

Asymptotic Equivalence of Regularization Methods in Thresholded Parameter Space

Yingying FAN and Jinchi LV

High-dimensional data analysis has motivated a spectrum of regularization methods for variable selection and sparse modeling, with two popular methods being convex and concave ones. A long debate has taken place on whether one class dominates the other, an important question both in theory and to practitioners. In this article, we characterize the asymptotic equivalence of regularization methods, with general penalty functions, in a thresholded parameter space under the generalized linear model setting, where the dimensionality can grow exponentially with the sample size. To assess their performance, we establish the oracle inequalities—as in Bickel, Ritov, and Tsybakov (2009)—of the global minimizer for these methods under various prediction and variable selection losses. These results reveal an interesting phase transition phenomenon. For polynomially growing dimensionality, the L_1 -regularization method of Lasso and concave methods are asymptotically equivalent, having the same convergence rates in the oracle inequalities. For exponentially growing dimensionality, concave methods are asymptotically equivalent but have faster convergence rates than the Lasso. We also establish a stronger property of the oracle risk inequalities of the regularization methods, as well as the sampling properties of computable solutions. Our new theoretical results are illustrated and justified by simulation and real data examples.

KEY WORDS: General penalty functions; Global minimizer; High-dimensional prediction and variable selection.

1. INTRODUCTION

Among all efforts on high-dimensional inference in the last decade, regularization methods have received much attention due to their ability to simultaneously conduct variable selection and estimation. The idea of regularization is to add a penalty term on model complexity to some model fitting loss measure. Then minimizing the penalized model fitting loss measure yields an estimate of the model parameters. Various penalty functions have been proposed in the literature. Broadly speaking, they can be classified into two classes: convex ones and concave ones. The former class is most popularly represented by the Lasso with the L_1 -penalty (Tibshirani 1996), and the latter class includes the smoothly clipped absolute deviation (SCAD; Fan and Li 2001), minimax concave penalty (MCP; Zhang 2010), and smooth integration of counting and absolute deviation (SICA; Lv and Fan 2009), among others.

There has been a long debate on which class of regularization methods one should use. Convex regularization methods enjoy nice computational properties and can be efficiently implemented with algorithms such as the least-angle regression (LARS; Efron et al. 2004) and coordinate optimization (Friedman et al. 2007; Wu and Lange 2008). On the theoretical side, Zhao and Yu (2006) introduced the irrepresentable conditions to characterize the model selection consistency of Lasso. See also, for example, Donoho, Elad, and Temlyakov (2006); Bunea, Tsybakov, and Wegkamp (2007); van de Geer (2008); and Bickel, Ritov, and Tsybakov (2009) for the properties of the L_1 -regularization method of Lasso. Despite its appealing properties, the Lasso suffers from an intrinsic bias issue (Fan

and Li 2001; Zou 2006; Zhang and Huang 2008). The irrepresentable conditions ensuring the model selection consistency of Lasso become stringent in high dimensions due to increased collinearity among predictors (Lv and Fan 2009; Fan and Lv 2011).

On the other hand, concave regularization methods, initiated in Fan and Li (2001), ameliorate the bias issue of Lasso and enjoy the model selection consistency property under much weaker conditions. Fan and Li (2001) proposed nonconcave penalized likelihood methods including the use of the SCAD penalty and established their oracle properties in the finite-dimensional setting. Their results were later extended by Fan and Peng (2004) to the moderate-dimensional setting with $p = o(n^{1/5})$ or $o(n^{1/3})$, where p is the dimensionality and n is the sample size. Recently, Lv and Fan (2009) established the weak oracle properties for regularization methods with general concave penalties in linear regression model, where p is allowed to grow exponentially with sample size n . Fan and Lv (2011) extended these results to generalized linear models (GLMs) and further proved the oracle properties of nonconcave penalized likelihood estimators. Despite all these theoretical developments, most existing studies on nonconvex regularization methods have focused on some appealing local minimizers. The global properties of these methods are still largely unknown and the theoretical characterizations of the global minimizers pose challenges.

The aforementioned advantages and potential issues of the two classes of regularization methods make it difficult for practitioners to decide which one to use. Understanding the connections and differences between different regularization methods is important both theoretically and empirically. An important question that has long puzzled researchers is: What are the connections and differences of all regularization methods? We intend to provide some answer to this question in this article. To

Yingying Fan is Assistant Professor (E-mail: fanyingy@marshall.usc.edu) and Jinchi Lv is Assistant Professor (E-mail: jinchilv@marshall.usc.edu), Data Sciences and Operations Department, Marshall School of Business, University of Southern California, Los Angeles, CA 90089. Fan's research was supported by NSF (National Science Foundation) CAREER Award DMS-1150318 and Grant DMS-0906784, and University of Southern California Marshall summer research funding. Lv's research was supported by NSF CAREER Award DMS-0955316 and Grant DMS-0806030. The authors sincerely thank the Co-Editor, Associate Editor, and two referees for their valuable comments that helped improve the article substantially.

© 2013 American Statistical Association
Journal of the American Statistical Association
September 2013, Vol. 108, No. 503, Theory and Methods
DOI: 10.1080/01621459.2013.803972

characterize the performance of different regularization methods, we establish the oracle inequalities and a stronger property of oracle risk inequalities of the global minimizer for regularization methods with general penalty functions, including both convex and concave ones.

The oracle inequalities have been frequently exploited to provide theoretical insights into high-dimensional inference methods and show how closely a sparse modeling method can mimic the oracle procedure. For example, Candes and Tao (2007) proved the oracle inequalities for the Dantzig selector, showing that the resulting estimator can achieve a loss within a logarithmic factor of the dimensionality for the oracle estimator. In a seminal article, Bickel, Ritov, and Tsybakov (2009) established the oracle inequalities simultaneously for two well-known L_1 -regularization methods, the Lasso and Dantzig selector. These oracle inequalities show that the two methods are asymptotically equivalent under certain regularity conditions. Extensive results on the oracle inequalities for general regularization methods were obtained in Antoniadis and Fan (2001) for the wavelets setting.

Our theoretical analysis reveals the asymptotic equivalence of regularization methods in a thresholded parameter space, in the sense of having the same convergence rates in the oracle inequalities and oracle risk inequalities. The introduction of the thresholded parameter space is motivated by the goal of distinguishing between important predictors and noise predictors in variable selection. The new results on oracle inequalities are parallel to those in Bickel, Ritov, and Tsybakov (2009) for the Lasso, but with improved sparsity bound. Our results on the oracle risk inequalities are stronger theoretical developments than those on the oracle inequalities. Specifically, in the case of polynomially growing dimensionality p , all regularization methods under consideration including the Lasso and concave ones have the same convergence rates, within a factor of $\log n$ of the oracle rates, in the oracle inequalities and oracle risk inequalities, leading to their asymptotic equivalence. In the case of exponentially growing dimensionality p , all concave regularization methods under consideration have the same convergence rates as in the previous case for both oracle inequalities and oracle risk inequalities, but the rates are faster than those of the Lasso, which are within a factor of $\log p$ of the oracle rates.

