j \in G} \frac{|W_j|}{\sigma_\varepsilon \sqrt{\hat{\Omega}_{j,j}}} \leq z \mid \mathbf{X} \right] = o_{\mathbb{P}}(1).$$

Therefore, conditionally on \mathbf{X} , the asymptotic distribution of

$$\max_{j \in G} n |\hat{b}_{\text{Lasso};j}|^2 / \sigma_\varepsilon^2 \hat{\Omega}_{j,j}$$

under the null-hypothesis $H_{0,G}; \beta_j^0 = 0 \forall j \in G$ is asymptotically equal to the maximum of dependent $\chi^2(1)$ variables $\max_{j \in G} |W_j|^2 / \sigma_\varepsilon^2 \hat{\Omega}_{j,j}$ whose distribution can be easily simulated since $\hat{\Omega}$ is known. The unknown σ_ε^2 may be replaced by a consistent estimator. For example, the scaled lasso [42] yields a consistent estimator for σ_ε^2 under the assumptions made for Theorem 2.2.

Theorem 2.2 is extended in Theorem 2.4 to the case of non-Gaussian errors and non-Gaussian design.

2.3.1. *Uniform convergence.* The statements of Theorem 2.2 also hold in a uniform sense, and thus the confidence intervals and tests based on these statements are honest [27]. In particular, the estimator \hat{b}_{Lasso} does not suffer the problems arising from the nonuniformity of limit theory for penalized estimators (described in, e.g., [37] or [38]). Such uniformity problems are also taken care of in [5] using an alternative procedure. However, using $\hat{b}_{\text{Lasso}} - \beta^0$ as pivot is asymptotically less conservative in general.

We consider the set of parameters

$$\mathcal{B}(s) = \{ \beta \in \mathbb{R}^p; |\{j : \beta_j \neq 0\}| \leq s \}.$$

We let \mathbb{P}_{β^0} be the distribution of the data under the linear model (1). Then the following for \hat{b}_{Lasso} in (9) holds.

COROLLARY 2.1. *Consider the linear model (1) with Gaussian error $\varepsilon \sim \mathcal{N}_n(0, \sigma_\varepsilon^2 I)$ where $\sigma_\varepsilon^2 = \mathcal{O}(1)$. Assume (A2) and the sparsity assumption $\beta^0 \in \mathcal{B}(s_0)$ with $s_0 = o(\sqrt{n}/\log(p))$. Suppose that $\max_j s_j = o(n/\log(p))$. Then, when using suitable choices with $\lambda \asymp \sqrt{\log(p)/n}$ for the lasso in (2), and $\lambda_j \asymp \sqrt{\log(p)/n}$ uniformly j for the lasso for nodewise regression in (8)*

$$\begin{aligned} \sqrt{n}(\hat{b}_{\text{Lasso}} - \beta^0) &= W + \Delta, \\ W \mid \mathbf{X} &\sim \mathcal{N}_p(0, \sigma_\varepsilon^2 \hat{\Omega}), \quad \hat{\Omega} := \hat{\Theta} \hat{\Sigma} \hat{\Theta}^T, \\ \|\Delta\|_\infty &= o_{\mathbb{P}_{\beta^0}}(1) \quad \text{uniformly in } \beta^0 \in \mathcal{B}(s_0). \end{aligned}$$

Moreover, since $\hat{\Omega}$ does not depend on β^0 we have as in Theorem 2.2, $\|\hat{\Omega} - \Sigma^{-1}\|_\infty = o_{\mathbb{P}}(1)$.

The proof is exactly the same as for Theorem 2.2 by simply noting that $\|\hat{\beta} - \beta^0\|_1 = \mathcal{O}_{\mathbb{P}_{\beta^0}}(s_0\sqrt{\log(p)/n})$ uniformly in $\beta^0 \in \mathcal{B}(s_0)$ [with high probability, the compatibility constant is bounded away from zero uniformly in all subsets S_0 with $|S_0| = o(\sqrt{n/\log(p)})$].

Corollary 2.1 implies that for $j \in \{1, \dots, p\}$ and all $z \in \mathbb{R}$,

$$\sup_{\beta^0 \in \mathcal{B}(s_0)} \left| \mathbb{P}_{\beta^0} \left[\frac{\sqrt{n}(\hat{b}_{\text{Lasso};j} - \beta_j^0)}{\sigma_\varepsilon \sqrt{\hat{\Omega}_{j,j}}} \leq z \mid \mathbf{X} \right] - \Phi(z) \right| = o_{\mathbb{P}}(1).$$

Thus one can construct p -values for each component. Based on many single p -values, we can use standard procedures for multiple testing adjustment to control for various type I error measures. The representation from Theorems 2.1 or 2.2 with $\|\Delta\|_\infty$ being sufficiently small allows to construct a multiple testing adjustment which takes the dependence in terms of the covariance $\hat{\Omega}$ (see Theorem 2.2) into account: the exact procedure is described in [8]. Especially when having strong dependence among the p -values, the method is much less conservative than the Bonferroni–Holm procedure for strongly controlling the family-wise error rate.

2.3.2. *Discussion of the assumptions.* The compatibility condition in (A1) is weaker than many others which have been proposed such as assumptions on restricted or sparse eigenvalues [48]: a relaxation by a constant factor has recently been given in [42]. Assumption (A2) is rather weak in the sense that it concerns the population inner product matrix. It implies condition (A1) with $1/\phi_0 = \mathcal{O}(1)$ (see Lemma 5.2) and $M = \mathcal{O}(1)$.

Regarding the sparsity assumption for s_0 in Theorem 2.1, our technique crucially uses the ℓ_1 -norm bound $\|\hat{\beta} - \beta^0\|_1 = \mathcal{O}_{\mathbb{P}}(s_0\sqrt{\log(p)/n})$; see Lemma 5.1. In order that this ℓ_1 -norm converges to zero, the sparsity constraint $s_0 = o(\sqrt{n/\log(p)})$ is usually required. Our sparsity assumption is slightly stricter by the factor $1/\sqrt{\log(p)}$ (because the normalization factor is \sqrt{n}), namely $s_0 = o(\sqrt{n}/\log(p))$.

2.3.3. *Optimality and semiparametric efficiency.* Corollary 2.1 establishes, in fact, that for any j , $\hat{b}_{\text{Lasso};j}$ is an asymptotically efficient estimator of β_j^0 , in the sense that it is asymptotically normal with asymptotic variance converging, as $n \rightarrow \infty$ to the variance of the best estimator. Consider, the *one*-dimensional sub-model,

$$(12) \quad Y = \beta_j^0(X_j - \mathbf{X}_{-j}\gamma_j) + \mathbf{X}_{-j}(\beta_{-j}^0 + \beta_j^0\mathbf{X}_{-j}\gamma_j) + \varepsilon,$$

where $X_j - \mathbf{X}_{-j}\gamma_j$ is the projection in $L_2(\mathbb{P})$ of X_j to the subspace orthogonal to \mathbf{X}_{-j} . Clearly, this is a linear submodel of the general model (1), passing through the true point. The Gauss–Markov theorem argues that the best variance of an unbiased estimator of β_j^0 in (12) is given by $\sigma_\varepsilon^2/(n \text{Var}(X_{1,j} - \mathbf{X}_{1,-j}\gamma_j))$.

Corollary 2.1 shows that $\sigma_\varepsilon^2 / \text{Var}(X_{1,j} - \mathbf{X}_{1,-j}\gamma_j)$ this is the asymptotic variance of $\sqrt{n}(\hat{b}_{\text{Lasso},j} - \beta_j^0)$. Thus, $\sqrt{n}(\hat{b}_{\text{Lasso},j} - \beta_j^0)$ is asymptotically normal, with the variance of the best possible unbiased estimator. Note, that any regular estimator (regular at least on parametric sub-models) must be asymptotically unbiased.

The main difference between this and most of the other papers on complex models is that usually the lasso is considered as solving a nonparametric model with parameter whose dimension p is increasing to infinity, while we consider the problem as a semiparametric model in which we concentrate on a low-dimensional model of interest, for example, β_1^0 , while the rest of the parameters, β_{-j}^0 , are considered as nuisance parameters. That is, we consider the problem as a semiparametric one.

In the rest of this discussion, we put the model in a standard semiparametric framework in which there is an infinite-dimensional population model. Without loss of generality, the parameter of interest is β_1^0 , that is, the first component (extension to more than one but finitely many parameters of interest is straightforward). Consider the random design model where the sequence $\{(Y_i, X_{i,1}, Z_i)\}_{i=1}^\infty$ is i.i.d. with

$$(13) \quad Y_1 = \beta_1^0 X_{1,1} + K(Z_1) + \varepsilon_1, \quad \varepsilon_1 \sim \mathcal{N}(0, \sigma_\varepsilon^2),$$

where $\beta_1^0 \in \mathbb{R}$ is an unknown parameter and $K(\cdot)$ is an unknown function. When observing $\{(Y_i, X_{i,1}, Z_i)\}_{i=1}^n$ this is the partially linear regression model, where \sqrt{n} -consistency for the parametric part β_1^0 can be achieved [40]. We observe the i.i.d. sequence $\{(Y_i, X_{i,1}, \{X_{i,j}^n\}_{j=2}^{p_n})\}_{i=1}^n$ such that

$$(14) \quad \begin{aligned} Y_1 &= \beta_1^0 X_{1,1} + \sum_{j=2}^{p_n} \beta_j^n X_{1,j}^n + \varepsilon_1^n, \\ &\varepsilon_1^n \text{ independent of } X_{1,1}, X_{1,2}^n, \dots, X_{1,p_n}^n, \\ \mathbb{E} \left[K(Z_1) - \sum_{j \in S_n \cap \{2, \dots, p_n\}} \beta_j^n X_{1,j}^n \right]^2 &\rightarrow 0, \quad |S_n| = o(\sqrt{n} / \log(p)), \\ \mathbb{E} \left[\mathbb{E}[X_{1,1} | Z_1] - \sum_{j=2}^{p_n} \gamma_{1,j}^n X_{1,j}^n \right]^2 &\rightarrow 0, \\ \left(K(Z_1) - \sum_{j=2}^{p_n} \beta_j^n X_{1,j}^n \right) \left(\mathbb{E}[X_{1,1} | Z_1] - \sum_{j=2}^{p_n} \gamma_{1,j}^n X_{1,j}^n \right) &= o_{\mathbb{P}}(n^{-1/2}). \end{aligned}$$

THEOREM 2.3. *Suppose (14) and the conditions of Theorem 2.2 are satisfied, then*

$$\hat{b}_{\text{Lasso};1} = \beta_1^0 + \frac{1}{n} \sum_{i=1}^n (X_{i,1} - \mathbb{E}[X_{i,1} | Z_i]) \varepsilon_i + o_{\mathbb{P}}(n^{-1/2}).$$

In particular, the limiting variance of $\sqrt{n}(\hat{b}_{\text{Lasso};1} - \beta_1^0)$ reaches the information bound $\sigma_\varepsilon^2 / \mathbb{E}(X_{1,1} - \mathbb{E}[X_{1,1}|Z_1])^2$. Furthermore, $\hat{b}_{\text{Lasso};1}$ is regular at the one-dimensional parametric sub-model with component β_1^0 , and hence, $\hat{b}_{\text{Lasso};1}$ is asymptotically efficient for estimating β_1^0 .

A proof is given in the supplemental article [45].

As a concrete example consider the following situation:

$$(15) \quad \begin{aligned} K(Z_1) &= \sum_{j=2}^{\infty} \beta_j^0 X_{1,j}, \\ X_{1,j}^n &\equiv X_{1,j} \quad \forall j = 1, \dots, p_n, \end{aligned}$$

where

$$\begin{aligned} \beta^0 \in \mathcal{B}(s_0) &:= \{(\beta_j)_{j \in \mathbb{N}}; |\{j : \beta_j \neq 0\}| \leq s_0\}, \\ s_0 < \infty \text{ fixed, } \mathbb{E}[X_{1,j}] &= 0 \ \forall j \text{ and } \max_{j \in \mathbb{N}} \text{var}(X_{1,j}) < \infty, \end{aligned}$$

$$\min_{S \subset \mathbb{N}} \Lambda_{\min}^2(S) > 0,$$

where $\Lambda_{\min}^2(S)$ is the smallest eigenvalue of the covariance matrix of $\{X_{1,j} : j \in S\}$,

$$|\{k : \gamma_{1,k} \neq 0\}| < \infty,$$

$$\text{where } \gamma_1 := \arg \min_{\gamma \in \mathbb{R}_{\infty}} \mathbb{E}_{\mathbb{P}} \left[\left(X_{1,1} - \sum_{k=2}^{\infty} \gamma_k X_{1,k} \right)^2 \right].$$

Note that the assumption about the minimal eigenvalues $\{\Lambda_{\min}^2(S) : S \subset \mathbb{N}\}$ is equivalent to saying that $\{X_{1,j}\}_{j \in \mathbb{N}}$ has a positive definite covariance function.

