

**n° 2014-05**  
**On the Prediction  
Performance of the Lasso**

**A. S. DALALYAN<sup>1</sup>**  
**M. HEBIRI<sup>2</sup>**  
**J. LEDERER<sup>3</sup>**

February 2014

Les documents de travail ne reflètent pas la position du CREST et n'engagent que leurs auteurs.  
Working papers do not reflect the position of CREST but only the views of the authors.

---

<sup>1</sup> CREST - ENSAE

<sup>2</sup> Université Paris Est

<sup>3</sup> Cornell University, USA.

# ON THE PREDICTION PERFORMANCE OF THE LASSO

BY ARNAK S. DALALYAN, MOHAMED HEBIRI, AND JOHANNES LEDERER

*ENSAE-CREST, Université Paris Est, and Cornell University*

Although the Lasso has been extensively studied, the relationship between its prediction performance and the correlations of the covariates is not fully understood. In this paper, we give new insights into this relationship in the context of multiple linear regression. We show, in particular, that the incorporation of a simple correlation measure into the tuning parameter leads to a nearly optimal prediction performance of the Lasso even for highly correlated covariates. However, we also reveal that for moderately correlated covariates, the prediction performance of the Lasso can be mediocre irrespective of the choice of the tuning parameter. For the illustration of our approach with an important application, we deduce nearly optimal rates for the least-squares estimator with total variation penalty.

---

*AMS 2000 subject classifications:* Primary 62G08, ; secondary 62C20,62G05,62G20

*Keywords and phrases:* multiple linear regression, sparse recovery, total variation penalty, oracle inequalities

**1. Introduction.** In recent years, considerable effort has been devoted to establishing sharp theoretical guarantees for the prediction performance<sup>1</sup> of the Lasso [Tib96]. Although there are already risk bounds for a variety of settings [Kol11, BvdG11], the prediction performance of the Lasso is still not completely understood. In this paper, we review and improve the sharpest known risk bounds to gain new insight into the prediction performance of the Lasso.

Our approach is valid for a broad class of models, but to avoid digression, we study the prediction performance of the Lasso only for Gaussian linear regression models with deterministic design. More specifically, we consider data consisting of  $n$  random observations  $y_1, \dots, y_n \in \mathbb{R}$  and  $p$  fixed covariates  $\mathbf{x}^1, \dots, \mathbf{x}^p \in \mathbb{R}^n$ . We further assume that there is a regression vector  $\boldsymbol{\beta}^* \in \mathbb{R}^p$  and a noise level  $\sigma^* > 0$  such that the residuals  $y_i - \beta_1^*(\mathbf{x}^1)_i - \dots - \beta_p^*(\mathbf{x}^p)_i$  are identically and independently distributed according to a centered Gaussian distribution with variance  $\sigma^{*2}$ . In vector notation, this reads

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\xi}, \quad \boldsymbol{\xi} \sim \sigma^* \mathcal{N}_n(0, \mathbf{I}_n), \quad (1)$$

where  $\mathbf{y} := (y_1, \dots, y_n)^\top \in \mathbb{R}^n$  is the response vector,  $\mathbf{X} := (\mathbf{x}^1, \dots, \mathbf{x}^p) \in \mathbb{R}^{n \times p}$  the design matrix (for which we assume, without loss of generality, that  $\|\mathbf{x}^j\|_2^2 \leq n$  for all  $j \in \{1, \dots, p\}$ ),  $\boldsymbol{\xi} \in \mathbb{R}^n$  the noise vector, and  $\mathbf{I}_n$  denotes the identity matrix in  $\mathbb{R}^{n \times n}$ . To keep the exposition simple, we restrict ourselves to Gaussian distributions for the noise vector and to fixed covariates, but our results extend to more general classes of distributions and, if the results are understood conditionally on the covariates, also hold for random covariates. Next, we recall that the Lasso is any solution of the convex optimization problem

$$\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}} \in \arg \min_{\boldsymbol{\beta}} \left\{ \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1 \right\}, \quad (2)$$

that can be efficiently solved even for very large values of  $p$  and  $n$  [EHJT04, BJMO12]. The magnitude of the tuning parameter  $\lambda > 0$  determines the amount of penalization and, therefore, has a crucial influence on the performance of the Lasso. Moreover, in particular for high dimensional models where  $p > n$ , the Lasso is typically not unique. This is irrelevant for the following, however: due to the (strict) convexity of the two terms in (2), any two minima  $\widehat{\boldsymbol{\beta}}_\lambda, \widehat{\boldsymbol{\beta}}'_\lambda$  fulfill  $\mathbf{X}\widehat{\boldsymbol{\beta}}_\lambda = \mathbf{X}\widehat{\boldsymbol{\beta}}'_\lambda$  and therefore yield equal prediction results.

In this paper, we study the prediction performance of the Lasso to answer the following four questions.

1. Numerous empirical results observed on synthetic data lead to the conjecture that the Lasso estimator used with the universal choice of the tuning parameter  $\lambda = \sqrt{2 \log(p)/n}$  has a prediction loss at least proportional to  $\frac{\log(p)}{n} \times \text{rank}(\mathbf{X})$ , where  $\text{rank}(\mathbf{X})$  is the rank of  $\mathbf{X}$ . However, as far as we know, no theoretical result corroborates this conjecture.
2. For sparse vectors  $\boldsymbol{\beta}^*$  with support  $J^* = \{j \in \{1, \dots, p\} : \beta_j^* \neq 0\}$  and for covariates that are strongly correlated in the sense that all irrelevant covariates  $\{\mathbf{x}^j : j \notin J^*\}$  are close to the linear span of relevant covariates  $\{\mathbf{x}^j : j \in J^*\}$ , empirical results

---

<sup>1</sup>Following [BRT09], the term prediction performance is understood here as the magnitude of the risk measured by the prediction loss  $\frac{1}{n} \|\mathbf{X}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|_2^2$ . This is not to be confused with the more common understanding of this term in machine learning literature, where the aim is to predict the label of a new unlabeled feature.

suggest that the smallest prediction loss is obtained choosing a tuning parameter  $\lambda$  that is substantially smaller than the universal one. The influence of correlations on the prediction performance of the Lasso was first considered in [vdGL13, HL13], where tuning parameters smaller than the classical ones are suggested if the covariates are correlated. In particular, for rates of convergence substantially faster than the slow rate  $(s^*/n)^{1/2}$  (where  $s^* = |J^*|$ ), their results suggest to incorporate the geometry of the covariates via a function of the entropy numbers of the symmetric convex hull of the covariates into the tuning parameters. These entropy numbers may be difficult to compute, however. Therefore, we aim at finding tuning parameters that incorporate the geometry of the covariates via an easily computable quantity and that lead to fast rates for prediction when the covariates are strongly correlated.

3. For really sparse vectors, that is, for  $s^*$  considerably smaller than  $n$  (for example,  $s^*$  is fixed and  $n \rightarrow \infty$ ), there are methods that satisfy fast rate bounds for prediction irrespective of the correlations of the covariates [BTW07a, DT07, RT11, DT12b, DT12a]. Fast rate bounds for Lasso prediction, in contrast, usually rely on assumptions on the correlations of the covariates such as low coherence [CP09], restricted eigenvalues [BRT09, RWY10], restricted isometry [CT07], compatibility [vdG07], *etc.*. For Lasso prediction, it is therefore not known whether fast rate bounds are available irrespective of the correlations of the covariates. This question is open even if we allow for oracle choices of the tuning parameter  $\lambda$ , that is, if we allow for  $\lambda$  that depend on the true regression vector  $\beta^*$ , the noise vector  $\xi$ , and the noise level  $\sigma^*$ .
4. Finally, known results imply fast rates for prediction with the Lasso in the following two extreme cases: First, when the covariates are mutually orthogonal, and second, when the covariates are all collinear. But how far from these two extreme cases can a design be such that it still permits fast rates for prediction with the Lasso? For the first case, the case of mutually orthogonal covariates, this question has been thoroughly studied [BRT09, BTW07b, Zha09, vdGB09, Wai09, CWX10, JN11a]. For the second case, the case of collinear covariates, this question has received much less attention and is therefore one of our main topics.

Let us stress that we focus only on the behavior of the Lasso in terms of the prediction loss and do not explore here such important aspects of the Lasso as variable selection and estimation.

1.1. *Notation.* Throughout the paper, for every integer  $k \in \mathbb{N}$ , we set  $[k] = \{1, \dots, k\}$ . For every  $q \in [0, \infty]$ , we denote by  $\|\mathbf{u}\|_q$  the usual  $\ell_q$ -(quasi)norm of a vector  $\mathbf{u} \in \mathbb{R}^k$ , that is

$$\|\mathbf{u}\|_q = \begin{cases} \text{Card}(\{j : u_j \neq 0\}), & q = 0, \\ (\sum_{j \in [k]} |u_j|^q)^{1/q}, & 0 < q < \infty, \\ \max_{j \in [k]} |u_j|, & q = \infty. \end{cases}$$

For any set  $T \subset [p]$ , we denote by  $T^c$  and  $|T|$  the complementary set  $[p] \setminus T$  and the cardinality of  $T$ , respectively. For every matrix  $\mathbf{A} \in \mathbb{R}^{p \times q}$  and any subset  $T$  of  $[q]$ , we denote by  $\mathbf{A}_T$  the matrix obtained from  $\mathbf{A}$  by removing all the columns belonging to  $T^c$ . For a vector  $\mathbf{u} \in \mathbb{R}^p$  and a set  $T \subset [p]$ ,  $\mathbf{u}_T$  is the vector obtained from  $\mathbf{u}$  by removing all the coordinates belonging to  $T^c$ . The transpose and the Moore-Penrose pseudoinverse of a matrix  $\mathbf{A}$  are denoted by  $\mathbf{A}^\top$  and  $\mathbf{A}^\dagger$ , respectively. For two vectors  $\mathbf{u}$  and  $\mathbf{u}'$  of the same dimension  $p$ , we define  $\odot$

as the coordinatewise product, that is  $\mathbf{u} \odot \mathbf{u}' = (u_1 u'_1, \dots, u_p u'_p)^\top$ . We write  $\mathbf{1}_p$  (resp.  $\mathbf{0}_p$ ) for the vector of  $\mathbb{R}^p$  having all coordinates equal to one (resp. zero). For the design matrix  $\mathbf{X}$  and any subset  $T$  of  $[p]$ , we denote by  $V_T$  the linear subspace of  $\mathbb{R}^n$  spanned by the columns of  $\mathbf{X}_T$ . Further, we denote by  $\Pi_T$  the orthogonal projector onto  $V_T$  and by  $\rho_T$  the maximal Euclidean distance between the normalized columns of  $\mathbf{X}$  and the set  $V_T$ , that is  $\rho_T = \max_{j \in [p]} \min_{\mathbf{v} \in V_T} \|\mathbf{v} - n^{-1/2} \mathbf{x}^j\|_2 = n^{-1/2} \max_{j \in [p]} \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2$ . For two vectors  $\boldsymbol{\beta}, \boldsymbol{\beta}' \in \mathbb{R}^p$ , we denote by  $\ell_n(\boldsymbol{\beta}, \boldsymbol{\beta}')$  the prediction loss  $\frac{1}{n} \|\mathbf{X}(\boldsymbol{\beta} - \boldsymbol{\beta}')\|_2^2$ . In all the asymptotic considerations, we will write  $a_n \lesssim b_n$  for two positive sequences  $(a_n)$  and  $(b_n)$  when for some  $c \in (0, \infty)$  it holds that  $\overline{\lim}_{n \rightarrow \infty} (a_n/b_n) \leq c$ . Further, we write  $a_n \asymp b_n$  for two sequences  $(a_n)$  and  $(b_n)$  which are of the same order, that is  $a_n \lesssim b_n \lesssim a_n$ .

**1.2. Outline of the paper.** The rest of this work is organized as follows. The next section presents some new risk bounds for the prediction risk of the Lasso under no condition on the covariates. These results provide an answer to the first question above. We also present in Section 2 an example showing that for some particularly unfavorable design matrices it is impossible to get rates faster than  $1/\sqrt{n}$ , even if  $|J^*|$  is very small. Section 3 is devoted to some refinements of the sharp sparsity oracle inequalities with fast rates based on compatibility factors. They imply, in particular, that the total variation estimator of piecewise constant signals is nearly rate optimal. A bunch of results, establishing “slow” rates that involve the quantity  $\rho_T$ , accounting for the severity of the correlations within covariates, are developed in Section 4. In particular, they allow us to answer the second and the fourth questions raised in the Introduction. Discussion with related work and some remarks on the computational complexity of the compatibility factors are placed in Section 5. We summarize the contributions of this work and outline some open questions in Section 6. The proofs of all the results stated in the paper are deferred to Section 7.

**2. Statistical complexity of the Lasso without any condition on the design.** The goal here is to present some new results concerning the accuracy of the Lasso in terms of the prediction loss when almost no assumption on the relationship between the covariates is required. In particular, we will show that the estimator  $\mathbf{X} \widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}$  of the mean  $\mathbf{X} \boldsymbol{\beta}^*$  of the vector  $\mathbf{y}$ , when projected on a subspace of  $\mathbb{R}^n$  spanned by a small number of columns of  $\mathbf{X}$ , achieves fast rates of convergence provided that  $\lambda$  is of the order of  $n^{-1/2}$ . This will be complemented in Section 4, where we establish new results characterizing the so called slow rates for the Lasso and show that, in some circumstances, these rates may be significantly faster than  $(s^*/n)^{1/2}$ .

**2.1. Fast rates for Lasso projections.** We begin by discussing one of the main points that contrasts our approach with the previous ones used in the literature. Let  $T$  be a subset of  $[p]$  which may be the set of relevant covariates or any other set. Let  $\Pi_T = \mathbf{X}_T (\mathbf{X}_T^\top \mathbf{X}_T)^\dagger \mathbf{X}_T^\top$  be the orthogonal projector onto the subspace of  $\mathbb{R}^n$  spanned by the columns of  $\mathbf{X}_T$ . An idea underpinning our results below is that when only noisy observations of the vector  $\mathbf{X} \boldsymbol{\beta}^*$  are available, it is practically impossible to make the difference between the true vector  $\boldsymbol{\beta}^*$  and the vector  $\boldsymbol{\beta}^{*,T}$  defined by the relations

$$\boldsymbol{\beta}_T^{*,T} = \boldsymbol{\beta}_T^* + (\mathbf{X}_T^\top \mathbf{X}_T)^\dagger \mathbf{X}_T^\top \boldsymbol{\xi} \quad \text{and} \quad \boldsymbol{\beta}_{T^c}^{*,T} = \boldsymbol{\beta}_{T^c}^*.$$

In fact, one easily checks that

$$\mathbf{y} = \mathbf{X}_T \boldsymbol{\beta}_T^* + \mathbf{X}_{T^c} \boldsymbol{\beta}_{T^c}^* + \boldsymbol{\xi} = \mathbf{X} \boldsymbol{\beta}^{*,T} + (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}. \quad (3)$$

When the rank of  $\Pi_T$  is much smaller than the sample size  $n$ , the noise vectors  $\boldsymbol{\xi}$  and  $(\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}$  exhibit similar behavior. Therefore, both  $\boldsymbol{\beta}^*$  and  $\boldsymbol{\beta}^{*,T}$  may be seen as the signal part of the noisy observation  $\mathbf{y}$ . In what follows, we exploit this idea in order to establish oracle inequalities<sup>2</sup> on the prediction error  $\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^{*,T})$  and some other related quantities. Since  $\boldsymbol{\beta}^{*,T}$  is merely a perturbation of  $\boldsymbol{\beta}^*$ , all the bounds proved for  $\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^{*,T})$  carry over similar bounds on the conventional prediction loss  $\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*)$ .

**THEOREM 1.** *Let  $T$  be any subset of  $[p]$  and let<sup>3</sup>  $\nu_T = \inf_{\mathbf{u} \in \mathbb{R}^{|T|}} \frac{\sqrt{|T|} \cdot \|\mathbf{X}_T \mathbf{u}\|_2}{\sqrt{n} \|\mathbf{u}\|_1}$ . For every  $\lambda > 0$ , it holds that*

$$\frac{1}{n} \|\Pi_T \mathbf{X} (\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}} - \boldsymbol{\beta}^{*,T})\|_2^2 \leq \frac{\lambda^2 |T|}{\nu_T^2}.$$

A remarkable fact is that the claim of the foregoing theorem is valid under very weak assumptions on the design matrix  $\mathbf{X}$ , for every value of the tuning parameter  $\lambda > 0$  and whatever the noise vector  $\boldsymbol{\xi}$  is. An immediate consequence of this result that follows from the triangle inequality is

$$\frac{1}{\sqrt{n}} \|\Pi_T \mathbf{X} (\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}} - \boldsymbol{\beta}^*)\|_2 \leq \left( \frac{\lambda \sqrt{|T|}}{\nu_T} + \frac{\|\Pi_T \boldsymbol{\xi}\|_2}{\sqrt{n}} \right), \quad \forall T \subset [p]. \quad (4)$$

Note that the vector  $\Pi_T \boldsymbol{\xi} / \sqrt{n}$  appearing in the last term in this inequality is exactly equal to the stochastic error of the least squares estimator when only the covariates  $\{\mathbf{x}^j : j \in T\}$  are considered as relevant. The Euclidean norm of this vector is typically of the order of  $\sigma^* \sqrt{|T|/n}$  and represents a lower bound on the risk when no information other than  $|T|$ -sparsity of  $\boldsymbol{\beta}^*$  is available. Since it is usually recommended to choose  $\lambda$  not larger than  $\sigma^* \sqrt{2 \log(p/\delta)/n}$ , for some prescribed tolerance level  $\delta \in (0, 1)$ , we conclude that  $\Pi_T \mathbf{X} \widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}$  estimates the vector  $\Pi_T \mathbf{X} \boldsymbol{\beta}^*$  with the fast rate of convergence  $\sigma^* \sqrt{|T| \log(p)/n}$ .

Relation (4) also demonstrates that the prediction loss of the Lasso decreases to zero at the fast rate of convergence  $s \log(p)/n$  in some particular cases with strongly correlated covariates. This result is stated in the following proposition.