The connections and differences between the two classes of regularization methods revealed by our study provide an interesting phase transition of how different regularization methods perform as the dimensionality grows with the sample size. To the best of our knowledge, the results and phase transition phenomenon shown in this article are new to the literature. In addition, our theoretical results are for the global minimizers of the regularization methods, which is different from most studies in the literature.

The rest of the article is organized as follows. Section 2 introduces the regularization methods in the thresholded parameter space. We present the sampling properties of the concave regularization methods in a thresholded parameter space in ultrahigh-dimensional GLMs, as well as the sampling properties of computable solutions, in Section 3. We discuss the implementation of the methods and present several simulation and real data examples in Section 4. Section 5 provides some

discussions of our results and their implications. All technical details are relegated to Appendices A and B.

2. REGULARIZATION METHODS IN THRESHOLDED PARAMETER SPACE

Let $(\mathbf{x}_i, y_i)_{i=1}^n$ be a sample of n independent observations from (\mathbf{x}, Y) in the GLM linking a p -dimensional predictor vector \mathbf{x} to a scalar response variable Y . The GLM assumes that with a canonical link, the conditional distribution of Y given the predictor vector \mathbf{x} belongs to the exponential family, with a density function taking the form

$$f(y; \theta, \phi) = \exp\{y\theta - b(\theta) + c(y, \phi)\}, \quad (1)$$

where $\theta = \mathbf{x}^T \boldsymbol{\beta}$ with $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T \in \mathbb{R}^p$ a regression coefficient vector, $b(\cdot)$ and $c(\cdot, \cdot)$ are some suitably chosen known functions, and ϕ is some positive dispersion parameter. The function $b(\cdot)$ is assumed to be smooth and convex and gives rise to the link function $g(\mu) = \theta$ with $\mu = E(Y|\mathbf{x}) = b'(\theta)$. Thus the log-likelihood function given by the sample is

$$\ell_n(\boldsymbol{\beta}) = \sum_{i=1}^n \{y_i \mathbf{x}_i^T \boldsymbol{\beta} - b(\mathbf{x}_i^T \boldsymbol{\beta}) + c(y_i, \phi)\}. \quad (2)$$

To ensure model identifiability and improve model interpretability in high dimensions, it is common to assume that only a portion of all predictors contributes to the response, that is, the true regression coefficient vector $\boldsymbol{\beta}_0 = (\beta_{0,1}, \dots, \beta_{0,p})^T$ is sparse with many components being zero. We refer to predictors with nonzero coefficients $\beta_{0,j}$ as true covariates and the remaining ones as noise covariates. Without loss of generality, we write $\boldsymbol{\beta}_0 = (\boldsymbol{\beta}_1^T, \mathbf{0}^T)^T$ with $\boldsymbol{\beta}_1$ consisting of all s nonzero coefficients. To ease the presentation, we suppress the dependence of all parameters such as s and p on n whenever there is no confusion.

In the GLM setting, the regularization method minimizes the penalized negative log-likelihood function

$$Q_n(\boldsymbol{\beta}) = -n^{-1} \{ \mathbf{y}^T \mathbf{X} \boldsymbol{\beta} - \mathbf{1}^T \mathbf{b}(\mathbf{X} \boldsymbol{\beta}) \} + \|p_\lambda(\boldsymbol{\beta})\|_1, \quad (3)$$

where $\mathbf{y} = (y_1, \dots, y_n)^T$ is an n -dimensional response vector, $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$ is an $n \times p$ deterministic design matrix, $\mathbf{b}(\boldsymbol{\theta}) = (b(\theta_1), \dots, b(\theta_n))^T$ is a vector-valued function with $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)^T$ and $\theta_i = \mathbf{x}_i^T \boldsymbol{\beta}$, and $\|p_\lambda(\boldsymbol{\beta})\|_1 = \sum_{j=1}^p p_\lambda(|\beta_j|)$ is a separable penalty term on model parameters with $p_\lambda(t)$ a penalty function defined on $t \in [0, \infty)$ and indexed by a nonnegative regularization parameter λ . The last term in the log-likelihood function (2) involving the dispersion parameter ϕ is dropped for simplicity. Here we use a compact notation $p_\lambda(\boldsymbol{\beta}) = p_\lambda(|\boldsymbol{\beta}|) = (p_\lambda(|\beta_1|), \dots, p_\lambda(|\beta_p|))^T$ with the penalty function applied componentwise and $|\boldsymbol{\beta}| = (|\beta_1|, \dots, |\beta_p|)^T$. To align all covariates to a common scale, we rescale each column vector of the $n \times p$ design matrix \mathbf{X} for each covariate to have L_2 -norm $n^{1/2}$. As mentioned in the Introduction, many penalty functions have been proposed for variable selection and sparse modeling; see the references therein for their specific forms.

The level of collinearity among the covariates typically increases with the dimensionality. When this level is high, the estimation can become unstable and the model identifiability may not be guaranteed. We consider the idea of bounding the

sparse model size to control the collinearity for sparse models and ensure identifiability and stability of model for reliable prediction and variable selection. A natural bound is given by the following concept of robust spark on the design matrix \mathbf{X} , as introduced in Zheng, Fan, and Lv (2014).

Definition 1 (Robust spark). The robust spark κ_c of the $n \times p$ design matrix \mathbf{X} is defined as the smallest possible positive integer such that there exists an $n \times \kappa_c$ submatrix of $n^{-1/2}\mathbf{X}$ having a singular value less than a given positive constant c .

The above concept of robust spark generalizes that of spark in Donoho and Elad (2003), which plays an important role in the problem of sparse recovery; see also Lv and Fan (2009). As $c \rightarrow 0+$, the robust spark κ_c approaches the spark of \mathbf{X} . For each sparse model with size $m < \kappa_c$, the corresponding $n \times m$ submatrix of $n^{-1/2}\mathbf{X}$ have all singular values bounded from below by c . The robust spark κ_c is always a positive integer no larger than $n + 1$ and can be some large number diverging with n . Although we consider the case of deterministic design matrix, the following proposition formally characterizes the order of κ_c when the design matrix \mathbf{X} is generated from Gaussian distribution.

Proposition 1. Assume $\log p = o(n)$ and that the rows of the $n \times p$ random design matrix \mathbf{X} are independent and identically distributed (iid) as $N(\mathbf{0}, \Sigma)$, where Σ has smallest eigenvalue bounded from below by some positive constant. Then there exist positive constants c and \tilde{c} such that with asymptotic probability one, $\kappa_c \geq \tilde{c}n/(\log p)$.

To compare different regularization methods in (3), we introduce the thresholded parameter space

$$\mathcal{B}_{\tau,c} = \{\boldsymbol{\beta} \in \mathbb{R}^p : \|\boldsymbol{\beta}\|_0 < \kappa_c/2 \text{ and for each } j, \beta_j = 0 \text{ or } |\beta_j| \geq \tau\}, \quad (4)$$

where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ and τ is some positive threshold on parameter magnitude. The threshold τ is key to distinguishing between important covariates and noise covariates for the purpose of variable selection. As shown in Theorem 1 in Section 3.2, the threshold τ is needed to satisfy $\tau\sqrt{n}/(\log p) \rightarrow \infty$ as $n \rightarrow \infty$, indicating that the threshold level should dominate the maximum noise level of p independent standard Gaussian errors asymptotically.