LEMMA 2.1. *Condition (14) is satisfied in the above example.*

A proof of this lemma is given in the supplemental article [45].

2.3.4. *Non-Gaussian design and non-Gaussian errors.* We extend Theorem 2.2 to allow for non-Gaussian designs and non-Gaussian errors. Besides covering a broader range for linear models, the result is important for the treatment of generalized linear models in Section 3.

Consider a random design matrix \mathbf{X} with i.i.d. rows having inner product matrix Σ with its inverse (assumed to exist) $\Theta = \Sigma^{-1}$. For $j = 1, \dots, p$, denote by $\gamma_j := \arg \min_{\gamma \in \mathbb{R}^{p-1}} \mathbb{E}[\|X_j - \mathbf{X}_{-j}\gamma\|_2^2]$. Define the error $\eta_j := X_j - \mathbf{X}_{-j}\gamma_j$ with variance $\tau_j^2 = \mathbb{E}[\|\eta_j\|_2^2/n] = 1/\Theta_{j,j}$, $j = 1, \dots, p$. We make the following assumptions:

(B1) The design \mathbf{X} has either i.i.d. sub-Gaussian rows (i.e., $\max_i \sup_{\|v\|_2 \leq 1} \mathbb{E} \exp[|\sum_{j=1}^p v_j \mathbf{X}_{i,j}|^2 / L^2] = \mathcal{O}(1)$ for some fixed constant $L > 0$) or i.i.d. rows and for some $K \geq 1$, $\|\mathbf{X}\|_\infty = \max_{i,j} |\mathbf{X}_{i,j}| = \mathcal{O}(K)$. The latter we call the bounded case. The strongly bounded case assumes in addition that $\max_j \|\mathbf{X}_{-j} \gamma_j\|_\infty = \mathcal{O}(K)$.

(B2) In the sub-Gaussian case, it holds that $\max_j \sqrt{s_j \log(p)/n} = o(1)$. In the (strongly) bounded case, we assume that $\max_j K^2 s_j \sqrt{\log(p)/n} = o(1)$.

(B3) The smallest eigenvalue Λ_{\min}^2 of Σ is strictly positive and $1/\Lambda_{\min}^2 = \mathcal{O}(1)$. Moreover, $\max_j \Sigma_{j,j} = \mathcal{O}(1)$.

(B4) In the bounded case, it holds that $\max_j \mathbb{E} \eta_{1,j}^4 = \mathcal{O}(K^4)$.

We note that the strongly bounded case in (B1) follows from the bounded case if $\|\gamma_j\|_1 = \mathcal{O}(1)$. Assumption (B2) is a standard sparsity assumption for Θ . Finally, assumption (B3) implies that $\|\Theta_j\|_2 \leq \Lambda_{\min}^{-2} = \mathcal{O}(1)$ uniformly in j so that in particular $\tau_j^2 = 1/\Theta_{j,j}$ stays away from zero. Note that (B3) also implies $\tau_j^2 \leq \Sigma_{j,j} = \mathcal{O}(1)$ uniformly in j .

To streamline the statement of the results, we write $K_0 = 1$ in the sub-Gaussian case and $K_0 = K$ in the (strongly) bounded case.

THEOREM 2.4. *Suppose the conditions (B1)–(B4) hold. Denote by $\hat{\Theta} := \hat{\Theta}_{\text{Lasso}}$ and $\hat{\tau}_j^2$, $j = 1, \dots, p$ the estimates from the nodewise lasso in (8). Then for suitable tuning parameters $\lambda_j \asymp K_0 \sqrt{\log(p)/n}$ uniformly in j , we have*

$$\begin{aligned} \|\hat{\Theta}_j - \Theta_j\|_1 &= \mathcal{O}_{\mathbb{P}}\left(K_0 s_j \sqrt{\frac{\log(p)}{n}}\right), \\ \|\hat{\Theta}_j - \Theta_j\|_2 &= \mathcal{O}_{\mathbb{P}}\left(K_0 \sqrt{\frac{s_j \log(p)}{n}}\right), \\ |\hat{\tau}_j^2 - \tau_j^2| &= \mathcal{O}_{\mathbb{P}}\left(K_0 \sqrt{\frac{s_j \log(p)}{n}}\right), \quad j = 1, \dots, p. \end{aligned}$$

Furthermore,

$$|\hat{\Theta}_j \Sigma \hat{\Theta}_j^T - \Theta_{j,j}| \leq \|\Sigma\|_\infty \|\hat{\Theta}_j - \Theta_j\|_1^2 \wedge \Lambda_{\max}^2 \|\hat{\Theta}_j - \Theta_j\|_2^2 + 2|\hat{\tau}_j^2 - \tau_j^2|, \quad j = 1, \dots, p,$$

where Λ_{\max}^2 is the maximal eigenvalue of Σ . In the sub-Gaussian or strongly bounded case the results are uniform in j .

Finally, assume model (1) but assume instead of Gaussian errors that $\{\varepsilon_i\}_{i=1}^n$ are i.i.d. with variance $\sigma_\varepsilon^2 = \mathcal{O}(1)$. Assume moreover in the sub-Gaussian case for \mathbf{X} that the errors are subexponential, that is, that $\mathbb{E} \exp[|\varepsilon_1|/L] = \mathcal{O}(1)$ for some

fixed L . Apply the estimator (2) with $\lambda \asymp K_0 \sqrt{\log(p)/n}$ suitably chosen. Assume that $K_0 s_0 \log(p)/\sqrt{n} = o(1)$ and $\max_j K_0 s_j \sqrt{\log(p)/n} = o(1)$. Then we have

$$\begin{aligned} \sqrt{n}(\hat{b}_{\text{Lasso}} - \beta^0) &= W + \Delta, \\ W &= \hat{\Theta} \mathbf{X}^T \varepsilon / \sqrt{n}, \\ |\Delta_j| &= o_{\mathbb{P}}(1) \quad \forall j \end{aligned}$$

and in the sub-Gaussian or strongly bounded case

$$\|\Delta\|_{\infty} = o_{\mathbb{P}}(1).$$

A proof is given in Section 5.6.

Note that the result is as in Theorem 2.2 except that $W|\mathbf{X}$ is not necessarily normally distributed. A central limit theorem argument can be used to obtain approximate Gaussianity of components of $W|\mathbf{X}$ of fixed dimension. This can also be done for moderately growing dimensions (see, e.g., [36]), which is useful for testing with large groups G .

3. Generalized linear models and general convex loss functions. We show here that the idea of de-sparsifying ℓ_1 -norm penalized estimators and corresponding theory from Section 2 carries over to models with convex loss functions such as generalized linear models (GLMs).

3.1. *The setting and de-sparsifying the ℓ_1 -norm regularized estimator.* We consider the following framework with $1 \times p$ vectors of covariables $x_i \in \mathcal{X} \subseteq \mathbb{R}^p$ and univariate responses $y_i \in \mathcal{Y} \subseteq \mathbb{R}$ for $i = 1, \dots, n$. As before, we denote by \mathbf{X} the design matrix with i th row equal to x_i . At the moment, we do not distinguish whether \mathbf{X} is random or fixed (e.g., when conditioning on \mathbf{X}).

For $y \in \mathcal{Y}$ and $x \in \mathcal{X}$ being a $1 \times p$ vector, we have a loss function

$$\rho_{\beta}(y, x) = \rho(y, x\beta) \quad (\beta \in \mathbb{R}^p),$$

which is assumed to be a strictly convex function in $\beta \in \mathbb{R}^p$. We now define

$$\dot{\rho}_{\beta} := \frac{\partial}{\partial \beta} \rho_{\beta}, \quad \ddot{\rho}_{\beta} := \frac{\partial}{\partial \beta} \frac{\partial}{\partial \beta^T} \rho_{\beta},$$

where we implicitly assume that the derivatives exist. For a function $g : \mathcal{Y} \times \mathcal{X} \rightarrow \mathbb{R}$, we write $P_n g := \sum_{i=1}^n g(y_i, x_i)/n$ and $P g := \mathbb{E} P_n g$. Moreover, we let $\|g\|_n^2 := P_n g^2$ and $\|g\|^2 := P g^2$.

The ℓ_1 -norm regularized estimator is

$$(16) \quad \hat{\beta} = \arg \min_{\beta} (P_n \rho_{\beta} + \lambda \|\beta\|_1).$$

As in Section 2.1, we desparsify the estimator. For this purpose, define

$$(17) \quad \hat{\Sigma} := P_n \ddot{\rho}_{\hat{\beta}}.$$

Note that in general, $\hat{\Sigma}$ depends on $\hat{\beta}$ (an exception being the squared error loss). We construct $\hat{\Theta} = \hat{\Theta}_{\text{Lasso}}$ by doing a nodewise lasso with $\hat{\Sigma}$ as input as detailed below in (21). We then define

$$(18) \quad \hat{b} := \hat{\beta} - \hat{\Theta} P_n \dot{\rho}_{\hat{\beta}}.$$

The estimator in (5) is a special case of (18) with squared error loss.

3.1.1. *Lasso for nodewise regression with matrix input.* Denote by $\hat{\Sigma}$ a matrix which we want to approximately invert using the nodewise lasso. For every row j , we consider the optimization

$$(19) \quad \hat{\gamma}_j := \arg \min_{\gamma \in \mathbb{R}^{p-1}} (\hat{\Sigma}_{j,j} - 2\hat{\Sigma}_{j,\setminus j} \gamma + \gamma^T \hat{\Sigma}_{\setminus j,\setminus j} \gamma + 2\lambda_j \|\gamma\|_1),$$

where $\hat{\Sigma}_{j,\setminus j}$ denotes the j th row of $\hat{\Sigma}$ without the diagonal element (j, j) , and $\hat{\Sigma}_{\setminus j,\setminus j}$ is the submatrix without the j th row and j th column. We note that for the case where $\hat{\Sigma} = \mathbf{X}^T \mathbf{X} / n$, $\hat{\gamma}_j$ is the same as in (7).

Based on $\hat{\gamma}_j$ from (19), we compute

$$(20) \quad \hat{\tau}_j^2 = \hat{\Sigma}_{j,j} - \hat{\Sigma}_{j,\setminus j} \hat{\gamma}_j.$$

Having $\hat{\gamma}_j$ and $\hat{\tau}_j^2$ from (19) and (20), we define the nodewise lasso as

$$(21) \quad \hat{\Theta}_{\text{Lasso}} \text{ as in (8) using (19)–(20) from matrix input } \hat{\Sigma} \text{ in (17).}$$

Moreover, we denote by

$$\hat{b}_{\text{Lasso}} := \hat{b} \text{ from (18) using the nodewise lasso from (21).}$$

Computation of (19), and hence of $\hat{\Theta}$ can be done efficiently via coordinate descent using the KKT conditions to characterize the zeroes. Furthermore, an active set strategy leads to additional speed-up. See, for example, [20] and [28].

For standard GLMs, the matrix input $\hat{\Sigma} = P_n \ddot{\rho}_{\hat{\beta}}$ in (17) can be written as $\hat{\Sigma} = \hat{\Sigma}_{\hat{\beta}} = \mathbf{X}_{\hat{\beta}}^T \mathbf{X}_{\hat{\beta}} / n$ with $\mathbf{X}_{\hat{\beta}} := W_{\hat{\beta}} \mathbf{X}$ and $W_{\hat{\beta}} = \text{diag}(w_{\hat{\beta}})$ for some weights $w_{i,\hat{\beta}} = w_{\hat{\beta}}(y_i, x_i)$ ($i = 1, \dots, n$). Then we can simply use the nodewise lasso as in (8) but based on the design matrix $\mathbf{X}_{\hat{\beta}}$: in particular, we can use the standard lasso algorithm.