**PROPOSITION 1.** *If there is a subset  $T$  of  $[p]$  of cardinality  $s$  such that all the covariates  $\{\mathbf{x}^j : j \in T^c\}$  belong to the linear span of  $\{\mathbf{x}^j : j \in T\}$ , then for every  $\lambda > 0$*

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*)^{1/2} \leq \frac{\lambda \sqrt{s}}{\nu_T} + \frac{\|\Pi_T \boldsymbol{\xi}\|_2}{\sqrt{n}}. \quad (5)$$

*In particular, for every random vector  $\boldsymbol{\xi}$  such that  $\mathbf{E}[\boldsymbol{\xi}] = 0$  and  $\max_i \mathbf{E}[\xi_i^2] \leq \sigma^{*2}$ ,*

$$\mathbf{E}[\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*)] \leq \frac{2\lambda^2 s}{\nu_T^2} + \frac{2\sigma^{*2} s}{n}. \quad (6)$$

<sup>2</sup>We refer the reader to [BTW07b] for an introduction to sparsity oracle inequalities.

<sup>3</sup>It follows from the Cauchy-Schwarz inequality that  $\nu_T$  is not smaller than the smallest singular value of the matrix  $\frac{1}{\sqrt{n}} \mathbf{X}_T$ .

If, in addition,  $\boldsymbol{\xi} \sim \sigma^* \mathcal{N}_n(0, \mathbf{I}_n)$ , then with probability at least  $1 - \delta$

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq \frac{2\lambda^2 s}{\nu_T^2} + \frac{4\sigma^{*2}(s + 2\log(1/\delta))}{n}. \quad (7)$$

The first two claims of this proposition trivially follow from (4), while the third claim follows from (5) using the fact that  $\|\Pi_T \boldsymbol{\xi}\|_2^2$  is drawn from the chi-squared distribution  $\chi_s^2$  in conjunction with the well-known results on the tails of the latter.

This proposition answers to the first question raised in the introduction concerning the performance of the Lasso as a function of the rank of  $\mathbf{X}$  when the latter is small as compared to  $n$ . In fact, let us denote by  $\bar{\nu}_r$  the maximal value of  $\nu_T$  over all possible subsets of  $[p]$  of cardinality  $r = \text{rank}(\mathbf{X})$ :  $\bar{\nu}_r = \max_{T:|T|=r} \nu_T$ . It follows from (7) that when  $\boldsymbol{\xi}$  is Gaussian and  $\lambda = \sqrt{2\log(p)/n}$ , for every  $\delta \in (0, 1)$ , with probability  $1 - \delta$ ,

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq \frac{4\log(p)\text{rank}(\mathbf{X})}{n\bar{\nu}_r^2} + \frac{4\sigma^{*2}(\text{rank}(\mathbf{X}) + 2\log(1/\delta))}{n}. \quad (8)$$

**2.2. Limits of fast rates: an example.** Unfortunately, in sparsity oracle inequalities with fast rates similar to (4) one cannot replace  $\Pi_T$  by  $\mathbf{I}_n$  without further assumptions on  $\mathbf{X}$ . To the best of our knowledge, this fact—although not so surprising—has never been formally proved in the literature. Here, we provide an example of matrix  $\mathbf{X}$  and a 2-sparse vector  $\boldsymbol{\beta}^*$  for which the prediction loss  $\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*)$  of the Lasso is at best of the order of  $n^{-1/2}$  with probability  $1/2$ , whatever the tuning parameter is.

**EXAMPLE 1.** Let  $n \geq 2$  be an integer. We set  $m$  to be the largest integer less than  $\sqrt{2n}$  and define the design matrix  $\mathbf{X} \in \mathbb{R}^{n \times 2m}$  by

$$\mathbf{X} = \sqrt{\frac{n}{2}} \begin{pmatrix} \mathbf{1}_m^\top & \mathbf{1}_m^\top \\ \mathbf{I}_m & -\mathbf{I}_m \\ \mathbf{0}_{(n-m-1) \times m} & \mathbf{0}_{(n-m-1) \times m} \end{pmatrix}.$$

If we denote by  $\{\mathbf{e}_j : j \in [n]\}$  the canonical basis of  $\mathbb{R}^n$ , the columns of this matrix are of the form  $\mathbf{x}^j = \sqrt{n/2}(\mathbf{e}_1 + \mathbf{e}_{j+1})$  for  $j = 1, \dots, m$  and  $\mathbf{x}^j = \sqrt{n/2}(\mathbf{e}_1 - \mathbf{e}_{j-m+1})$  for  $j = m+1, \dots, 2m$ . To avoid unnecessary technicalities, we assume in this example that the noise vector is composed of i.i.d. Rademacher random variables, that is  $\mathbf{P}(\boldsymbol{\xi} = \mathbf{s}) = 2^{-n}$  for every  $\mathbf{s} \in \{\pm 1\}^n$  (thus  $\sigma^* = 1$ ). Let the true regression vector be  $\boldsymbol{\beta}^* \in \mathbb{R}^{2m}$  such that  $\beta_1^* = \beta_{m+1}^* = 1$  and  $\beta_j^* = 0$  for every  $j \in [2m] \setminus \{1, m+1\}$ .

**PROPOSITION 2.** For any  $\lambda > 0$ , the prediction loss of the Lasso  $\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}$  satisfies the inequality

$$\mathbf{P}\left(\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \geq \frac{1}{2\sqrt{2n}}\right) \geq \frac{1}{2}.$$

There are at least three reasons that make this example particularly instructive. First, it shows that the correlations between the covariates need not be close to  $\pm 1$  to cause the failure of the fast rates. Even in the case of small fixed correlations the rate of convergence of the Lasso in prediction loss may be not smaller than  $Cn^{-1/2}$ . Second, the foregoing result is true for

every  $\lambda > 0$ . Thus, even an oracle choice of  $\lambda$  cannot prevent slow rates. Third, it is valid for a small value of sparsity index: the  $\ell_0$ -norm of  $\beta^*$  is equal to 2. In the literature, other examples on which the Lasso fails to achieve fast rates have been proposed (see Section 2 in [CP09]), however, to the best of our knowledge, this is the first counter-example in which such a result is analytically proved for fixed sparsity, fixed correlations, any value of  $\lambda$  and a  $\beta^*$  independent of  $n$ .

This example clearly demonstrates the limits of the Lasso as a method of prediction. While for several other prediction procedures [DT07, RT11, DT12b, DT12a] fast rates are valid without any condition on the correlations between the predictors, some relatively strong assumptions are necessary for the Lasso to achieve fast rates. It should be noted in defense of the Lasso that it presents major advantages in terms of computational complexity.

**3. Fast rates under relaxed compatibility assumptions on the design matrix.** To the best of our knowledge, the sharpest oracle inequality for the Lasso available in the literature is the one presented in [SZ12]. We begin by stating their result<sup>4</sup> in order to discuss what can be learnt from it concerning the questions presented in the introduction. Then, we state a new oracle inequality that combines the proof of [SZ12] and the idea of estimating  $\beta^{*,T}$  instead of  $\beta^*$  in order to get some improvements.

For every set  $T \subset [p]$  and any  $\bar{c} > 0$  we recall the definition of the compatibility factor  $\kappa_{T,\bar{c}} \geq 0$ :

$$\kappa_{T,\bar{c}} = \inf_{\delta \in \mathbb{R}^p: \|\delta_{T^c}\|_1 < \bar{c}\|\delta_T\|_1} \frac{|T| \cdot \|\mathbf{X}\delta\|_2^2}{n(\|\delta_T\|_1 - \bar{c}^{-1}\|\delta_{T^c}\|_1)^2}. \quad (9)$$

**THEOREM 2** ([SZ12], Theorem 4). *Let  $\delta \in (0, 1)$  be a fixed tolerance level. If for some  $\gamma > 1$ , the tuning parameter of the Lasso satisfies  $\lambda = \gamma\sigma^* \left(\frac{2}{n} \log(p/\delta)\right)^{1/2}$ , then with probability at least  $1 - \delta$ ,*

$$\ell_n(\widehat{\beta}_\lambda^{\text{Lasso}}, \beta^*) \leq \inf_{\bar{\beta} \in \mathbb{R}^p, T \subset [p]} \left\{ \ell_n(\bar{\beta}, \beta^*) + 4\lambda \|\bar{\beta}_{T^c}\|_1 + \frac{2(1+\gamma)^2 \sigma^{*2} |T| \log(p/\delta)}{n\kappa_{T,(\gamma+1)/(\gamma-1)}} \right\}.$$

An important feature of this inequality is its sharpness, reflected by the fact that the constant in front of the infimum, often referred to as the leading constant of an oracle inequality (OI), is equal to one. The first sharp OI with fast rate of convergence of the remainder term has been proved in [KLT11]. It was then refined and extended to the procedure square-root Lasso (also known as the scaled Lasso) in [SZ12].

Two particularly interesting consequences of this inequality can be obtained as follows. First, by choosing  $\bar{\beta}$  to be any sparse vector and  $T$  to be the support  $J(\bar{\beta})$  of  $\bar{\beta}$ , we get

$$\ell_n(\widehat{\beta}_\lambda^{\text{Lasso}}, \beta^*) \leq \min_{\bar{\beta} \in \mathbb{R}^p} \left\{ \ell_n(\bar{\beta}, \beta^*) + \frac{2(1+\gamma)^2 \sigma^{*2} \|\bar{\beta}\|_0 \log(p/\delta)}{n\kappa_{J(\bar{\beta}),(\gamma+1)/(\gamma-1)}} \right\}. \quad (10)$$

<sup>4</sup>The results stated below do not match exactly with those stated in [SZ12], but they can be easily deduced from the proofs in [SZ12]

Second, choosing  $\bar{\boldsymbol{\beta}} = \boldsymbol{\beta}^*$  and  $T = J_s^*$ , the set of  $s$  largest in absolute value coefficients of  $\boldsymbol{\beta}^*$ , we get

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq 4\gamma\sigma^* \left(\frac{2}{n} \log(p/\delta)\right)^{1/2} \|\boldsymbol{\beta}_{(J_s^*)^c}^*\|_1 + \frac{2(1+\gamma)^2\sigma^{*2}s \log(p/\delta)}{n\kappa_{J_s^*, (\gamma+1)/(\gamma-1)}}. \quad (11)$$

Finally, let us mention that an intermediate result in [SZ12] implies that if the tuning parameter of the Lasso satisfies  $\lambda \geq \sigma^* \left(\frac{2}{n} \log(p/\delta)\right)^{1/2}$ , then

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq \min_{\bar{\boldsymbol{\beta}}} \left\{ \ell_n(\bar{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) + 4\lambda \|\bar{\boldsymbol{\beta}}\|_1 \right\}. \quad (12)$$

This kind of inequalities are frequently referred to as sparsity oracle inequalities with *slow rates*. They have the advantage of requiring no assumption on the relationship between the covariates but provides rates of convergence of the order of  $\sigma^* \left(\frac{s \log p}{n}\right)^{1/2}$  which in typical situations are slower than those available for orthogonal design matrices.

Let us state now some refinements of Theorem 2. For any subset  $T$  of  $[p]$ , let us introduce the weights<sup>5</sup>

$$\omega_j(T, \mathbf{X}) = \frac{1}{\sqrt{n}} \|(\mathbf{I}_n - \Pi_T)\mathbf{x}^j\|_2, \quad \bar{\omega}_j(T, \mathbf{X}) = \frac{\omega_j(T, \mathbf{X})}{\max_{\ell \in [p]} \omega_\ell(T, \mathbf{X})}, \quad \forall j \in [p]. \quad (13)$$

Since  $\mathbf{x}^j$  are normalized to have an  $\ell_2$  norm at most equal to  $\sqrt{n}$ , the weights  $\omega_j(T, \mathbf{X})$  are all between zero and one. Furthermore, they vanish whenever  $\mathbf{x}^j$  belongs to the linear span of  $\{\mathbf{x}^\ell, \ell \in T\}$ . In particular,  $\omega_j(T, \mathbf{X}) = 0$  for every  $j \in T$ . Using these weights and any  $\gamma > 0$ , we define the sets

$$\mathcal{C}_0(T, \gamma, \boldsymbol{\omega}) = \left\{ \boldsymbol{\delta} \in \mathbb{R}^p : \|(\mathbf{1}_p - \gamma^{-1}\boldsymbol{\omega})_{T^c} \odot \boldsymbol{\delta}_{T^c}\|_1 < \|\boldsymbol{\delta}_T\|_1 \right\}.$$

When  $\boldsymbol{\omega} = \mathbf{1}_p$ , we write  $\mathcal{C}_0(T, \gamma)$  instead of  $\mathcal{C}_0(T, \gamma, \boldsymbol{\omega})$ .

**DEFINITION 1** (Compatibility factors). For every vector  $\boldsymbol{\omega} \in \mathbb{R}^p$  with nonnegative entries, we call the weighted compatibility factor the quantity

$$\bar{\kappa}_{T, \gamma, \boldsymbol{\omega}} = \inf_{\boldsymbol{\delta} \in \mathcal{C}_0(T, \gamma, \boldsymbol{\omega})} \frac{|T| \cdot \|\mathbf{X}\boldsymbol{\delta}\|_2^2}{n \left\{ \|\boldsymbol{\delta}_T\|_1 - \|(\mathbf{1}_p - \gamma^{-1}\boldsymbol{\omega})_{T^c} \odot \boldsymbol{\delta}_{T^c}\|_1 \right\}^2}.$$

The weighted compatibility factors with weights  $\boldsymbol{\omega}$  and  $\bar{\boldsymbol{\omega}}$  defined in (13) are particularly useful for explaining the accuracy of the Lasso as measured by the prediction loss. They relax the assumptions previously known in the literature that lead to fast rates.

**THEOREM 3.** *Let  $\delta \in (0, 1)$  be a fixed tolerance level. If for some value  $\gamma > 1$ , the tuning parameter of the Lasso satisfies  $\lambda = \gamma\sigma^* \sqrt{2 \log(p/\delta)/n}$ , then on an event of probability at least  $1 - 2\delta$ , the following bound holds:*

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq \inf_{\bar{\boldsymbol{\beta}} \in \mathbb{R}^p, T \subset [p]} \left\{ \ell_n(\bar{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) + 4\lambda \|\bar{\boldsymbol{\beta}}_{T^c}\|_1 + \frac{4\sigma^{*2}|T| \log(p/\delta)}{n} \cdot r_{n,p,T} \right\}, \quad (14)$$

<sup>5</sup>In the definition of  $\bar{\boldsymbol{\omega}}$ , we use the convention  $0/0 = 0$ .

where the remainder term is given by  $r_{n,p,T} = \log^{-1}(p/\delta) + 2|T|^{-1} + \gamma^2 \bar{\kappa}_{T,\gamma,\omega}^{-1}$ . Furthermore, if for some  $T \subset [p]$  and some  $\gamma > 1$ ,  $\lambda = \gamma \sigma^* \rho_T \sqrt{2 \log(p/\delta)/n}$ , then with probability at least  $1 - 2\delta$ , the following bound holds

$$\ell_n(\hat{\beta}_\lambda^{\text{Lasso}}, \beta^*) \leq \inf_{\beta} \left\{ \ell_n(\beta, \beta^*) + 4\lambda \|\beta_{T^c}\|_1 \right\} + \frac{4\sigma^{*2} \rho_T^2 |T| \log(p/\delta)}{n} \cdot \bar{r}_{n,p,T}, \quad (15)$$

where the remainder term is given by  $\bar{r}_{n,p,T} = \frac{(1+2|T|^{-1} \log(1/\delta))}{\rho_T^2 \log(p/\delta)} + \frac{\gamma^2}{\bar{\kappa}_{T,\gamma,\omega}}$ .

The main difference between inequalities (14) and (15) is the presence of the factor  $\rho_T^2$  in the numerator of the last term. This factor is always not larger than 1. However, in order to introduce it we needed to replace the compatibility factor  $\bar{\kappa}_{T,\gamma,\omega}$  by  $\bar{\kappa}_{T,\gamma,\bar{\omega}}$  and to deflate  $\lambda$  by the factor  $\rho_T$ . From a practical point of view, this last modification is not always easy to implement, since the quantity  $\rho_T$  depends on the set  $T$  which can be thought of as the best possible set of covariates. This set being unknown, the claim of (15) is to be interpreted as a theoretical justification for choosing the tuning parameter smaller than the universal value. Such a choice can be made, for instance, by cross validation. It is also possible to perform a sparse PCA on the set of covariates in order to choose a suitable value of  $\lambda$  (cf. Section 5.2 for more details).

**EXAMPLE 2** (Total variation penalty for piecewise constant functions). In image denoising and signal processing, total variation type penalties are often employed to enforce similarity between neighboring pixels or values of the signal. In the one-dimensional setting, the problem may be formulated as follows. Assume that a piecewise constant function  $f^* : [0, 1] \rightarrow \mathbb{R}$  is observed on the regular grid in a noisy environment:  $y_i = f^*(i/n) + \xi_i$ , for  $i = 1, \dots, n$ . Let us denote the unknown vector of values of  $f^*$  on the grid by  $\mathbf{f}^* = (f^*(1/n), \dots, f^*(1))^\top$  and define the total variation penalty of a vector  $\mathbf{f} \in \mathbb{R}^n$  by  $\|\mathbf{f}\|_{\text{TV}} = \sum_{i=1}^n |f_i - f_{i-1}|$  with the convention that  $f_0 = 0$ . Then, the TV-penalized least squares estimator of  $\mathbf{f}^*$  is defined as

$$\hat{\mathbf{f}}^{\text{TV}} \in \arg \min_{\mathbf{f} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{f}\|_2^2 + \lambda \|\mathbf{f}\|_{\text{TV}} \right\}, \quad (16)$$

where  $\lambda > 0$  is a tuning parameter. This estimator, hereafter referred to as TV-estimator, has been shown to be closely related to the Lasso [HLL07, HLL10]. More precisely, if we define the vector of differences  $\beta \in \mathbb{R}^n$  by  $\beta_j = f_j - f_{j-1}$ , then (16) is equivalent to (2) with the  $n \times n$  design matrix  $\mathbf{X} = (\mathbf{1}(i \geq j))_{i,j}$ . Despite its popularity in applications, it is very surprising that the TV-estimator and, more precisely, its prediction accuracy is not yet completely understood. In a purely asymptotic setting  $n \rightarrow \infty$ , [MvdG87] establish that the  $\ell_2$ -loss of  $\hat{\mathbf{f}}^{\text{TV}}$ , defined by  $\frac{1}{n} \|\hat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2$ , is of the order of  $n^{-2/3}$ . This is, however, just an upper bound on the risk of  $\hat{\mathbf{f}}^{\text{TV}}$ , and this upper bound is much worse than the optimal rate of convergence, known to be  $n^{-1}$  in the problem of estimating piecewise constant functions. This optimal rate is achieved, for instance, by the penalized least squares with a penalty proportional to the number of jumps, *i.e.*, the  $\ell_0$ -norm of the difference vector (see [BKL<sup>+</sup>09] and the references therein). The question whether it is possible or not to improve the rate  $n^{-2/3}$  for the TV-penalized estimator and, eventually, to achieve the minimax rate, has remained open so far.

