

On the prediction performance of the Lasso

ARNAK S. DALALYAN¹, MOHAMED HEBIRI² and JOHANNES LEDERER³

¹ENSAE-CREST, 3 Avenue Pierre Larousse, 92240 Malakoff, France. E-mail: arnak.dalalyan@ensae.fr

²Université 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

³Cornell University, 1188 Comstock Hall, Ithaca, NY 14853-2601. E-mail: johannesleder@cornell.edu

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 can lead 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. We finally show that our results also lead to near-optimal rates for the least-squares estimator with total variation penalty.

Keywords: multiple linear regression; oracle inequalities; sparse recovery; total variation penalty

1. Introduction

In recent years, considerable effort has been devoted to establishing sharp theoretical guarantees for the prediction performance¹ of the Lasso [33]. Although there are already risk bounds for a variety of settings [7,25], 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 σ^{*2} . In vector notation, this reads

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

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 $j \in \{1, \dots, p\}$),

¹Following [5], the term prediction performance is understood here as the magnitude of the risk measured by the prediction loss $\frac{1}{n} \|\mathbf{X}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|_2^2$. This is not to be confused with the understanding of this term more common in machine learning literature, where the aim is to predict the label of a new unlabeled feature.

$\xi \in \mathbb{R}^n$ 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

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

that can be efficiently solved even for very large values of p and n [2,19]. 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. Note that, in particular for high dimensional models where $p > n$, there are typically multiple solutions of (2); however, since $\mathbf{X}\widehat{\boldsymbol{\beta}}_\lambda = \mathbf{X}\widehat{\boldsymbol{\beta}}'_\lambda$ for any two solution $\widehat{\boldsymbol{\beta}}_\lambda, \widehat{\boldsymbol{\beta}}'_\lambda$, all solutions have the same prediction performance.

In this paper, we study the prediction performance of the Lasso and make, in particular, the following five contributions:

1. Numerous empirical results indicate that the prediction error of the Lasso with the universal tuning parameter $\lambda = \sqrt{2 \log(p)/n}$ is at most proportional to $\frac{\log(p)}{n} \times \text{rank}(\mathbf{X})$. Equation (9) of this paper is the first theoretical confirmation of 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 suggest that the smallest prediction loss is obtained choosing a tuning parameter λ that is substantially smaller than the universal one. The influence of correlations on the prediction performance of the Lasso was first considered in [23,35], 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. One contribution of this paper is to advance and complement these results: First, we introduce a new measure for the geometry of the covariates. In contrast to earlier measures, the new measure is computable and allows for an exhaustive characterization of Lasso prediction as a function of the correlations. Second, Corollary 1 establishes that even with the universal tuning parameter (which does not incorporate the geometry of the covariates), Lasso can have fast rates of convergence in highly correlated settings.

These results relate Lasso prediction with the choice of the tuning parameter and, in particular, provide a description of the Lasso prediction performance for optimal tuning parameters. We expect, therefore, that our results could be complemented by the current research on tuning parameter calibration [3,14].

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 [9,16–18,30]. Fast rate bounds for Lasso prediction, in contrast, usually rely on assumptions on the correlations of the covariates such as low coherence [12], restricted eigenvalues [5,29], restricted isometry [11], compatibility [34], cone invertibility [39],

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 λ , that is, if we allow for λ that depend on the true regression vector β^* , the noise vector ξ , and the noise level σ^* . In the present work, we give a negative response to this question in Example 2.

4. 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 [5,8,10,24,36,38,42]. For the second case, the case of collinear covariates, this question has received much less attention. Therefore, this question is one of our main topics, and we give answers in Corollary 1 and Proposition 5.

5. We finally show that our new Lasso prediction guarantees also lead to optimal guarantees for prediction with total variation penalties. Total variation penalties can enforce similarities between neighboring pixels or between values of signals and are therefore popular for image denoising and signal processing. However, the known theoretical results for prediction with total variation penalties are fragmentary. Using completely new probabilistic approaches, we relate prediction with total variation penalties with our prediction bounds for the Lasso. This allows us to state in Propositions 3, 6, and 7 and in equation (16) near-optimal guarantees for prediction with total variation penalties for a large variety of settings.

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. Moreover, while we give some insights into computational aspects, this work is essentially a theoretical contribution.

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 ℓ_q -(quasi)norm of a vector $\mathbf{u} \in \mathbb{R}^k$, that is

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

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 Π_T the orthogonal projector onto V_T and by ρ_T the maximal Euclidean distance between the normalized columns of \mathbf{X} and the set V_T , that is,

$$\rho_T := \max_{j \in [p]} \min_{\mathbf{v} \in V_T} \|\mathbf{v} - n^{-1/2} \mathbf{x}^j\|_2 = n^{-1/2} \max_{j \in [p]} \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2. \tag{3}$$

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. Section 3 is devoted to some refinements of the sharp sparsity oracle inequalities with fast rates based on compatibility factors [32]. They imply, in particular, that the total variation estimator of piecewise constant signals is nearly rate optimal. We present in Section 4 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. ‘‘Slow’’ rates that involve the quantity ρ_T , accounting for the severity of the correlations within covariates, are developed in Section 5. In particular, they allow us to answer the second and the fourth questions raised in the Introduction. 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. Fast rates for Lasso projections

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 $\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}$ 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 λ is of the order of $n^{-1/2}$. This will be complemented in Section 5, 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 $(s^*/n)^{1/2}$.

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

$$\mathbf{y} = \mathbf{X}_T \boldsymbol{\beta}_T^* + \mathbf{X}_{T^c} \boldsymbol{\beta}_{T^c}^* + \boldsymbol{\xi} = \mathbf{X} \boldsymbol{\beta}^{*,T} + (\mathbf{I}_n - \Pi_T) \boldsymbol{\xi}. \tag{4}$$

When the rank of Π_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 inequalities² 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 let³ $v_T = \inf_{\mathbf{u} \in \mathbb{R}^{|T|}} \frac{\sqrt{|T|} \cdot \|\mathbf{X}_T \mathbf{u}\|_2}{\sqrt{n} \|\mathbf{u}\|_1}$. For every $\lambda > 0$, it holds that*

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

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

$$\frac{1}{\sqrt{n}} \|\Pi_T \mathbf{X} (\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}} - \boldsymbol{\beta}^*)\|_2 \leq \left(\frac{\lambda \sqrt{|T|}}{v_T} + \frac{\|\Pi_T \boldsymbol{\xi}\|_2}{\sqrt{n}} \right) \quad \forall T \subset [p]. \tag{5}$$

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 λ 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 $\sigma^{*2} |T| \log(p)/n$.