3.2. *Theoretical results.* We show here that the components of the estimator \hat{b} in (18), when normalized with the easily computable standard error, converge to a standard Gaussian distribution. Based on such a result, the construction of confidence intervals and tests is straightforward.

Let $\beta^0 \in \mathbb{R}^p$ be the unique minimizer of $P\rho_{\beta}$ with s_0 denoting the number of nonzero coefficients. We use analogous notation as in Section 2.3 but with modifications for the current context. The asymptotic framework, which allows for

Gaussian approximation of averages, is as in Section 2.3 for $p = p_n \geq n \rightarrow \infty$, and thus, $Y := (y_1, \dots, y_n)^T = Y_n$, $\mathbf{X} = \mathbf{X}_n$, $\beta^0 = \beta_n^0$ and underlying parameters are all (potentially) depending on n . As before, we usually suppress the corresponding index n .

We make the following assumptions which are discussed in Section 3.3.1. Thereby, we assume (C3), (C5), (C6) and (C8) for some constant $K \geq 1$ and positive constants λ_* and s_* . The constant λ is the tuning parameter in (16). In Section 3.3.1, we will discuss the conditions with $\lambda \asymp \sqrt{\log p/n}$ and for all j , $\lambda_* \asymp \lambda_j \asymp \sqrt{\log(p)/n}$ where λ_j is the tuning parameter in (19). Moreover, there we will assume $s_* \geq s_j$ for all j . Here, $s_j = |\{k \neq j : \Theta_{\beta^0, j, k} \neq 0\}|$, $j = 1, \dots, p$ with $\Theta_{\beta^0} := (P\ddot{\rho}_{\beta^0})^{-1}$ (assumed to exist).

(C1) The derivatives

$$\dot{\rho}(y, a) := \frac{d}{da}\rho(y, a), \quad \ddot{\rho}(y, a) := \frac{d^2}{da^2}\rho(y, a),$$

exist for all y, a , and for some δ -neighborhood ($\delta > 0$), $\ddot{\rho}(y, a)$ is Lipschitz:

$$\max_{a_0 \in \{x_i \beta^0\}} \sup_{|a-a_0| \vee |\hat{a}-a_0| \leq \delta} \sup_{y \in \mathcal{Y}} \frac{|\ddot{\rho}(y, a) - \ddot{\rho}(y, \hat{a})|}{|a - \hat{a}|} \leq 1.$$

Moreover,

$$\max_{a_0 \in \{x_i \beta^0\}} \sup_{y \in \mathcal{Y}} |\dot{\rho}(y, a_0)| = \mathcal{O}(1), \quad \max_{a_0 \in \{x_i \beta^0\}} \sup_{|a-a_0| \leq \delta} \sup_{y \in \mathcal{Y}} |\ddot{\rho}(y, a)| = \mathcal{O}(1).$$

(C2) It holds that $\|\hat{\beta} - \beta^0\|_1 = \mathcal{O}_{\mathbb{P}}(s_0\lambda)$, $\|\mathbf{X}(\hat{\beta} - \beta^0)\|^2 = \mathcal{O}_{\mathbb{P}}(s_0\lambda^2)$, and $\|\mathbf{X}(\hat{\beta} - \beta^0)\|_n^2 = \mathcal{O}_{\mathbb{P}}(s_0\lambda^2)$.

(C3) It holds that $\|\mathbf{X}\|_{\infty} := \max_{i,j} |\mathbf{X}_{i,j}| = \mathcal{O}(K)$.

(C4) It holds that $\|P_n \ddot{\rho}_{\hat{\beta}} \hat{\Theta}_j^T - e_j\|_{\infty} = \mathcal{O}_{\mathbb{P}}(\lambda_*)$.

(C5) It holds that $\|\mathbf{X} \hat{\Theta}_j^T\|_{\infty} = \mathcal{O}_{\mathbb{P}}(K)$ and $\|\hat{\Theta}_j\|_1 = \mathcal{O}_{\mathbb{P}}(\sqrt{s_*})$.

(C6) It holds that $\|(P_n - P)\dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T\|_{\infty} = \mathcal{O}_{\mathbb{P}}(K^2\lambda)$ and moreover

$$\max_j 1/(\hat{\Theta} P \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \hat{\Theta}^T)_{j,j} = \mathcal{O}(1).$$

(C7) For every j , the random variable

$$\frac{\sqrt{n}(\hat{\Theta} P_n \dot{\rho}_{\beta^0})_j}{\sqrt{(\hat{\Theta} P \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \hat{\Theta}^T)_{j,j}}}$$

converges weakly to a $\mathcal{N}(0, 1)$ -distribution.

(C8) It holds that

$$Ks_0\lambda^2 = o(n^{-1/2}), \quad \lambda_*\lambda s_0 = o(n^{-1/2}) \quad \text{and} \quad K^2s_*\lambda + K^2\sqrt{s_0}\lambda = o(1).$$

The following main result holds for fixed or random design according to whether the assumptions hold for one or the other case.

THEOREM 3.1. *Assume (C1)–(C8). For the estimator in (18), we have for each $j \in \{1, \dots, p\}$:*

$$\sqrt{n}(\hat{b}_j - \beta_j^0)/\hat{\sigma}_j = V_j + o_{\mathbf{P}}(1),$$

where V_j converges weakly to a $\mathcal{N}(0, 1)$ -distribution and where

$$\hat{\sigma}_j^2 := (\hat{\Theta} P_n \dot{\rho}_{\hat{\beta}} \dot{\rho}_{\hat{\beta}}^T \hat{\Theta}^T)_{j,j}.$$

A proof is given in Section 5.7. Assumption (C1) of Theorem 3.1 means that we regress to the classical conditions for asymptotic normality in the one-dimensional case as in, for example, [15]. Assumption (C8) is a sparsity assumption: for $K = O(1)$ and choosing $\lambda_* \asymp \lambda \asymp \sqrt{\log(p)/n}$ the condition reads as $s_0 = o(\sqrt{n}/\log(p))$ (as in Theorem 2.2) and $s_* = o(\sqrt{n}/\log(p))$. All the other assumptions (C2)–(C7) follow essentially from the conditions of Corollary 3.1 presented later, with the exception that (C3) is straightforward to understand. For more details, see Section 3.3.1.

3.3. About nodewise regression with certain random matrices. We justify in this section most of the assumptions for Theorem 3.1 when using the node-wise lasso estimator $\hat{\Theta} = \hat{\Theta}_{\text{Lasso}}$ as in (21) and when the matrix input is parameterized by $\hat{\beta}$ as for standard generalized linear models. For notational simplicity, we drop the subscript “lasso” in $\hat{\Theta}$. Let w_β be an n -vector with entries $w_{i,\beta} = w_\beta(y_i, x_i)$. We consider the matrix $\mathbf{X}_\beta := W_\beta \mathbf{X}$ where $W_\beta = \text{diag}(w_\beta)$. We define $\hat{\Sigma}_\beta := \mathbf{X}_\beta^T \mathbf{X}_\beta/n$. We fix some j and consider $\hat{\Theta}_{\hat{\beta},j}$ as the j th row of the nodewise regression $\hat{\Theta} = \hat{\Theta}_{\hat{\beta}}$ in (21) based on the matrix input $\hat{\Sigma}_{\hat{\beta}}$.

We let $\Sigma_\beta = \mathbb{E}[\mathbf{X}_\beta^T \mathbf{X}_\beta/n]$ and define $\Theta := \Theta_{\beta^0} := \Sigma_{\beta^0}^{-1}$ (assumed to exist). Let $s_j := s_{\beta^0,j}$ be the number of off-diagonal zeros of the j th row of Θ_{β^0} . Analogous to Section 2.3.4, we let $\mathbf{X}_{\beta^0,-j} \gamma_{\beta^0,j}$ be the projection of $\mathbf{X}_{\beta^0,j}$ on $\mathbf{X}_{\beta^0,-j}$ using the inner products in the matrix Σ_{β^0} and let $\eta_{\beta^0,j} := \mathbf{X}_{\beta^0,j} - \mathbf{X}_{\beta^0,-j} \gamma_{\beta^0,j}$. We then make the following assumptions:

(D1) The pairs of random variables $\{(y_i, x_i)\}_{i=1}^n$ are i.i.d. and $\|\mathbf{X}\|_\infty = \max_{i,j} |\mathbf{X}_{i,j}| = O(K)$ and $\|\mathbf{X}_{\beta^0,-j} \gamma_{\beta^0,j}\|_\infty = O(K)$ for some $K \geq 1$.

(D2) It holds that $K^2 s_j \sqrt{\log(p)/n} = o(1)$.

(D3) The smallest eigenvalue of Σ_{β^0} is bounded away from zero, and moreover, $\|\Sigma_{\beta^0}\|_\infty = O(1)$.

(D4) For some $\delta > 0$ and all $\|\beta - \beta^0\|_1 \leq \delta$, it holds that w_β stays away from zero and that $\|w_\beta\|_\infty = O(1)$. We further require that for all such β and all x and y

$$|w_\beta(y, x) - w_{\beta^0}(y, x)| \leq |x(\beta - \beta^0)|.$$

(D5) It holds that

$$\|\mathbf{X}(\hat{\beta} - \beta^0)\|_n = \mathcal{O}_{\mathbb{P}}(\lambda\sqrt{s_0}), \quad \|\hat{\beta} - \beta^0\|_1 = \mathcal{O}_{\mathbb{P}}(\lambda s_0).$$

Condition (D5) and (C2) typically hold when $\lambda\sqrt{s_0} = o(1)$ with tuning parameter $\lambda \asymp \sqrt{\log(p)/n}$ since the compatibility condition is then inherited from (D3) (see also Section 3.3.1). We have the following result.

THEOREM 3.2. *Assume the conditions (D1)–(D5). Then, using $\lambda_j \asymp K\sqrt{\log(p)/n}$ for the nodewise lasso $\hat{\Theta}_{\hat{\beta},j}$.*

$$\begin{aligned} \|\hat{\Theta}_{\hat{\beta},j} - \Theta_{\beta^0,j}\|_1 &= \mathcal{O}_{\mathbb{P}}(Ks_j\sqrt{\log(p)/n}) + \mathcal{O}_{\mathbb{P}}(K^2s_0((\lambda^2/\sqrt{\log(p)/n}) \vee \lambda)), \\ \|\hat{\Theta}_{\hat{\beta},j} - \Theta_{\beta^0,j}\|_2 &= \mathcal{O}_{\mathbb{P}}(K\sqrt{s_j \log(p)/n}) + \mathcal{O}_{\mathbb{P}}(K^2\sqrt{s_0}\lambda), \end{aligned}$$

and for $\tau_{\beta^0,j}^2 := \Theta_{\beta^0,j,j}$

$$|\hat{\tau}_{\hat{\beta},j}^2 - \tau_{\beta^0,j}^2| = \mathcal{O}_{\mathbb{P}}(K\sqrt{s_j \log(p)/n}) + \mathcal{O}_{\mathbb{P}}(K^2\sqrt{s_0}\lambda).$$

Moreover,

$$\begin{aligned} &|\hat{\Theta}_{\hat{\beta},j} \Sigma_{\beta^0} \hat{\Theta}_{\hat{\beta},j}^T - \Theta_{\beta^0,j,j}| \\ &\leq \|\Sigma_{\beta^0}\|_{\infty} \|\hat{\Theta}_{\hat{\beta},j} - \Theta_{\beta^0,j}\|_1^2 \wedge \Lambda_{\max}^2 \|\hat{\Theta}_{\hat{\beta},j} - \Theta_{\beta^0,j}\|_2^2 + 2|\hat{\tau}_{\hat{\beta},j}^2 - \tau_{\beta^0,j}^2|, \end{aligned}$$

where Λ_{\max}^2 is the maximal eigenvalue of Σ_{β^0} .

A proof using ideas for establishing Theorem 2.4 is given in the supplemental article [45].