On the other hand, more recent papers [HLL07, HLL10] propose nonasymptotic risk bounds for  $\widehat{\mathbf{f}}^{\text{TV}}$ . Without any assumption, they show that for  $\lambda \asymp n^{-1/2}$  their risk bound is of the order of  $n^{-1/2}$ . They also notice that if the TV-estimator with  $\lambda \asymp n^{-1}$  has only a few jumps, then its  $\ell_2$ -loss is of the optimal order  $n^{-1}$ . This result is, however, not very satisfactory since diminishing  $\lambda$  up to the order  $n^{-1}$  is quite likely to significantly increase the number of jumps in the TV-estimator. Put differently, there is no theoretical result assessing the probability of getting only a few jumps when  $\lambda \asymp n^{-1}$ . This raises some new questions: (a) Is it possible to establish sharp oracle inequalities for TV-estimator with optimal rate of convergence? (b) Is it really necessary to choose  $\lambda$  very small for achieving the optimal rate? (c) What is the rate of convergence in terms of the number of jumps, when the latter is allowed to increase with  $n$ ? In order to show that the theoretical tools developed in this section provide almost exhaustive answers to these questions, we need the following result.

**PROPOSITION 3.** *Let  $\mathbf{X}$  be the  $n \times n$  matrix with entries  $x_{ij} = \mathbf{1}(i \geq j)$  and let  $\mathbf{a} \in \mathbb{R}_+^n$  be a given vector of “weights”. For every  $T = \{j_1, \dots, j_s\} \subset [n]$  and for every  $\mathbf{u} \in \mathbb{R}^n$ , we have*

$$\|\mathbf{u}_T \odot \mathbf{a}_T\|_1 - \|\mathbf{u}_{T^c} \odot \mathbf{a}_{T^c}\|_1 \leq 4\|\mathbf{X}\mathbf{u}\|_2 \left( 2 \sum_{j \in [n]} |a_j - a_{j+1}|^2 + 2(s+1)\|\mathbf{a}\|_\infty^2 \Delta_{\min, T}^{-1} \right)^{1/2},$$

where  $\Delta_{\min, T} = \min_{\ell \in [s+1]} |j_{\ell+1} - j_\ell|$  with the convention  $j_0 = 1$  and  $j_{s+1} = 1$ .

The proof of this result, deferred to Section 7, is carried out using a completely new approach based on a probabilistic argument. We believe that this argument may be used in other situations for evaluating the compatibility factors theoretically. All the previous efforts for evaluating the compatibility and restricted eigenvalue constants, focused on weakly correlated designs (see, for instance, [BRT09]). In contrast with this, our approach provides bounds on compatibility factors even for strongly correlated designs. Indeed, many pairs of columns of matrix  $\mathbf{X}$  corresponding to the TV-estimator have correlation of the order of  $1 - n^{-1/2}$ .

For applying Proposition 3 to the TV-estimator, we choose  $a_j = 1$  for every  $j$  belonging to the set  $T$ , which presumably contains the jumps of  $f^*$ , and  $a_j = 1 - \frac{1}{2\sqrt{n}}\|(\mathbf{I}_n - \Pi_T)\mathbf{x}^j\|_2$ ,  $j \in T^c$ . In what follows, we denote by  $\Delta_{\min, T}$  the smallest distance between two jumps, that is  $\Delta_{\min, T} = \min_{\ell \in [s+1]} |j_\ell - j_{\ell-1}|$  with the convention that  $j_0 = 1$  and  $j_{s+1} = n + 1$ .

**PROPOSITION 4.** *Let  $\mathbf{f}^*$  be a piecewise constant vector and  $J^* = \{j \in [n] : f_j^* \neq f_{j+1}^*\}$ . If the tuning parameter satisfies  $\lambda = 2\sigma^* \{(2/n) \log(n/\delta)\}^{1/2}$ , then on an event of probability at least  $1 - 2\delta$ , the following bound holds for every nonempty  $T \subset [n]$ :*

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \inf_{\bar{\mathbf{f}} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\bar{\mathbf{f}} - \mathbf{f}^*\|_2^2 + 4\lambda \|\bar{\mathbf{f}}_{T^c}\|_{\text{TV}} \right\} + \frac{4\sigma^{*2}|T| \log(n/\delta)}{n} \cdot r_{n, T}, \quad (17)$$

where the bounded remainder term is given by  $r_{n, T} = 3 + 256(\log(n) + (n/\Delta_{\min, T}))$ .

The risk bound (17) drastically improves the results on the  $\ell_2$ -loss of the TV-estimator ever proved in the literature. Not only it holds for finite samples, is with a leading constant one and provides a risk bound valid with high probability, but, more importantly, it has a remainder term of the order of  $|T|(\log(n))^2/n$ . The rate of decay to zero of this term when  $n \rightarrow \infty$  is

much faster than what was known before and should be contrasted with  $n^{-2/3}$  established in [MvdG87]. More precisely, when the true function  $\mathbf{f}^*$  is piecewise constant on a partition of  $s$  intervals, taking in (17)  $\widehat{\mathbf{f}} = \mathbf{f}^*$  and  $T = \{j \in [n] : f_j^* \neq f_{j-1}^*\}$ , the terms in accolades at the right hand-side vanish and one gets the inequality

$$\mathbf{P}\left(\frac{1}{n}\|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \frac{4\sigma^{*2}|J^*|\log(n/\delta)}{n} \cdot (3 + 256(\log(n) + (n/\Delta_{\min, J^*})))\right) \geq 1 - 2\delta. \quad (18)$$

When the vector  $\mathbf{f} \in \mathbb{R}^n$  consists of the values of a piecewise constant function  $f$  at the points  $\{i/n, i \in [n]\}$ , of the regular grid, the term  $n/\Delta_{\min, J^*}$  is bounded by a constant (for  $n \rightarrow \infty$  and fixed  $f$ ). In this case, the upper bound in (18) is of the nearly optimal order. Furthermore, risk bound (18) holds for every  $|J^*|$ , even if it tends to infinity with  $n$ . To the best of our knowledge, this is the first result of this type. All the previous asymptotic results considered the number of jumps  $|J^*|$  as fixed. Moreover, our result is valid for the universal choice of the tuning parameter and not the very small one evoked in [HLL10]. To complete this discussion, let us mention that the constant 256 in (18) is definitely sub-optimal and it is out of scope of this work to look for the best possible constants.

**4. “Slow” rates accounting for high correlations.** In the preceding section, we have discussed fast rate bounds, that is, bounds that contain the tuning parameters to the power two. In this section, we turn to slow rate bounds, that is, bounds that contain the tuning parameters to the power one. We present slow rate bounds that entail—in contrast to what the nomenclature suggests—fast rates if the correlations are properly incorporated into the tuning parameters. These results corroborate and refine results in [HL13, vdGL13] and are of particular interest for the Least-Squares estimator with total variation penalty (TV-estimator). We can deduce, in particular, that the TV-estimator is almost minimax for the estimation of monotone or Hölder continuous signals and, therefore, improve on results in [MvdG87], for example.

The following slow rate bound is the main result of this section.

**THEOREM 4.** *Let  $T \subset [p]$  be a set of indices and let  $\delta > 0$ ,  $\gamma \geq 1$  be constants. Then, if the tuning parameter  $\lambda$  is not smaller than  $\gamma\sigma^*\rho_T\sqrt{2\log(p/\delta)}/n$ , the Lasso (2) fulfills*

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) + \frac{2(\gamma-1)\lambda}{\gamma}\|\widehat{\boldsymbol{\beta}}\|_1 \leq \inf_{\boldsymbol{\beta} \in \mathbb{R}^p} \left\{ \ell_n(\boldsymbol{\beta}, \boldsymbol{\beta}^*) + \frac{2(\gamma+1)\lambda}{\gamma}\|\boldsymbol{\beta}\|_1 \right\} + \frac{2\sigma^{*2}(|T| + 2\log(1/\delta))}{n}$$

with probability at least  $1 - 2\delta$ .

The factor  $\rho_T$  in the lower bound for the tuning parameter  $\lambda$  makes this theorem particularly interesting. Slow rate bounds can be derived using the proofs in [SZ12] (see Equation (12)), but they involve tuning parameters of order  $\sqrt{\log(p)/n}$ . Theorem 4 allows for considerably smaller tuning parameters if the variables are correlated; this then leads to rates in between the classical slow rates and fast rates of order (neglecting logarithmic factors)  $\sqrt{s/n}$  and  $s/n$ , respectively. Theorem 4 implies in particular fast rates for highly correlated covariates:

**COROLLARY 1.** *Assume that  $T_n \subset [p]$  is a set of indices (that may depend on the sample size  $n$ ) such that all covariates  $\{\mathbf{x}^j : j \in [p]\}$  are very close to the linear span of the set of*

vectors  $\{\mathbf{x}^j : j \in T_n\}$  in the sense that  $\rho_{T_n} \lesssim n^{-r}$  for a positive constant  $r > 0$ . Then, if the tuning parameter satisfies  $\lambda \geq c\sigma^* \sqrt{\log(p)/n^{2r+1}}$  for a sufficiently large constant  $c > 0$ , the Lasso (2) fulfills

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \lesssim \left( \sqrt{\frac{\log(p)}{n^{2r+1}}} \|\boldsymbol{\beta}^*\|_1 \right) \vee \frac{|T_n|}{n} \quad (19)$$

with high probability.

If, in particular, the irrelevant covariates  $\{\mathbf{x}^j : j \notin J^*\}$  are within Euclidean distance 1 of the linear space spanned by the relevant covariates  $\{\mathbf{x}^j : j \in J^*\}$ , it holds that  $r = 1/2$  and, therefore, the Lasso achieves the fast rate  $s/n$  up to logarithmic factors, provided that  $\lambda$  is chosen of order  $\sqrt{\log(p)}/n$  (with sufficiently large constants).

REMARK 1 (Effective number of parameters). The bound (19) can be further refined replacing the number of parameters  $p$  by an effective number of parameters as described in [HL13, Section 3.2.2]. This effective number of parameters can be considerably smaller than  $p$  if the correlations are high, therefore reducing the bound by a factor up to  $\sqrt{\log(p)}$ .

Corollary 1 exhibits fast rates for highly correlated but not necessarily perfectly collinear designs. We call a design perfectly collinear if all covariates belong to the linear space spanned by the relevant covariates, that is,  $\{\mathbf{x}^j : j \in [p]\} \subset \text{Span}\{\mathbf{x}^j : j \in J^*\}$ . For these designs, fast rates can be deduced from known results. Corollary 1, in contrast, exhibits fast rates even for designs that differ from perfectly collinear designs by an order of  $n^{-1/2}$  (as measured by the maximal distance  $\rho_{J^*}$ ). Thus, Corollary 1 answers the fourth question raised in the Introduction and, for example, corroborates the findings in [vdGL13] (see Section 5.2 for a more detailed comparison with [vdGL13]).

The dependence of the tuning parameters on the set  $T$  in Theorem 4 and Corollary 1 can lead to additional computational costs. For some applications, such as the total variation penalization discussed below, the set  $T$  is completely predetermined. For some other applications, however, the set  $T$  is not completely predetermined, promoting the minimization of  $\rho_T$  over a class of sets (for example, all sets with a fixed cardinality), which can be computationally expensive. Proposition 5 below provides another risk bound that helps to evade minimizations with respect to  $T$  under favorable circumstances.

PROPOSITION 5. Let  $T \subset [p]$  be a set of indices and  $\delta > 0$ ,  $\gamma > 1$  be constants. If the tuning parameter satisfies  $\lambda \geq \gamma\sigma^* \rho_T \sqrt{2 \log(p/\delta)/n}$ , the Lasso (2) fulfills

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq \frac{4\rho_T^2 \gamma^2}{(\gamma - 1)^2} \|\boldsymbol{\beta}^*\|_1^2 + \frac{4\sigma^{*2}(|T| + 2 \log(1/\delta))}{n} + \frac{2|T|\lambda^2}{\nu_T^2}$$

with probability at least  $1 - 2\delta$ .

As before, Proposition 5 shows that the correlations can be exploited adapting the tuning parameter to the design via the measure  $\rho_T$ , but, in strong contrast to the above results, Proposition 5 provides fast rate bounds for strongly correlated covariates even for standard, non-adapted tuning parameters of the order  $\sqrt{\log(p)}/n$ . For example, if  $\rho_{J^*} \asymp n^{-1/2}$  and

$|J^*| \asymp 1$ , fast rates can be deduced from Proposition 5 even with universal tuning parameters of order  $\sqrt{\log(p)/n}$ ; in strong contrast, considerably smaller tuning parameters of order  $\sqrt{\log(p)/n}$  are required to deduce fast rates from Theorem 4 for this example. Note, however, that Proposition 5 does not supersede Theorem 4 in general: For moderate correlations, the first term of the bound in Proposition 5 is large, and Theorem 4 is then considerably more beneficial.

Theorem 4 provides, in particular, minimax rates for signal denoising with total variation penalties. In the previous section, we have studied the TV-estimator for piecewise constant signals. In the remainder of this section, we study the TV-estimator for monotone signals (or slightly more general, signals with bounded variation) and for Hölder continuous signals. A recent review on this topic and a detailed analysis of the maximum likelihood estimator in this context can be found in [SCS13] and an earlier risk bound can be found in [Zha02].

**EXAMPLE 3** (Predicting monotone functions with the TV-estimator). In this example, we derive an almost minimax risk bound that is particularly interesting for signals with bounded variation. For this, we apply Theorem 4 exploiting that the TV-estimator can be considered as a special case of the Lasso. As mentioned earlier, the TV-estimator (16) corresponds to the Lasso (2) with the design matrix  $\mathbf{X} \in \mathbb{R}^{n \times n}$  with entries  $x_{ij} = \mathbf{1}(i \geq j)$ . To transfer the results for the Lasso to the TV-estimator, we assume that  $n \geq 3$ , fix a positive integer  $h \in [n-1]$ , denote by  $k \geq 2$  the largest integer such that  $(k-1)h < n$ , and finally set  $T = \{1, h+1, 2h+1, 3h+1, \dots, (k-1)h+1\}$ . The set  $T$  induces the partition <sup>6</sup> $\{\llbracket 1, h+1 \llbracket, \llbracket h+1, 2h+1 \llbracket, \dots, \llbracket (k-1)h+1, n+1 \llbracket\}$  of  $[n]$  with at most  $h$  points in the last interval and exactly  $h$  points in all other intervals. Moreover,  $\Pi_T$  is the orthogonal projection onto the subspace spanned by the vectors that are constant on each of the elements of this partition. This implies  $\rho_T = n^{-1/2} \max_j \|(\mathbf{I}_n - \Pi_T)\mathbf{x}^j\|_2 = \max_{j \in [h]} \sqrt{(j-1)(h-j+1)/(nh)} \leq \sqrt{h/(4n)}$ . Using  $h \leq 2n/k$ , we then obtain  $\rho_T \leq 1/\sqrt{2k}$  so that we can deduce from Theorem 4 the following risk bound for the TV-estimator.

**PROPOSITION 6.** *Assume that we observe the random vector  $\mathbf{y} = \mathbf{f}^* + \boldsymbol{\xi}$ , where  $\mathbf{f}^* \in \mathbb{R}^n$  is the fixed but unknown vector of interest obscured by Gaussian noise  $\boldsymbol{\xi} \sim \sigma^* \mathcal{N}(0, \mathbf{I}_n)$ . Let  $\delta > 0$  be a constant and  $k$  be the smallest integer larger than  $(\|\mathbf{f}^\dagger\|_{\text{TV}}^2 n \log(n/\delta)/\sigma^{*2})^{1/3}$ , where  $\mathbf{f}^\dagger$  is the orthogonal projection of  $\mathbf{f}^*$  on the convex polyhedral cone of vectors with nondecreasing entries. Then, for the tuning parameter  $\lambda = \sigma^* \sqrt{\log(n/\delta)/(kn)}$ , the TV-estimator (16) fulfills*

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \frac{1}{n} \|\mathbf{f}^\dagger - \mathbf{f}^*\|_2^2 + \frac{2\sigma^{*2}(1 + 2\log(1/\delta))}{n} + 6 \left( \frac{\sigma^{*4} \|\mathbf{f}^\dagger\|_{\text{TV}}^2 \log(n/\delta)}{n^2} \right)^{1/3}$$

with probability at least  $1 - 2\delta$ .

Proposition 6 has four crucial features. First, it provides nearly minimax rates for the TV-estimator: The dominating term is typically the last term, which is minimax up to the logarithmic factor [Zha02, Eq. (1.4)]. We conjecture that such logarithmic factors are always required in bounds that hold with high probability (note that the bounds in [Zha02], in contrast, are in expectation). Second, the unknown quantities in the tuning parameter of Proposition 6

<sup>6</sup>Here and in the sequel, we use the notation  $\llbracket a, b \llbracket := [a, b] \cap \mathbb{N}$ .

can be avoided or readily estimated: The noise variance  $\sigma^*$  can most likely be avoided using scaled versions of the Lasso [BCW13, SZ12]; the factor  $\|\mathbf{f}^\dagger\|_{\text{TV}}$  measuring the total variation of the function  $\mathbf{f}^\dagger$  can be roughly estimated<sup>7</sup> by  $\max_{i,j}(y_i - y_j)$ . Third, Proposition 6 allows for model mis-specifications with respect to monotone functions. Finally, it is nonasymptotic holding for any sample size  $n \geq 3$ .

**EXAMPLE 4** (TV-penalty for Hölder continuous functions). In this example, we derive an almost minimax risk bound that is particularly interesting for Hölder continuous signals. For this, we again apply Theorem 4 exploiting that the TV-estimator can be considered as a special case of the Lasso. It is known that the least squares estimator with an  $\ell_0$ -norm penalty of the increments can achieve the minimax rate  $n^{-2\alpha/(2\alpha+1)}$  up to logarithmic factors over the set of Hölder continuous functions  $\{f : [0, 1] \rightarrow \mathbb{R} : |f(x) - f(x')| \leq L|x - x'|^\alpha\}$  with parameters  $\alpha \in (0, 1]$  and  $L > 0$  [BKL<sup>+</sup>09]. In contrast, the best known rate for the TV-estimator over these sets is  $n^{-2\alpha/3}$  and therefore clearly suboptimal [MvdG87]. Using Theorem 4, we can improve on this bound and demonstrate that the TV-estimator can also achieve the minimax rate  $n^{-2\alpha/(2\alpha+1)}$  up to logarithmic factors over these sets if the tuning parameter is appropriately chosen.