Relation (5) 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 $\lambda > 0$*

$$\ell_n(\widehat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*)^{1/2} \leq \frac{\lambda \sqrt{s}}{v_T} + \frac{\|\Pi_T \boldsymbol{\xi}\|_2}{\sqrt{n}}. \tag{6}$$

²We refer the reader to [8] for an introduction to sparsity oracle inequalities.

³It follows from the Cauchy–Schwarz inequality that v_T is not smaller than the smallest singular value of the matrix $\frac{1}{\sqrt{n}} \mathbf{X}_T$.

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

$$\mathbf{E}[\ell_n(\widehat{\beta}_\lambda^{\text{Lasso}}, \beta^*)] \leq \frac{2\lambda^2 s}{v_T^2} + \frac{2\sigma^{*2} s}{n}. \tag{7}$$

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

$$\ell_n(\widehat{\beta}_\lambda^{\text{Lasso}}, \beta^*) \leq \frac{2\lambda^2 s}{v_T^2} + \frac{4\sigma^{*2}(s + 2\log(1/\delta))}{n}. \tag{8}$$

The first two claims of this proposition trivially follow from (5), while the third claim follows from (6) using the fact that $\|\Pi_T \xi\|_2^2$ is drawn from the chi-squared distribution χ_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{v}_r the maximal value of v_T over all possible subsets of $[p]$ of cardinality $r = \text{rank}(\mathbf{X})$: $\bar{v}_r = \max_{T:|T|=r} v_T$. It follows from (8) that when ξ is Gaussian and $\lambda = \sqrt{2\log(p)/n}$, for every $\delta \in (0, 1)$, with probability $1 - \delta$,

$$\ell_n(\widehat{\beta}_\lambda^{\text{Lasso}}, \beta^*) \leq \frac{4\log(p)\text{rank}(\mathbf{X})}{n\bar{v}_r^2} + \frac{4\sigma^{*2}(\text{rank}(\mathbf{X}) + 2\log(1/\delta))}{n}. \tag{9}$$

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 [32]. We begin by stating their result⁴ 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 [32] and the idea of estimating $\beta^{*,T}$ instead of β^* 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 $\kappa_{T,\bar{c}} \geq 0$:

$$\kappa_{T,\bar{c}} = \inf_{\delta \in \mathbb{R}^p: \|\delta_{T^c}\|_1 < \bar{c}\|\delta_T\|_1} \frac{|T| \cdot \|\mathbf{X}\delta\|_2^2}{n(\|\delta_T\|_1 - \bar{c}^{-1}\|\delta_{T^c}\|_1)^2}. \tag{10}$$

Theorem 2 ([32], Theorem 4). *Let $\delta \in (0, 1)$ be a fixed tolerance level. If for some $\gamma > 1$, the tuning parameter of the Lasso satisfies $\lambda = \gamma\sigma^*(\frac{2}{n}\log(p/\delta))^{1/2}$, then with probability at least*

⁴The results stated below do not match exactly with those stated in [32], but they can be easily deduced from the proofs in [32].

$1 - \delta$,

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

This inequality is sharp in the sense 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 [26]. It was then refined and extended to the procedure square-root Lasso (also known as the scaled Lasso) in [32]. We also refer to [4], which appeared online between the online appearances of the two mentioned papers. [4] contains bounds that are not sharp in the above sense but, in contrast to [32], allow for more general design matrices that include, for example, replicas of predictors.

Let us state now some refinements of Theorem 2. For any subset T of $[p]$, let us introduce the weights⁵

$$\omega_j(T, \mathbf{X}) = \frac{1}{\sqrt{n}} \|\mathbf{I}_n - \Pi_T\| \mathbf{x}^j\|_2, \quad \bar{\omega}_j(T, \mathbf{X}) = \frac{\omega_j(T, \mathbf{X})}{\max_{\ell \in [p]} \omega_{\ell}(T, \mathbf{X})} \quad \forall j \in [p]. \quad (11)$$

Since \mathbf{x}^j are normalized to have an ℓ_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

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

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

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

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

The weighted compatibility factors with weights $\boldsymbol{\omega}$ and $\bar{\boldsymbol{\omega}}$ defined in (11) 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:

$$\ell_n(\widehat{\boldsymbol{\beta}}_{\lambda}^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq \inf_{\bar{\boldsymbol{\beta}} \in \mathbb{R}^p, T \subset [p]} \left\{ \ell_n(\bar{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) + 4\lambda \|\bar{\boldsymbol{\beta}}_{T^c}\|_1 + \frac{4\sigma^{*2} |T| \log(p/\delta)}{n} \cdot r_{n, p, T} \right\}, \quad (12)$$

⁵In 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

$$\ell_n(\hat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \leq \inf_{\boldsymbol{\beta}} \{ \ell_n(\bar{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) + 4\lambda \|\bar{\boldsymbol{\beta}}_{T^c}\|_1 \} + \frac{4\sigma^{*2} \rho_T^2 |T| \log(p/\delta)}{n} \cdot \bar{r}_{n,p,T}, \quad (13)$$

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 (12) and (13) is the presence of the factor ρ_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 weighted compatibility factor $\bar{\kappa}_{T,\gamma,\omega}$ by $\bar{\kappa}_{T,\gamma,\bar{\omega}}$ and to deflate λ by the factor ρ_T . From a practical point of view, this last modification is not always easy to implement, since the quantity ρ_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 (13) is to be interpreted as a theoretical justification for choosing the tuning parameter smaller than the universal value. Tuning parameter selection schemes, such as cross validation, might provide such a smaller value.