COROLLARY 3.1. *Assume the conditions of Theorem 3.2, with tuning parameter $\lambda \asymp \sqrt{\log(p)/n}$, $K \asymp 1$, $s_j = o(\sqrt{n}/\log(p))$ and $s_0 = o(\sqrt{n}/\log(p))$. Then*

$$\begin{aligned} \|\hat{\Theta}_{\hat{\beta},j} - \Theta_{\beta^0,j}\|_1 &= o_{\mathbb{P}}(1/\sqrt{\log(p)}), \\ \|\hat{\Theta}_{\hat{\beta},j} - \Theta_{\beta^0,j}\|_2 &= o_{\mathbb{P}}(n^{-1/4}) \end{aligned}$$

and

$$|\hat{\Theta}_{\hat{\beta},j} \Sigma_{\beta^0} \hat{\Theta}_{\hat{\beta},j}^T - \Theta_{\beta^0,j,j}| = o_{\mathbb{P}}(1/\log(p)).$$

The next lemma is useful when estimating the asymptotic variance.

LEMMA 3.1. *Assume the conditions of Corollary 3.1. Let for $i = 1, \dots, n$, ξ_i be a real-valued random variable and $x_i^T \in \mathbb{R}^p$, and let $(x_i, \xi_i)_{i=1}^n$ be i.i.d. Assume $\mathbb{E}x_i^T \xi_i = 0$ and that $|\xi_i| \leq 1$. Then*

$$\hat{\Theta}_{\hat{\beta},j} \sum_{i=1}^n x_i^T \xi_i / n = \Theta_{\beta^0,j} \sum_{i=1}^n x_i^T \xi_i / n + o_{\mathbb{P}}(n^{-1/2}).$$

Let $A := \mathbb{E}x_i^T x_i \xi_i^2$ (assumed to exist). Assume that $\|A\Theta_j^T\|_{\infty} = \mathcal{O}(1)$ and that $1/(\Theta_j A \Theta_j^T) = \mathcal{O}(1)$. Then

$$\hat{\Theta}_{\hat{\beta},j} A \hat{\Theta}_{\hat{\beta},j}^T = \Theta_{\beta^0,j} A \Theta_{\beta^0,j}^T + o_{\mathbb{P}}(1).$$

Moreover, then

$$\frac{\hat{\Theta}_{\hat{\beta},j} \sum_{i=1}^n x_i^T \xi_i / \sqrt{n}}{\sqrt{\hat{\Theta}_{\hat{\beta},j} A \hat{\Theta}_{\hat{\beta},j}^T}}$$

converges weakly to a $\mathcal{N}(0, 1)$ -distribution.

A proof is given in the supplemental article [45].

3.3.1. *Consequence for GLMs.* Consider the case where $a \mapsto \rho(y, a)$ is convex for all y . We let $\{(y_i, x_i)\}_{i=1}^n \sim \mathbb{P}$ be i.i.d. random variables. We denote by \mathbf{X}_{β^0} the weighted design matrix $W_{\beta^0} \mathbf{X}$ with W_{β^0} the diagonal matrix with elements $\{\sqrt{\ddot{\rho}(y_i, x_i \beta^0)}\}_{i=1}^n$. We further let $\mathbf{X}_{\beta^0, -j} \gamma_{\beta^0, j}^0$ be the projection in $L_2(\mathbb{P})$ of $\mathbf{X}_{\beta^0, j}$ on $\mathbf{X}_{\beta^0, -j}$, $j = 1, \dots, p$. We write $\Sigma_{\beta^0} := \mathbb{E} \mathbf{X}_{\beta^0}^T \mathbf{X}_{\beta^0} / n$ and let s_j be the number of nonzero lower-diagonal elements of the j th column of Σ_{β^0} ($j = 1, \dots, p$).

THEOREM 3.3. *Let $\{(y_i, x_i)\}_{i=1}^n \sim \mathbb{P}$ be i.i.d. random variables. Assume:*

- (i) *Condition (C1),*
- (ii) $\|1/\ddot{\rho}_{\beta^0}\|_{\infty} = \mathcal{O}(1)$,
- (iii) $\|\mathbf{X}\|_{\infty} = \mathcal{O}(1)$,
- (iv) $\|\mathbf{X}\beta^0\|_{\infty} = \mathcal{O}(1)$ and $\|\mathbf{X}_{\beta^0, -j} \gamma_{\beta^0, j}^0\|_{\infty} = \mathcal{O}(1)$ for each j ,
- (v) *the smallest eigenvalue of Σ_{β^0} stays away from zero,*
- (vi) $1/(\Theta_{\beta^0, j} P \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \Theta_{\beta^0, j}^T) = \mathcal{O}(1) \forall j$,
- (vii) $s_0 = o(\sqrt{n}/\log(p))$ and $s_j = \sqrt{n/\log(p)}$ for all j .

Take $\hat{\Theta}$ equal to $\hat{\Theta}_{\text{Lasso}}$ given in (21) with $\lambda_j \asymp \sqrt{\log(p)/n}$ ($j = 1, \dots, p$) suitably chosen. For the estimator in (18), with suitable $\lambda \asymp \sqrt{\log(p)/n}$, we have for each j

$$\sqrt{n}(\hat{b}_j - \beta_j^0) / \hat{\sigma}_j = V_j + o_{\mathbb{P}}(1),$$

where V_j converges weakly to a $\mathcal{N}(0, 1)$ -distribution and where

$$\hat{\sigma}_j^2 := (\hat{\Theta} P_n \hat{\rho}_{\hat{\beta}} \hat{\rho}_{\hat{\beta}}^T \hat{\Theta}^T)_{j,j}.$$

A proof is given in Section 5.8.

Note that for the case where ρ_{β} is the minus log-likelihood, $P \hat{\rho}_{\beta^0} \hat{\rho}_{\beta^0}^T = \Sigma_{\beta^0}$, and hence $\Theta_{\beta^0,j} P \hat{\rho}_{\beta^0} \hat{\rho}_{\beta^0}^T \Theta_{\beta^0,j}^T = \Theta_{\beta^0,j,j}$. Assumption (vi) then follows from assumptions (i)–(iii) since $1/\Theta_{\beta^0,j,j} \leq \Sigma_{\beta^0,j,j}$.

4. Empirical results. We consider finite sample behavior for inference of individual regression coefficients β_j^0 , including adjustment for the case of multiple hypothesis testing.

4.1. *Methods and models.* We compare our method based on \hat{b}_{Lasso} with a procedure based on multiple sample splitting [32] (for multiple hypothesis testing only) and with a residual bootstrap method proposed by [14].

The implementational details for inference based on \hat{b}_{Lasso} are as follows. For the linear regression of the response Y versus the design \mathbf{X} , we use the scaled lasso [42] with its universal regularization parameter, and we use its estimate $\hat{\sigma}_{\varepsilon}^2$ of the error variance. For logistic regression, we use the corresponding lasso estimator with tuning parameter from 10-fold cross-validation. Regarding the nodewise lasso (for linear and logistic regression), we choose the same tuning parameter $\lambda_j \equiv \lambda_X$ by 10-fold cross-validation among all nodewise regressions. An alternative method which we did not yet examine in the simulations would be to do nodewise regression with square-root lasso using a universal choice for the tuning parameter (see Remark 2.1). For the bootstrap method from [14], we use 10-fold cross-validation to sequentially select the tuning parameter for lasso and subsequently for adaptive lasso. For multiple sample splitting [32], we do variable screening with the lasso whose regularization parameter is chosen by 10-fold cross-validation.

The construction of confidence intervals and hypothesis tests for individual parameters β_j^0 based on \hat{b}_{Lasso} is straightforward, as described in Section 2.1. Adjustment for multiple testing of hypotheses $H_{0,j}$ over all $j = 1, \dots, p$ is done using the Bonferroni–Holm procedure for controlling the family-wise error rate (FWER). For the bootstrap procedure from [14], the Bonferroni–Holm adjustment is not sensible, unless we would draw very many bootstrap resamples (e.g., 10,000 or more): with fewer resamples, we cannot reliably estimate the distribution in the tails needed for Bonferroni–Holm correction. Thus, for this bootstrap method, we only consider construction of confidence intervals. Finally, the multiple sample splitting method [32] is directly giving p -values which control the FWER.

For our simulation study, we consider (logistic) linear models where the rows of \mathbf{X} are fixed i.i.d. realizations from $\mathcal{N}_p(0, \Sigma)$. We specify two different covariance

matrices:

$$\begin{aligned} \text{Toeplitz:} \quad & \Sigma_{j,k} = 0.9^{|j-k|}, \\ \text{Equi corr:} \quad & \Sigma_{j,k} \equiv 0.8 \quad \text{for all } j \neq k, \quad \Sigma_{j,j} \equiv 1 \quad \text{for all } j. \end{aligned}$$

The active set has either cardinality $s_0 = |S_0| = 3$ or $s_0 = 15$, and each of it is of one of the following forms:

$$S_0 = \{1, 2, \dots, s_0\}, \text{ or: realization of random support } S_0 = \{u_1, \dots, u_{s_0}\},$$

where u_1, \dots, u_{s_0} is a fixed realization of s_0 draws without replacement from $\{1, \dots, p\}$. The regression coefficients are from a fixed realization of s_0 i.i.d. Uniform $U[0, c]$ variables with $c \in \{1, 2, 4\}$. For linear models, the distribution of the errors is always $\varepsilon_1, \dots, \varepsilon_n \sim \mathcal{N}(0, 1)$; see comment below regarding t -distributed errors. We also consider logistic regression models with binary response and

$$\log(\pi(x)/(1 - \pi(x))) = x\beta^0, \quad \pi(x) = \mathbb{P}[y_1 = 1|x_1 = x].$$

Sample size is always $n = 100$ (with some exceptions in the supplemental article [45]) and the number of variables is $p = 500$. We then consider many combinations of the different specifications above. All our results are based on 100 independent simulations of the model with fixed design and fixed regression coefficients (i.e., repeating over 100 independent simulations of the errors in a linear model).

4.2. Results for simulated data.

4.2.1. *Linear model: Confidence intervals.* We consider average coverage and average length of the intervals for individual coefficients corresponding to variables in either S_0 or S_0^c : denoting by CI_j a two-sided confidence interval for β_j^0 , we report empirical versions of

$$\text{Avcov } S_0 = s_0^{-1} \sum_{j \in S_0} \mathbb{P}[\beta_j^0 \in CI_j],$$

$$\text{Avcov } S_0^c = (p - s_0)^{-1} \sum_{j \in S_0^c} \mathbb{P}[0 \in CI_j],$$

$$\text{Avglength } S_0 = s_0^{-1} \sum_{j \in S_0} \text{length}(CI_j); \quad \text{and analogously for Avglength } S_0^c.$$

The following Tables 1–4 are for different active sets.

Discussion. As the main finding, we summarize that the desparsified lasso estimator is clearly better for the variables in S_0 than the residual based bootstrap. For the variables in S_0^c with regression coefficients equal to zero, the residual bootstrap exhibits the super-efficiency phenomenon: the average length of the interval

TABLE 1

Linear model: average coverage and length of confidence intervals, for nominal coverage equal to 0.95. “Lasso-Pro” (lasso-projection) denotes the procedure based on our desparsified estimator \hat{b}_{Lasso} ; “Res-Boot” is the residual based bootstrap from [14]

Measure	Method	Toeplitz		Equi corr	
		$U(\{0, 2\})$	$U(\{0, 4\})$	$U(\{0, 2\})$	$U(\{0, 4\})$
Active set $S_0 = \{1, 2, 3\}$					
Avcov S_0	Lasso-Pro	0.86	0.84	0.90	0.89
	Res-Boot	0.66	0.85	0.45	0.57
Avglength S_0	Lasso-Pro	0.786	0.787	0.762	0.760
	Res-Boot	0.698	0.918	0.498	0.670
Avcov S_0^c	Lasso-Pro	0.95	0.95	0.95	0.95
	Res-Boot	1.00	1.00	1.00	1.00
Avglength S_0^c	Lasso-Pro	0.786	0.787	0.811	0.808
	Res-Boot	0.000	0.000	0.006	0.007

is often very close to zero while coverage equals one. This cannot happen with the desparsified lasso estimator: in contrast to the residual based bootstrap, the desparsified lasso estimator allows for a convergence result which is uniform for a large class of parameters, and hence leading to honest confidence intervals; see Section 2.3.1. Furthermore, our empirical results for active sets with $s_0 = 15$ indicate that inference with the desparsified lasso has its limit when the problem is not sufficiently sparse, especially for the case with equi-correlated design: this is in line with our theoretical results.