**PROPOSITION 7.** *Assume that we observe the random vector  $\mathbf{y} = \mathbf{f}^* + \boldsymbol{\xi}$ , where  $\mathbf{f}^* \in \mathbb{R}^n$  is the fixed but unknown vector of interest obscured by Gaussian noise  $\boldsymbol{\xi} \sim \sigma^* \mathcal{N}(0, \mathbf{I}_n)$ . Let  $\delta, L > 0$  and  $\alpha \in (0, 1]$  be constants and let  $k$  be the smallest integer larger than  $(L^2 n / (\sigma^{*2} \log(n/\delta)))^{1/(2\alpha+1)}$ . Moreover, let  $\mathcal{H}_{\alpha, L}^n = \{\mathbf{f} \in \mathbb{R}^n : |f_i - f_j| \leq L n^{-\alpha} |i - j|^\alpha \forall i, j \in [n]\}$  be the Hölder class with parameters  $\alpha$  and  $L$ . Then, for the tuning parameter  $\lambda = \sigma^* \sqrt{\log(n/\delta)/(kn)}$ , the TV-estimator (16) fulfills*

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \inf_{\mathbf{f} \in \mathcal{H}_{\alpha, L}^n} \left\{ \frac{1}{n} \|\mathbf{f} - \mathbf{f}^*\|_2^2 \right\} + \frac{8\sigma^{*2} \log(n/\delta)}{n} + 16L^2 \left( \frac{\sigma^{*2} \log(n/\delta)}{nL^2} \right)^{2\alpha/(2\alpha+1)}$$

with probability at least  $1 - 2\delta$ .

Proposition 7 for the TV-estimator and the risk bounds in [BKL<sup>+</sup>09] for the  $\ell_0$ -penalized least-squares estimator provide exactly the same, nearly minimax rates  $(n/\log(n))^{-2\alpha/(2\alpha+1)}$ . The results differ, however, in other important aspects. Benefits of Proposition 7, on the one hand, are its finite sample bounds and the inclusion of model mis-specifications; the risk bounds in [BKL<sup>+</sup>09], in contrast, are purely asymptotic and do not take model mis-specifications into account. A deficiency of Proposition 7, on the other hand, is the dependence of the tuning parameter on the constants  $\alpha$  and  $L$ .

## 5. Discussion.

5.1. *Conditions of Belloni, Chernozhukov and Wang [BCW13].* Using an intelligent trick, the authors of [BCW13] managed to replace  $\kappa_{T, \bar{c}}$  by  $\kappa_{T, 1}$  by means of introducing a new constant  $\varrho_{T, \gamma}$  which is the  $1 - \delta$  quantile of the stochastic term  $\max_{\mathbf{u} \in \mathcal{C}_0(T, \gamma)} |(\boldsymbol{\xi}/\sigma^*)^\top \mathbf{X}\mathbf{u}| / \|\mathbf{X}\mathbf{u}\|_2$ .

<sup>7</sup>One may also expect that choosing  $\lambda$  by cross validation or by minimizing an unbiased estimator of the risk would lead to bounds similar to that of Proposition 5. However, we have no theoretical result corroborating this claim.

The fact of being able to replace  $\bar{c}$  by 1 leads to a qualitative enlargement of the set of matrices satisfying the condition  $\kappa_{T,\bar{c}} > 0$ . In fact, as proved in [BCW13], while  $\kappa_{T,1}$  is invariant by including identical columns in  $\mathbf{X}$ ,  $\kappa_{T,\bar{c}}$  with any  $\bar{c} > 1$  vanishes if we copy a column of  $\mathbf{X}_T$  in  $\mathbf{X}_{T^c}$  or vice-versa. Note that this property of invariance by copying covariates from  $T$  to  $T^c$  and, reciprocally, from  $T^c$  to  $T$  holds true for the weighted compatibility factors defined in Section 3 as well.

Combining our approach with the idea of [BCW13], it is possible to replace  $\bar{\kappa}_{T,\gamma,\omega}^{-1}$  in (14) by  $\frac{1}{|T|}\bar{\varrho}_{T,\gamma,\omega}^2 + \kappa_{T,1}^{-1}$ , where  $\bar{\varrho}_{T,\gamma,\omega}$  is the  $(1-\delta)$ -quantile of the random variable  $\max_{\mathbf{u} \in \mathcal{C}_0(T,\gamma,\omega)} |\boldsymbol{\xi}^\top (\mathbf{I}_n - \Pi_T)\mathbf{X}_{T^c}\mathbf{u}_{T^c}| / \|\mathbf{X}\mathbf{u}\|_2$ . At the first sight this upper bound is tighter than the one of Theorem 3, but it is less interpretable because of the presence of  $\bar{\varrho}_{T,\gamma,\omega}$ . However, as we prove below in Proposition 8, quantities like  $\bar{\varrho}_{T,\gamma,\omega}$  do not really lead to a substantially smaller risk bound than the one expressed in terms of compatibility factors. To complete the comparison of our results with those in [BCW13], let us simply remark that since our model is simpler than the one considered in [BCW13], the results we get are sharper. Indeed, we get a leading constant one in the oracle inequality, while the proof technique of [BCW13] would produce a leading constant strictly larger than one in the mis-specified case.

**PROPOSITION 8.** *Let  $\delta \in (0, 1/2)$  and  $\varrho_{T,\bar{c}}$  be the  $1 - \delta$  quantile of the random variable  $\bar{\eta} = \sup_{\mathbf{u} \in \mathbb{R}^p: \|\mathbf{u}_{T^c}\|_1 \leq \bar{c}\|\mathbf{u}_T\|_1} \frac{|\boldsymbol{\xi}^\top \mathbf{X}\mathbf{u}|}{\sigma^* \|\mathbf{X}\mathbf{u}\|_2}$ . For every  $10 \leq |J| \leq n$  we have*

$$\varrho_{T,\bar{c}} \geq \frac{\lambda_{\min,J}|T|^{1/2}}{16\kappa_{T,\bar{c}}^{1/2}} \bigwedge \frac{|J|^{1/2}}{4}, \quad (20)$$

where  $\lambda_{\min,J}$  is the smallest singular value of the matrix  $\frac{1}{\sqrt{n}}\mathbf{X}_J$ .

In many concrete examples of design matrices  $\mathbf{X}$ , there is a set  $J \subset T^c$  of cardinality of the same order as  $n$  such that  $\mathbf{X}_J$  is of full rank. For such matrices,  $\lambda_{\min,J}$  is a constant and the proposition tells us that  $|T|^{-1}\varrho_{T,\bar{c}}^2$  is of the order of  $\kappa_{T,\bar{c}}^{-1} \wedge (n/|T|)$ . As we already mentioned, the risk bound in [BCW13] is proportional to  $|T|^{-1}\varrho_{T,\bar{c}}^2 + \kappa_{T,1}^{-1} \asymp \kappa_{T,\bar{c}}^{-1}$ . Therefore, according to the result of the last proposition, there is no significant gain in the rate of convergence nor in the severity of the assumptions imposed on  $\mathbf{X}$  when using  $\varrho_{T,\bar{c}}$  instead of  $\kappa_{T,\bar{c}}$ .

**REMARK 2.** The constant  $\bar{\varrho}_{T,\bar{c}}$  slightly differs from the one used in [BCW13], where the additional constraint  $\|\boldsymbol{\beta}^* + \mathbf{u}\|_1 \leq \bar{c}\|\boldsymbol{\beta}^*\|_1$  is included in the definition of  $\mathcal{C}_0(T, \bar{c})$ . On the one hand, more generally, one can define the set  $\mathcal{C}_0(T, \bar{c}, \boldsymbol{\beta}) = \{\mathbf{u} \in \mathbb{R}^p : \|\boldsymbol{\beta} + \mathbf{u}\|_1 - \|\boldsymbol{\beta}_T\|_1 < (\bar{c} - 1)(\bar{c} + 1)^{-1}\|\mathbf{u}\|_1\}$ , contained in the one used in [BCW13], and let  $\varrho_{T,\bar{c},\boldsymbol{\beta}}$  be the  $1 - \delta$  quantile of

$$\bar{\eta} = \sup_{\mathbf{u} \in \mathcal{C}_0(T,\bar{c},\boldsymbol{\beta})} \frac{|\boldsymbol{\xi}^\top \mathbf{X}\mathbf{u}|}{\sigma^* \|\mathbf{X}\mathbf{u}\|_2}.$$

Then the risk bound of [BCW13] holds true with  $\bar{\varrho}_{T,\bar{c}}$  replaced by  $\bar{\varrho}_{T,\bar{c},\boldsymbol{\beta}^*}$ . On the other hand, if one looks for a characteristic independent of  $\boldsymbol{\beta}^*$ , then it is necessary to take a supremum over all  $\boldsymbol{\beta}^*$  that are zero outside  $T$ . This amounts to taking the supremum over  $\cup_{\boldsymbol{\beta}^*: \|\boldsymbol{\beta}_{T^c}^*\|_1=0} \mathcal{C}_0(T, \bar{c}, \boldsymbol{\beta}^*) = \mathcal{C}_0(T, \bar{c})$ , that is to considering the quantity  $\bar{\varrho}_{T,\bar{c}}$  of Proposition 8. Finally, one can repeat the arguments of the proof of Proposition 8 to check that for every

$J \subset T^c$  satisfying  $10 \leq |J| \leq n$ , we have  $\varrho_{T,\bar{c},\beta^*} \geq \frac{\lambda_{\min,J|T|^{1/2}}}{16\tilde{\kappa}_{T,\bar{c},\beta^*}^{1/2}} \wedge \frac{|J|^{1/2}}{4}$ , where  $\tilde{\kappa}_{T,\bar{c},\beta^*}$  is the inf over all  $\mathbf{u} \in \mathcal{C}(T, \bar{c}, \beta^*)$  of the ratio  $\frac{1}{n} \|\mathbf{X}\mathbf{u}\|_2^2 / (\|\beta_T^*\|_1 + (\bar{c} - 1)(\bar{c} + 1)^{-1} \|\mathbf{u}\|_1 - \|\beta^* + \mathbf{u}\|_1)$ .

5.2. *Relation to the previous work on the Lasso with correlated covariates.* Our bound (19) is close in spirit to the bound in [HL13, Theorem 3.1] which is a consequence of results in [vdGL13]. Both bounds demonstrate that the Lasso can achieve fast rates for prediction even for highly correlated design matrices if the tuning parameter is chosen appropriately. In order to make the comparison with our results easier, let us state the main result of [vdGL13] using the notation of the present work. In fact, Theorem 4.1 in [vdGL13] establishes that for any  $\alpha \in [0, 1)$ ,  $\lambda > 0$ , and  $\lambda_0 > 0$ , on the event  $\mathcal{B}_{\alpha,\lambda_0} = \{\sup_{\beta:\|\beta\|_1=1} 4|\xi^\top \mathbf{X}\beta| / \|\mathbf{X}\beta\|_2^{1-\alpha} \leq n^{(1+\alpha)/2} \lambda_0\}$ , it holds that

$$\ell_n(\hat{\beta}_\lambda^{\text{Lasso}}, \beta^*) \leq 7 \inf_{\bar{\beta} \in \mathbb{R}^p, T \subset [p]} \left\{ \ell_n(\bar{\beta}, \beta^*) + \frac{8\lambda}{3} \|\bar{\beta}_{T^c}\|_1 + \frac{7}{6} \left( \frac{\lambda_0}{\lambda^\alpha} \right)^{\frac{2}{1-\alpha}} + \frac{224\lambda^2|T|}{\kappa_{T,6}} \right\}, \quad (21)$$

where  $\kappa_{T,6}$  is the compatibility constant (9). Furthermore, the authors of [vdGL13] provide sufficient conditions in terms of the entropy of the set  $\mathcal{F} = \{\beta \in \mathbb{R}^p : \|\mathbf{X}\beta\|_2^2 \leq n; \|\beta\|_1 \leq 1\}$  ensuring that the probability of the event  $\mathcal{B}_{\alpha,\lambda_0}$  is close to one for some  $\alpha \in (0, 1)$  and  $\lambda_0 \asymp (\log(n)/n)^{1/2}$ . The main advantages of the results stated in the present work as compared to (21) are that (a) the risk bounds are with leading constant one, (b) the quantity  $\rho_T$  governing the choice of  $\lambda$  and the rate of convergence of the prediction risk is, in general, easier to compute than the entropy, and (c) the compatibility factor is replaced by the weighted compatibility factor that is strictly positive in several important cases where the compatibility factor vanishes (*e.g.*, for total variation penalization). Moreover (d), the benefits of our results do not necessarily rely on a high overall correlation: If, for example, the covariates can be clustered into (a reasonably small number of) sets of highly correlated covariates, one can find a small set  $T$  such that the measure  $\rho_T$  is small. In contrast, the entropy measure in [vdGL13] is not necessarily small for this example, because it measures the symmetric convex hull of all variables, which is still a large set if the clusters are not highly correlated with each other (we expect, however, that the approach in [vdGL13] can be refined in this respect). On the other hand, risk bound (21) may potentially offer more flexibility due to the presence of the parameter  $\alpha$ . Besides, it is very likely that the proof technique used in this work allow for removing the factor 7 in front of the inf at the right hand-side of (21).

There is another direction of research, explored in the recent paper [BRvdGZ13] and in discussions [BW13, SS13], that replaces the original design matrix by a new one with weaker correlations between the covariates. This is achieved by clustering the columns of  $\mathbf{X}$  and replacing the groups of strongly correlated covariates by one representer (CRL), or by gathering strongly correlated covariates in disjoint groups for prediction with the group-Lasso procedure (CGL). While the theoretical results developed in [BRvdGZ13] demonstrate advantages with respect to those available for the Lasso, the experimental results reported in Tables 2-5 show that the Lasso remains perfectly competitive with the new procedures CRL and CGL. The results of this work explain, at least partially, these empirical results, in that we proved that the prediction loss of the standard Lasso is small even if the design matrix contains strongly correlated columns.

Furthermore, we believe that the clustering strategy developed in [BRvdGZ13] may be benefi-

cial in conjunction with the Lasso, without any modification. In fact, the result of the clustering can be used for tuning the penalty level  $\lambda$ . More precisely, let  $G_1, \dots, G_M$  be the clusters we get, forming a partition of  $[p]$ . Theoretical results developed in previous sections suggest to choose one representer per cluster by setting  $j_m = \arg \min_{i \in G_m} \max_{j \in G_m} \|(\mathbf{I}_n - \Pi_i)\mathbf{x}^j\|_2$  and to define  $T = \{j_1, \dots, j_M\}$  together with  $\lambda = 2\sigma^* \rho_T \sqrt{2 \log(p/\delta)/n}$ . If the clusters are tight—the vectors within each cluster are very close to one another—then  $\rho_T$  will be small, since  $\rho_T = n^{-1/2} \max_j \|(\mathbf{I}_n - \Pi_T)\mathbf{x}^j\|_2 \leq n^{-1/2} \max_m \max_{j \in G_m} \|(\mathbf{I}_n - \Pi_{j_m})\mathbf{x}^j\|_2 \leq n^{-1/2} \max_m \max_{j \in G_m} \|(\mathbf{I}_n - \Pi_{j_m})\mathbf{x}^j\|_2$ . Note that once the clusters are found, the aforementioned computation of  $j_m$ 's and of  $\lambda$  is not time consuming. Another compelling alternative is to replace the clustering step by sparse PCA. Indeed, our theoretical findings advocate for choosing as  $T$  a subset of  $[p]$  which is simultaneously of small cardinality and such that all  $\mathbf{x}^j$ 's are close to the linear space spanned by  $\{\mathbf{x}^j : j \in T\}$ . It is precisely the task of the sparse PCA to find such a subset  $T$  (cf. [BJNP13] and the references therein).

5.3. *Computation of the compatibility factors by sequential convex programming.* It may be of practical interest to evaluate the risk bounds presented in our main results in order to understand how accurate the Lasso prediction is. To this end, one may often need to compute, at least approximately, the compatibility factor  $\kappa_{T, \bar{c}}$  or its weighted counterpart  $\bar{\kappa}_{T, \gamma, \omega}$ . In general, this task is difficult to accomplish since the computation of the aforementioned factors amounts to minimizing a nonconvex function over a nonconvex subset of  $\mathbb{R}^p$  with large dimensionality  $p$ . (For a different procedure, efficiently verifiable conditions entailing theoretical guarantees for sparse recovery are proposed in [JN11b].) There is, however, a particular case where this task may be solved with a reasonable computational complexity. This corresponds to subsets  $T$  of small or moderately large cardinalities. In fact, in the remaining of this subsection we will show that for every  $T \subset [p]$  and every  $\bar{c}, \epsilon > 0$ , one can find an interval  $[\kappa_*, \kappa^*]$  of length at most  $\epsilon$  and containing  $\bar{\kappa}_{T, \bar{c}}$  by solving at most  $2^{|T|} \log_2(|T|/\epsilon)$  convex programs. Furthermore, each of these convex programs is a second-order cone program (SOCP) and the global computation can be split into  $2^{|T|}$  parallel programs, each program involving  $\log_2(|T|/\epsilon)$  SOCPs.

The procedure we propose relies on the well-known bisection algorithm. Since we know that  $\bar{\kappa}_{T, \bar{c}}$  always belong to the interval  $[0, |T|]$ , we start by setting  $\kappa_* = 0$  and  $\kappa^* = |T|$ . Then, at each step of iteration, we set  $\kappa = (\kappa_* + \kappa^*)/2$  and  $\mu = (n\kappa/|T|)^{1/2} \bar{c}^{-1}$  and solve the problem

$$\begin{aligned} & \text{minimize} && \|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu \|\boldsymbol{\delta}_{T^c}\|_1 \\ & \text{subject to} && \|\boldsymbol{\delta}_T\|_1 = 1. \end{aligned} \tag{22}$$

If the solution  $\boldsymbol{\delta}^\mu$  of this problem satisfies  $\|\mathbf{X}\boldsymbol{\delta}^\mu\|_2 + \mu \|\boldsymbol{\delta}_{T^c}^\mu\|_1 \leq \mu \bar{c}$ , then we leave  $\kappa_*$  unchanged and decrease  $\kappa^*$  by setting

$$\kappa^* = \frac{|T| \cdot \|\mathbf{X}\boldsymbol{\delta}^\mu\|_2^2}{n(1 - \bar{c}^{-1} \|\boldsymbol{\delta}_{T^c}^\mu\|_1)^2}.$$

(Note that the right hand-side is always less than or equal to  $\kappa$ ). If, in contrast, the solution  $\boldsymbol{\delta}^\mu$  of problem (22) satisfies  $\|\mathbf{X}\boldsymbol{\delta}^\mu\|_2 + \mu \|\boldsymbol{\delta}_{T^c}^\mu\|_1 > \mu \bar{c}$ , then we leave  $\kappa^*$  unchanged and increase  $\kappa_*$  by setting  $\kappa_* = \kappa$ . We iterate this process until we get  $\kappa^* - \kappa_* \leq \epsilon$ . Since at the first step the gap  $\kappa^* - \kappa_*$  equals  $|T|$  and at each step this gap is divided by at least a factor 2, the total number of iterations to get precision  $\epsilon$  is not larger than  $\log_2(|T|/\epsilon)$ .