Example 1 (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 $\mathbf{f}^* = (f^*(1/n), \dots, f^*(1))^T$ 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

$$\hat{\mathbf{f}}^{\text{TV}} \in \arg \min_{\mathbf{f} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{f}\|_2^2 + \lambda \|\mathbf{f}\|_{\text{TV}} \right\}, \quad (14)$$

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 [21,22]. More precisely, if we define the vector of differences $\boldsymbol{\beta} \in \mathbb{R}^n$ by $\beta_j = f_j - f_{j-1}$, then (14) is equivalent to (2) with the $n \times n$ design matrix $\mathbf{X} = (\mathbb{1}(i \geq j))_{i,j}$. 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 $n \rightarrow \infty$, [28] establish that the ℓ_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, *that is*, the ℓ_0 -norm of the difference vector (see [6] 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 [21,22] 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 ℓ_2 -loss is of the optimal order n^{-1} . This result is, however, not very satisfactory since reducing λ down 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 λ 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 2. *Let \mathbf{X} be the $n \times n$ matrix with entries $x_{ij} = \mathbb{1}(i \geq j)$ 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 compatibility factors theoretically. All the previous efforts for evaluating compatibility and restricted eigenvalue constants, focused on weakly correlated designs (see, for instance, [5]). 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 2 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 3. *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 $T \subset [n]$:*

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \inf_{\bar{\mathbf{f}} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\bar{\mathbf{f}} - \mathbf{f}^*\|_2^2 + 4\lambda \|\bar{\mathbf{f}}_{T^c}\|_{\text{TV}} \right\} + \frac{4\sigma^{*2} |T| \log(n/\delta)}{n} \cdot r_{n, T}, \quad (15)$$

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

The risk bound (15) drastically improves the results on the ℓ_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 [28]. More precisely, when the true function \mathbf{f}^* is piecewise constant on a partition of s intervals, taking in (15) $\mathbf{f} = \mathbf{f}^*$ 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

$$\mathbf{P}\left(\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \frac{4\sigma^{*2}|J^*| \log(n/\delta)}{n} \cdot (3 + 256(\log(n) + (n/\Delta_{\min, J^*})))\right) \geq 1 - 2\delta. \quad (16)$$

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 (16) is of the nearly optimal order. Furthermore, risk bound (16) 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 [22]. To complete this discussion, let us mention that the constant 256 in (16) is definitely sub-optimal and it is out of scope of this work to look for the best possible constants.

4. Limits of fast rates: An example

In this section, we show that the prediction loss of the Lasso is in some cases at best of the order of $n^{-1/2}$, whatever the tuning parameter is. This example provides – to the best of our knowledge – the first proof that in some cases, the Lasso can not achieve fast rates even if the regression vector β^* has fixed length.

Example 2. 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 $\beta^* \in \mathbb{R}^{2m}$ such that $\beta_1^* = \beta_{m+1}^* = 1$ and $\beta_j^* = 0$ for every $j \in [2m] \setminus \{1, m + 1\}$.

Proposition 4. For any $\lambda > 0$, the prediction loss of the Lasso $\widehat{\beta}_\lambda^{\text{Lasso}}$ satisfies the inequality

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

There are at least three reasons that make this example particularly instructive. First, it shows that the correlations between the covariates need not to be close to ± 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 λ cannot prevent slow rates. Third, it is valid for a small value of sparsity index: the ℓ_0 -norm of β^* 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 [12]), 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 λ and a β^* independent of n .

This example clearly demonstrates the limits of the Lasso as a method of prediction. While for several other prediction procedures [16–18,30] 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.

5. “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 considerably extend and sharpen insights obtained in [23,35] 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 [28], 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 λ is not smaller than $\gamma\sigma^*\rho_T\sqrt{2\log(p/\delta)/n}$, the Lasso (2) fulfills*

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

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

The factor ρ_T (defined in equation (3)) in the lower bound for the tuning parameter λ makes this theorem particularly interesting. Slow rate bounds can be derived using the proofs in [32], 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*

$$\ell_n(\hat{\boldsymbol{\beta}}_\lambda^{\text{Lasso}}, \boldsymbol{\beta}^*) \lesssim \left(\sqrt{\frac{\log(p)}{n^{2r+1}}} \|\boldsymbol{\beta}^*\|_1 \right) \vee \frac{|T_n|}{n} \tag{17}$$

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 λ is chosen of order $\sqrt{\log(p)}/n$ (with sufficiently large constants).

Remark 1 (Effective number of parameters). The bound (17) can be further refined replacing the number of parameters p by an effective number of parameters as described in [23], 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 ρ_{J^*} , see equation (3)). Thus, Corollary 1 is the forth key contribution mentioned in the Introduction.

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 ρ_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*

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

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 ρ_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 [13] and an earlier risk bound can be found in [41].

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 (14) corresponds to the Lasso (2) with the design matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ with entries $x_{ij} = \mathbb{1}(i \geq j)$. 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 partition⁶ $\{\llbracket 1, h+1 \rrbracket, \llbracket h+1, 2h+1 \rrbracket, \dots, \llbracket (k-1)h+1, n+1 \rrbracket\}$ of $[n]$ with at most h points in the last interval and exactly h points in all other intervals. Moreover, Π_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 (14) 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 $1 - 2\delta$.

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

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 [41], equation (1.4). We conjecture that such logarithmic factors are always required in bounds that hold with high probability (note that the bounds in [41], in contrast, are in expectation). Second, the unknown quantities in the tuning parameter of Proposition 6 can be avoided or readily estimated: The noise variance σ^* can most likely be avoided using scaled versions of the Lasso [4,32]; the factor $\|\mathbf{f}^\uparrow\|_{\text{TV}}$ measuring the total variation of the function \mathbf{f}^\uparrow can be roughly estimated⁷ by $\max_{i,j}(y_i - y_j)$. Third, Proposition 6 allows for model misspecifications 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 ℓ_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 $L > 0$ [6]. In contrast, the best known rate for the TV-estimator over these sets is $n^{-2\alpha/3}$ and therefore clearly suboptimal [28]. 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 α and L . Then, for the tuning parameter $\lambda = \sigma^* \sqrt{\log(n/\delta)/(kn)}$, the TV-estimator (14) 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 [6] for the ℓ_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 misspecifications; the risk bounds in [6], in contrast, are purely asymptotic and do not take model misspecifications into account. A deficiency of Proposition 7, on the other hand, is the dependence of the tuning parameter on the constants α and L .

⁷One may also expect that choosing λ by cross validation or by minimizing an unbiased estimator of the risk would lead to bounds similar to that of Proposition 6. However, there are no theoretical result corroborating this claim.

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](#)⁸ of ρ_T , a 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, see [Example 2](#). 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 cannot 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 2](#), 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 [[27](#), [40](#)] 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 [[37](#)]. 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 [[4](#)], scaled Lasso [[31](#)] or scaled Dantzig selector [[15](#)]. We also believe that the geometry of the design, measured by the quantities ρ_T , may lead to better recommendations for the tuning parameter in the transductive setting [[1](#)]. 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 [[28](#)], but we believe that the results of the present work may lead to improved risk bounds.