The use of the thresholded parameter space $\mathcal{B}_{\tau,c}$ in (4) is motivated by the approach of the best subset regression with the L_0 -regularization, which was proved in Barron, Birge, and Massart (1999) to enjoy the oracle risk inequalities under the prediction loss. The following proposition is satisfied by any global minimizer of the regularization problem (3) when the L_0 -penalty $p_\lambda(t) = \lambda 1_{\{t \neq 0\}}$ is used.

Proposition 2 (Hard-thresholding property). For the L_0 -penalty $p_\lambda(t) = \lambda 1_{\{t \neq 0\}}$, the global minimizer $\hat{\boldsymbol{\beta}} = (\hat{\beta}_1, \dots, \hat{\beta}_p)^T$ of the regularization problem (3) over \mathbb{R}^p satisfies that each component $\hat{\beta}_j$ is either 0 or has magnitude larger than some positive threshold.

The above hard-thresholding property is shared by many other penalty functions. For example, Zheng, Fan, and Lv (2014) and Fan and Lv (2012) proved such a property in the setting of

penalized least squares for the hard-thresholding penalty (Hard) and SICA penalty, respectively. These continuous concave penalties are also considered in our study. Intuitively, if some covariates have weak effects, that is, having regression coefficients with magnitude below certain threshold, we can keep these variables out of the model to improve the prediction accuracy with reduced estimation variability because they may have negligible effects on prediction. Moreover, these weak signals are generally difficult to stand out compared with some noise variables due to the impact of high dimensionality.

3. ASYMPTOTIC EQUIVALENCE OF REGULARIZATION METHODS

In this section, we establish the asymptotic equivalence of the regularization methods (3) in the thresholded parameter space $\mathcal{B}_{\tau,c}$, with various penalty functions, in the sense of having the same convergence rates in the oracle inequalities and oracle risk inequalities.

3.1 Technical Conditions

We first introduce some notation and two key events to facilitate our technical presentation. Denote by $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_n)^T = \mathbf{Y} - E\mathbf{Y}$ the n -dimensional random model error vector with \mathbf{Y} the n -dimensional random response vector and $\boldsymbol{\alpha}_0 = \text{supp}(\boldsymbol{\beta}_0) = \{1, \dots, s\}$ the support of the true regression coefficient vector $\boldsymbol{\beta}_0$, that is, the true underlying sparse model. Throughout the article, we consider a universal choice of the regularization parameter $\lambda = c_0\sqrt{(\log p)/n}$ with some positive constant c_0 , where p is implicitly understood as $n \vee p$ in all bounds. Define two events

$$\begin{aligned} \mathcal{E} &= \{\|n^{-1}\mathbf{X}^T\boldsymbol{\varepsilon}\|_\infty \leq \lambda/2\} \quad \text{and} \\ \mathcal{E}_0 &= \{\|n^{-1}\mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon}\|_\infty \leq c_0\sqrt{(\log n)/n}\}, \end{aligned} \quad (5)$$

where \mathbf{X}_α denotes a submatrix of the design matrix \mathbf{X} consisting of columns with indices in a given set $\alpha \subset \{1, \dots, p\}$.

Condition 1 (Error tail distribution). The complements of the two events in (5) satisfy $P(\mathcal{E}^c) = O(p^{-c_1})$ and $P(\mathcal{E}_0^c) = O(n^{-c_1})$ for some positive constant c_1 that can be sufficiently large for large enough c_0 .

Condition 2 (Bounded variance). The function $b(\theta)$ satisfies that $c_2 \leq b''(\theta) \leq c_2^{-1}$ in its domain, where c_2 is some positive constant.

Condition 3 (Concave penalty function). The penalty function $p_\lambda(t)$ is increasing and concave in $t \in [0, \infty)$ with $p_\lambda(0) = 0$ and is differentiable with $p'_\lambda(0+) = c_3\lambda$ for some positive constant c_3 .

Condition 4 (Ultrahigh dimensionality). It holds that $\log p = O(n^a)$ for some constant $a \in (0, 1)$.

Condition 5 (True parameter vector). It holds that $s = o(n^{1-a})$ and there exists a constant $c > 0$ such that the robust spark $\kappa_c > 2s$. Moreover, $\min_{1 \leq j \leq s} |\beta_{0,j}| \gg \sqrt{(\log p)/n}$.

Condition 1 puts a constraint on the error tail distribution. The same event \mathcal{E} was considered in Bickel, Ritov, and Tsybakov (2009) for Gaussian error, and the probability bound on $P(\mathcal{E}^c)$ can be easily derived using the classical Gaussian tail

probability bound. We introduce a second event \mathcal{E}_0 to derive improved estimation and prediction bounds for the regularized estimator. The probability bound on $P(\mathcal{E}_0^c)$ holds similarly for Gaussian error. Condition 1 also holds for error distributions other than Gaussian, including bounded or light-tailed error, with no or mild condition on the design matrix \mathbf{X} . We discuss some technical details of this condition in Appendix A.

Condition 2 is a mild condition that is commonly assumed in the GLM setting and requires that the variances of all responses are bounded away from zero and infinity. Condition 3 is a common, mild assumption on the penalty function for studying regularization methods; see also Lv and Fan (2009) and Fan and Lv (2011). It requires that the penalty function $p_\lambda(t)$ is concave on the positive half axis $[0, \infty)$. In this context, a wide class of penalty functions, including the L_1 -penalty in Lasso, SCAD, MCP, and SICA, satisfy Condition 3 and belong to the class of concave penalty functions.

Condition 4 allows the dimensionality p to increase up to exponentially fast with the sample size n . Condition 5 puts constraints on the design matrix \mathbf{X} , the model sparsity, and the minimum signal strength. If τ is chosen such that $\tau\sqrt{n}/(\log p) \rightarrow \infty$ and $\tau < \min_{1 \leq j \leq s} |\beta_{0,j}|$, and Condition 5 is satisfied, then it is seen that $\beta_0 \in \mathcal{B}_{\tau,c}$ with $\mathcal{B}_{\tau,c}$ defined in (4). For the reason presented above, in the future presentation, we only consider appropriately chosen τ such that $\beta_0 \in \mathcal{B}_{\tau,c}$. In addition, since we only need the existence of a constant c satisfying Condition 5 and its exact value is not needed in implementation, we will suppress the dependence of $\mathcal{B}_{\tau,c}$ on c and write it as \mathcal{B}_τ hereafter.