Let us now analyze the complexity of the optimization problem (22). The objective function of this problem is convex but the set of feasible solutions is not. Interestingly, for every  $\mathbf{s} \in \{\pm 1\}^{|T|}$ , if we restrict the optimization to the orthant  $\{\boldsymbol{\delta} : s_j \delta_j \geq 0, \forall j \in T\}$ , the constraints become convex as well. Thus, our proposal consists in replacing (22) by  $2^{|T|}$  optimization problems

$$\begin{aligned} & \text{minimize} && \|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu\|\boldsymbol{\delta}_{T^c}\|_1 \\ & \text{subject to} && \mathbf{s}^\top \boldsymbol{\delta}_T = 1, \text{ and } s_j \delta_j \geq 0, \forall j \in T. \end{aligned} \tag{23}$$

For every  $\mathbf{s} \in \{\pm 1\}^{|T|}$ , (23) can be rewritten as a SOCP using standard arguments. Denoting by  $\boldsymbol{\delta}^{\mu, \mathbf{s}}$  any solution of (23), we determine a solution of (22) by minimizing the common objective function of the above optimization problems over the finite set  $\{\boldsymbol{\delta}^{\mu, \mathbf{s}} : \mathbf{s} \in \{\pm 1\}^{|T|}\}$ , that is

$$\boldsymbol{\delta}^\mu = \arg \min_{\boldsymbol{\delta} \in \{\boldsymbol{\delta}^{\mu, \mathbf{s}} : \mathbf{s} \in \{\pm 1\}^{|T|}\}} (\|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu\|\boldsymbol{\delta}_{T^c}\|_1).$$

These consideration lead to the procedure summarized in Algorithm 1.

---

**Algorithm 1** Pseudo-code for computing the compatibility factor

---

**Require:**  $n \times p$  matrix  $\mathbf{X}$ , set  $T \subset [p]$ , constant  $\bar{c} > 0$ , precision  $\epsilon > 0$

**Ensure:** interval  $[\kappa_*, \kappa^*]$  containing the compatibility factor  $\kappa_{T, \bar{c}}$

```

1:  $\kappa_* \leftarrow 0$  and  $\kappa^* \leftarrow |T|$ 
2: while  $\kappa^* - \kappa_* < \epsilon$  do
3:    $\kappa \leftarrow (\kappa^* + \kappa_*)/2$ 
4:    $\mu \leftarrow (n\kappa/|T|)^{1/2} \bar{c}^{-1}$ 
5:   for  $\mathbf{s} \in \{\pm 1\}^{|T|}$  do
6:      $v^{\mathbf{s}} \leftarrow \min\{\|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu\|\boldsymbol{\delta}_{T^c}\|_1\}$  subject to  $\mathbf{s}^\top \boldsymbol{\delta}_T = 1$  and  $s_j \delta_j \geq 0, \forall j \in T$ 
7:      $\boldsymbol{\delta}^{\mathbf{s}} \leftarrow \arg \min\{\|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu\|\boldsymbol{\delta}_{T^c}\|_1\}$  subject to  $\mathbf{s}^\top \boldsymbol{\delta}_T = 1$  and  $s_j \delta_j \geq 0, \forall j \in T$ 
8:   end for
9:   if  $\min_{\mathbf{s}} v^{\mathbf{s}} > \mu \bar{c}$  then
10:     $\kappa_* \leftarrow \kappa$ 
11:   else
12:     $\boldsymbol{\delta}^\mu \leftarrow \arg \min_{\boldsymbol{\delta} \in \{\boldsymbol{\delta}^{\mathbf{s}} : \mathbf{s} \in \{\pm 1\}^{|T|}\}} (\|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu\|\boldsymbol{\delta}_{T^c}\|_1)$ 
13:     $\kappa^* \leftarrow |T| \cdot \|\mathbf{X}\boldsymbol{\delta}^\mu\|_2^2 / (n\{1 - \bar{c}^{-1}\|\boldsymbol{\delta}_{T^c}^\mu\|_1\}^2)$ 
14:   end if
15: end while

```

---

**6. Conclusions.** Our results lead to a better understanding of the prediction performance of the Lasso, as they demonstrate that correlations are not necessarily obstructive but even helpful in some cases. This permits more accurate comparisons of the Lasso with its many competitors. Our results are based on the introduction of  $\rho_T$ , a new, simple measure of the correlations of the covariates. If this measure is incorporated in the choice of the tuning parameter, the Lasso prediction risk decays at a fast rate for a broad variety of settings including settings with strongly correlated covariates. To derive this, we did not invoke the usual assumptions such as restricted isometry, restricted eigenvalues, etc., but rather relied on the slow rate bounds that hold for arbitrary designs. Consequences of our results are then substantially improved risk bounds for the least squares estimator with total variation penalty.

We also introduce compatibility factors that are not only abstract concepts but can be both evaluated numerically and bounded theoretically and, therefore, implicate many applications. We introduce in particular a new, weighted compatibility factor  $\bar{\kappa}$ , which—in contrast to its original version—may be bounded away from zero even for strongly correlated covariates. This allows us to apply the corresponding results to the least squares estimator with total variation penalty, for example, where the correlations between the covariates are up to  $1 - (1/n)$ .

Our results finally indicate that the prediction performance of the Lasso can not be characterized by only the maximal correlation between covariates: On the one hand, as described above, the Lasso can provide accurate prediction even if the covariates are highly correlated. On the other hand, as indicated by Example 1, the Lasso can perform poorly in prediction even for moderately correlated covariates.

Future research directions include developing our approach in the case of the group Lasso [YL06, LPvdGT11] in order to understand how to optimally exploit the (correlation) structure of the Gram matrix for defining the groups. This problem is also interesting for applications to the total-variation penalization as discussed in [VB10]. Another relevant question is how the refinements proposed in the present work may be adapted to scale invariant versions of the Lasso, such as the square-root Lasso [BCW13], scaled Lasso [SBvdG10] or scaled Dantzig selector [DC12]. We also believe that the geometry of the design, measured by the quantities  $\rho_T$ , may lead to better recommendations for the tuning parameter in the transductive setting [AH12]. Finally, we would like to explore the consequences of our results when applied to the nonparametric estimation of a regression function  $f$  by penalized least squares with a penalty proportional to the discrete counterpart of the  $L_1$ -norm of the  $k$ th derivative of  $f$ . This problem has been studied in [MvdG87], but we believe that the results of the present work may lead to improved risk bounds.

**7. Proofs.** We gather in this sections the proofs of all the theorems and propositions stated in previous sections. To ease notation, throughout these proofs we write  $\hat{\beta}$  instead of  $\hat{\beta}_\lambda^{\text{Lasso}}$ . In the sequel, we denote by  $\text{sgn}(x)$  the sub-differential of the function  $x \mapsto |x|$ , that is

$$\text{sgn}(x) = \begin{cases} \{1\}, & x > 0, \\ [-1, 1], & x = 0, \\ \{-1\}, & x < 0. \end{cases}$$

PROOF OF THEOREM 1. We first use the Karush-Kuhn-Tucker conditions to infer that

$$\frac{1}{n} \mathbf{X}^\top (\mathbf{y} - \mathbf{X} \hat{\beta}) \in \lambda \text{sgn}(\hat{\beta}).$$

This implies that for every vector  $\bar{\beta} \in \mathbb{R}^p$

$$\begin{aligned} \frac{1}{n} \hat{\beta}_T^\top \mathbf{X}_T^\top (\mathbf{y} - \mathbf{X} \hat{\beta}) &= \lambda \|\hat{\beta}_T\|_1, \\ \frac{1}{n} \bar{\beta}_T^\top \mathbf{X}_T^\top (\mathbf{y} - \mathbf{X} \hat{\beta}) &\leq \lambda \|\bar{\beta}_T\|_1. \end{aligned}$$

Subtracting the first relation from the second one, we get

$$\frac{1}{n} (\bar{\beta} - \hat{\beta})_T^\top \mathbf{X}_T^\top (\mathbf{y} - \mathbf{X} \hat{\beta}) \leq \lambda (\|\bar{\beta}_T\|_1 - \|\hat{\beta}_T\|_1). \quad (24)$$

We define now the vector  $\bar{\boldsymbol{\beta}}$  by the relations

$$\bar{\boldsymbol{\beta}}_T = \boldsymbol{\beta}_T^{*,T} + (\mathbf{X}_T^\top \mathbf{X}_T)^\dagger \mathbf{X}_T^\top \mathbf{X}_{T^c} (\boldsymbol{\beta}^{*,T} - \hat{\boldsymbol{\beta}})_{T^c} \quad \text{and} \quad \bar{\boldsymbol{\beta}}_{T^c} = 0.$$

This choice of  $\bar{\boldsymbol{\beta}}$  may appear somewhat strange and complicated, but is made in order that the relation  $\mathbf{X}_T(\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_T = \Pi_T \mathbf{X}(\boldsymbol{\beta}^{*,T} - \hat{\boldsymbol{\beta}})$  be satisfied. On the other hand, one easily checks that  $\mathbf{y} = \mathbf{X}_T \boldsymbol{\beta}_T^* + \mathbf{X}_{T^c} \boldsymbol{\beta}_{T^c}^* + \boldsymbol{\xi} = \mathbf{X} \boldsymbol{\beta}^{*,T} + (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}$ . Replacing these expressions of  $\bar{\boldsymbol{\beta}}$  and  $\mathbf{y}$  in (24), we find that for every  $T \subset [p]$ ,

$$\frac{1}{n} (\boldsymbol{\beta}^{*,T} - \hat{\boldsymbol{\beta}})^\top \mathbf{X}^\top \Pi_T \mathbf{X} (\boldsymbol{\beta}^{*,T} - \hat{\boldsymbol{\beta}}) \leq \lambda (\|\bar{\boldsymbol{\beta}}_T\|_1 - \|\hat{\boldsymbol{\beta}}_T\|_1) \leq \lambda \|\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_1.$$

Equivalently, this relation may be written as

$$\frac{1}{n} \|\Pi_T \mathbf{X}(\boldsymbol{\beta}^{*,T} - \hat{\boldsymbol{\beta}})\|_2^2 = \frac{1}{n} \|\mathbf{X}_T(\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_T\|_2^2 \leq \lambda \|\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_1.$$

In view of the fact that  $\sqrt{|T|/n} \|\mathbf{X}_T \mathbf{u}\|_2 \geq \nu_T \|\mathbf{u}\|_1$  for every  $\mathbf{u} \in \mathbb{R}^{|T|}$ , we get

$$\|\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_1 \leq \frac{\sqrt{|T|} \|\mathbf{X}_T(\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_T\|_2}{\sqrt{n} \nu_T}.$$

Combining the last two displays, we obtain

$$\frac{1}{n} \|\Pi_T \mathbf{X}(\boldsymbol{\beta}^{*,T} - \hat{\boldsymbol{\beta}})\|_2^2 = \frac{1}{n} \|\mathbf{X}_T(\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_T\|_2^2 \leq \lambda \frac{\sqrt{|T|} \|\mathbf{X}_T(\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_T\|_2}{\sqrt{n} \nu_T}.$$

Dividing the both sides of the last inequality by  $\|\mathbf{X}_T(\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_T\|_2^2$  the desired result follows.  $\square$

PROOF OF THEOREM 3. Recall that according to (24), for every  $J \subset [p]$  and  $\bar{\boldsymbol{\beta}} \in \mathbb{R}^p$ ,

$$\frac{1}{n} (\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_J^\top \mathbf{X}_J^\top (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}}) \leq \lambda (\|\bar{\boldsymbol{\beta}}_J\|_1 - \|\hat{\boldsymbol{\beta}}_J\|_1).$$

Replacing the expression of  $\mathbf{y}$  in this inequality, we find that

$$\frac{1}{n} (\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}})_J^\top \mathbf{X}_J^\top (\mathbf{X} \boldsymbol{\beta}^* + \boldsymbol{\xi} - \mathbf{X} \hat{\boldsymbol{\beta}}) \leq \lambda (\|\bar{\boldsymbol{\beta}}_J\|_1 - \|\hat{\boldsymbol{\beta}}_J\|_1).$$

Let us introduce the two difference vectors  $\boldsymbol{\delta} = \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$  and  $\bar{\boldsymbol{\delta}} = \hat{\boldsymbol{\beta}} - \bar{\boldsymbol{\beta}}$ . The last display combined with the decomposition  $\boldsymbol{\xi} = \Pi_T \boldsymbol{\xi} + (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}$ , for every  $T \subset [p]$ , yields

$$\frac{1}{n} \bar{\boldsymbol{\delta}}_J^\top \mathbf{X}_J^\top \mathbf{X} \boldsymbol{\delta} \leq \frac{1}{n} \bar{\boldsymbol{\delta}}_J^\top \mathbf{X}_J^\top \Pi_T \boldsymbol{\xi} + \frac{1}{n} \bar{\boldsymbol{\delta}}_J^\top \mathbf{X}_J^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi} + \lambda (\|\bar{\boldsymbol{\beta}}_J\|_1 - \|\hat{\boldsymbol{\beta}}_J\|_1).$$

Using the identity  $\mathbf{u}^\top \mathbf{u}' = \frac{1}{2} (\|\mathbf{u}\|_2^2 + \|\mathbf{u}'\|_2^2 - \|\mathbf{u} - \mathbf{u}'\|_2^2)$  we get for every  $J \subset [p]$ ,

$$\begin{aligned} \frac{\|\mathbf{X}_J \bar{\boldsymbol{\delta}}_J\|_2^2 + \|\mathbf{X} \boldsymbol{\delta}\|_2^2}{n} &\leq \frac{1}{n} \|\mathbf{X}_J \bar{\boldsymbol{\delta}}_J - \mathbf{X} \boldsymbol{\delta}\|_2^2 + \frac{2}{n} \bar{\boldsymbol{\delta}}_J^\top \mathbf{X}_J^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi} + \frac{2}{n} \bar{\boldsymbol{\delta}}_J^\top \mathbf{X}_J^\top \Pi_T \boldsymbol{\xi} \\ &\quad + 2\lambda (\|\bar{\boldsymbol{\beta}}_J\|_1 - \|\hat{\boldsymbol{\beta}}_J\|_1). \end{aligned} \tag{25}$$

To prove (14) we choose  $J = [p]$ , for which (25) becomes

$$\frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2 + \frac{1}{n} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq \frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}} - \mathbf{X}\boldsymbol{\delta}\|_2^2 + \frac{2}{n} \bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \boldsymbol{\xi} + 2\lambda(\|\bar{\boldsymbol{\beta}}\|_1 - \|\widehat{\boldsymbol{\beta}}\|_1) \quad (26)$$

$$\begin{aligned} &\leq \frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}} - \mathbf{X}\boldsymbol{\delta}\|_2^2 + \frac{2}{n} \bar{\boldsymbol{\delta}}_{T^c}^\top \mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi} + \frac{2}{n} \bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \Pi_T \boldsymbol{\xi} \\ &\quad + 2\lambda(\|\bar{\boldsymbol{\delta}}_T\|_1 - \|\bar{\boldsymbol{\delta}}_{T^c}\|_1) + 4\lambda\|\bar{\boldsymbol{\beta}}_{T^c}\|_1, \quad \forall T \subset [p]. \end{aligned} \quad (27)$$

Notice that  $\bar{\boldsymbol{\delta}}_{T^c}^\top \mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi} \leq \sum_{j \in T^c} |\mathbf{x}^j{}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}| \cdot |\bar{\delta}_j|$  and  $|\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \Pi_T \boldsymbol{\xi}| \leq \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2 \|\Pi_T \boldsymbol{\xi}\|_2$ . Replacing  $\lambda$  by its value  $\gamma\sigma^* \left(\frac{2}{n} \log(p/\delta)\right)^{1/2}$  and restricting our attention to the event  $\mathcal{B}_T = \left\{ \max_{j \in T^c} |\boldsymbol{\xi}^\top (\mathbf{I}_n - \Pi_T) \mathbf{x}^j| / \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2 \leq \sigma^* (2 \log(p/\delta))^{1/2} \right\} \cap \left\{ \|\Pi_T \boldsymbol{\xi}\|_2 \leq \sigma^* (\sqrt{|T|} + \sqrt{2 \log(1/\delta)}) \right\}$  we obtain

$$\begin{aligned} \frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2 + \frac{1}{n} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 &\leq \frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}} - \mathbf{X}\boldsymbol{\delta}\|_2^2 + 4\lambda\|\bar{\boldsymbol{\beta}}_{T^c}\|_1 + \frac{2}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2 \|\Pi_T \boldsymbol{\xi}\|_2 \\ &\quad + 2\lambda \left( \|\bar{\boldsymbol{\delta}}_T\|_1 - \|\bar{\boldsymbol{\delta}}_{T^c}\|_1 + \gamma^{-1} \|(\bar{\boldsymbol{\delta}} \odot \boldsymbol{\omega})_{T^c}\|_1 \right). \end{aligned} \quad (28)$$

The definition of  $\bar{\kappa}_{T,\gamma,\boldsymbol{\omega}}$  implies that  $\|\bar{\boldsymbol{\delta}}_T\|_1 - \|\bar{\boldsymbol{\delta}}_{T^c}\|_1 + \gamma^{-1} \|(\bar{\boldsymbol{\delta}} \odot \boldsymbol{\omega})_{T^c}\|_1 \leq \frac{|T|^{1/2} \cdot \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2}{(n\bar{\kappa}_{T,\gamma,\boldsymbol{\omega}})^{1/2}}$  (note that this inequality is trivial when the left hand-side is negative), therefore

$$\begin{aligned} &\frac{2}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2 \|\Pi_T \boldsymbol{\xi}\|_2 + 2\lambda \left( \|\bar{\boldsymbol{\delta}}_T\|_1 - \|\bar{\boldsymbol{\delta}}_{T^c}\|_1 + \gamma^{-1} \|(\bar{\boldsymbol{\delta}} \odot \boldsymbol{\omega})_{T^c}\|_1 \right) \\ &\leq 2 \frac{\|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2}{\sqrt{n}} \left( \frac{1}{\sqrt{n}} \|\Pi_T \boldsymbol{\xi}\|_2 + \lambda(|T|/\bar{\kappa}_{T,\gamma,\boldsymbol{\omega}})^{1/2} \right) \\ &\leq \frac{\|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2}{n} + \left( \frac{1}{\sqrt{n}} \|\Pi_T \boldsymbol{\xi}\|_2 + \lambda(|T|/\bar{\kappa}_{T,\gamma,\boldsymbol{\omega}})^{1/2} \right)^2 \\ &\leq \frac{\|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2}{n} + \frac{4\sigma^{*2}(|T| + 2 \log(1/\delta))}{n} + \frac{4\gamma^2 \sigma^{*2} |T| \log(p/\delta)}{n\bar{\kappa}_{T,\gamma,\boldsymbol{\omega}}}. \end{aligned} \quad (29)$$

After replacing (29) in (28), we remark that the terms  $\frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2$  cancel out and we get inequality (14). Classical results on the tails of Gaussian and  $\chi^2$  distributions imply that  $\mathcal{B}_T$  is at least of probability  $1 - 2\delta$ , for every  $T$ . The assertion of (14) follows by choosing  $T$  to be a subset of  $[p]$  minimizing the right hand-side of (14).