7. Proofs

In this section, we gather the proofs of all the theorems and propositions stated in previous sections. For ease of notation, we write $\widehat{\beta}$ instead of $\widehat{\beta}_\lambda^{\text{Lasso}}$ throughout these proofs. In the sequel,

⁸A very similar quantity appears also in the risk bound provided by [Theorem 8.2](#) in [[25](#)], but the result therein does not suggest the incorporation of this quantity into the tuning parameter, and the established risk bounds are less accurate than those of the present work.

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

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

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

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

Subtracting the first relation from the second one, we get

$$\frac{1}{n} (\bar{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}})_T^\top \mathbf{X}_T^\top (\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}}) \leq \lambda (\|\bar{\boldsymbol{\beta}}_T\|_1 - \|\widehat{\boldsymbol{\beta}}_T\|_1). \quad (18)$$

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

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

This choice of $\bar{\boldsymbol{\beta}}$ may appear somewhat strange and complicated, but is made in order that the relation $\mathbf{X}_T (\bar{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}})_T = \Pi_T \mathbf{X} (\boldsymbol{\beta}^{*,T} - \widehat{\boldsymbol{\beta}})$ 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 (18), we find that for every $T \subset [p]$,

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

Equivalently, this relation may be written as

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

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

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

Combining the last two displays, we obtain

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

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

Proof of Theorem 3. Recall that according to (18), for every $J \subset [p]$ and $\bar{\boldsymbol{\beta}} \in \mathbb{R}^p$,

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

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

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

Let us introduce the two difference vectors $\boldsymbol{\delta} = \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$ and $\bar{\boldsymbol{\delta}} = \widehat{\boldsymbol{\beta}} - \bar{\boldsymbol{\beta}}$. 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

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

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 $J \subset [p]$,

$$\begin{aligned} \frac{\|\mathbf{X}_J \bar{\boldsymbol{\delta}}_J\|_2^2 + \|\mathbf{X}\boldsymbol{\delta}\|_2^2}{n} &\leq \frac{1}{n}\|\mathbf{X}_J \bar{\boldsymbol{\delta}}_J - \mathbf{X}\boldsymbol{\delta}\|_2^2 + \frac{2}{n}\bar{\boldsymbol{\delta}}_J^\top \mathbf{X}_J^\top (\mathbf{I}_n - \Pi_T)\boldsymbol{\xi} + \frac{2}{n}\bar{\boldsymbol{\delta}}_J^\top \mathbf{X}_J^\top \Pi_T \boldsymbol{\xi} \\ &\quad + 2\lambda(\|\bar{\boldsymbol{\beta}}_J\|_1 - \|\widehat{\boldsymbol{\beta}}_J\|_1). \end{aligned} \quad (19)$$

To prove (12), we choose $J = [p]$, for which (19) becomes

$$\begin{aligned} \frac{1}{n}\|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2 + \frac{1}{n}\|\mathbf{X}\boldsymbol{\delta}\|_2^2 &\leq \frac{1}{n}\|\mathbf{X}\bar{\boldsymbol{\delta}} - \mathbf{X}\boldsymbol{\delta}\|_2^2 + \frac{2}{n}\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \boldsymbol{\xi} + 2\lambda(\|\bar{\boldsymbol{\beta}}\|_1 - \|\widehat{\boldsymbol{\beta}}\|_1) \\ &\leq \frac{1}{n}\|\mathbf{X}\bar{\boldsymbol{\delta}} - \mathbf{X}\boldsymbol{\delta}\|_2^2 + \frac{2}{n}\bar{\boldsymbol{\delta}}_{T^c}^\top \mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T)\boldsymbol{\xi} + \frac{2}{n}\bar{\boldsymbol{\delta}}^\top \mathbf{X}^\top \Pi_T \boldsymbol{\xi} \\ &\quad + 2\lambda(\|\bar{\boldsymbol{\delta}}_T\|_1 - \|\bar{\boldsymbol{\delta}}_{T^c}\|_1) + 4\lambda\|\bar{\boldsymbol{\beta}}_{T^c}\|_1 \quad \forall T \subset [p]. \end{aligned} \quad (21)$$

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 λ by its value $\gamma\sigma^*(\frac{2}{n}\log(p/\delta))^{1/2}$ and restricting our attention to the event $\mathcal{B}_T = \{\max_{j \in T^c} |\boldsymbol{\xi}^\top (\mathbf{I}_n - \Pi_T)\mathbf{x}^j| / \|(\mathbf{I}_n - \Pi_T)\mathbf{x}^j\|_2 \leq \sigma^*(2\log(p/\delta))^{1/2}\} \cap \{\|\Pi_T \boldsymbol{\xi}\|_2 \leq \sigma^*(\sqrt{|T|} + \sqrt{2\log(1/\delta)})\}$ we obtain

$$\begin{aligned} \frac{1}{n}\|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2^2 + \frac{1}{n}\|\mathbf{X}\boldsymbol{\delta}\|_2^2 &\leq \frac{1}{n}\|\mathbf{X}\bar{\boldsymbol{\delta}} - \mathbf{X}\boldsymbol{\delta}\|_2^2 + 4\lambda\|\bar{\boldsymbol{\beta}}_{T^c}\|_1 + \frac{2}{n}\|\mathbf{X}\bar{\boldsymbol{\delta}}\|_2 \|\Pi_T \boldsymbol{\xi}\|_2 \\ &\quad + 2\lambda(\|\bar{\boldsymbol{\delta}}_T\|_1 - \|\bar{\boldsymbol{\delta}}_{T^c}\|_1 + \gamma^{-1}\|(\bar{\boldsymbol{\delta}} \odot \boldsymbol{\omega})_{T^c}\|_1). \end{aligned} \quad (22)$$

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

$$\begin{aligned} & \frac{2}{n} \|\mathbf{X}\bar{\delta}\|_2 \|\Pi_T \xi\|_2 + 2\lambda (\|\bar{\delta}_T\|_1 - \|\bar{\delta}_{T^c}\|_1 + \gamma^{-1} \|(\bar{\delta} \odot \omega)_{T^c}\|_1) \\ & \leq 2 \frac{\|\mathbf{X}\bar{\delta}\|_2}{\sqrt{n}} \left(\frac{1}{\sqrt{n}} \|\Pi_T \xi\|_2 + \lambda (|T|/\bar{\kappa}_{T,\gamma,\omega})^{1/2} \right) \\ & \leq \frac{\|\mathbf{X}\bar{\delta}\|_2^2}{n} + \left(\frac{1}{\sqrt{n}} \|\Pi_T \xi\|_2 + \lambda (|T|/\bar{\kappa}_{T,\gamma,\omega})^{1/2} \right)^2 \\ & \leq \frac{\|\mathbf{X}\bar{\delta}\|_2^2}{n} + \frac{4\sigma^{*2}(|T| + 2\log(1/\delta))}{n} + \frac{4\gamma^2\sigma^{*2}|T|\log(p/\delta)}{n\bar{\kappa}_{T,\gamma,\omega}}. \end{aligned} \quad (23)$$

After replacing (23) in (22), we remark that the terms $\frac{1}{n}\|\mathbf{X}\bar{\delta}\|_2^2$ cancel out and we get inequality (12). Classical results on the tails of Gaussian and χ^2 distributions imply that \mathcal{B}_T is at least of probability $1 - 2\delta$, for every T . The assertion of (12) follows by choosing T to be a subset of $[p]$ minimizing the right-hand side of (12).