Finally, we have also looked at non-Gaussian models where the error terms are from a scaled t_5 distribution (Student distribution with 5 degrees of freedom) with

TABLE 2
See caption of Table 1

Measure	Method	Toeplitz		Equi corr	
		$U(\{0, 2\})$	$U(\{0, 4\})$	$U(\{0, 2\})$	$U(\{0, 4\})$
Active set with $s_0 = 3$ and support from fixed random realization					
Avcov S_0	Lasso-Pro	0.95	0.94	0.89	0.87
	Res-Boot	0.58	0.73	0.31	0.51
Avglength S_0	Lasso-Pro	0.890	0.934	0.822	0.821
	Res-Boot	0.336	0.463	0.500	0.743
Avcov S_0^c	Lasso-Pro	0.95	0.96	0.95	0.95
	Res-Boot	1.00	1.00	1.00	1.00
Avglength S_0^c	Lasso-Pro	0.879	0.923	0.805	0.804
	Res-Boot	0.002	0.003	0.008	0.009

TABLE 3
See caption of Table 1

Measure	Method	Toeplitz		Equi corr	
		$U([0, 2])$	$U([0, 4])$	$U([0, 2])$	$U([0, 4])$
Active set $S_0 = \{1, 2, \dots, 15\}$					
Avgcov S_0	Lasso-Pro	0.76	0.73	0.56	0.53
	Res-Boot	0.79	0.87	0.54	0.63
Avglength S_0	Lasso-Pro	0.813	0.814	0.559	0.554
	Res-Boot	1.012	1.138	0.746	0.844
Avgcov S_0^c	Lasso-Pro	0.96	0.96	0.93	0.93
	Res-Boot	1.00	1.00	0.99	0.99
Avglength S_0^c	Lasso-Pro	0.788	0.789	0.568	0.562
	Res-Boot	0.001	0.000	0.047	0.044

variance equal to one. The results (not reported here) look essentially identical as in Tables 1–4.

4.2.2. *Linear model: Multiple testing.* We consider multiple two-sided testing of hypotheses $H_{0,j}; \beta_j^0 = 0$ among all $j = 1, \dots, p$. We correct the p -values based on our \hat{b}_{Lasso} with the Bonferroni–Holm procedure to control the familywise error rate (FWER). The method based on multiple sample splitting [32] automatically yields p -values for controlling the FWER. For measuring power, we report on the empirical version of

$$\text{Power} = s_0^{-1} \sum_{j \in S_0} \mathbb{P}[H_{0,j} \text{ is rejected}].$$

TABLE 4
See caption of Table 1

Measure	Method	Toeplitz		Equi corr	
		$U([0, 2])$	$U([0, 4])$	$U([0, 2])$	$U([0, 4])$
Active set with $s_0 = 15$ and support from fixed random realization					
Avgcov S_0	Lasso-Pro	0.93	0.94	0.55	0.44
	Res-Boot	0.45	0.54	0.48	0.55
Avglength S_0	Lasso-Pro	2.391	4.354	0.572	0.552
	Res-Boot	0.480	0.599	0.675	0.809
Avgcov S_0^c	Lasso-Pro	0.95	0.95	0.93	0.92
	Res-Boot	0.98	0.97	0.99	0.99
Avglength S_0^c	Lasso-Pro	2.370	4.317	0.570	0.550
	Res-Boot	0.029	0.035	0.048	0.050

TABLE 5

Linear model: family-wise error rate (FWER) and power of multiple testing, for nominal FWER equal to 0.05. “Lasso-Pro” (lasso-projection) denotes the procedure based on our de-sparsified estimator \hat{b}_{Lasso} with Bonferroni–Holm adjustment for multiple testing; “MS-Split” is the multiple sample splitting method from [32]

Measure	Method	Toeplitz		Equi corr	
		$U([0, 2])$	$U([0, 4])$	$U([0, 2])$	$U([0, 4])$
		Active set $S_0 = \{1, 2, 3\}$			
Power	Lasso-Pro	0.42	0.69	0.48	0.82
	MS-Split	0.60	0.83	0.35	0.63
FWER	Lasso-Pro	0.03	0.05	0.13	0.13
	MS-Split	0.16	0.25	0.00	0.00

The following Tables 5–8 are for different active sets.

Discussion. Similarly to what we found for confidence intervals above, multiple testing with the desparsified lasso estimator is reliable and works well for sparse problems (i.e., $s_0 = 3$). For less sparse problems (i.e., $s_0 = 15$), the error control is less reliable, especially for equi-correlated designs. For sparse Toeplitz designs, the lasso-projection method has more power than multiple sample splitting, a finding which is in line with our established optimality theory.

4.2.3. *Logistic regression: Multiple testing.* The residual bootstrap method [14] cannot be used in a straightforward way for logistic regression. As for linear models, we compare our desparsified lasso estimator with the multiple sample splitting procedure, in the context of multiple testing for controlling the FWER.

For the case of logistic regression shown in Tables 9–10, inference with the de-sparsified lasso method is not very reliable with respect to the FWER. The multiple sample splitting method is found to perform better. We present in the

TABLE 6
See caption of Table 5

Measure	Method	Toeplitz		Equi corr	
		$U([0, 2])$	$U([0, 4])$	$U([0, 2])$	$U([0, 4])$
		Active set with $s_0 = 3$ and support from fixed random realization			
Power	Lasso-Pro	0.54	0.81	0.56	0.79
	MS-Split	0.44	0.71	0.40	0.69
FWER	Lasso-Pro	0.00	0.00	0.10	0.11
	MS-Split	0.00	0.00	0.00	0.00

TABLE 7
See caption of Table 5

Measure	Method	Toeplitz		Equi corr	
		$U([0, 2])$	$U([0, 4])$	$U([0, 2])$	$U([0, 4])$
Active set $S_0 = \{1, 2, \dots, 15\}$					
Power	Lasso-Pro	0.73	0.89	0.70	0.92
	MS-Split	0.23	0.67	0.00	0.00
FWER	Lasso-Pro	0.03	0.02	1.00	1.00
	MS-Split	0.00	0.00	0.00	0.00

supplemental article [45] some additional results for sample sizes $n = 200$ and $n = 400$, illustrating that the FWER control as well as the power for the desparsified lasso improve.

4.3. *Real data analysis.* We consider a dataset about riboflavin (vitamin B_2) production by bacillus subtilis. The data has been kindly provided by DSM (Switzerland) and is publicly available [9]. The real-valued response variable is the logarithm of the riboflavin production rate and there are $p = 4088$ covariates (genes) measuring the logarithm of the expression level of 4088 genes. These measurements are from $n = 71$ samples of genetically engineered mutants of bacillus subtilis. We model the data with a high-dimensional linear model and obtain the following results for significance. The desparsified lasso procedure finds no significant coefficient while the multiple sample splitting method claims significance of one variable at the 5% significance level for the FWER. Such low power is to be expected in presence of thousands of variables: finding significant groups of highly correlated variables would seem substantially easier, at the price of not being able to infer significant of variables at the individual level.

TABLE 8
See caption of Table 5

Measure	Method	Toeplitz		Equi corr	
		$U([0, 2])$	$U([0, 4])$	$U([0, 2])$	$U([0, 4])$
Active set with $s_0 = 15$ and support from fixed random realization					
Power	Lasso-Pro	0.06	0.07	0.65	0.86
	MS-Split	0.07	0.14	0.00	0.00
FWER	Lasso-Pro	0.02	0.00	0.96	0.98
	MS-Split	0.02	0.13	0.00	0.00

TABLE 9
Logistic regression: All other specifications as in Table 5

Measure	Method	Toeplitz		
		$U(\{0, 1\})$	$U(\{0, 2\})$	$U(\{0, 4\})$
Power	Lasso-ProG	0.06	0.27	0.50
	MS-Split	0.07	0.37	0.08
FWER	Lasso-ProG	0.03	0.08	0.23
	MS-Split	0.01	0.00	0.00

5. Proofs and materials needed.

5.1. *Bounds for $\|\hat{\beta} - \beta^0\|_1$ with fixed design.* The following known result gives a bound for the ℓ_1 -norm estimation accuracy.

LEMMA 5.1. *Assume a linear model as in (1) with Gaussian error and fixed design \mathbf{X} which satisfies the compatibility condition with compatibility constant ϕ_0^2 and with $\hat{\Sigma}_{j,j} \leq M^2 < \infty$ for all j . Consider the lasso with regularization parameter $\lambda \geq 2M\sigma_\varepsilon\sqrt{\frac{2(t^2+\log(p))}{n}}$. Then, with probability at least $1 - 2\exp(-t^2)$,*

$$\|\hat{\beta} - \beta^0\|_1 \leq 8\lambda \frac{s_0}{\phi_0^2} \quad \text{and} \quad \|\mathbf{X}(\hat{\beta} - \beta^0)\|_2^2/n \leq 8\lambda^2 \frac{s_0}{\phi_0^2}.$$

A proof follows directly from the arguments in [10], Theorem 6.1, which can be modified to treat the case with unequal values of $\hat{\Sigma}_{j,j}$ for various j .

5.2. *Proof of Theorem 2.1.* It is straightforward to see that

$$(22) \quad \begin{aligned} \|\Delta\|_\infty/\sqrt{n} &= \|(\hat{\Theta}_{\text{Lasso}}\hat{\Sigma} - I)(\hat{\beta} - \beta^0)\|_\infty \\ &\leq \|(\hat{\Theta}_{\text{Lasso}}\hat{\Sigma} - I)\|_\infty \|\hat{\beta} - \beta^0\|_1. \end{aligned}$$

TABLE 10
Logistic regression: All other specifications as in Table 5

Measure	Method	Toeplitz		
		$U(\{0, 1\})$	$U(\{0, 2\})$	$U(\{0, 4\})$
Power	Lasso-ProG	0.02	0.16	0.35
	MS-Split	0.00	0.17	0.27
FWER	Lasso-ProG	0.08	0.16	0.27
	MS-Split	0.00	0.03	0.01

Therefore, by (10) we have that $\|\Delta\|_\infty \leq \sqrt{n}\|\hat{\beta} - \beta^0\|_1 \max_j \lambda_j / \hat{\tau}_j^2$, and using the bound from Lemma 5.1 completes the proof.

5.3. *Random design: Bounds for compatibility constant and $\|\hat{T}^{-2}\|_\infty$.* The compatibility condition with constant ϕ_0^2 being bounded away from zero is ensured by a rather natural condition about sparsity. We have the following result.

LEMMA 5.2. *Assume (A2). Furthermore, assume that $s_0 = o(n/\log(p))$. Then there is a constant $L = \mathcal{O}(1)$ depending on Λ_{\min} only such that with probability tending to one the compatibility condition holds with compatibility constant $\phi_0^2 \geq 1/L^2$.*

A proof follows directly as in [39], Theorem 1. Lemmas 5.1 and 5.2 say that we have a bound

$$(23) \quad \begin{aligned} \|\hat{\beta} - \beta^0\|_1 &= \mathcal{O}_{\mathbb{P}}\left(s_0 \sqrt{\frac{\log(p)}{n}}\right), \\ \|\mathbf{X}(\hat{\beta} - \beta^0)\|_2^2/n &= \mathcal{O}_{\mathbb{P}}\left(\frac{s_0 \log(p)}{n}\right), \end{aligned}$$

when assuming (A2) and sparsity $s_0 = o(n/\log(p))$.

When using the lasso for nodewise regression in (8), we would like to have a bound for $\|\hat{T}_{\text{Lasso}}^{-2}\|_\infty$ appearing in Theorem 2.1.