3.2 Oracle Inequalities of Global Minimizer

In this section, we aim to establish the oracle inequalities for the global minimizer of the penalized negative log-likelihood (3) in the thresholded parameter space \mathcal{B}_τ , that is,

$$\hat{\beta} = \arg \min_{\beta \in \mathcal{B}_\tau} Q_n(\beta). \tag{6}$$

In general, there may exist multiple global minimizers of $Q_n(\beta)$. Our theoretical results are satisfied by any of these global minimizers. Throughout the article, we refer to any global minimizer as the regularized estimator. The oracle inequalities for the Lasso estimator under estimation and prediction losses were established in Bickel, Ritov, and Tsybakov (2009) to study the asymptotic equivalence of the Lasso estimator and Dantzig selector. In addition to common estimation and prediction losses, we introduce a variable selection loss defined as the total number of falsely discovered signs of covariates by an estimator $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_p)^T$,

$$\text{FS}(\hat{\beta}) = |\{j : \text{sgn}(\hat{\beta}_j) \neq \text{sgn}(\beta_{0,j}), 1 \leq j \leq p\}|. \tag{7}$$

This loss of false signs $\text{FS}(\hat{\beta})$ is a stronger measure than commonly used ones such as the number of false positives and the number of false negatives. We will use this measure to study the sign consistency property of the regularized estimator $\hat{\beta}$ (Zhao and Yu 2006).

Theorem 1 (Oracle inequalities). Assume that Conditions 1–5 hold and τ is chosen such that $\tau < \min_{1 \leq j \leq s} |\beta_{0,j}|$ and $\lambda = c_0\sqrt{(\log p)/n} = o(\tau)$. Then the global minimizer defined

in (6) exists, and any such global minimizer satisfies that with probability at least $1 - O(p^{-c_1})$, it holds simultaneously that

- (a) (*False signs*). $\text{FS}(\hat{\beta}) \leq Cs\lambda^2\tau^{-2}/(1 - C\lambda^2\tau^{-2})$;
- (b) (*Estimation losses*). $\|\hat{\beta} - \beta_0\|_q \leq C\lambda s^{1/q}(1 - C\lambda^2\tau^{-2})^{-1/q}$ for each $q \in [1, 2]$ and $\|\hat{\beta} - \beta_0\|_\infty \leq C\lambda s^{1/2}(1 - C\lambda^2\tau^{-2})^{-1/2}$;
- (c) (*Prediction loss*). $n^{-1/2}\|\mathbf{X}(\hat{\beta} - \beta_0)\|_2 \leq C\lambda s^{1/2}(1 - C\lambda^2\tau^{-2})^{-1/2}$,

where C is some positive constant.

Theorem 1 shows the existence of the global minimizer defined in (6) and presents the oracle inequalities for the regularized estimator for a wide class of penalty functions characterized by Condition 3. All theoretical results in the article hold uniformly over the set of all possible global minimizers.

Since the regularization parameter λ represents the minimum regularization level needed to suppress the noise covariates, and the thresholding level τ is just below the minimum signal strength, a valid thresholding level requires $\lambda = o(\tau)$ to ensure that all true covariates will not be screened out asymptotically. Since $\lambda\tau^{-1} \rightarrow 0$, the above bound on false signs $\text{FS}(\hat{\beta})$ is of a smaller order than the true model size s , meaning that the proportion of missed signs for signals, that is, $\text{FS}(\hat{\beta})/s$, vanishes asymptotically. This tight bound on false signs is a unique feature of introducing the thresholded parameter space. In contrast, the bound on estimated model size $\|\hat{\beta}\|_0$ for the ordinary Lasso estimator is of order $O(\phi_{\max}s)$ with ϕ_{\max} the largest eigenvalue of the Gram matrix $n^{-1}\mathbf{X}^T\mathbf{X}$ (Bickel, Ritov, and Tsybakov 2009), and thus the proportion of missed signs $\text{FS}(\hat{\beta})/s$ in this estimator can be of order $O(\phi_{\max})$, which does not vanish asymptotically. In view of $\lambda\tau^{-1} \rightarrow 0$ and $\lambda = c_0\sqrt{(\log p)/n}$, the bounds on the estimation and prediction losses in Theorem 1 satisfy that for each $q \in [1, 2]$,

$$\begin{aligned} \|\hat{\beta} - \beta_0\|_q &= O\{s^{1/q}\sqrt{(\log p)/n}\} \quad \text{and} \\ n^{-1/2}\|\mathbf{X}(\hat{\beta} - \beta_0)\|_2 &= O(\sqrt{s(\log p)/n}), \end{aligned}$$

whose convergence rates are within a logarithmic factor of $\log p$ of the oracle rates. The above convergence rates in these oracle inequalities are consistent with those in Bickel, Ritov, and Tsybakov (2009) for the Lasso estimator.

We next show that under some additional conditions, the sign consistency of the regularized estimator $\hat{\beta}$ can be obtained and the convergence rates in Theorem 1 can be further improved. Define a small neighborhood of β_0 in the thresholded parameter space as

$$\begin{aligned} \mathcal{B}_1^* &= \{\beta \in \mathcal{B}_\tau : \text{supp}(\beta) = \text{supp}(\beta_0) \text{ and} \\ &\|\beta - \beta_0\|_2 \leq 2Cs^{1/2}\lambda\} \end{aligned} \tag{8}$$

with constant C given in Theorem 1. Note that this neighborhood is asymptotically shrinking since $s^{1/2}\lambda \rightarrow 0$ as guaranteed by Conditions 4 and 5. We introduce two important constants

$$\gamma_n^* = \sup_{\beta_i \in \mathcal{B}_1^*, i=1, \dots, n} \left\| \left\{ \frac{1}{n} \mathbf{X}_{\alpha_0}^T \mathbf{H}(\beta_1, \dots, \beta_n) \mathbf{X}_{\alpha_0} \right\}^{-1} \right\|_\infty, \tag{9}$$

$$\gamma_n = \sup_{\beta \in \mathcal{B}_1^*, \alpha \in \{s+1, \dots, p\} \text{ and } |\alpha| \leq s} \left\| \frac{1}{n} \mathbf{X}_{\alpha_0}^T \mathbf{H}(\beta) \mathbf{X}_\alpha \right\|_\infty, \tag{10}$$

where $\mathbf{H}(\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_n) = \text{diag}\{b''(\mathbf{x}_1^T \boldsymbol{\beta}_1), \dots, b''(\mathbf{x}_n^T \boldsymbol{\beta}_n)\}$ and $\mathbf{H}(\boldsymbol{\beta}) = \text{diag}\{b''(\mathbf{x}_1^T \boldsymbol{\beta}), \dots, b''(\mathbf{x}_n^T \boldsymbol{\beta})\}$ are diagonal matrices of variances. To get some intuition on the constants γ_n^* and γ_n , let us consider the special case of Gaussian linear model with $b''(\theta) \equiv 1$. In such case, we have

$$\begin{aligned} \gamma_n^* &= \left\| \left(\frac{1}{n} \mathbf{X}_{\alpha_0}^T \mathbf{X}_{\alpha_0} \right)^{-1} \right\|_{\infty} \quad \text{and} \\ \gamma_n &= \sup_{\alpha \subset \{s+1, \dots, p\} \text{ and } |\alpha| \leq s} \left\| \frac{1}{n} \mathbf{X}_{\alpha_0}^T \mathbf{X}_{\alpha} \right\|_{\infty}. \end{aligned} \quad (11)$$

Since each column of \mathbf{X} is rescaled to have L_2 -norm $n^{1/2}$, it is seen that γ_n^* is only associated with the design matrix of the true model α_0 , while γ_n is related to the correlation between true covariates and noise covariates.