The proof of the second claim of the theorem is identical to that of (12) and, therefore, is left to the reader. \square

Proof of Theorem 4. According to (20), for every $\bar{\beta} \in \mathbb{R}^p$ and for $\delta = \hat{\beta} - \beta^*$ and $\bar{\delta} = \hat{\beta} - \bar{\beta}$, the inequality

$$\frac{1}{n} \|\mathbf{X}\bar{\delta}\|_2^2 + \frac{1}{n} \|\mathbf{X}\delta\|_2^2 \leq \frac{1}{n} \|\mathbf{X}\bar{\delta} - \mathbf{X}\delta\|_2^2 + \frac{2}{n} \bar{\delta}^\top \mathbf{X}^\top \xi + 2\lambda (\|\bar{\beta}\|_1 - \|\hat{\beta}\|_1) \quad (24)$$

holds for every $\lambda > 0$. We split the stochastic term $\bar{\delta}^\top \mathbf{X}^\top \xi$ into two terms $\bar{\delta}^\top \mathbf{X}^\top (\mathbf{I}_n - \Pi_T)\xi$ and $\bar{\delta}^\top \mathbf{X}^\top \Pi_T \xi$. The first one can be bounded using the duality inequality $\bar{\delta}^\top \mathbf{X}^\top (\mathbf{I}_n - \Pi_T)\xi = \bar{\delta}_{T^c}^\top \mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T)\xi \leq \|\bar{\delta}_{T^c}\|_1 \|\mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T)\xi\|_\infty \leq (\|\bar{\beta}_{T^c}\|_1 + \|\hat{\beta}_{T^c}\|_1) \|\mathbf{X}_{T^c}^\top (\mathbf{I}_n - \Pi_T)\xi\|_\infty$, while the second one, in view of the Cauchy–Schwarz inequality, satisfies $\bar{\delta}^\top \mathbf{X}^\top \Pi_T \xi \leq \|\mathbf{X}\bar{\delta}\|_2 \|\Pi_T \xi\|_2 \leq (\|\mathbf{X}\bar{\delta}\|_2^2 + \|\Pi_T \xi\|_2^2)/2$. As in the proof of fast rates, we restrict our attention to the event $\mathcal{B}_T = \{\max_{j \in T^c} |\xi^\top (\mathbf{I}_n - \Pi_T)\mathbf{x}^j| / \|(\mathbf{I}_n - \Pi_T)\mathbf{x}^j\|_2 \leq \sigma^*(2\log(p/\delta))^{1/2}\} \cap \{\|\Pi_T \xi\|_2 \leq \sigma^*(\sqrt{|T|} + \sqrt{2\log(1/\delta)})\}$. On this event, we get

$$2\bar{\delta}^\top \mathbf{X}^\top \xi \leq \|\mathbf{X}\bar{\delta}\|_2^2 + 2\sigma^{*2}(|T| + 2\log(1/\delta)) + 2\lambda\gamma^{-1}n(\|\bar{\beta}_{T^c}\|_1 + \|\hat{\beta}_{T^c}\|_1). \quad (25)$$

Combining this inequality with (24), we get the desired result. \square

Proof of Proposition 4. We first recall that given data (\mathbf{y}, \mathbf{X}) and tuning parameter λ , all solutions (2) provide the same prediction, which means that it is sufficient to consider only one of the solutions. For completeness, let us briefly derive this fact. We do this by contradiction,

considering two solutions $\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\beta}}'$

$$\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\beta}}' \in \arg \min_{\boldsymbol{\beta}} g(\boldsymbol{\beta}), \quad g(\boldsymbol{\beta}) := \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1,$$

and assuming that these solutions lead to different predictions, that is, $\mathbf{X}\widehat{\boldsymbol{\beta}} \neq \mathbf{X}\widehat{\boldsymbol{\beta}}'$. Fix now a number $u \in (0, 1)$ and introduce the linear combination $\widehat{\boldsymbol{\beta}}'' := u\widehat{\boldsymbol{\beta}} + (1-u)\widehat{\boldsymbol{\beta}}'$. The definition of $\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\beta}}'$ as minimizers of the objective function g implies $g(\widehat{\boldsymbol{\beta}}) = g(\widehat{\boldsymbol{\beta}}') \leq g(\widehat{\boldsymbol{\beta}}'')$. On the other hand, we can deduce the strict inequality $g(\widehat{\boldsymbol{\beta}}'') < ug(\widehat{\boldsymbol{\beta}}) + (1-u)g(\widehat{\boldsymbol{\beta}}')$ from the strict convexity of the function $\boldsymbol{y} \mapsto \|\mathbf{y} - \boldsymbol{y}\|_2^2$. Combining these findings leads to $g(\widehat{\boldsymbol{\beta}}'') < ug(\widehat{\boldsymbol{\beta}}'') + (1-u)g(\widehat{\boldsymbol{\beta}}'') = g(\widehat{\boldsymbol{\beta}}'')$. This is a contradiction, so that we can conclude that for given (\mathbf{y}, \mathbf{X}) and λ , all Lasso solutions provide the same prediction.