LEMMA 5.3. *Assume (A2) with row-sparsity for $\Theta := \Sigma^{-1}$ bounded by*

$$\max_j s_j = o(n/\log(p)).$$

Then, when suitably choosing the regularization parameters $\lambda_j \asymp \sqrt{\log(p)/n}$ uniformly in j ,

$$\max_j 1/\hat{\tau}_j^2 = \mathcal{O}_{\mathbb{P}}(1).$$

PROOF. A proof follows using standard arguments. With probability tending to one the compatibility assumption holds uniformly for all nodewise regressions with compatibility constant bounded away from zero uniformly in j , as in Lemma 5.2 and invoking the union bound. Furthermore, the population error variance $\tau_j^2 = \mathbb{E}[(X_{1,j} - \sum_{k \neq j} \gamma_{j,k} X_{1,k})^2]$, where $\gamma_{j,k}$ are the population regression coefficients of $X_{1,j}$ versus $\{X_{1,k}; k \neq j\}$ satisfy: uniformly in j , $\tau_j^2 = 1/\Theta_{j,j} \geq \Lambda_{\min}^2 > 0$ and $\tau_j^2 \leq \mathbb{E}[X_{1,j}^2] = \Sigma_{j,j} = \mathcal{O}(1)$, thereby invoking assumption (A2). Thus, all the error variances behave nicely. Recall that

$$\hat{\tau}_j^2 := \|X_j - \mathbf{X}_{-j}\hat{\gamma}_j\|_2^2/n + \lambda_j \|\hat{\gamma}_j\|_1.$$

In the following, the probability statements are again uniformly in j by the union bound for suitable tuning parameters $\lambda_j \asymp \sqrt{\log(p)/n}$ uniformly in j . Each node-

wise regression satisfies $\|\mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j)\|_2^2/n = \mathcal{O}_{\mathbb{P}}(s_j \log(p)/n)$ [see Lemma 5.1 or (23), now applied to the lasso estimator for the regression of X_j on \mathbf{X}_{-j}]. It follows that

$$\begin{aligned} \|X_j - \mathbf{X}_{-j}\hat{\gamma}_j\|_2^2/n &= \|X_j - \mathbf{X}_{-j}\gamma_j\|_2^2/n + \|\mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j)\|_2^2/n \\ &\quad + 2(X_j - \mathbf{X}_{-j}\gamma_j)^T \mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j) \\ &= \tau_j^2 + \mathcal{O}_{\mathbb{P}}(n^{-1/2}) + \mathcal{O}_{\mathbb{P}}(\lambda_j^2 s_j) + \mathcal{O}_{\mathbb{P}}(\lambda_j \sqrt{s_j}) = \tau_j^2 + o_{\mathbb{P}}(1). \end{aligned}$$

Note further that

$$\|\gamma_j\|_1 \leq \sqrt{s_j} \|\gamma_j\|_2 \leq \sqrt{s_j} \Sigma_{j,j} / \Lambda_{\min}.$$

Moreover, by the same arguments giving the bounds in (23), $\|\hat{\gamma}_j - \gamma_j\|_1 = \mathcal{O}_{\mathbb{P}}(s_j \lambda_j)$ so that

$$\lambda_j \|\hat{\gamma}_j\|_1 \leq \lambda_j \|\gamma_j\|_1 + \lambda_j \|\hat{\gamma}_j - \gamma_j\|_1 = \lambda_j \mathcal{O}(\sqrt{s_j}) + \lambda_j \mathcal{O}_{\mathbb{P}}(\lambda_j s_j) = o_{\mathbb{P}}(1).$$

Hence, the statement of the lemma follows. \square

5.4. *Bounds for $\|\hat{\beta} - \beta^0\|_2$ with random design.* Note that $\|\mathbf{X}(\hat{\beta} - \beta^0)\|_2^2/n = (\hat{\beta} - \beta^0)^T \hat{\Sigma}(\hat{\beta} - \beta^0)$. Lemma 5.2 uses [39], Theorem 1. The same result can be invoked to conclude that when (A2) holds and when $\lambda \asymp \sqrt{\log(p)/n}$ is suitably chosen, then for a suitably chosen fixed C , with probability tending to one

$$\begin{aligned} &(\hat{\beta} - \beta^0)^T \Sigma(\hat{\beta} - \beta^0) \\ &\leq (\hat{\beta} - \beta^0)^T \hat{\Sigma}(\hat{\beta} - \beta^0) C + \sqrt{\frac{\log(p)}{n}} \|\hat{\beta} - \beta^0\|_1 C. \end{aligned}$$

Hence,

$$(\hat{\beta} - \beta^0)^T \Sigma(\hat{\beta} - \beta^0) = \mathcal{O}_{\mathbb{P}}\left(\frac{s_0 \log(p)}{n}\right).$$

So under (A2) for suitable $\lambda \asymp \sqrt{\log(p)/n}$

$$(24) \quad \|\hat{\beta} - \beta^0\|_2 = \mathcal{O}_{\mathbb{P}}(\sqrt{s_0 \log(p)/n})$$

(see also [6]). This result will be applied in the next subsection, albeit to the lasso for node wise regression instead of for the original linear model.

5.5. *Proof of Theorem 2.2.* Invoking Theorem 2.1 and Lemma 5.3, we have that

$$\|\Delta\|_{\infty} \leq \mathcal{O}_{\mathbb{P}}(s_0 \log(p)/\sqrt{n}) = o_{\mathbb{P}}(1),$$

where the last bound follows by the sparsity assumption on s_0 .

What remains to be shown is that $\|\hat{\Omega} - \Theta\|_{\infty} = o_{\mathbb{P}}(1)$, as detailed by the following lemma.

LEMMA 5.4. *Let $\hat{\Theta} := \hat{\Theta}_{\text{Lasso}}$ with suitable tuning parameters λ_j satisfying $\lambda_j \asymp \sqrt{\log(p)/n}$ uniformly in j . Assume the conditions of Lemma 5.3. Suppose that $\max_j \lambda_j^2 s_j = o(1)$. Then*

$$\|\hat{\Omega} - \Theta\|_\infty = o_{\mathbb{P}}(1).$$

PROOF. By the same arguments as in the proof of Lemma 5.3, uniformly in j ,

$$\|\hat{\Theta}_j\|_1 = O_{\mathbb{P}}(\sqrt{s_j}).$$

Furthermore, we have

$$(25) \quad \hat{\Omega} = \hat{\Theta} \hat{\Sigma} \hat{\Theta}^T = (\hat{\Theta} \hat{\Sigma} - I) \hat{\Theta}^T + \hat{\Theta}^T$$

and

$$(26) \quad \|(\hat{\Theta} \hat{\Sigma} - I) \hat{\Theta}^T\|_\infty \leq \max_j \lambda_j \|\hat{\Theta}_j\|_1 / \hat{\tau}_j^2 = o_{\mathbb{P}}(1),$$

which follows from Lemma 5.3. Finally, we have using standard arguments for the ℓ_2 -norm bounds [see also (24)]

$$(27) \quad \|\hat{\Theta} - \Theta\|_\infty \leq \max_j \|\hat{\Theta}_j - \Theta_j\|_2 \leq \max_j \lambda_j \sqrt{s_j} = o_{\mathbb{P}}(1).$$

Using (25)–(27), we complete the proof. \square

The proof of Theorem 2.2 is now complete.

5.6. *Proof of Theorem 2.4.* Under the sub-Gaussian assumption we know that η_j is also sub-Gaussian. So then $\|\eta_j^T \mathbf{X}_{-j}/n\|_\infty = O_{\mathbb{P}}(\sqrt{\log(p)/n})$. If $\|\mathbf{X}\|_\infty = O(K)$, we can use the work in [16] to conclude that

$$\|\eta_j^T \mathbf{X}_{-j}\|_\infty/n = O_{\mathbb{P}}(K \sqrt{\log(p)/n}).$$

However, this result does not hold uniformly in j . Otherwise, in the strongly bounded case, we have

$$\|\eta_j\|_\infty \leq \|X_j\|_\infty + \|\mathbf{X}_{-j} \gamma_j\|_\infty = O(K).$$

So then $\|\eta_j^T \mathbf{X}_{-j}/n\|_\infty = O_{\mathbb{P}}(K \sqrt{\log(p)/n}) + O_{\mathbb{P}}(K^2 \log(p)/n)$, which is uniform in j .

Then by standard arguments (see, e.g., [6], and see [10] which complements the concentration results in [26] for the case of errors with only second moments) for $\lambda_j \asymp K_0 \sqrt{\log(p)/n}$ [recall that $K_0 = 1$ in the sub-Gaussian case and $K_0 = K$ in the (strongly) bounded case]

$$\|\mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j)\|_n^2 = O_{\mathbb{P}}(s_j \lambda_j^2), \quad \|\hat{\gamma}_j - \gamma_j\|_1 = O_{\mathbb{P}}(s_j \lambda_j).$$

The condition $K^2 s_j \sqrt{\log(p)/n}$ is used in the (strongly) bounded case to be able to conclude that the empirical compatibility condition holds (see [10], Section 6.12).

In the sub-Gaussian case, we use that $\sqrt{s_j \log(p)/n} = o(1)$ and an extension of Theorem 1 in [39] from the Gaussian case to the sub-Gaussian case. This gives again that the empirical compatibility condition holds.

We further find that

$$\|\hat{\gamma}_j - \gamma_j\|_2 = \mathcal{O}_{\mathbb{P}}(K_0 \sqrt{s_j \log(p)/n}).$$

To show this, we first introduce the notation $v^T \Sigma v := \|\mathbf{X}v\|^2$. Then in the (strongly) bounded case

$$|\|\mathbf{X}v\|_n^2 - \|\mathbf{X}v\|^2| \leq \|\hat{\Sigma} - \Sigma\|_{\infty} \|v\|_1^2 = \mathcal{O}_{\mathbb{P}}(K^2 \sqrt{\log(p)/n}) \|v\|_1^2.$$

Since $\|\hat{\gamma}_j - \gamma_j\|_1 = \mathcal{O}_{\mathbb{P}}(K_0 s_j \sqrt{\log(p)/n})$ and the smallest eigenvalue Λ_{\min}^2 of Σ stays away from zero, this gives

$$\begin{aligned} \mathcal{O}_{\mathbb{P}}(K_0^2 s_j \log(p)/n) &= \|\mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j)\|_n^2 \\ &\geq \Lambda_{\min}^2 \|\hat{\gamma}_j - \gamma_j\|_2^2 - \mathcal{O}_{\mathbb{P}}(K_0^4 s_j^2 (\log(p)/n)^{3/2}) \\ &\geq \Lambda_{\min}^2 \|\hat{\gamma}_j - \gamma_j\|_2^2 - o_{\mathbb{P}}(K_0^2 \log(p)/n), \end{aligned}$$

where we again used that $K_0^2 s_j \sqrt{\log(p)/n} = o(1)$. In the sub-Gaussian case, the result for the $\|\cdot\|_2$ -estimation error follows by similar arguments invoking again a sub-Gaussian extension of Theorem 1 in [39].