To evaluate the prediction property, we consider the Kullback–Leibler divergence of the fitted model from the true model given by

$$D(\widehat{\boldsymbol{\beta}}) = -(E\mathbf{Y})^T \mathbf{X}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) + \mathbf{1}^T [\mathbf{b}(\mathbf{X}\widehat{\boldsymbol{\beta}}) - \mathbf{b}(\mathbf{X}\boldsymbol{\beta}_0)],$$

where $E\mathbf{Y} = (b'(\mathbf{x}_1^T \boldsymbol{\beta}_0), \dots, b'(\mathbf{x}_n^T \boldsymbol{\beta}_0))^T$ is the true mean response vector for the GLM.

Theorem 2 (Sign consistency and oracle inequalities). Assume that conditions of Theorem 1 hold with $\min_{1 \leq j \leq s} |\beta_{0,j}| \geq 2\tau$, $\lambda = c_0 \sqrt{(\log p)/n} = o(s^{-1/2}\tau)$, and $\gamma_n = o\{\tau \sqrt{n}/(s \log n)\}$. Then any global minimizer $\widehat{\boldsymbol{\beta}}$ in (6) satisfies that with probability at least $1 - O(n^{-c_1})$, it holds simultaneously that

- (a) (*Sign consistency*). $\text{sgn}(\widehat{\boldsymbol{\beta}}) = \text{sgn}(\boldsymbol{\beta}_0)$;
- (b) (*Estimation and prediction losses*). If the penalty function further satisfies $p'_\lambda(\tau) = O\{\sqrt{(\log n)/n}\}$, then we have for each $q \in [1, 2]$,

$$\begin{aligned} \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_q &\leq Cs^{1/q} \sqrt{(\log n)/n}, \\ \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_{\infty} &\leq C\gamma_n^* \sqrt{(\log n)/n}, \end{aligned}$$

$$\text{and } n^{-1} D(\widehat{\boldsymbol{\beta}}) \leq Cs(\log n)/n,$$

where C is some positive constant.

In comparison with Theorem 1(a), we obtain in Theorem 2(a) a stronger property of sign consistency of the regularized estimator. The additional condition on the penalty function $p'_\lambda(\tau) = O\{\sqrt{(\log n)/n}\}$ can be easily satisfied by concave penalties such as the SCAD and SICA, with appropriately chosen λ . For penalty functions satisfying this additional condition, the convergence rates of the regularized estimator are improved with the $\log p$ term (see Theorem 1) replaced with $\log n$ (see Theorem 2). In this sense, our study provides a setting showing the general nonoptimality of the logarithmic factor of the dimensionality $\log p$ in oracle inequalities.

To gain more insights into Theorem 2, we consider again the case of Gaussian linear model. In view of (11) and the robust spark condition in (4), we have an upper bound on γ_n^* given by

$$\gamma_n^* \leq s^{1/2} \left\| \left(\frac{1}{n} \mathbf{X}_{\alpha_0}^T \mathbf{X}_{\alpha_0} \right)^{-1} \right\|_2 \leq c^{-1} s^{1/2}.$$

Observing that γ_n in (11) measures the correlation between noise covariates and true covariates, the condition $\gamma_n = o\{\tau \sqrt{n}/(s \log n)\}$ in Theorem 2 essentially requires that the

noise covariates and true covariates should not be too highly correlated with each other. Note that each column of \mathbf{X}_{α_0} is rescaled to have L_2 -norm $n^{1/2}$. When all true covariates are orthogonal to each other, we have $\gamma_n^* = 1$ and thus the bound on the L_{∞} -estimation loss in Theorem 2 becomes

$$\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_{\infty} \leq C\sqrt{(\log n)/n},$$

whose convergence rate is within a logarithmic factor of $\log n$ of the oracle rate.

Combining Theorems 1 and 2 shows that for polynomially growing dimensionality with $p = O(n^a)$ for some positive constant a , the L_1 -regularization method of Lasso and concave regularization methods with penalties satisfying Condition 3 are asymptotically equivalent in the thresholded parameter space, meaning that all methods have the same convergence rates in the oracle inequalities, with a logarithmic factor of $\log n$. For exponentially growing dimensionality with $\log p = O(n^a)$ for some positive constant a less than 1, the concave regularization methods satisfying the additional condition $p'_\lambda(\tau) = O\{\sqrt{(\log n)/n}\}$ are asymptotically equivalent and still enjoy the same convergence rates in the oracle inequalities, with a logarithmic factor of $\log n$. For the L_1 -penalty used in Lasso, the condition $p'_\lambda(\tau) = O\{\sqrt{(\log n)/n}\}$ and the choice of the regularization parameter $\lambda = c_0 \sqrt{(\log p)/n}$ are, however, incompatible with each other in the case of $\log p = O(n^a)$. Thus in the ultrahigh-dimensional case, the convergence rates in the oracle inequalities for Lasso, which have a logarithmic factor of $\log p$, are slower than those for concave regularization methods. These results reveal an interesting phase diagram on how the performance of regularization methods, in the thresholded parameter space, evolves with the dimensionality and the penalty function, in terms of convergence rates in the oracle inequalities.

Among different approaches to alleviating the bias issue of the Lasso, the adaptive Lasso (Zou 2006) exploits the weighted L_1 -penalty $\lambda \|\mathbf{w} \circ \boldsymbol{\beta}\|_1$ with weight vector $\mathbf{w} = (w_1, \dots, w_p)^T$, where $w_j = |\beta_{\text{ini},j}|^{-\gamma}$ for some $\gamma > 0$, $1 \leq j \leq p$, with $\boldsymbol{\beta}_{\text{ini}} = (\beta_{\text{ini},1}, \dots, \beta_{\text{ini},p})^T$ an initial estimator, and \circ denotes the componentwise product. Under some particular choices of the initial estimator, the adaptive Lasso can enjoy the properties established in Theorems 1 and 2, similarly as the Lasso. For instance, the choice of the trivial initial estimator $\boldsymbol{\beta}_{\text{ini}} = \mathbf{1}$ gives the Lasso estimator. How to choose other nontrivial initial estimators is crucial to ensuring that the adaptive Lasso has improved convergence rates as concave methods in ultrahigh dimensions. Another popular method, the bridge regression in Frank and Friedman (1993), uses the L_q -penalty $p_\lambda(t) = \lambda t^q$ for $0 < q \leq 2$. When $0 < q < 1$, the bridge regression is also a concave regularization method since $p_\lambda(t)$ is concave on $[0, \infty)$. However, such a method falls outside the class of regularization methods in our framework, since $p'_\lambda(0+) = \infty$ in this case, which violates Condition 3. As a consequence, a key inequality (B.7) in our technical analysis does not hold in general for the bridge estimator with $0 < q < 1$. It is yet unclear whether similar results to those in Theorems 1 and 2 would also hold for the bridge estimator in the case of $0 < q < 1$.