We can now look at specific solutions to (2) for Example 2. To this end, we note that the Karush–Kuhn–Tucker conditions for Example 2 are the following:

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

After simplification, we get

$$1 + \frac{1}{\sqrt{2n}}(\xi_1 + \xi_{j+1}) - \frac{1}{2} \sum_{k=1}^{2m} \widehat{\beta}_k - \frac{1}{2}(\widehat{\beta}_j - \widehat{\beta}_{m+j}) \in \lambda \operatorname{sgn}(\widehat{\beta}_j), \quad (26)$$

$$1 + \frac{1}{\sqrt{2n}}(\xi_1 - \xi_{j+1}) - \frac{1}{2} \sum_{k=1}^{2m} \widehat{\beta}_k + \frac{1}{2}(\widehat{\beta}_j - \widehat{\beta}_{m+j}) \in \lambda \operatorname{sgn}(\widehat{\beta}_{m+j}). \quad (27)$$

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 $\widehat{\boldsymbol{\beta}} = \mathbf{0}$ is a solution to the system (26)–(27). Therefore, $\widehat{\boldsymbol{\beta}} = \mathbf{0}$ and hence $\ell_n(\widehat{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) = 2$. Thus, in the case $\lambda \geq 1$ the Lasso is not consistent, which is not a surprise since the theory recommends to always choose λ 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

$$\widehat{\beta}_j = \begin{cases} \frac{2(1-\lambda)}{m+1}, & \xi_{j+1} > 0, \\ 0, & \xi_{j+1} < 0, \end{cases} \quad \widehat{\beta}_{m+j} = \begin{cases} 0, & \xi_{j+1} > 0, \\ \frac{2(1-\lambda)}{m+1}, & \xi_{j+1} < 0. \end{cases} \quad (28)$$

Indeed, for instance if $\xi_{j+1} > 0$, replacing these values of $\widehat{\beta}$ in (26) and (27) we get on the event \mathcal{B} :

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

where the last inclusion follows from the relation $m + 1 \geq \sqrt{2n}$. For the vector (28), we check that the prediction loss

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

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

$$\begin{aligned}
 \widehat{\beta}_j &= \begin{cases} \frac{\sqrt{2n} + m - 1}{m\sqrt{2n}} - \lambda, & \xi_{j+1} > 0, \\ \frac{\sqrt{2n} - m - 1}{m\sqrt{2n}} + \lambda, & \xi_{j+1} < 0, \end{cases} \\
 \widehat{\beta}_{m+j} &= \begin{cases} \frac{\sqrt{2n} - m - 1}{m\sqrt{2n}} + \lambda, & \xi_{j+1} > 0, \\ \frac{\sqrt{2n} + m - 1}{m\sqrt{2n}} - \lambda, & \xi_{j+1} < 0, \end{cases}
 \end{aligned} \tag{29}$$

for every $j \in [m]$. Indeed, for instance if $\xi_{j+1} > 0$, replacing these values of $\widehat{\beta}$ in (26) and (27) we get on the event \mathcal{B} :

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

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

The prediction loss of this estimator is

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

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 2. 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 that⁹

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

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

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

⁹For notational convenience, we assume hereafter that T is augmented by the elements $\{1, n + 1\}$.

where

$$\begin{aligned} \Psi_\ell &= \min_{\substack{\mathbf{t} \in \{\pm 1\}^{n+1} \\ (t_{j\ell-1}, t_{j\ell}) = (s_{j\ell-1}, s_{j\ell})}} \sum_{j=j\ell-1}^{j\ell-1} f_j(t_j a_j - t_{j+1} a_{j+1}) \\ &= \min_{\substack{\bar{\mathbf{t}} \in \{\pm 1\}^{\Delta_\ell+1} \\ (\bar{t}_1, \bar{t}_{\Delta_\ell+1}) = (s_{j\ell-1}, s_{j\ell})}} \sum_{j=1}^{\Delta_\ell} f_{j\ell-1+j-1}(\bar{t}_j a_{j\ell-1+j-1} - \bar{t}_{j+1} a_{j\ell-1+j}). \end{aligned}$$

In what follows, we propose a bound on Ψ_1 . The other Ψ_ℓ 's can be evaluated similarly. Let $\varepsilon_1, \dots, \varepsilon_n$ be independent random variables with values in ± 1 such that $p := \mathbf{P}(\varepsilon_j = 1) = (1 - (2\Delta_1)^{-1})$ and $\mathbf{P}(\varepsilon_j = -1) = (2\Delta_1)^{-1}$. Further, we define $\bar{\varepsilon}_1 = s_1$ and $\bar{\varepsilon}_{j+1} = \bar{\varepsilon}_1 \times \varepsilon_1 \times \dots \times \varepsilon_j$ for every $j \in B_1$. Since ε_j 's are independent, $\{\bar{\varepsilon}_j\}_{j \in [j_1]}$ is a Markov chain with values in $\{\pm 1\}$. For this Markov chain, we first check that

$$\mathbf{P}(\bar{\varepsilon}_{j_1} = s_{j_1}) \geq 1/4. \tag{30}$$

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 $\{\mathbf{e}_{B_1} = \mathbf{1}_{B_1}\} \subset \{\bar{\varepsilon}_{j_1} = s_{j_1}\}$ to infer that

$$\mathbf{P}(\bar{\varepsilon}_{j_1} = s_{j_1}) \geq \mathbf{P}(\mathbf{e}_{B_1} = \mathbf{1}_{B_1}) = \prod_{j=1}^{j_1-1} \mathbf{P}(\varepsilon_j = 1) = (1 - (2\Delta_1)^{-1})^{\Delta_1}.$$

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

$$\mathbf{P}(\bar{\varepsilon}_{j_1} = s_{j_1}) \geq \sum_{j=1}^{\Delta_1} \mathbf{P}(\mathbf{e}_{-j} = \mathbf{1}, \varepsilon_j = -1) = \sum_{j=1}^{\Delta_1} (1 - (2\Delta_1)^{-1})^{\Delta_1-1} \times (2\Delta_1)^{-1} \geq \frac{1}{4}.$$

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

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

We need now to evaluate the expectation of the random variable $\Upsilon = |\sum_{j=1}^{\Delta_1} f_j \bar{\varepsilon}_j (a_j - \varepsilon_{j+1} a_{j+1})|$. To this end, since all the a_j are nonnegative, we remark that $\mathbf{E}(|a_j - \varepsilon_{j+1} a_{j+1}|) = |a_j - a_{j+1}| + (a_j \wedge a_{j+1})/\Delta_1$ 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 $A = 4(\sum_{j=1}^{\Delta_1} f_j^2)^{1/2}(\sum_{j=1}^{\Delta_1} (2|a_j - a_{j+1}|^2 + 2a_j^2 \Delta_1^{-2}))^{1/2}$, we get

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

Combined with inequality (30), this entails that

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

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. □

Proof of Proposition 3. We apply Theorem 3 with $\gamma = 2$. This leads to (12) with