We moreover have

$$\begin{aligned} |\hat{\tau}_j^2 - \tau_j^2| &= \underbrace{|\eta_j^T \eta_j/n - \tau_j^2|}_I + \underbrace{|\eta_j^T \mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j)/n|}_II \\ &\quad + \underbrace{|\eta_j^T \mathbf{X}_{-j} \gamma_j/n|}_III + \underbrace{|(\gamma_j)^T \mathbf{X}_{-j}^T \mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j)/n|}_IV. \end{aligned}$$

Now, since we assume fourth moments of the errors,

$$I = \mathcal{O}_{\mathbb{P}}(K_0^2 n^{-1/2}).$$

Moreover,

$$II = \mathcal{O}_{\mathbb{P}}(K_0 \sqrt{\log(p)/n}) \|\hat{\gamma}_j - \gamma_j\|_1 = \mathcal{O}_{\mathbb{P}}(K_0^2 s_j \log(p)/n).$$

As for III, we have

$$III = \mathcal{O}_{\mathbb{P}}(K_0 \sqrt{\log(p)/n}) \|\gamma_j\|_1 = \mathcal{O}_{\mathbb{P}}(K_0 \sqrt{s_j \log(p)/n})$$

since $\|\gamma_j\|_1 \leq \sqrt{s_j} \|\gamma_j\|_2 = \mathcal{O}(\sqrt{s_j})$. Finally, by the KKT conditions,

$$\|\mathbf{X}_{-j}^T \mathbf{X}_{-j}(\hat{\gamma}_j - \gamma_j)\|_{\infty}/n = \mathcal{O}_{\mathbb{P}}(K_0 \sqrt{\log(p)/n}),$$

and hence

$$IV = \mathcal{O}_{\mathbb{P}}(K_0\sqrt{\log(p)/n})\|\gamma_j\|_1 = \mathcal{O}_{\mathbb{P}}(K_0\sqrt{s_j \log(p)/n}).$$

So now we have shown that

$$|\hat{\tau}_j^2 - \tau_j^2| = \mathcal{O}_{\mathbb{P}}(K_0\sqrt{s_j \log(p)/n}).$$

Since $1/\tau_j^2 = \mathcal{O}(1)$, this implies that also

$$1/\hat{\tau}_j^2 - 1/\tau_j^2 = \mathcal{O}_{\mathbb{P}}(K_0\sqrt{s_j \log(p)/n}).$$

We conclude that

$$\begin{aligned} \|\hat{\Theta}_j - \Theta_j\|_1 &= \|\hat{C}_j/\hat{\tau}_j^2 - C_j/\tau_j^2\|_1 \\ &\leq \underbrace{\|\hat{\gamma}_j - \gamma_j\|_1/\hat{\tau}_j^2}_i + \underbrace{\|\gamma_j\|_1(1/\hat{\tau}_j^2 - 1/\tau_j^2)}_{ii}, \end{aligned}$$

where

$$i = \mathcal{O}_{\mathbb{P}}(K_0s_j\sqrt{\log(p)/n})$$

since $\hat{\tau}_j^2$ is a consistent estimator of τ_j^2 and $1/\tau_j^2 = \mathcal{O}(1)$, and also

$$ii = \mathcal{O}_{\mathbb{P}}(K_0s_j\sqrt{\log(p)/n}),$$

since $\|\gamma_j\|_1 = \mathcal{O}(\sqrt{s_j})$.

Recall that

$$\|\hat{\gamma}_j - \gamma_j\|_2 = \mathcal{O}_{\mathbb{P}}(K_0\sqrt{s_j \log(p)/n}).$$

But then

$$\begin{aligned} \|\hat{\Theta}_j - \Theta_j\|_2 &\leq \|\hat{\gamma}_j - \gamma_j\|_2/\hat{\tau}_j^2 + \|\gamma_j\|_2(1/\hat{\tau}_j^2 - 1/\tau_j^2) \\ &= \mathcal{O}_{\mathbb{P}}(K_0\sqrt{s_j \log(p)/n}). \end{aligned}$$

For the last part, we write

$$\begin{aligned} \hat{\Theta}_j \Sigma \hat{\Theta}_j^T - \Theta_{j,j} &= (\hat{\Theta}_j - \Theta_j) \Sigma (\hat{\Theta}_j - \Theta_j)^T + \Theta_j \Sigma (\hat{\Theta}_j - \Theta_j)^T + \Theta_j \Sigma \Theta_j^T - \Theta_{j,j} \\ &= (\hat{\Theta}_j - \Theta_j) \Sigma (\hat{\Theta}_j - \Theta_j)^T + 2(1/\hat{\tau}_j^2 - 1/\tau_j^2), \end{aligned}$$

since $\Theta_j \Sigma = e_j^T$, $\Theta_j \Sigma \Theta_j^T = \Theta_{j,j}$, $\hat{\Theta}_{j,j} = 1/\hat{\tau}_j^2$, and $\Theta_{j,j} = 1/\tau_j^2$. But

$$(\hat{\Theta}_j - \Theta_j) \Sigma (\hat{\Theta}_j - \Theta_j)^T \leq \|\Sigma\|_{\infty} \|\hat{\Theta}_j - \Theta_j\|_1.$$

We may also use

$$(\hat{\Theta}_j - \Theta_j) \Sigma (\hat{\Theta}_j - \Theta_j)^T \leq \Lambda_{\max}^2 \|\hat{\Theta}_j - \Theta_j\|_2^2.$$

The last statement of the theorem follows as in Theorem 2.1, as $\sqrt{n}(\hat{b}_{\text{Lasso},j} - \beta_j^0) = W_j + \Delta_j$, with $\Delta_j \leq \sqrt{n} \lambda_j / \hat{\tau}_j^2 \|\hat{\beta} - \beta^0\|_1$, and $\lambda_j / \hat{\tau}_j^2 \asymp \lambda_j \asymp \sqrt{\log(p)/n}$, the latter being uniformly in j in the sub-Gaussian or strongly bounded case.

5.7. *Proof of Theorem 3.1.* Note that

$$\dot{\rho}(y, x_i \hat{\beta}) = \dot{\rho}(y, x_i \beta^0) + \ddot{\rho}(y, \tilde{a}_i) x_i (\hat{\beta} - \beta^0),$$

where \tilde{a}_i is a point intermediating $x_i \hat{\beta}$ and $x_i \beta^0$, so that $|\tilde{a}_i - x_i \hat{\beta}| \leq |x_i (\hat{\beta} - \beta^0)|$.

We find by the Lipschitz condition on $\ddot{\rho}$ [condition (C1)]

$$\begin{aligned} & |\ddot{\rho}(y, \tilde{a}_i) x_i (\hat{\beta} - \beta^0) - \ddot{\rho}(y, x_i \hat{\beta}) x_i (\hat{\beta} - \beta^0)| \\ & \leq |\tilde{a}_i - x_i \hat{\beta}| |x_i (\hat{\beta} - \beta^0)| \leq |x_i (\hat{\beta} - \beta^0)|^2. \end{aligned}$$

Thus, using that by condition (C5) $|x_i \hat{\Theta}_j^T| = \mathcal{O}_{\mathbb{P}}(K)$ uniformly in j ,

$$\hat{\Theta}_j P_n \dot{\rho}_{\hat{\beta}} = \hat{\Theta}_j P_n \dot{\rho}_{\beta^0} + \hat{\Theta}_j P_n \ddot{\rho}_{\hat{\beta}} (\hat{\beta} - \beta^0) + \text{Rem}_1,$$

where

$$\begin{aligned} \text{Rem}_1 &= \mathcal{O}_{\mathbb{P}}(K) \sum_{i=1}^n |x_i (\hat{\beta} - \beta^0)|^2 / n = \mathcal{O}(K) \|\mathbf{X}(\hat{\beta} - \beta^0)\|_n^2 \\ &= \mathcal{O}_{\mathbb{P}}(K s_0 \lambda^2) = o_{\mathbb{P}}(1), \end{aligned}$$

where we used condition (C2) and in the last step condition (C8).

We know that by condition (C4)

$$\|\hat{\Theta}_j P_n \ddot{\rho}_{\hat{\beta}} - e_j^T\|_{\infty} = \mathcal{O}(\lambda_*).$$

It follows that

$$\begin{aligned} b_j - \beta_j^0 &= \hat{\beta}_j - \beta_j^0 - \hat{\Theta}_j P_n \dot{\rho}_{\hat{\beta}} \\ &= \hat{\beta}_j - \beta_j^0 - \hat{\Theta}_j P_n \dot{\rho}_{\beta^0} - \hat{\Theta}_j P_n \ddot{\rho}_{\hat{\beta}} (\hat{\beta} - \beta^0) - \text{Rem}_1 \\ &= -\hat{\Theta}_j P_n \dot{\rho}_{\beta^0} - (\hat{\Theta}_j P_n \ddot{\rho}_{\hat{\beta}} - e_j^T) (\hat{\beta} - \beta^0) - \text{Rem}_1 \\ &= -\hat{\Theta}_j P_n \dot{\rho}_{\beta^0} - \text{Rem}_2, \end{aligned}$$

where

$$|\text{Rem}_2| \leq |\text{Rem}_1| + \mathcal{O}(\lambda_*) \|\hat{\beta} - \beta^0\|_1 = o_{\mathbb{P}}(n^{-1/2}) + \mathcal{O}_{\mathbb{P}}(s_0 \lambda \lambda_*) = o_{\mathbb{P}}(n^{-1/2})$$

since by condition (C2) $\|\hat{\beta} - \beta^0\|_1 = \mathcal{O}_{\mathbb{P}}(\lambda s_0)$, and by the second part of condition (C8) also $\lambda_* \lambda s_0 = o(n^{-1/2})$.

We now have to show that our estimator of the variance is consistent. We find

$$\begin{aligned} & |(\hat{\Theta} P \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \hat{\Theta}^T)_{j,j} - (\hat{\Theta} P_n \dot{\rho}_{\hat{\beta}} \dot{\rho}_{\hat{\beta}}^T \hat{\Theta}^T)_{j,j}| \\ & \leq \underbrace{|(\hat{\Theta} (P_n - P) \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \hat{\Theta}^T)_{j,j}|}_I + \underbrace{|(\hat{\Theta} P \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \hat{\Theta}^T)_{j,j} - (\hat{\Theta} P \dot{\rho}_{\hat{\beta}} \dot{\rho}_{\hat{\beta}}^T \hat{\Theta}^T)_{j,j}|}_{II}. \end{aligned}$$

But, writing $\varepsilon_{k,l} := (P_n - P) \dot{\rho}_{k,\beta^0} \dot{\rho}_{l,\beta^0}$, we see that

$$\begin{aligned} I &= |(\hat{\Theta} (P_n - P) \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \hat{\Theta}^T)_{j,j}| = \left| \sum_{k,l} \hat{\Theta}_{j,k} \hat{\Theta}_{j,l} \varepsilon_{k,l} \right| \leq \|\hat{\Theta}_j\|_1^2 \|\varepsilon\|_\infty \\ &= \mathcal{O}_{\mathbb{P}}(s_* K^2 \lambda), \end{aligned}$$

where we used conditions (C5) and (C6).

Next, we will handle *II*. We have

$$\begin{aligned} \dot{\rho}_{\hat{\beta}}(y, x) \dot{\rho}_{\hat{\beta}}^T(y, x) - \dot{\rho}_{\beta^0}(y, x) \dot{\rho}_{\beta^0}^T(y, x) &= [\dot{\rho}^2(y - x \hat{\beta}) - \dot{\rho}^2(y - x \beta^0)] x^T x \\ &:= v(y, x) x^T x, \end{aligned}$$

with

$$|v(y, x)| := |\dot{\rho}^2(y - x \hat{\beta}) - \dot{\rho}^2(y - x \beta^0)| = \mathcal{O}_{\mathbb{P}}(1) |x(\hat{\beta} - \beta^0)|,$$

where we use that $\dot{\rho}_{\beta^0}$ is bounded and $\dot{\rho}$ is locally bounded [condition (C1)]. It follows from condition (C2) that

$$P|v| \leq \sqrt{P|v|^2} = \|\mathbf{X}(\hat{\beta} - \beta^0)\| = \mathcal{O}_{\mathbb{P}}(\lambda \sqrt{s_0}).$$

Moreover, by condition (C5),

$$\|\hat{\Theta}_j x^T\|_\infty = \mathcal{O}_{\mathbb{P}}(K)$$

so that

$$|(\hat{\Theta} v(x, y) x^T x \hat{\Theta}^T)_{j,j}| \leq \mathcal{O}(K^2) |v(y, x)|.$$

Thus,

$$|(\hat{\Theta} P \dot{\rho}_{\beta^0} \dot{\rho}_{\beta^0}^T \hat{\Theta}^T)_{j,j} - (\hat{\Theta} P \dot{\rho}_{\hat{\beta}} \dot{\rho}_{\hat{\beta}}^T \hat{\Theta}^T)_{j,j}| = \mathcal{O}_{\mathbb{P}}(K^2 \sqrt{s_0} \lambda).$$

It follows that

$$I + II = \mathcal{O}_{\mathbb{P}}(K^2 s_* \lambda) + \mathcal{O}_{\mathbb{P}}(K^2 \sqrt{s_0} \lambda) = o_{\mathbb{P}}(1)$$

by the last part of condition (C8).