3.3 Oracle Risk Inequalities of Global Minimizer

The oracle inequalities presented in Section 3.2 are derived by conditioning on the event \mathcal{E} (Theorem 1) or $\mathcal{E} \cap \mathcal{E}_0$ (Theorem

2) defined in (5), and thus they may not hold on the complement \mathcal{E}^c or $\mathcal{E}^c \cup \mathcal{E}_0^c$. We now establish a stronger property of the oracle risk inequalities for the regularized estimator $\widehat{\beta}$ in (6), which gives upper bounds on the expectations of various variable selection, estimation, and prediction losses.

Theorem 3 (Oracle risk inequalities). Assume that conditions of Theorem 2 hold and the fourth moments of errors $E\varepsilon_i^4$ are uniformly bounded. Then any global minimizer $\widehat{\beta}$ in (6) satisfies that

- (a) (*Sign risk*). $E\{\text{FS}(\widehat{\beta})\} = \frac{1}{p_\lambda(\tau)}\{\|p_\lambda(\beta_0)\|_1 + s\lambda^2\} O(n^{-c_1}) + O(p^{-c_1/2}\kappa_c)$;
- (b) (*Estimation and prediction risks*). If the penalty function further satisfies $p'_\lambda(\tau) = O\{\sqrt{(\log n)/n}\}$, then we have for each $q \in [1, 2]$,

$$\begin{aligned} E\|\widehat{\beta} - \beta_0\|_q^q &\leq Cs[(\log n)/n]^{q/2}, \\ E\|\widehat{\beta} - \beta_0\|_\infty &\leq C\gamma_n^* \sqrt{(\log n)/n}, \text{ and} \\ E\{n^{-1}D(\widehat{\beta})\} &\leq Cs(\log n)/n, \end{aligned}$$

where C is some positive constant.

The expectation of the number of falsely discovered signs converges to zero at a polynomial rate of n . In the wavelet setting of Gaussian linear model with $p = n$ and orthogonal design matrix \mathbf{X} , it has been proved in Antoniadis and Fan (2001) that the risks of the regularized estimators under the L_2 -loss are bounded by $O\{s(\log n)/n\}$, which is consistent with our aforementioned results. This indicates that there is no additional cost in risk bounds for generalizing to the ultrahigh-dimensional nonlinear model setting of the GLM.

3.4 Sampling Properties of Computable Solutions

The theoretical results presented in previous sections are on any global minimizer of the penalized negative log-likelihood $Q_n(\beta)$ in the thresholded parameter space \mathcal{B}_τ . The global minimizer may not be guaranteed to be found by a computational algorithm. Therefore, it is also important to study the sampling properties of the computable solution produced by any algorithm. Define a vector-valued function $\mu(\theta) = (b'(\theta_1), \dots, b'(\theta_n))^T$ for $\theta = (\theta_1, \dots, \theta_n)^T$, which is the mean function in the GLM.

Theorem 4. Let $\widehat{\beta} \in \mathcal{B}_\tau$ be a computable solution to the minimization problem (6) produced by any algorithm that is the global minimizer when constrained on the subspace given by $\text{supp}(\widehat{\beta})$, and $\eta_n = \|n^{-1}\mathbf{X}^T[\mathbf{y} - \mu(\mathbf{X}\widehat{\beta})]\|_\infty$. Assume in addition that there exists some positive constant c_4 such that $\|n^{-1}\mathbf{X}_\alpha^T[\mu(\mathbf{X}\beta) - \mu(\mathbf{X}\beta_0)]\|_2 \geq c_4\|\beta - \beta_0\|_2$ for any $\beta \in \mathcal{B}_\tau$ and $\alpha = \text{supp}(\beta) \cup \text{supp}(\beta_0)$, if model (1) is nonlinear. If $\eta_n + \lambda = o(\tau)$ and $\min_{1 \leq j \leq s} |\beta_{0,j}| > c_5s^{1/2}(\eta_n + \lambda)$ with c_5 some sufficiently large positive constant, then $\widehat{\beta}$ enjoys the same asymptotic properties as for any global minimizer in Theorems 1–3 under the same conditions therein.

The condition that $\widehat{\beta}$ is the global minimizer of problem (6) when constrained on the subspace given by its support can hold under some mild condition on the penalty function. Such a property has been formally characterized in proposition 1 of Fan and Lv (2011). For example, when condition (12) in

Section 4.1 is satisfied, the penalized negative log-likelihood $Q_n(\beta)$ in (3) is strictly convex on the above subspace, which entails that the local minimizer found by any algorithm will necessarily be the global minimizer over this subspace.

As shown in the proof of Theorem 4, the above additional condition on the mean deviation vector $\mu(\mathbf{X}\beta) - \mu(\mathbf{X}\beta_0)$ always holds for linear model with $c_4 = c^2$. In nonlinear models, such a condition requires that a deviation from the true mean vector $\mu(\mathbf{X}\beta_0)$ can be captured by the covariates involved. Theorem 4 shows that a computable solution produced by any algorithm can share the same nice asymptotic properties as for any global minimizer, when the maximum correlation between the covariates and the residual vector $\mathbf{y} - \mu(\mathbf{X}\widehat{\beta})$ is a smaller order of the threshold τ . Such a solution needs not to be the global minimizer.

4. NUMERICAL STUDIES

4.1 Implementation

Algorithms for implementing regularization methods include those mentioned in the Introduction, the local quadratic approximation algorithm (Fan and Li 2001), and local linear approximation algorithm (Zou and Li 2008). In particular, the coordinate optimization algorithm, which solves the problem one coordinate a time and cycles through all coordinates, has received much recent attention for solving large-scale problems thanks to its very low computational cost for each coordinate. For example, the iterative coordinate ascent (ICA) algorithm (Fan and Lv 2011) implements regularization methods by combining the ideas of second-order quadratic approximation of likelihood function and coordinate optimization. For each coordinate within each iteration, the quadratic approximation of the likelihood function at the p -vector from the previous step along that coordinate reduces the problem to a univariate penalized least squares, which admits a closed-form solution for many commonly used penalty functions. See, for example, Lin and Lv (2013) for an analysis of convergence properties of this algorithm.

In this article, we apply the ICA algorithm to implement concave regularization methods in the thresholded parameter space. A key ingredient of these methods is the use of the thresholded parameter space, which naturally puts an additional constraint on each component of the parameter vector. For each coordinate within each iteration, we solve the univariate penalized least-squares problem with the corresponding quadratic approximation of the likelihood function and update this coordinate only when the global minimizer has magnitude above the given threshold τ . We found that this optimization algorithm works well for producing the solution paths for concave regularization methods in the thresholded parameter space. The thresholding also induces additional sparsity of the regularized estimate and thus makes the algorithm converge faster.