The proof of the second claim of the theorem is identical to that of (14) and, therefore, is left to the reader.  $\square$

**PROOF OF THEOREM 4.** According to (26), for every  $\bar{\boldsymbol{\beta}} \in \mathbb{R}^p$  and for  $\boldsymbol{\delta} = \bar{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$  and  $\widehat{\boldsymbol{\beta}} = \bar{\boldsymbol{\beta}} - \boldsymbol{\beta}$ , the inequality

$$\frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2 + \frac{1}{n} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq \frac{1}{n} \|\mathbf{X}\bar{\boldsymbol{\delta}} - \mathbf{X}\boldsymbol{\delta}\|_2^2 + \frac{2}{n} \bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \boldsymbol{\xi} + 2\lambda(\|\bar{\boldsymbol{\beta}}\|_1 - \|\widehat{\boldsymbol{\beta}}\|_1) \quad (30)$$

holds for every  $\lambda > 0$ . We split the stochastic term  $\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \boldsymbol{\xi}$  into two terms  $\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}$  and  $\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \Pi_T \boldsymbol{\xi}$ . The first one can be bounded using the duality inequality  $\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi} = \bar{\boldsymbol{\delta}}_{T^c}^\top \mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi} \leq \|\bar{\boldsymbol{\delta}}_{T^c}\|_1 \|\mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}\|_\infty \leq (\|\bar{\boldsymbol{\beta}}_{T^c}\|_1 + \|\widehat{\boldsymbol{\beta}}_{T^c}\|_1) \|\mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}\|_\infty$ , while

the second one, in view of the Cauchy-Schwarz inequality, satisfies  $\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \Pi_T \boldsymbol{\xi} \leq \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2 \|\Pi_T \boldsymbol{\xi}\|_2 \leq (\|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2 + \|\Pi_T \boldsymbol{\xi}\|_2^2)/2$ . As in the proof of fast rates, we restrict our attention to the event  $\mathcal{B}_T = \{\max_{j \in T^c} |\boldsymbol{\xi}^\top (\mathbf{I}_n - \Pi_T) \mathbf{x}^j| / \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2 \leq \sigma^* (2 \log(p/\delta))^{1/2}\} \cap \{\|\Pi_T \boldsymbol{\xi}\|_2 \leq \sigma^* (\sqrt{|T|} + \sqrt{2 \log(1/\delta)})\}$ . On this event, we get

$$2\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \boldsymbol{\xi} \leq \|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2 + 2\sigma^{*2}(|T| + 2 \log(1/\delta)) + 2\lambda\gamma^{-1}n(\|\bar{\boldsymbol{\beta}}_{T^c}\|_1 + \|\hat{\boldsymbol{\beta}}_{T^c}\|_1). \quad (31)$$

Combining this inequality with (30), we get the desired result.  $\square$

**PROOF OF PROPOSITION 2.** The Karush-Kuhn-Tucker conditions for this example are the following:

$$\begin{aligned} \frac{1}{\sqrt{2}}(\mathbf{e}_1 + \mathbf{e}_{j+1})^\top (\sqrt{2}\mathbf{e}_1 + \frac{1}{\sqrt{n}}\boldsymbol{\xi} - \frac{1}{\sqrt{2}}\sum_{k=1}^{2m}\hat{\beta}_k\mathbf{e}_1 - \frac{1}{\sqrt{2}}(\hat{\beta}_j - \hat{\beta}_{m+j})\mathbf{e}_{j+1}) &\in \lambda \operatorname{sgn}(\hat{\beta}_j), \\ \frac{1}{\sqrt{2}}(\mathbf{e}_1 - \mathbf{e}_{j+1})^\top (\sqrt{2}\mathbf{e}_1 + \frac{1}{\sqrt{n}}\boldsymbol{\xi} - \frac{1}{\sqrt{2}}\sum_{k=1}^{2m}\hat{\beta}_k\mathbf{e}_1 - \frac{1}{\sqrt{2}}(\hat{\beta}_j - \hat{\beta}_{m+j})\mathbf{e}_{j+1}) &\in \lambda \operatorname{sgn}(\hat{\beta}_{m+j}). \end{aligned}$$

After simplification, we get

$$1 + \frac{1}{\sqrt{2n}}(\xi_1 + \xi_{j+1}) - \frac{1}{2}\sum_{k=1}^{2m}\hat{\beta}_k - \frac{1}{2}(\hat{\beta}_j - \hat{\beta}_{m+j}) \in \lambda \operatorname{sgn}(\hat{\beta}_j), \quad (32)$$

$$1 + \frac{1}{\sqrt{2n}}(\xi_1 - \xi_{j+1}) - \frac{1}{2}\sum_{k=1}^{2m}\hat{\beta}_k + \frac{1}{2}(\hat{\beta}_j - \hat{\beta}_{m+j}) \in \lambda \operatorname{sgn}(\hat{\beta}_{m+j}). \quad (33)$$

We will restrict our attention to the event  $\mathcal{B} = \{\xi_1 < 0\}$  which has a probability 1/2. First note that if  $\lambda \geq 1$ , then the vector  $\hat{\boldsymbol{\beta}} = \mathbf{0}$  is a solution to the system (32)-(33). Therefore,  $\hat{\boldsymbol{\beta}} = \mathbf{0}$  and hence  $\ell_n(\hat{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) = 2$ . Thus, in the case  $\lambda \geq 1$  the Lasso is not consistent, which is not a surprise since the theory recommends to always choose  $\lambda$  of order  $O((\log(p)/n)^{1/2})$ .

In the more interesting case  $\lambda \in (\frac{m+1-\sqrt{2n}}{m\sqrt{2n}}, 1)$ , a Lasso solution is given by

$$\hat{\beta}_j = \begin{cases} \frac{2(1-\lambda)}{m+1}, & \xi_{j+1} > 0, \\ 0, & \xi_{j+1} < 0, \end{cases} \quad \hat{\beta}_{m+j} = \begin{cases} 0, & \xi_{j+1} > 0, \\ \frac{2(1-\lambda)}{m+1}, & \xi_{j+1} < 0. \end{cases} \quad (34)$$

Indeed, for instance if  $\xi_{j+1} > 0$ , replacing these values of  $\hat{\boldsymbol{\beta}}$  in (32) and (33) we get on the event  $\mathcal{B}$ :

$$\begin{aligned} 1 + \frac{\xi_1 + \xi_{j+1}}{\sqrt{2n}} - \frac{1}{2}\sum_{k=1}^{2m}\hat{\beta}_k - \frac{1}{2}(\hat{\beta}_j - \hat{\beta}_{m+j}) &= 1 - \frac{m}{2} \times \frac{2(1-\lambda)}{m+1} - \frac{1}{2} \times \frac{2(1-\lambda)}{m+1} = \lambda \\ 1 + \frac{\xi_1 - \xi_{j+1}}{\sqrt{2n}} - \frac{1}{2}\sum_{k=1}^{2m}\hat{\beta}_k + \frac{1}{2}(\hat{\beta}_j - \hat{\beta}_{m+j}) &= 1 - \frac{2}{\sqrt{2n}} - \frac{m}{2} \times \frac{2(1-\lambda)}{m+1} + \frac{1}{2} \times \frac{2(1-\lambda)}{m+1} \\ &= \lambda \left(1 - \frac{2}{m+1}\right) - \frac{\sqrt{2}}{\sqrt{n}} + \frac{2}{m+1} \in [-\lambda, \lambda], \end{aligned}$$

where the last inclusion follows from the relation  $m+1 \geq \sqrt{2n}$ . For the vector (34), we check that the prediction loss

$$\begin{aligned} \ell_n(\widehat{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) &= \left( \sqrt{2} - \frac{\sqrt{2}m(1-\lambda)}{m+1} \right)^2 + \frac{2m(1-\lambda)^2}{(m+1)^2} \\ &= 2 \left( \frac{1+m\lambda}{m+1} \right)^2 + \frac{2m(1-\lambda)^2}{(m+1)^2} \\ &= \frac{2+2m\lambda^2}{m+1} \geq \frac{1}{m} \geq \frac{1}{\sqrt{2n}}. \end{aligned}$$

Finally, in the case  $\lambda \in [0, \frac{m+1-\sqrt{2n}}{m\sqrt{2n}}]$ , a Lasso solution on the event  $\mathcal{B}$  is given by

$$\widehat{\beta}_j = \begin{cases} \frac{\sqrt{2n+m-1}}{m\sqrt{2n}} - \lambda, & \xi_{j+1} > 0, \\ \frac{\sqrt{2n-m-1}}{m\sqrt{2n}} + \lambda, & \xi_{j+1} < 0, \end{cases} \quad \widehat{\beta}_{m+j} = \begin{cases} \frac{\sqrt{2n-m-1}}{m\sqrt{2n}} + \lambda, & \xi_{j+1} > 0, \\ \frac{\sqrt{2n+m-1}}{m\sqrt{2n}} - \lambda, & \xi_{j+1} < 0, \end{cases} \quad (35)$$

for every  $j \in [m]$ . Indeed, for instance if  $\xi_{j+1} > 0$ , replacing these values of  $\widehat{\boldsymbol{\beta}}$  in (32) and (33) we get on the event  $\mathcal{B}$ :

$$\begin{aligned} 1 + \frac{\xi_1 + \xi_{j+1}}{\sqrt{2n}} - \frac{1}{2} \sum_{k=1}^{2m} \widehat{\beta}_k - \frac{1}{2} (\widehat{\beta}_j - \widehat{\beta}_{m+j}) &= 1 - \frac{m}{2} \times \frac{2(\sqrt{2n}-1)}{m\sqrt{2n}} - \frac{1}{2} \times \left( \frac{2m}{m\sqrt{2n}} - 2\lambda \right) = \lambda \\ 1 + \frac{\xi_1 - \xi_{j+1}}{\sqrt{2n}} - \frac{1}{2} \sum_{k=1}^{2m} \widehat{\beta}_k + \frac{1}{2} (\widehat{\beta}_j - \widehat{\beta}_{m+j}) &= 1 - \frac{2}{\sqrt{2n}} - \frac{(\sqrt{2n}-1)}{\sqrt{2n}} + \left( \frac{1}{\sqrt{2n}} - \lambda \right) = -\lambda. \end{aligned}$$

The prediction loss of this estimator is

$$\begin{aligned} \ell_n(\widehat{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) &\geq m \left( \frac{1}{\sqrt{n}} - \sqrt{2}\lambda \right)^2 \\ &\geq 2m \left( \frac{1}{\sqrt{2n}} - \frac{m+1-\sqrt{2n}}{m\sqrt{2n}} \right)^2 \\ &= 2m \left( \frac{\sqrt{2n}-1}{m\sqrt{2n}} \right)^2 \geq \frac{1}{2\sqrt{2n}}, \end{aligned}$$

where for the last inequality we have used the facts that  $\sqrt{2n}-1 \geq \sqrt{n/2}$ ,  $\forall n \geq 2$ , and  $m \leq \sqrt{2n}$ . This completes the proof of the proposition.  $\square$

**PROOF OF PROPOSITION 3.** We will use a probabilistic argument. Set  $\mathbf{f} = \mathbf{X}\mathbf{u}$ . Denoting  $s_j = \text{sgn}(u_j) = \text{sgn}(f_j - f_{j-1})$ ,  $j \in T$ , and  $s_j = -\text{sgn}(u_j) = -\text{sgn}(f_j - f_{j-1})$ ,  $j \in T^c$ , with the convention that  $f_0 = 0$ ,  $a_{n+1} = a_n$  and  $s_{n+1} = -\text{sgn}(f_n)$ , we get

$$\begin{aligned} \|\mathbf{u}_T \odot \mathbf{a}_T\|_1 - \|\mathbf{u}_{T^c} \odot \mathbf{a}_{T^c}\|_1 &= \sum_{j \in T} a_j |f_j - f_{j-1}| - \sum_{j \in T^c} a_j |f_j - f_{j-1}| \\ &\leq \sum_{j \in T} a_j |f_j - f_{j-1}| - \sum_{j \in T^c} a_j |f_j - f_{j-1}| + a_n |f_n| \\ &= \sum_{j \in [n]} s_j a_j (f_j - f_{j-1}) - a_n s_{n+1} f_n = \sum_{j \in [n]} f_j (s_j a_j - s_{j+1} a_{j+1}). \end{aligned}$$

More interestingly, since  $|x| = \max_{t \in \{\pm 1\}} tx$  for every  $x \in \mathbb{R}$ , one easily checks that<sup>8</sup>

$$\|\mathbf{u}_T \odot \mathbf{a}_T\|_1 - \|\mathbf{u}_{T^c} \odot \mathbf{a}_{T^c}\|_1 = \min_{\substack{\mathbf{t} \in \{\pm 1\}^{n+1} \\ \mathbf{t}_T = \mathbf{s}_T}} \sum_{j \in [n]} f_j(t_j a_j - t_{j+1} a_{j+1}).$$

Let us set  $j_0 = 1$ ,  $j_{s+1} = n + 1$  and consider the partition  $\{B_\ell = \llbracket j_{\ell-1}, j_\ell \llbracket, \ell \in [s+1]\}$  of  $[n]$ . Let  $\Delta_\ell = |B_\ell|$  for every  $\ell \in [s+1]$ . Permuting the minimum and the summation, we get

$$\|\mathbf{u}_T \odot \mathbf{a}_T\|_1 - \|\mathbf{u}_{T^c} \odot \mathbf{a}_{T^c}\|_1 = \sum_{\ell=1}^{s+1} \Psi_\ell$$

where

$$\begin{aligned} \Psi_\ell &= \min_{\substack{\mathbf{t} \in \{\pm 1\}^{n+1} \\ (t_{j_{\ell-1}}, t_{j_\ell}) = (s_{j_{\ell-1}}, s_{j_\ell})}} \sum_{j=j_{\ell-1}}^{j_\ell-1} f_j(t_j a_j - t_{j+1} a_{j+1}) \\ &= \min_{\substack{\bar{\mathbf{t}} \in \{\pm 1\}^{\Delta_\ell+1} \\ (t_1, t_{\Delta_\ell+1}) = (s_{j_{\ell-1}}, s_{j_\ell})}} \sum_{j=1}^{\Delta_\ell} f_{j_{\ell-1}+j-1}(\bar{t}_j a_{j_{\ell-1}+j-1} - \bar{t}_{j+1} a_{j_{\ell-1}+j}). \end{aligned}$$

In what follows, we propose a bound on  $\Psi_1$ . The other  $\Psi_\ell$ 's can be evaluated similarly. Let  $\epsilon_1, \dots, \epsilon_n$  be independent random variables with values in  $\pm 1$  such that  $p := \mathbf{P}(\epsilon_j = 1) = (1 - (2\Delta_1)^{-1})$  and  $\mathbf{P}(\epsilon_j = -1) = (2\Delta_1)^{-1}$ . Further, we define  $\bar{\epsilon}_1 = s_1$  and  $\bar{\epsilon}_{j+1} = \epsilon_1 \times \dots \times \epsilon_j$ , for every  $j \in B_1$ . Since  $\epsilon_j$ 's are independent,  $\{\bar{\epsilon}_j\}_{j \in [j_1]}$  is a Markov chain with values in  $\{\pm 1\}$ . For this Markov chain, we first check that  $\mathbf{P}(\bar{\epsilon}_{j_1} = s_{j_1}) \geq 1/4$ . Indeed, by symmetry, it suffices to consider the two cases  $(s_1, s_{j_1}) = (1, 1)$  and  $(s_1, s_{j_1}) = (1, -1)$ . In the first case, we use the inclusion  $\{\boldsymbol{\epsilon}_{B_1} = \mathbf{1}_{B_1}\} \subset \{\bar{\epsilon}_{j_1} = s_{j_1}\}$  to infer that

$$\mathbf{P}(\bar{\epsilon}_{j_1} = s_{j_1}) \geq \mathbf{P}(\boldsymbol{\epsilon}_{B_1} = \mathbf{1}_{B_1}) = \prod_{j=1}^{j_1-1} \mathbf{P}(\epsilon_j = 1) = (1 - (2\Delta_1)^{-1})^{\Delta_1}.$$

For  $\Delta_1 \geq 1$ , one checks that  $(1 - (2\Delta_1)^{-1})^{\Delta_1} \geq 1/2$ , which yields  $\mathbf{P}(\bar{\epsilon}_{j_1} = s_{j_1}) \geq \frac{1}{2}$ . In the second case,  $s_1 = -s_{j_1}$ , we use the inclusion  $\cup_{j \in B_1} \{\boldsymbol{\epsilon}_{-j} = \mathbf{1}, \epsilon_j = -1\} \subset \{\bar{\epsilon}_{j_1} = s_{j_1} = -s_1\}$ , where  $\boldsymbol{\epsilon}_{-j}$  is the vector obtained from  $\boldsymbol{\epsilon}_{B_1}$  by removing the  $j$ th entry. This inclusion yields

$$\mathbf{P}(\bar{\epsilon}_{j_1} = s_{j_1}) \geq \sum_{j=1}^{\Delta_1} \mathbf{P}(\boldsymbol{\epsilon}_{-j} = \mathbf{1}, \epsilon_j = -1) = \sum_{j=1}^{\Delta_1} (1 - (2\Delta_1)^{-1})^{\Delta_1-1} \times (2\Delta_1)^{-1} \geq \frac{1}{4}. \quad (36)$$

On the other hand, for every  $A > 0$ , we have

$$\mathbf{P}\left(\sum_{j=1}^{\Delta_1} f_j(\bar{\epsilon}_j a_j - \bar{\epsilon}_{j+1} a_{j+1}) > A\right) \leq \frac{1}{A} \mathbf{E}\left[\left|\sum_{j=1}^{\Delta_1} f_j(\bar{\epsilon}_j a_j - \bar{\epsilon}_{j+1} a_{j+1})\right|\right].$$

<sup>8</sup>For notational convenience, we assume hereafter that  $T$  is augmented by the elements  $\{1, n+1\}$ .