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

It remains to find a lower bound for $\bar{\kappa}_{T,2,\omega}$. To this end, we resort to Proposition 2 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 $B_\ell = \llbracket j_{\ell-1}, j_\ell \llbracket$ 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 $\mathbf{x}^j = [\mathbb{1}(i \geq j)]_{i \in [n]}$, the projector Π_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 2 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 $\delta = \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$. We consider two cases separately. The first case is when the inequality $\|\widehat{\boldsymbol{\beta}}\|_1 \leq \frac{\gamma+1}{\gamma-1} \|\boldsymbol{\beta}^*\|_1$ is satisfied. Then, using the Parseval identity, we have $\ell_n(\widehat{\boldsymbol{\beta}}, \boldsymbol{\beta}^*) = \frac{1}{n} \|\Pi_T \mathbf{X} \delta\|_2^2 + \frac{1}{n} \|(\mathbf{I}_n - \Pi_T) \mathbf{X} \delta\|_2^2$. The first summand can be bounded using inequality (5), so we focus on the second summand. Using the triangle inequality, we get $\|(\mathbf{I}_n - \Pi_T) \mathbf{X} \delta\|_2 = \|\sum_j (\mathbf{I}_n - \Pi_T) \mathbf{x}^j \delta_j\|_2 \leq \sum_j \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j \delta_j\|_2 \leq \|\delta\|_1 \max_{j \in T^c} \|(\mathbf{I}_n - \Pi_T) \mathbf{x}^j\|_2$. Furthermore, we have $\|\delta\|_1 \leq \|\widehat{\boldsymbol{\beta}}\|_1 + \|\boldsymbol{\beta}^*\|_1 \leq \frac{2\gamma}{\gamma-1} \|\boldsymbol{\beta}^*\|_1$. Hence, putting these bounds together, we find

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

with a probability at least $1 - \delta$. In the second case, $\|\widehat{\boldsymbol{\beta}}\|_1 > \frac{\gamma+1}{\gamma-1} \|\boldsymbol{\beta}^*\|_1$, according to inequalities (24) and (25) applied to $\bar{\boldsymbol{\delta}} = \boldsymbol{\delta}$, we have

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

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

$$\frac{1}{n} \|\widehat{\mathbf{f}}^{\text{TV}} - \mathbf{f}^*\|_2^2 \leq \inf_{\bar{\mathbf{f}} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\bar{\mathbf{f}} - \mathbf{f}^*\|_2^2 + 4\lambda \|\bar{\mathbf{f}}\|_{\text{TV}} \right\} + \frac{2\sigma^{*2}(k + 2\log(1/\delta))}{n} \quad (31)$$

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

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

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

$$\|\bar{\mathbf{f}} - \mathbf{f}^*\|_2^2 \leq \|\mathbf{g}^* - \mathbf{f}^*\|_2^2 + \|\bar{\mathbf{f}} - \mathbf{g}^*\|_2^2. \quad (32)$$

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 (32). In conjunction

with (31), 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

Appendix: Computation of the compatibility factors by sequential convex programming

The contribution of the present paper is theoretical. However, 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 [24].) 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 $\bar{c}, \varepsilon > 0$, one can find an interval $[\kappa_*, \kappa^*]$ of length at most ε and containing $\bar{\kappa}_{T, \bar{c}}$ by solving at most $2^{|T|} \log_2(|T|/\varepsilon)$ 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 $\log_2(|T|/\varepsilon)$ SOCPs. Note that an alternative approach proposed in [20] is to lower-bound the compatibility factor by a quantity that can be efficiently computed by linear programming.

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

$$\begin{aligned} & \text{minimize} && \|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu \|\boldsymbol{\delta}_{T^c}\|_1 \\ & \text{subject to} && \|\boldsymbol{\delta}_T\|_1 = 1. \end{aligned} \tag{33}$$

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 κ_* unchanged and decrease κ^* 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 κ .) If, in contrast, the solution $\boldsymbol{\delta}^\mu$ of problem (33) satisfies $\|\mathbf{X}\boldsymbol{\delta}^\mu\|_2 + \mu \|\boldsymbol{\delta}_{T^c}^\mu\|_1 > \mu \bar{c}$, then we leave κ^* unchanged and increase κ_* by setting $\kappa_* = \kappa$. We iterate this process until we get $\kappa^* - \kappa_* \leq \varepsilon$. 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 ε is not larger than $\log_2(|T|/\varepsilon)$.

Let us now analyze the complexity of the optimization problem (33). The objective function of this problem is convex but the set of feasible solutions is not. Interestingly, for every $\mathbf{s} \in \{\pm 1\}^{|T|}$,

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 $\varepsilon > 0$

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

```

1:  $\kappa_* \leftarrow 0$  and  $\kappa^* \leftarrow |T|$ 
2: while  $\kappa^* - \kappa_* > \varepsilon$  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

```

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 (33) by $2^{|T|}$ optimization problems

$$\begin{aligned} & \text{minimize} \quad \|\mathbf{X}\boldsymbol{\delta}\|_2 + \mu\|\boldsymbol{\delta}_{T^c}\|_1 \\ & \text{subject to} \quad \mathbf{s}^\top \boldsymbol{\delta}_T = 1, \text{ and } s_j \delta_j \geq 0, \forall j \in T. \end{aligned} \tag{34}$$

For every $\mathbf{s} \in \{\pm 1\}^{|T|}$, (34) can be rewritten as a SOCP using standard arguments. Denoting by $\boldsymbol{\delta}^{\mu, \mathbf{s}}$ any solution of (34), we determine a solution of (33) 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

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

lead to the procedure summarized in Algorithm 1.