To gain some insight into the stability of the computational algorithm, assume that the penalty function $p_\lambda(t)$ has maximum concavity

$$\rho(p_\lambda) = \sup_{0 < t_1 < t_2 < \infty} \left\{ -\frac{p'_\lambda(t_2) - p'_\lambda(t_1)}{t_2 - t_1} \right\} < cc_2, \quad (12)$$

where constants c and c_2 are given in Definition 1 and Condition 2, respectively. This condition holds for penalties satisfying Condition 3 with suitably chosen regularization parameter λ and shape parameter. For example, the L_1 -penalty $p_\lambda(t) = \lambda t$ in Lasso has maximum concavity 0; the SCAD penalty $p_\lambda(t)$ having derivative $p'_\lambda(t) = \lambda I(t \leq \lambda) + (a-1)^{-1}(a\lambda - t)_+ I(t > \lambda)$, with shape parameter $a > 2$, has maximum concavity $\rho(p_\lambda) = (a-1)^{-1}$; and the SICA penalty $p_\lambda(t; a) = \lambda(a+1)t/(a+t)$ with shape parameter a has maximum concavity $2\lambda(a^{-1} + a^{-2})$. Condition (12) on the maximum concavity of penalty function ensures that the penalized negative log-likelihood $Q_n(\beta)$ in (3) is strictly convex on a union of coordinate subspaces $\{\beta \in \mathbb{R}^p : \|\beta\|_0 < \kappa_c\}$, which is key to the stability of the sparse solution found by any algorithm.

In implementation, we need to select two tuning parameters: the threshold τ for the thresholded parameter space \mathcal{B}_τ and the regularization parameter λ for the penalty function $p_\lambda(t)$. As shown in the theoretical results, the threshold τ should be larger than the regularization parameter $\lambda = c_0 \sqrt{(\log p)/n}$ to filter the noise. Thus we choose τ as $\tau = c_6 (\log n)^{1/2} \sqrt{(\log p)/n}$ for some positive constant c_6 . As for the regularization parameter λ , we use the validation set or cross-validation (CV) to select τ .

4.2 Simulation Studies

In this section, we investigate the finite-sample properties of several concave regularization methods in the thresholded parameter space, in three commonly used GLMs: the linear regression model, the logistic regression model, and the Poisson regression model, as well as in a real data example. Since the main purpose of our simulation study is to justify the theoretical results, we select the tuning parameters by minimizing the prediction error calculated using an independent validation set, with size equal to the sample size in the study. This tuning parameter selection criterion reduces additional estimation variability incurred by the CV. Fivefold CV was used for tuning parameter selection in real data analysis.

4.2.1 Linear Regression. We start with the linear regression model (1) written in the matrix form

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}. \quad (13)$$

We generated 100 datasets from this model with error $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I}_n)$ independent of the design matrix \mathbf{X} . The sample size n and error standard deviation σ were chosen to be 100 and 0.4, respectively. For each dataset, the rows of the design matrix \mathbf{X} were sampled as iid copies of random p -vector from $N(\mathbf{0}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = (r^{|j-k|})_{1 \leq j, k \leq p}$ for some number r . We considered three settings for the pair (p, r) of dimensionality and population collinearity level: (1000, 0.25), (1000, 0.5), and (5000, 0.25). In addition to the population collinearity, the sample collinearity among the covariates can be of a much higher level due to the high dimensionality. The true regression coefficient vector $\boldsymbol{\beta}$ was set to be $\boldsymbol{\beta}_0 = (1, -0.5, 0.7, -1.2, -0.9, 0.5, 0.55, 0, \dots, 0)^T$. We take the oracle procedure, using the information of the true underlying sparse model, as the benchmark variable selection method, and compare the Lasso, SCAD, MCP, Hard, and SICA in the thresholded parameter space, which are referred to as Lasso_t, SCAD_t, MCP_t, Hard_t, and SICA_t for simplicity, respectively. We also include the original SCAD in the comparison.

Simulation results show that SCAD_t, MCP_t, and Hard_t had very similar performance, so we omit the results on MCP_t and Hard_t to save space. The shape parameter a of the SCAD and SICA penalties was chosen to be 3.7, and 10^{-4} or 10^{-2} , respectively.

To evaluate the selected models, we consider several performance measures for prediction and variable selection. The first measure is the prediction error (PE) defined as $E(Y - \mathbf{x}^T \widehat{\boldsymbol{\beta}})^2$ with $\widehat{\boldsymbol{\beta}}$ an estimate and (\mathbf{x}^T, Y) an independent observation for the p covariates and response. An independent test sample of size 10,000 was generated to calculate the PE. The second to fourth measures are the L_q -estimation losses $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_q$ with $q = 2, 1$, and ∞ , respectively. The fifth and sixth measures are variable selection losses of false positives (FP) and false negatives (FN), where a false positive represents a falsely selected noise covariate in the model and a false negative represents a missed true covariate. The seventh measure is the model selection consistency probability of each method based on 100 simulations. We also compare the estimate $\widehat{\sigma}$ of the error standard deviation σ in linear model for all methods.

The model selection consistency results are summarized in Table 2 and all other results are summarized in Table 1. We see that across all settings and over all performance measures through their means and standard errors, all concave regularization methods in the thresholded parameter space mimicked very closely the oracle procedure. In particular, the model selection consistency probability for each of these methods was very close to one and its estimated error standard deviation followed very closely that by the oracle procedure. Figure 1 further shows that the regularized estimators given by these methods had almost identical sampling distributions. These numerical results are in line with our theory presented in Section 3. We also observe that SCAD_t improved over original SCAD in both prediction and variable selection. The model selection consistency probability of SCAD was particularly improved when considering the thresholded parameter space.

We also consider three additional settings for (n, p, r) : (100, 5000, 0.5), (200, 5000, 0.5), and (400, 5000, 0.5). The comparison results for all methods are presented in Tables 3 and 4. Due to the high collinearity in the setting of $(n, p, r) = (100, 5000, 0.5)$, all methods performed worse than the oracle procedure. As sample size increases, these methods followed more closely the oracle procedure, which are consistent with our theoretical results.

4.2.2 Logistic Regression. We consider the logistic regression model (1) with the parameter θ_i for the response Y_i given by

$$\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)^T = \mathbf{X}\boldsymbol{\beta}. \quad (14)$$

We generated 100 datasets from this model, each of which contains an n -dimensional response vector \mathbf{y} sampled from the Bernoulli distribution with success probability vector $(e^{\theta_1}/(1 + e^{\theta_1}), \dots, e^{\theta_n}/(1 + e^{\theta_n}))^T$, where $\boldsymbol{\theta}$ is given in (14). The sample size n and the true regression coefficient vector $\boldsymbol{\beta}$ were set to be 200 and $\boldsymbol{\beta}_0 = (2, 0, -2.3, 0, 2.8, 0, -2.2, 0, 2.5, 0, \dots, 0)^T$, respectively. The rest of the setting is the same as that in Section 4.2.1. We compared the same concave regularization methods with the oracle procedure and used the same seven prediction and variable selection performance measures as in

Table 1. The means and standard errors (in parentheses) of various performance measures as well as the estimated error standard deviation for all methods in Section 4.2.1