We need now to evaluate the expectation of the random variable  $\Upsilon = |\sum_{j=1}^{\Delta_1} f_j \bar{\epsilon}_j (a_j - \epsilon_{j+1} a_{j+1})|$ . To this end, since all the  $a_j$  are nonnegative, we remark that  $\mathbf{E}(|a_j - \epsilon_{j+1} a_{j+1}|) = |a_j - a_{j+1}| + (a_j \wedge a_{j+1})/\Delta_1$  and, therefore,

$$\mathbf{E}[\Upsilon] \leq \sum_{j=1}^{\Delta_1} |f_j| \cdot (|a_j - a_{j+1}| + a_j \Delta_1^{-1}) \leq \left( \sum_{j=1}^{\Delta_1} f_j^2 \right)^{1/2} \left( \sum_{j=1}^{\Delta_1} (2|a_j - a_{j+1}|^2 + 2a_j^2 \Delta_1^{-2}) \right)^{1/2}.$$

Thus, taking  $A = 4 \left( \sum_{j=1}^{\Delta_1} f_j^2 \right)^{1/2} \left( \sum_{j=1}^{\Delta_1} (2|a_j - a_{j+1}|^2 + 2a_j^2 \Delta_1^{-2}) \right)^{1/2}$ , we get

$$\mathbf{P} \left( \sum_{j=1}^{\Delta_1} f_j (\bar{\epsilon}_j a_j - \bar{\epsilon}_{j+1} a_{j+1}) > A \right) \leq 1/4.$$

Combined with inequality (36), this entails that

$$\mathbf{P} \left( \sum_{j=1}^{\Delta_1} f_j (\bar{\epsilon}_j a_j - \bar{\epsilon}_{j+1} a_{j+1}) \leq A \text{ and } (\bar{\epsilon}_1, \bar{\epsilon}_{j_1}) = (s_1, s_{j_1}) \right) > 0.$$

Consequently,  $\Psi_1 \leq A$ . Applying the same argument to arbitrary  $\ell \in [s+1]$ , we get

$$\Psi_\ell \leq 4 \left( \sum_{j \in B_\ell} f_j^2 \right)^{1/2} \left( \sum_{j \in B_\ell} (2|a_j - a_{j+1}|^2 + 2a_j^2 \Delta_\ell^{-2}) \right)^{1/2}.$$

In view of the Cauchy-Schwarz inequality, we get

$$\begin{aligned} \|\mathbf{u}_T \odot \mathbf{a}_T\|_1 - \|\mathbf{u}_{T^c} \odot \mathbf{a}_{T^c}\|_1 &\leq 4 \left( \sum_{j \in [n]} f_j^2 \right)^{1/2} \left( \sum_{\ell=1}^{s+1} \sum_{j \in B_\ell} (2|a_j - a_{j+1}|^2 + 2a_j^2 \Delta_\ell^{-2}) \right)^{1/2} \\ &\leq 4 \|\mathbf{f}\|_2 \left( 2 \sum_{j \in [n]} |a_j - a_{j+1}|^2 + 2(s+1) \|\mathbf{a}\|_\infty^2 \max_\ell \Delta_\ell^{-1} \right)^{1/2}. \end{aligned}$$

This completes the proof.  $\square$

**PROOF OF PROPOSITION 4.** We apply Theorem 3 with  $\gamma = 2$ . This leads to (14) with

$$r_{n,n,T} = \frac{1 + 2|T|^{-1} \log(1/\delta)}{\log(p/\delta)} + \frac{4}{\bar{\kappa}_{T,2,\omega}} \leq 3 + \frac{4}{\bar{\kappa}_{T,2,\omega}}.$$

It remains to find a lower bound for  $\bar{\kappa}_{T,2,\omega}$ . To this end, we resort to Proposition 3 with  $a_j = 1$  for every  $j$  belonging to the set  $T$  and  $a_j = 1 - \frac{1}{2\sqrt{n}} \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2$ ,  $j \in T^c$ . Let  $T = \{j_1, \dots, j_s\}$  and  $B_\ell = \llbracket j_{\ell-1}, j_\ell \rrbracket$  for  $\ell \in [s+1]$  with the convention that  $j_0 = 1$  and  $j_{s+1} = n+1$ . Since the columns of  $\mathbf{X}$  are given by  $\mathbf{x}^j = [\mathbf{1}(i \geq j)]_{i \in [n]}$ , the projector  $\Pi_T$  projects onto the subspace of  $\mathbb{R}^n$  containing all the vectors that are constant on the partition  $\{B_\ell\}$ . Therefore, one easily checks that  $\|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2 = \sqrt{\frac{(j-j_{\ell-1})(j_\ell-j)}{j_\ell-j_{\ell-1}}}$ , for every  $j \in B_\ell$ .

This implies that  $\mathbf{a}_T = \mathbf{1}_T$  and  $\|\mathbf{a}\|_\infty = 1$ . Furthermore, we have

$$\begin{aligned} \sum_{j \in [n]} |a_j - a_{j+1}|^2 &\leq \frac{1}{4n} \sum_{\ell \in [s+1]} \sum_{j \in B_\ell} \frac{(\sqrt{(j-j_{\ell-1})(j_\ell-j)} - \sqrt{(j+1-j_{\ell-1})(j_\ell-j-1)})^2}{j_\ell - j_{\ell-1}} \\ &\leq \frac{1}{4n} \sum_{\ell \in [s+1]} \sum_{j=1}^{\Delta_\ell} \frac{((j-1)(\Delta_\ell - (j-1)) - j(\Delta_\ell - j))^2}{\Delta_\ell(\sqrt{(j-1)(\Delta_\ell - (j-1))} + \sqrt{j(\Delta_\ell - j)})^2} \\ &\leq \frac{1}{4n} \sum_{\ell \in [s+1]} \sum_{j=1}^{\Delta_\ell} \frac{(2j-3-\Delta_\ell)^2}{\Delta_\ell((j-1)(\Delta_\ell - (j-1)) + j(\Delta_\ell - j))} \\ &\leq \frac{1}{4n} \sum_{\ell \in [s+1]} \sum_{j \leq \Delta_\ell/2} \frac{\Delta_\ell^2}{\Delta_\ell \times j \times \Delta_\ell/2} = \frac{1}{2n} \sum_{\ell \in [s+1]} \sum_{j \leq \Delta_\ell/2} j^{-1}. \end{aligned}$$

Since obviously  $\Delta_\ell \leq n$ , we can bound the sum  $\sum_{j \leq \Delta_\ell/2} j^{-1}$  by  $1 + \log(n/2) \leq 2 \log(n)$ , provided that  $n \geq 3$ . This yields  $\sum_{j \in [n]} |a_j - a_{j+1}|^2 \leq (s+1) \log(n)/n$ . Therefore, Proposition 3 implies that

$$\|\mathbf{u}_T\|_1 - \|\mathbf{u}_{T^c} \odot \mathbf{a}_{T^c}\|_1 \leq 4\|\mathbf{X}\mathbf{u}\|_2 \left( 2(s+1) \left[ \frac{\log(n)}{n} + \frac{1}{\Delta_{\min}} \right] \right)^{1/2}.$$

Using the inequality  $s+1 \leq 2s$ , we infer from the above inequality that  $\bar{\kappa}_{T,2,\omega} \geq \{64(\log(n) + (n/\Delta_{\min}))\}^{-1}$ .  $\square$

**PROOF OF PROPOSITION 5.** Let us denote  $\boldsymbol{\delta} = \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$ . We consider two cases separately. The first case is when the inequality  $\|\hat{\boldsymbol{\beta}}\|_1 \leq \frac{\gamma+1}{\gamma-1} \|\boldsymbol{\beta}^*\|_1$  is satisfied. Then, using the Parseval identity, we have  $\ell_n(\hat{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) = \frac{1}{n} \|\Pi_T \mathbf{X} \boldsymbol{\delta}\|_2^2 + \frac{1}{n} \|(\mathbf{I}_n - \Pi_T) \mathbf{X} \boldsymbol{\delta}\|_2^2$ . The first summand can be bounded using inequality (4), so we focus on the second summand. Using the triangle inequality, we get  $\|(\mathbf{I}_n - \Pi_T) \mathbf{X} \boldsymbol{\delta}\|_2 = \|\sum_j (\mathbf{I}_n - \Pi_T) \mathbf{x}^j \delta_j\|_2 \leq \sum_j \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j \delta_j\|_2 \leq \|\boldsymbol{\delta}\|_1 \max_{j \in T^c} \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2$ . Furthermore, we have  $\|\boldsymbol{\delta}\|_1 \leq \|\hat{\boldsymbol{\beta}}\|_1 + \|\boldsymbol{\beta}^*\|_1 \leq \frac{2\gamma}{\gamma-1} \|\boldsymbol{\beta}^*\|_1$ . Hence, putting these bounds together, we find

$$\begin{aligned} \ell_n(\hat{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) &\leq \frac{1}{n} \|\Pi_T \mathbf{X} \boldsymbol{\delta}\|_2^2 + \frac{4\gamma^2 \|\boldsymbol{\beta}^*\|_1^2}{n(\gamma-1)^2} \max_{j \in T^c} \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2^2 \\ &\leq \frac{2\lambda^2 |T|}{\nu_T^2} + \frac{4\sigma^{*2} (|T| + 2 \log(1/\delta))}{n} + \frac{4\gamma^2 \|\boldsymbol{\beta}^*\|_1^2}{n(\gamma-1)^2} \max_{j \in T^c} \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2^2, \end{aligned}$$

with a probability at least  $1 - \delta$ . In the second case,  $\|\hat{\boldsymbol{\beta}}\|_1 > \frac{\gamma+1}{\gamma-1} \|\boldsymbol{\beta}^*\|_1$ , according to inequalities (30) and (31) applied to  $\bar{\boldsymbol{\delta}} = \boldsymbol{\delta}$ , we have

$$\begin{aligned} \frac{2}{n} \|\mathbf{X} \boldsymbol{\delta}\|_2^2 &\leq \frac{2}{n} \boldsymbol{\delta}^\top \mathbf{X}^\top \boldsymbol{\xi} + 2\lambda (\|\boldsymbol{\beta}^*\|_1 - \|\hat{\boldsymbol{\beta}}\|_1) \\ &\leq \frac{1}{n} \|\mathbf{X} \boldsymbol{\delta}\|_2^2 + \frac{2\sigma^{*2} (|T| + 2 \log(1/\delta))}{n} + 2\lambda \gamma^{-1} \underbrace{((\gamma+1)\|\boldsymbol{\beta}^*\|_1 - (\gamma-1)\|\hat{\boldsymbol{\beta}}\|_1)}_{\leq 0}, \end{aligned}$$

with probability at least  $1 - 2\delta$ . This completes the proof.  $\square$

PROOF OF PROPOSITION 6. Applying Theorem 4 to  $\lambda = \sigma^*(\log(n/\delta)/(nk))^{1/2}$ , in conjunction with the bound  $\rho_T \leq (2k)^{-1/2}$ , we get that the inequality

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \inf_{\bar{\mathbf{f}} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\bar{\mathbf{f}} - \mathbf{f}^*\|_2^2 + 4\lambda \|\bar{\mathbf{f}}\|_{\text{TV}} \right\} + \frac{2\sigma^{*2}(k + 2\log(1/\delta))}{n} \quad (37)$$

holds true with a probability at least  $1 - 2\delta$ . Replacing  $\bar{\mathbf{f}}$  by  $\mathbf{f}^\dagger$ , and replacing  $\lambda$  by its expression, we get

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \frac{1}{n} \|\mathbf{f}^\dagger - \mathbf{f}^*\|_2^2 + \frac{4\sigma^{*2} \log(1/\delta)}{n} + \underbrace{4\sigma^* \left( \frac{\log(n/\delta)}{nk} \right)^{1/2} \|\mathbf{f}^\dagger\|_{\text{TV}} + \frac{2\sigma^{*2}k}{n}}_{:=\Psi(k)},$$

with probability  $\geq 1 - 2\delta$ . One readily checks that  $x \mapsto \Psi(x)$  achieves its (global) minimum at  $x_{\min} = (\|\mathbf{f}^\dagger\|_{\text{TV}}/\sigma^*)^{2/3} (n \log(n/\delta))^{1/3}$ . Furthermore, the definition of  $k$  entails that  $1/k \leq 1/x_{\min}$  and  $k \leq x_{\min} + 1$ . This yields  $\Psi(k) \leq \Psi(x_{\min}) + 2\sigma^{*2}/n$  and the desired result follows.  $\square$

PROOF OF PROPOSITION 7. We start by observing that  $\mathcal{H}_{\alpha,L}^n$  is a closed convex subset of  $\mathbb{R}^n$ . Therefore, for every  $\bar{\mathbf{f}} \in \mathcal{H}_{\alpha,L}^n$  and for  $\mathbf{g}^* = \arg \min_{\mathbf{f} \in \mathcal{H}_{\alpha,L}^n} \|\mathbf{f} - \mathbf{f}^*\|_2$ , we have

$$\|\bar{\mathbf{f}} - \mathbf{f}^*\|_2^2 \leq \|\mathbf{g}^* - \mathbf{f}^*\|_2^2 + \|\bar{\mathbf{f}} - \mathbf{g}^*\|_2^2. \quad (38)$$

Now, let  $T = \{a_1, a_2, \dots, a_k\} \subset [n]$  be any set satisfying  $a_1 = 1$  and  $0 \leq a_{j+1} - a_j \leq 2n/k$ ,  $\forall j \in [k]$ . This set induces the partition  $\mathcal{T} = \{I_1, \dots, I_k\}$  where each  $I_j = \llbracket a_j, a_{j+1} \rrbracket$  is an interval of length smaller or equal to  $2n/k$ . We define  $\bar{\mathbf{f}}$  as an approximation of  $\mathbf{g}^*$  by a piecewise linear vector on the partition  $\mathcal{T}$ :  $\bar{f}_i = g_{a_j}^* + \frac{i-a_j}{a_{j+1}-a_j} \{g_{a_{j+1}}^* - g_{a_j}^*\}$ ,  $\forall i \in I_j$ . On the one hand, one can easily check that  $\bar{\mathbf{f}}$  belongs to the set  $\mathcal{H}_{\alpha,L}^n$  and, therefore, satisfies inequality (38). In conjunction with (37), this implies that for  $\lambda = \sigma^*(\log(n/\delta)/(nk))^{1/2}$ , the inequalities

$$\begin{aligned} \frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 &\leq \inf_{\bar{\mathbf{f}} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\bar{\mathbf{f}} - \mathbf{f}^*\|_2^2 + 4\lambda \|\bar{\mathbf{f}}\|_{\text{TV}} \right\} + \frac{2\sigma^{*2}(k + 2\log(1/\delta))}{n} \\ &\leq \frac{1}{n} \|\mathbf{g}^* - \mathbf{f}^*\|_2^2 + \frac{1}{n} \|\bar{\mathbf{f}} - \mathbf{g}^*\|_2^2 + 4\lambda \|\bar{\mathbf{f}}\|_{\text{TV}} + \frac{2\sigma^{*2}(k + 2\log(1/\delta))}{n} \end{aligned}$$

hold true with a probability at least  $1 - 2\delta$ . Since  $\mathbf{g}^* \in \mathcal{H}_{\alpha,L}^n$ , we have

$$\begin{aligned} \|\bar{\mathbf{f}} - \mathbf{g}^*\|_2^2 &\leq \sum_j \sum_{i \in I_j} \left| g_{a_j}^* + \frac{i-a_j}{a_{j+1}-a_j} \{g_{a_{j+1}}^* - g_{a_j}^*\} - g_i^* \right|^2 \\ &\leq \sum_j \sum_{i \in I_j} \left( \frac{a_{j+1}-i}{a_{j+1}-a_j} |g_{a_j}^* - g_i^*|^2 + \frac{i-a_j}{a_{j+1}-a_j} |g_{a_{j+1}}^* - g_i^*|^2 \right) \\ &\leq nL^2 n^{-2\alpha} (2n/k)^{2\alpha} \leq 4nL^2 k^{-2\alpha}. \end{aligned}$$

On the other hand, since  $\bar{\mathbf{f}}$  is piecewise constant, we have  $\|\bar{\mathbf{f}}\|_{\text{TV}} = \sum_j |g_{a_{j+1}}^* - g_{a_j}^*| \leq kLn^{-\alpha} (2n/k)^\alpha \leq 2Lk^{1-\alpha}$ . Combining all these bounds, we get that with probability at least

$1 - 2\delta$ ,

$$\begin{aligned} \frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 &\leq \frac{1}{n} \|\mathbf{g}^* - \mathbf{f}^*\|_2^2 + \frac{4\sigma^{*2} \log(1/\delta)}{n} \\ &\quad + 4L^2 k^{-2\alpha} + 8L\sigma^* \left( \frac{\log(n/\delta)}{nk} \right)^{1/2} k^{1-\alpha} + \frac{2\sigma^{*2} k}{n}. \end{aligned}$$

The inequality between the geometric and arithmetic means yields  $8L\sigma^* \left( \frac{\log(n/\delta)}{nk} \right)^{1/2} k^{1-\alpha} \leq 8L^2 k^{-2\alpha} + 2\sigma^{*2} \frac{k \log(n/\delta)}{n}$ . Thus, with probability at least  $1 - 2\delta$ , we have

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \frac{1}{n} \|\mathbf{g}^* - \mathbf{f}^*\|_2^2 + \frac{4\sigma^{*2} \log(1/\delta)}{n} + 12L^2 k^{-2\alpha} + \frac{4\sigma^{*2} k \log(n/\delta)}{n}.$$