5.8. *Proof of Theorem 3.3.* This follows from Theorem 3.1. The assumptions (C2), (C4)–(C8) follow from the conditions of Corollary 3.1 with $\Sigma_\beta := P\ddot{\rho}_\beta$ and $w_\beta^2(y, x) := \ddot{\rho}(y, x\beta)$, where we take $\hat{\Theta} = \hat{\Theta}_{\text{Lasso}}$ and $s_* = s_j$ and $\lambda_* = \lambda_j$. Condition (C2) holds because the compatibility condition is met as Σ_{β^0} is nonsingular and

$$\|\hat{\Sigma} - \Sigma_{\beta^0}\|_\infty = \mathcal{O}_{\mathbb{P}}(\lambda_*).$$

The condition that $\dot{\rho}(y, x\beta^0)$ is bounded ensures that $\rho(y, a)$ is locally Lipschitz, so that we can control the empirical process $(P_n - P)(\rho_{\hat{\beta}} - \rho_{\beta^0})$ as in [47] (see also [10] or [46]). [In the case of a GLM with canonical loss (e.g., least squares loss) we can relax the condition of a locally bounded derivative because the empirical process is then linear.] Condition (C3) is assumed to hold with $\|\mathbf{X}\|_\infty = \mathcal{O}(1)$, and condition (C4) holds with $\lambda_* \asymp \sqrt{\log p/n}$. This is because in the node-wise regression construction, the $1/\hat{\tau}_j^2$ are consistent estimators of $(\Sigma_{\beta^0}^{-1})_{jj}$ (see Theorem 3.2). Condition (C5) holds as well. Indeed, $\|\Theta_{\beta^0, j}\|_1 = \mathcal{O}(\sqrt{s_j})$, and $\|\hat{\Theta}_{\hat{\beta}, j} - \Theta_{\beta^0, j}\|_1 = \mathcal{O}_{\mathbb{P}}(\lambda_j s_j) = \mathcal{O}_{\mathbb{P}}(\sqrt{s_j})$. Condition (C6) holds, too, since we assume that $\|\dot{\rho}_{\beta^0}\|_\infty = \mathcal{O}(1)$ as well as $\|\mathbf{X}\|_\infty = \mathcal{O}(1)$. As for condition (C7), this follows from Lemma 3.1, since $|\Theta_{\beta^0, j}\dot{\rho}_{\beta^0}(y, x)| = |\Theta_{\beta^0, j}x^T\dot{\rho}(y, x\beta^0)| = \mathcal{O}(1)$, which implies for $A := P\dot{\rho}_{\beta^0}\dot{\rho}_{\beta^0}^T$ that $\|A\Theta_{\beta^0, j}^T\|_\infty = \mathcal{O}(1)$.

SUPPLEMENTARY MATERIAL

Supplement to “On asymptotically optimal confidence regions and tests for high-dimensional models” (DOI: [10.1214/14-AOS1221SUPP](https://doi.org/10.1214/14-AOS1221SUPP); .pdf). The supplementary article contains additional empirical results, as well as the proofs of Theorems 2.3 and 3.2, Lemmas 2.1 and 3.1.

REFERENCES

[1] BELLONI, A., CHERNOZHUKOV, V. and HANSEN, C. (2014). Inference on treatment effects after selection amongst high-dimensional controls. *Rev. Econ. Stud.* **81** 608–650.
 [2] BELLONI, A., CHERNOZHUKOV, V. and KATO, K. (2013). Uniform postselection inference for LAD regression models. Available at [arXiv:1306.0282](https://arxiv.org/abs/1306.0282).
 [3] BELLONI, A., CHERNOZHUKOV, V. and WANG, L. (2011). Square-root lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika* **98** 791–806. [MR2860324](https://doi.org/10.1093/biomet/asq024)
 [4] BELLONI, A., CHERNOZHUKOV, V. and WEI, Y. (2013). Honest confidence regions for logistic regression with a large number of controls. Available at [arXiv:1306.3969](https://arxiv.org/abs/1306.3969).
 [5] BERK, R., BROWN, L., BUJA, A., ZHANG, K. and ZHAO, L. (2013). Valid post-selection inference. *Ann. Statist.* **41** 802–837. [MR3099122](https://doi.org/10.1214/12-AOS1122)
 [6] BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. [MR2533469](https://doi.org/10.1214/07-AN116)
 [7] BÜHLMANN, P. (2006). Boosting for high-dimensional linear models. *Ann. Statist.* **34** 559–583. [MR2281878](https://doi.org/10.1214/05-AN1178)
 [8] BÜHLMANN, P. (2013). Statistical significance in high-dimensional linear models. *Bernoulli* **19** 1212–1242. [MR3102549](https://doi.org/10.1080/00036811307549)

- [9] BÜHLMANN, P., KALISCH, M. and MEIER, L. (2014). High-dimensional statistics with a view toward applications in biology. *Annual Review of Statistics and Its Applications* **1** 255–278.
- [10] BÜHLMANN, P. and VAN DE GEER, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, Heidelberg. [MR2807761](#)
- [11] BUNEA, F., TSYBAKOV, A. and WEGKAMP, M. (2007). Sparsity oracle inequalities for the Lasso. *Electron. J. Stat.* **1** 169–194. [MR2312149](#)
- [12] CANDÈS, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.* **35** 2313–2351. [MR2382644](#)
- [13] CHATTERJEE, A. and LAHIRI, S. N. (2011). Bootstrapping lasso estimators. *J. Amer. Statist. Assoc.* **106** 608–625. [MR2847974](#)
- [14] CHATTERJEE, A. and LAHIRI, S. N. (2013). Rates of convergence of the adaptive LASSO estimators to the oracle distribution and higher order refinements by the bootstrap. *Ann. Statist.* **41** 1232–1259. [MR3113809](#)
- [15] CRAMÉR, H. (1946). *Mathematical Methods of Statistics. Princeton Mathematical Series* **9**. Princeton Univ. Press, Princeton, NJ. [MR0016588](#)
- [16] DÜMBGEN, L., VAN DE GEER, S. A., VERAAR, M. C. and WELLNER, J. A. (2010). Nemirovski’s inequalities revisited. *Amer. Math. Monthly* **117** 138–160. [MR2590193](#)
- [17] FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 849–911. [MR2530322](#)
- [18] FAN, J. and LV, J. (2010). A selective overview of variable selection in high dimensional feature space. *Statist. Sinica* **20** 101–148. [MR2640659](#)
- [19] FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
- [20] FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2010). Regularized paths for generalized linear models via coordinate descent. *Journal of Statistical Software* **33** 1–22.
- [21] GREENSHTEIN, E. and RITOV, Y. (2004). Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli* **10** 971–988. [MR2108039](#)
- [22] JAVANMARD, A. and MONTANARI, A. (2013). Confidence intervals and hypothesis testing for high-dimensional regression. Available at [arXiv:1306.3171](#).
- [23] JAVANMARD, A. and MONTANARI, A. (2013). Hypothesis testing in high-dimensional regression under the Gaussian random design model: Asymptotic theory. Available at [arXiv:1301.4240v1](#).
- [24] JUDITSKY, A., KILIÇ KARZAN, F., NEMIROVSKI, A. and POLYAK, B. (2012). Accuracy guarantees for ℓ_1 recovery of block-sparse signals. *Ann. Statist.* **40** 3077–3107. [MR3097970](#)
- [25] KNIGHT, K. and FU, W. (2000). Asymptotics for lasso-type estimators. *Ann. Statist.* **28** 1356–1378. [MR1805787](#)
- [26] LEDERER, J. and VAN DE GEER, S. (2014). New concentration inequalities for suprema of empirical processes. *Bernoulli*. To appear. Available at [arXiv:1111.3486](#).
- [27] LI, K.-C. (1989). Honest confidence regions for nonparametric regression. *Ann. Statist.* **17** 1001–1008. [MR1015135](#)
- [28] MEIER, L., VAN DE GEER, S. and BÜHLMANN, P. (2008). The group Lasso for logistic regression. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 53–71. [MR2412631](#)
- [29] MEINSHAUSEN, N. (2013). Assumption-free confidence intervals for groups of variables in sparse high-dimensional regression. Available at [arXiv:1309.3489](#).
- [30] MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist.* **34** 1436–1462. [MR2278363](#)

- [31] MEINSHAUSEN, N. and BÜHLMANN, P. (2010). Stability selection. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **72** 417–473. [MR2758523](#)
- [32] MEINSHAUSEN, N., MEIER, L. and BÜHLMANN, P. (2009). p -values for high-dimensional regression. *J. Amer. Statist. Assoc.* **104** 1671–1681. [MR2750584](#)
- [33] MEINSHAUSEN, N. and YU, B. (2009). Lasso-type recovery of sparse representations for high-dimensional data. *Ann. Statist.* **37** 246–270. [MR2488351](#)
- [34] NEGAHBAN, S. N., RAVIKUMAR, P., WAINWRIGHT, M. J. and YU, B. (2012). A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers. *Statist. Sci.* **27** 538–557. [MR3025133](#)
- [35] NICKL, R. and VAN DE GEER, S. (2013). Confidence sets in sparse regression. *Ann. Statist.* **41** 2852–2876. [MR3161450](#)
- [36] PORTNOY, S. (1987). A central limit theorem applicable to robust regression estimators. *J. Multivariate Anal.* **22** 24–50. [MR0890880](#)
- [37] PÖTSCHER, B. M. (2009). Confidence sets based on sparse estimators are necessarily large. *Sankhyā* **71** 1–18. [MR2579644](#)
- [38] PÖTSCHER, B. M. and LEEB, H. (2009). On the distribution of penalized maximum likelihood estimators: The LASSO, SCAD, and thresholding. *J. Multivariate Anal.* **100** 2065–2082. [MR2543087](#)
- [39] RASKUTTI, G., WAINWRIGHT, M. J. and YU, B. (2010). Restricted eigenvalue properties for correlated Gaussian designs. *J. Mach. Learn. Res.* **11** 2241–2259. [MR2719855](#)
- [40] ROBINSON, P. M. (1988). Root- N -consistent semiparametric regression. *Econometrica* **56** 931–954. [MR0951762](#)
- [41] SHAH, R. D. and SAMWORTH, R. J. (2013). Variable selection with error control: Another look at stability selection. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **75** 55–80. [MR3008271](#)
- [42] SUN, T. and ZHANG, C.-H. (2012). Scaled sparse linear regression. *Biometrika* **99** 879–898. [MR2999166](#)
- [43] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **58** 267–288. [MR1379242](#)
- [44] VAN DE GEER, S. (2007). The deterministic Lasso. In *JSM Proceedings*, 2007, 140. Am. Statist. Assoc., Alexandria, VA.
- [45] VAN DE GEER, S., BÜHLMANN, P., RITOV, Y. and DEZEURE, R. (2014). Supplement to “On asymptotically optimal confidence regions and tests for high-dimensional models.” DOI:10.1214/14-AOS1221SUPP.
- [46] VAN DE GEER, S. and MÜLLER, P. (2012). Quasi-likelihood and/or robust estimation in high dimensions. *Statist. Sci.* **27** 469–480. [MR3025129](#)
- [47] VAN DE GEER, S. A. (2008). High-dimensional generalized linear models and the lasso. *Ann. Statist.* **36** 614–645. [MR2396809](#)
- [48] VAN DE GEER, S. A. and BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.* **3** 1360–1392. [MR2576316](#)
- [49] WAINWRIGHT, M. J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans. Inform. Theory* **55** 2183–2202. [MR2729873](#)
- [50] WASSERMAN, L. and ROEDER, K. (2009). High-dimensional variable selection. *Ann. Statist.* **37** 2178–2201. [MR2543689](#)
- [51] ZHANG, C.-H. and HUANG, J. (2008). The sparsity and bias of the LASSO selection in high-dimensional linear regression. *Ann. Statist.* **36** 1567–1594. [MR2435448](#)
- [52] ZHANG, C.-H. and ZHANG, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **76** 217–242. [MR3153940](#)

- [53] ZHAO, P. and YU, B. (2006). On model selection consistency of Lasso. *J. Mach. Learn. Res.* 7 2541–2563. MR2274449

S. VAN DE GEER
P. BÜHLMANN
R. DEZEURE
SEMINAR FÜR STATISTIK
ETH ZÜRICH
RÄMISTRASSE 101
8092 ZÜRICH
SWITZERLAND
E-MAIL: geer@stat.math.ethz.ch
buhlmann@stat.math.ethz.ch
dezeure@stat.math.ethz.ch

Y. RITOV
DEPARTMENT OF STATISTICS
THE HEBREW UNIVERSITY OF JERUSALEM
MOUNT SCOPUS, JERUSALEM 91905
ISRAEL
E-MAIL: yaacov.ritov@gmail.com