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

- [1] Alquier, P. and Hebiri, M. (2012). Transductive versions of the LASSO and the Dantzig selector. *J. Statist. Plann. Inference* **142** 2485–2500. [MR2922000](#)
- [2] Bach, F.R., Jenatton, R., Mairal, J. and Obozinski, G. (2012). Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning* **4** 1–106.
- [3] Belloni, A. and Chernozhukov, V. (2011). ℓ_1 -penalized quantile regression in high-dimensional sparse models. *Ann. Statist.* **39** 82–130. [MR2797841](#)
- [4] Belloni, A., Chernozhukov, V. and Wang, L. (2014). Pivotal estimation via square-root Lasso in non-parametric regression. *Ann. Statist.* **42** 757–788. [MR3210986](#)
- [5] Bickel, P.J., Ritov, Y. and Tsybakov, A.B. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. [MR2533469](#)
- [6] Boysen, L., Kempe, A., Liebscher, V., Munk, A. and Wittich, O. (2009). Consistencies and rates of convergence of jump-penalized least squares estimators. *Ann. Statist.* **37** 157–183. [MR2488348](#)
- [7] Bühlmann, P. and van de Geer, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Series in Statistics. Heidelberg: Springer. [MR2807761](#)
- [8] Bunea, F., Tsybakov, A. and Wegkamp, M. (2007). Sparsity oracle inequalities for the Lasso. *Electron. J. Stat.* **1** 169–194. [MR2312149](#)
- [9] Bunea, F., Tsybakov, A.B. and Wegkamp, M.H. (2007). Aggregation for Gaussian regression. *Ann. Statist.* **35** 1674–1697. [MR2351101](#)
- [10] Cai, T.T., Wang, L. and Xu, G. (2010). Shifting inequality and recovery of sparse signals. *IEEE Trans. Signal Process.* **58** 1300–1308. [MR2730209](#)
- [11] Candès, E. and Tao, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.* **35** 2313–2351. [MR2382644](#)
- [12] Candès, E.J. and Plan, Y. (2009). Near-ideal model selection by ℓ_1 minimization. *Ann. Statist.* **37** 2145–2177. [MR2543688](#)
- [13] Chatterjee, S., Guntuboyina, A. and Sen, B. (2013). Improved risk bounds in isotonic regression. Technical report. Available at [arxiv:1311.3765](#).
- [14] Chichignoud, M., Lederer, J. and Wainwright, M. (2014). Tuning Lasso for sup-norm optimality. Preprint. Available at [arXiv:1410.0247](#).
- [15] Dalalyan, A.S. and Chen, Y. (2012). Fused sparsity and robust estimation for linear models with unknown variance. In *Advances in Neural Information Processing Systems 25: NIPS* 1268–1276. Red Hook, NY.
- [16] Dalalyan, A.S. and Tsybakov, A.B. (2007). Aggregation by exponential weighting and sharp oracle inequalities. In *Learning Theory. Lecture Notes in Computer Science* **4539** 97–111. Springer, Berlin. [MR2397581](#)
- [17] Dalalyan, A.S. and Tsybakov, A.B. (2012). Mirror averaging with sparsity priors. *Bernoulli* **18** 914–944. [MR2948907](#)
- [18] Dalalyan, A.S. and Tsybakov, A.B. (2012). Sparse regression learning by aggregation and Langevin Monte-Carlo. *J. Comput. System Sci.* **78** 1423–1443. [MR2926142](#)
- [19] Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. (2004). Least angle regression. *Ann. Statist.* **32** 407–499. [MR2060166](#)
- [20] Gautier, E. and Tsybakov, A.E. (2011). High-dimensional instrumental variables regression and confidence sets. Technical report. Available at [arXiv:1105.2454](#). [MR2867623](#)
- [21] Harchaoui, Z. and Lévy-Leduc, C. (2007). Catching change-points with lasso. In *Annual Conference on Neural Information Processing Systems*.
- [22] Harchaoui, Z. and Lévy-Leduc, C. (2010). Multiple change-point estimation with a total variation penalty. *J. Amer. Statist. Assoc.* **105** 1480–1493. [MR2796565](#)

- [23] Hebiri, M. and Lederer, J. (2013). How correlations influence Lasso prediction. *IEEE Trans. Inform. Theory* **59** 1846–1854. [MR3030757](#)
- [24] Juditsky, A. and Nemirovski, A. (2011). Accuracy guarantees for ℓ_1 -recovery. *IEEE Trans. Inform. Theory* **57** 7818–7839. [MR2895363](#)
- [25] Koltchinskii, V. (2011). *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems. Lecture Notes in Math.* **2033**. Heidelberg: Springer. [MR2829871](#)
- [26] Koltchinskii, V., Lounici, K. and Tsybakov, A.B. (2011). Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.* **39** 2302–2329. [MR2906869](#)
- [27] Lounici, K., Pontil, M., van de Geer, S. and Tsybakov, A.B. (2011). Oracle inequalities and optimal inference under group sparsity. *Ann. Statist.* **39** 2164–2204. [MR2893865](#)
- [28] Mammen, E. and van de Geer, S. (1997). Locally adaptive regression splines. *Ann. Statist.* **25** 387–413. [MR1429931](#)
- [29] Raskutti, G., Wainwright, M.J. and Yu, B. (2010). Restricted eigenvalue properties for correlated Gaussian designs. *J. Mach. Learn. Res.* **11** 2241–2259. [MR2719855](#)
- [30] Rigollet, P. and Tsybakov, A. (2011). Exponential screening and optimal rates of sparse estimation. *Ann. Statist.* **39** 731–771. [MR2816337](#)
- [31] Städler, N., Bühlmann, P. and van de Geer, S. (2010). ℓ_1 -penalization for mixture regression models. *TEST* **19** 209–256. [MR2677722](#)
- [32] Sun, T. and Zhang, C.-H. (2012). Scaled sparse linear regression. *Biometrika* **99** 879–898. [MR2999166](#)
- [33] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **58** 267–288. [MR1379242](#)
- [34] van de Geer, S. (2007). The deterministic lasso. In *Proc. of Joint Statistical Meeting*. Amer. Statist. Assoc.
- [35] van de Geer, S. and Lederer, J. (2013). The Lasso, correlated design, and improved oracle inequalities. In *From Probability to Statistics and Back: High-Dimensional Models and Processes. Inst. Math. Stat. (IMS) Collect.* **9** 303–316. IMS, Beachwood, OH. [MR3202642](#)
- [36] van de Geer, S.A. and Bühlmann, P. (2009). On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.* **3** 1360–1392. [MR2576316](#)
- [37] Vert, J.-P. and Bleakley, K. (2010). Fast detection of multiple change-points shared by many signals using group lars. In *Advances in Neural Information Processing Systems* 2343–2351.
- [38] Wainwright, M.J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans. Inform. Theory* **55** 2183–2202. [MR2729873](#)
- [39] Ye, F. and Zhang, C.-H. (2010). Rate minimaxity of the Lasso and Dantzig selector for the ℓ_q loss in ℓ_r balls. *J. Mach. Learn. Res.* **11** 3519–3540. [MR2756192](#)
- [40] Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 49–67. [MR2212574](#)
- [41] Zhang, C.-H. (2002). Risk bounds in isotonic regression. *Ann. Statist.* **30** 528–555. [MR1902898](#)
- [42] Zhang, T. (2009). Some sharp performance bounds for least squares regression with L_1 regularization. *Ann. Statist.* **37** 2109–2144. [MR2543687](#)

Received October 2014 and revised July 2015