Using the inequalities  $k - 1 \leq (L^2 n / (\sigma^{*2} \log(n/\delta)))^{1/(2\alpha+1)} \leq k$  we get

$$12L^2 k^{-2\alpha} \leq 12L^2 \left( \frac{L^2 n}{\sigma^{*2} \log(n/\delta)} \right)^{-2\alpha/(2\alpha+1)} = 12L^2 \left( \frac{\sigma^{*2} \log(n/\delta)}{L^2 n} \right)^{2\alpha/(2\alpha+1)},$$

and

$$\begin{aligned} \frac{4\sigma^{*2} k \log(n/\delta)}{n} &\leq \frac{4\sigma^{*2} \log(n/\delta)}{n} \times \left( \frac{L^2 n}{\sigma^{*2} \log(n/\delta)} \right)^{1/(2\alpha+1)} + \frac{4\sigma^{*2} \log(n/\delta)}{n} \\ &= 4L^2 \left( \frac{\sigma^{*2} \log(n/\delta)}{L^2 n} \right)^{2\alpha/(2\alpha+1)} + \frac{4\sigma^{*2} \log(n/\delta)}{n}. \end{aligned}$$

To complete, we use the fact that  $\|\mathbf{g}^* - \mathbf{f}^*\|_2^2 = \inf_{\mathbf{f} \in \mathcal{H}_{\alpha, L}^n} \frac{1}{n} \|\mathbf{f} - \mathbf{f}^*\|_2^2$ .  $\square$

**PROOF OF PROPOSITION 8.** Let  $\mathbf{u}^*$  be any vector satisfying the cone constraint  $\|\mathbf{u}_{T^c}^*\|_1 \leq \bar{c} \|\mathbf{u}_T^*\|_1$ . Clearly, the vector  $-\mathbf{u}^*$  satisfies the same constraint. For any  $J \subset T^c$ , we define  $\boldsymbol{\zeta} = \sigma^{*-1} (\mathbf{X}_J^\top \mathbf{X}_J)^\dagger \mathbf{X}_J^\top \boldsymbol{\xi}$ . Let  $\mathbf{v} \in \mathbb{R}^p$  be the random vector defined by  $\mathbf{v}_J = \alpha \boldsymbol{\zeta}$  and  $\mathbf{v}_{J^c} = 0$ , where  $\alpha = (\bar{c} \|\mathbf{u}_T^*\|_1 - \|\mathbf{u}_{T^c}^*\|_1) / \|\boldsymbol{\zeta}\|_1$  is a positive number ensuring that the vectors  $\pm \mathbf{u}^* + \mathbf{v}$  satisfy the cone constraint. Indeed, the triangle inequality implies that

$$\begin{aligned} \|(\pm \mathbf{u}^* + \mathbf{v})_{T^c}\|_1 - \bar{c} \|(\pm \mathbf{u}^* + \mathbf{v})_T\|_1 &= \|(\pm \mathbf{u}^* + \mathbf{v})_{T^c}\|_1 - \bar{c} \|\mathbf{u}_T^*\|_1 \\ &\leq \|\mathbf{u}_{T^c}^*\|_1 + \|\mathbf{v}_J\|_1 - \bar{c} \|\mathbf{u}_T^*\|_1 = 0. \end{aligned}$$

Further, we remark that  $\boldsymbol{\xi}^\top \mathbf{X} \mathbf{v} = (\alpha / \sigma^*) \boldsymbol{\xi}^\top \Pi_J \boldsymbol{\xi} \geq 0$  and therefore

$$\begin{aligned} \bar{\eta} &\geq \max_{\mathbf{u} \in \{\pm \mathbf{u}^*\}} \frac{|\boldsymbol{\xi}^\top \mathbf{X}(\mathbf{u} + \mathbf{v})|}{\sigma^* \|\mathbf{X}(\mathbf{u} + \mathbf{v})\|_2} \geq \max_{\mathbf{u} \in \{\pm \mathbf{u}^*\}} \frac{|\boldsymbol{\xi}^\top \mathbf{X}(\mathbf{u} + \mathbf{v})|}{\sigma^* (\|\mathbf{X} \mathbf{u}^*\|_2 + \|\mathbf{X} \mathbf{v}\|_2)} \\ &= \frac{|\boldsymbol{\xi}^\top \mathbf{X} \mathbf{u}^*| + \boldsymbol{\xi}^\top \mathbf{X} \mathbf{v}}{\sigma^* (\|\mathbf{X} \mathbf{u}^*\|_2 + \|\mathbf{X} \mathbf{v}\|_2)} \geq \frac{\boldsymbol{\xi}^\top \mathbf{X} \mathbf{v}}{\sigma^* (\|\mathbf{X} \mathbf{u}^*\|_2 + \|\mathbf{X} \mathbf{v}\|_2)} := \eta_1. \end{aligned}$$

Using the definition of  $\mathbf{v}$ , we get the relations  $\boldsymbol{\xi}^\top \mathbf{X} \mathbf{v} = \boldsymbol{\xi}^\top \mathbf{X}_J \mathbf{v}_J = (\alpha / \sigma^*) \|\Pi_J \boldsymbol{\xi}\|_2^2$  and  $\|\mathbf{X} \mathbf{v}\|_2 = \|\mathbf{X}_J \mathbf{v}_J\|_2 = (\alpha / \sigma^*) \|\Pi_J \boldsymbol{\xi}\|_2$  implying that

$$\eta_1 = \frac{\alpha \|\Pi_J(\boldsymbol{\xi} / \sigma^*)\|_2^2}{\|\mathbf{X} \mathbf{u}^*\|_2 + \alpha \|\Pi_J(\boldsymbol{\xi} / \sigma^*)\|_2}.$$

To complete the proof, it remains to find appropriate lower bounds for the terms  $\|\Pi_J(\boldsymbol{\xi}/\sigma^*)\|_2^2$  and  $\alpha$ . Since  $\|\Pi_J(\boldsymbol{\xi}/\sigma^*)\|_2^2 \sim \chi_{|J|}^2$ , we have  $\mathbf{P}\{\|\Pi_J\boldsymbol{\xi}\|_2^2 \geq \sigma^{*2}(|J| - 2\sqrt{|J|\log(2/\delta')})\} \geq 1 - (\delta'/2)$ . Therefore, for every  $|J| \geq 10 > (8/3)^2 \log 4$ , we have  $\mathbf{P}\{\|\Pi_J\boldsymbol{\xi}\|_2^2 \geq \sigma^{*2}|J|/4\} \geq 1 - 1/4$ . Let  $\mathbf{X}_J = \mathbf{V}\boldsymbol{\Lambda}\mathbf{U}^\top$  be the singular value decomposition of  $\mathbf{X}_J$  and  $\tilde{\boldsymbol{\xi}} = \mathbf{V}^\top \boldsymbol{\xi}/\sigma^* \in \mathbb{R}^{|J|}$  with Gaussian distribution  $\mathcal{N}(0, \mathbf{I}_{|J|})$ . We remark that

$$\begin{aligned} \|(\mathbf{X}_J^\top \mathbf{X}_J)^\dagger \mathbf{X}_J^\top \boldsymbol{\xi}\|_1 &= \sigma^* \|\mathbf{U}\boldsymbol{\Lambda}^{-1} \tilde{\boldsymbol{\xi}}\|_1 \\ &\leq \sigma^* |J|^{1/2} \|\mathbf{U}\boldsymbol{\Lambda}^{-1} \tilde{\boldsymbol{\xi}}\|_2 \\ &= \sigma^* |J|^{1/2} \|\boldsymbol{\Lambda}^{-1} \tilde{\boldsymbol{\xi}}\|_2 \\ &\leq \sigma^* (|J|/n\lambda_{\min,J}^2)^{1/2} \|\tilde{\boldsymbol{\xi}}\|_2 \end{aligned}$$

Therefore, using tail bounds for the  $\chi^2$  distribution, we find

$$\mathbf{P}\left(\|\zeta\|_1 \leq (|J|/n\lambda_{\min,J}^2)^{1/2} \underbrace{(|J|^{1/2} + \sqrt{2\log(4)})}_{\leq 2|J|^{1/2}}\right) \geq 1 - \frac{1}{4}.$$

From this inequality we infer that  $\alpha \geq n^{1/2}\lambda_{\min,J}(\bar{c}\|\mathbf{u}_T^*\|_1 - \|\mathbf{u}_{T^c}^*\|_1)/(2|J|)$  with probability at least 3/4. Combining the lower bounds obtained for  $\alpha$  and  $\|\Pi_J\boldsymbol{\xi}\|_2^2$ , we get

$$\mathbf{P}\left(\eta_1 \geq \frac{|J|}{4} \left\{ \frac{2|J| \cdot \|\mathbf{X}\mathbf{u}^*\|_2}{n^{1/2}\lambda_{\min,J}\bar{c}(\|\mathbf{u}_T^*\|_1 - \bar{c}^{-1}\|\mathbf{u}_{T^c}^*\|_1)} + \frac{|J|^{1/2}}{2} \right\}^{-1}\right) \geq 1/2.$$

This implies that

$$\bar{\varrho}_{T,\bar{c}} \geq \left\{ \frac{8\|\mathbf{X}\mathbf{u}^*\|_2}{n^{1/2}\lambda_{\min,J}(\|\mathbf{u}_T^*\|_1 - \bar{c}^{-1}\|\mathbf{u}_{T^c}^*\|_1)} + 2|J|^{-1/2} \right\}^{-1}.$$

Since this is true for every  $\mathbf{u}^*$  belonging to the cone, we can take the supremum of the right hand-side to get

$$\bar{\varrho}_{T,\bar{c}} \geq \frac{1}{8\lambda_{\min,J}^{-1}(\kappa_{T,\bar{c}}/|T|)^{1/2} + 2|J|^{-1/2}},$$

and the desired result follows.  $\square$

**8. Acknowledgments.** Johannes Lederer acknowledges partial financial support as fellow of the Swiss National Science Foundation. The work of Arnak Dalalyan was partially supported by the grant Investissements d'Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047).

## References.

- [AH12] Pierre Alquier and Mohamed Hebiri. Transductive versions of the lasso and the dantzig selector. *Journal of Statistical Planning and Inference*, 142(9):2485–2500, 2012.
- [BCW13] Alexandre Belloni, Victor Chernozhukov, and Lie Wang. Pivotal estimation via square-root lasso in nonparametric regression. Technical Report arXiv:1105.1475, December 2013.
- [BJMO12] Francis R. Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.

- [BJNP13] Aharon Birnbaum, Iain M. Johnstone, Boaz Nadler, and Debashis Paul. Minimax bounds for sparse PCA with noisy high-dimensional data. *Ann. Statist.*, 41(3):1055–1084, 2013.
- [BKL<sup>+</sup>09] Leif Boysen, Angela Kempe, Volkmar Liebscher, Axel Munk, and Olaf Wittich. Consistencies and rates of convergence of jump-penalized least squares estimators. *The Annals of Statistics*, 37(1):157–183, 02 2009.
- [BRT09] Peter J. Bickel, Ya'acov Ritov, and Alexandre B. Tsybakov. Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.*, 37(4):1705–1732, 2009.
- [BRvdGZ13] Peter Bühlmann, Philipp Rütimann, Sara van de Geer, and Cun-Hui Zhang. Correlated variables in regression: clustering and sparse estimation. *J. Statist. Plann. Inference*, 143(11):1835–1858, 2013.
- [BTW07a] F. Bunea, A. Tsybakov, and M. Wegkamp. Aggregation for Gaussian regression. *Ann. Statist.*, 35(4):1674–1697, 2007.
- [BTW07b] Florentina Bunea, Alexandre Tsybakov, and Marten Wegkamp. Sparsity oracle inequalities for the Lasso. *Electron. J. Stat.*, 1:169–194, 2007.
- [BvdG11] Peter Bühlmann and Sara van de Geer. *Statistics for High Dimensional Data. Methods, Theory and Applications*. Springer, 2011.
- [BW13] Jacob Bien and Marten Wegkamp. Discussion of “Correlated variables in regression: clustering and sparse estimation” [mr3095072]. *J. Statist. Plann. Inference*, 143(11):1859–1862, 2013.
- [CP09] Emmanuel J. Candès and Yaniv Plan. Near-ideal model selection by  $\ell_1$  minimization. *Ann. Statist.*, 37(5A):2145–2177, 2009.
- [CT07] Emmanuel Candès and Terence Tao. The Dantzig selector: statistical estimation when  $p$  is much larger than  $n$ . *Ann. Statist.*, 35(6):2313–2351, 2007.
- [CWX10] T. Tony Cai, Lie Wang, and Guangwu Xu. Shifting inequality and recovery of sparse signals. *IEEE Trans. Signal Process.*, 58(3, part 1):1300–1308, 2010.
- [DC12] Arnak S. Dalalyan and Yin Chen. Fused sparsity and robust estimation for linear models with unknown variance. In *Advances in Neural Information Processing Systems 25: NIPS*, pages 1268–1276, 2012.
- [DT07] Arnak S. Dalalyan and Alexandre B. Tsybakov. Aggregation by exponential weighting and sharp oracle inequalities. In *Learning theory (COLT2007), Lecture Notes in Comput. Sci., Vol. 4539*, pages 97–111, 2007.
- [DT12a] Arnak S. Dalalyan and Alexandre B. Tsybakov. Mirror averaging with sparsity priors. *Bernoulli*, 18(3):914–944, 2012.
- [DT12b] Arnak S. Dalalyan and Alexandre B. Tsybakov. Sparse regression learning by aggregation and Langevin Monte-Carlo. *J. Comput. System Sci.*, 78(5):1423–1443, 2012.
- [EHJT04] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Statist.*, 32(2):407–499, 2004. With discussion, and a rejoinder by the authors.
- [HL13] Mohamed Hebiri and Johannes Lederer. How correlations influence Lasso prediction. *IEEE Trans. Inform. Theory*, 59(3):1846–1854, 2013.
- [HLL07] Zaïd Harchaoui and Céline Lévy-Leduc. Catching change-points with lasso. In *NIPS*, 2007.
- [HLL10] Zaïd Harchaoui and Céline Lévy-Leduc. Multiple change-point estimation with a total variation penalty. *Journal of the American Statistical Association*, 105(492):1480–1493, 2010.
- [JN11a] Anatoli Juditsky and Arkadi Nemirovski. Accuracy guarantees for  $\ell_1$ -recovery. *IEEE Trans. Inform. Theory*, 57(12):7818–7839, 2011.
- [JN11b] Anatoli Juditsky and Arkadi Nemirovski. Accuracy guarantees for  $\ell_1$ -recovery. *IEEE Trans. Inform. Theory*, 57(12):7818–7839, 2011.
- [KLT11] Vladimir Koltchinskii, Karim Lounici, and Alexandre B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.*, 39(5):2302–2329, 2011.
- [Kol11] Vladimir Koltchinskii. *Oracle inequalities in empirical risk minimization and sparse recovery problems*, volume 2033 of *Lecture Notes in Mathematics*. Springer, Heidelberg, 2011. Lectures from the 38th Probability Summer School held in Saint-Flour, 2008, École d’Été de Probabilités de Saint-Flour. [Saint-Flour Probability Summer School].
- [LPvdGT11] Karim Lounici, Massimiliano Pontil, Sara van de Geer, and Alexandre B. Tsybakov. Oracle inequalities and optimal inference under group sparsity. *Ann. Statist.*, 39(4):2164–2204, 2011.
- [MvdG87] Enno Mammen and Sara van de Geer. Locally adaptive regression splines. *The Annals of Statistics*, 25(1):387–413, 1987.
- [RT11] P. Rigollet and A. Tsybakov. Exponential Screening and optimal rates of sparse estimation.

- Ann. Statist.*, 39(2):731–771, 2011.
- [RWY10] Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated Gaussian designs. *J. Mach. Learn. Res.*, 11:2241–2259, 2010.
- [SBvdG10] Nicolas Städler, Peter Bühlmann, and Sara van de Geer.  $\ell_1$ -penalization for mixture regression models. *TEST*, 19(2):209–256, 2010.
- [SCS13] Adityanand Guntuboyina Sabyasachi Chatterjee and Bodhisattva Sen. Improved risk bounds in isotonic regression. Technical report, arxiv:1311.3765, 2013.
- [SS13] Rajen D. Shah and Richard J. Samworth. Discussion of ‘Correlated variables in regression: clustering and sparse estimation’ by Peter Bühlmann, Philipp Rütimann, Sara van de Geer and Cun-Hui Zhang [mr3095072]. *J. Statist. Plann. Inference*, 143(11):1866–1868, 2013.
- [SZ12] Tingni Sun and Cun-Hui Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- [Tib96] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996.
- [VB10] Jean-Philippe Vert and Kevin Bleakley. Fast detection of multiple change-points shared by many signals using group lars. In *Advances in Neural Information Processing Systems*, pages 2343–2351, 2010.
- [vdG07] Sara van de Geer. The deterministic lasso. In *Proc. of Joint Statistical Meeting*, 2007.
- [vdGB09] Sara van de Geer and Peter Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.*, 3:1360–1392, 2009.
- [vdGL13] Sara van de Geer and Johannes Lederer. The Lasso, correlated design, and improved oracle inequalities. *IMS Collections*, 9:303–316, 2013.
- [Wai09] Martin J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using  $\ell_1$ -constrained quadratic programming (Lasso). *IEEE Trans. Inform. Theory*, 55(5):2183–2202, 2009.
- [YL06] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(1):49–67, 2006.
- [Zha02] Cun-Hui Zhang. Risk bounds in isotonic regression. *Ann. Statist.*, 30(2):528–555, 2002.
- [Zha09] Tong Zhang. Some sharp performance bounds for least squares regression with  $L_1$  regularization. *Ann. Statist.*, 37(5A):2109–2144, 2009.

3, AVENUE PIERRE LAROUSSE,  
92240 MALAKOFF  
FRANCE  
E-MAIL: [arnak.dalalyan@ensae.fr](mailto:arnak.dalalyan@ensae.fr)  
URL: <http://arnak-dalalyan.fr/>

UNIVERSITÉ PARIS-EST – MARNE-LA-VALLÉE,  
5, BOULEVARD DESCARTES, CHAMPS SUR MARNE,  
77454 MARNE-LA-VALLÉE, CEDEX 2 FRANCE.  
E-MAIL: [Mohamed.Hebiri@univ-mlv.fr](mailto:Mohamed.Hebiri@univ-mlv.fr)  
URL: <http://perso-math.univ-mlv.fr/users/hebiri.mohamed/>

CORNELL UNIVERSITY  
1188 COMSTOCK HALL  
ITHACA, NY 14853-2601  
E-MAIL: [johannesleder@cornell.edu](mailto:johannesleder@cornell.edu)  
URL: <http://www.johannesleder.de/>