Measure	Method				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
Setting I					
PE (×0.1)	1.722 (0.007)	1.736 (0.007)	1.721 (0.007)	1.719 (0.007)	1.719 (0.007)
L ₂ -loss (×0.1)	1.122 (0.032)	1.184 (0.030)	1.115 (0.030)	1.106 (0.031)	1.106 (0.031)
L ₁ -loss (×0.1)	2.485 (0.077)	2.972 (0.100)	2.425 (0.071)	2.414 (0.071)	2.414 (0.071)
L _∞ -loss (×0.01)	7.48 (0.24)	7.67 (0.21)	7.61 (0.21)	7.55 (0.23)	7.55 (0.23)
FP	0.01 (0.01)	3.84 (0.47)	0 (0)	0 (0)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
$\hat{\sigma}$ (×0.1)	4.040 (0.035)	3.959 (0.034)	4.019 (0.034)	4.011 (0.034)	4.011 (0.034)
Setting II					
PE (×0.1)	1.789 (0.045)	1.741 (0.008)	1.735 (0.008)	1.738 (0.019)	1.719 (0.007)
L ₂ -loss (×0.1)	1.445 (0.100)	1.403 (0.039)	1.375 (0.040)	1.353 (0.062)	1.301 (0.038)
L ₁ -loss (×0.1)	3.360 (0.318)	3.558 (0.118)	3.180 (0.108)	2.957 (0.132)	2.862 (0.088)
L _∞ -loss (×0.01)	9.42 (0.65)	8.99 (0.28)	8.95 (0.28)	9.22 (0.49)	8.76 (0.26)
FP	0.22 (0.18)	4.11 (0.48)	0.56 (0.12)	0.01 (0.01)	0 (0)
FN	0.01 (0.01)	0 (0)	0 (0)	0.01 (0.01)	0 (0)
$\hat{\sigma}$ (×0.1)	4.023 (0.033)	3.937 (0.034)	3.963 (0.036)	4.016 (0.035)	4.010 (0.034)
Setting III					
PE (×0.1)	1.722 (0.008)	1.743 (0.008)	1.719 (0.007)	1.724 (0.008)	1.715 (0.006)
L ₂ -loss (×0.1)	1.133 (0.034)	1.228 (0.033)	1.123 (0.032)	1.138 (0.034)	1.104 (0.031)
L ₁ -loss (×0.1)	2.457 (0.074)	3.455 (0.139)	2.455 (0.071)	2.488 (0.075)	2.438 (0.070)
L _∞ -loss (×0.01)	7.77 (0.28)	7.79 (0.24)	7.61 (0.26)	7.80 (0.29)	7.43 (0.24)
FP	0.02 (0.01)	8.25 (0.84)	0.01 (0.01)	0.06 (0.02)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
$\hat{\sigma}$ (×0.1)	4.003 (0.032)	3.859 (0.034)	3.988 (0.031)	3.966 (0.032)	3.983 (0.031)

NOTE: Settings I, II, and III refer to cases of $(p, r) = (1000, 0.25)$, $(1000, 0.5)$, and $(5000, 0.25)$, respectively.

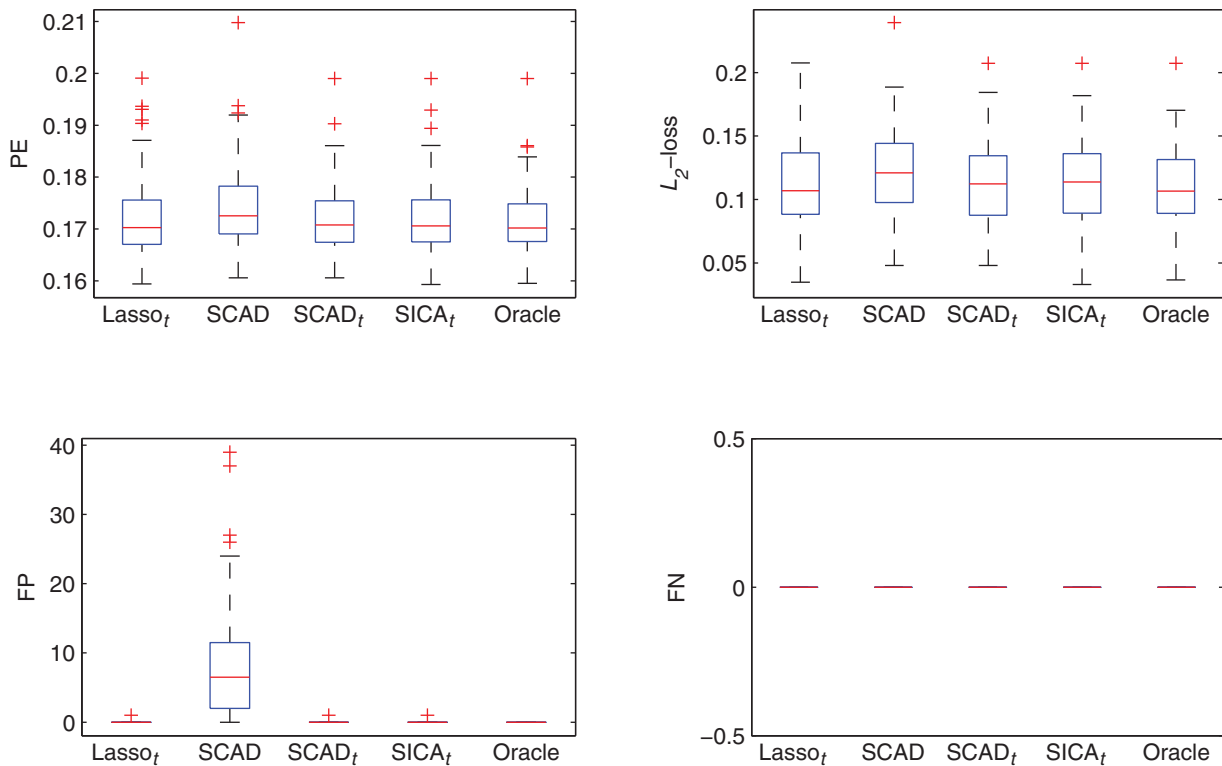


Figure 1. Boxplots of the PE, L₂-loss, FP, and FN over 100 simulations for all methods in Section 4.2.1, with $(p, r) = (5000, 0.25)$. The x-axis represents different methods.

Table 2. Model selection consistency probabilities of all methods in Section 4.2.1

Setting of (p, r)	Model selection consistency probability				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
(1000, 0.25)	0.99	0.26	1	1	1
(1000, 0.5)	0.96	0.26	0.71	0.98	1
(5000, 0.25)	0.98	0.14	0.99	0.94	1

Table 4. Model selection consistency probabilities of all methods in Section 4.2.1 with $(p, r) = (5000, 0.5)$

n	Model selection consistency probability				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
100	0.78	0.10	0.68	0.84	1
200	1	0.55	1	1	1
400	1	0.69	1	1	1

Section 4.2.1. The prediction error is defined as $E\{Y - \exp(\mathbf{x}^T \hat{\boldsymbol{\beta}}) / [1 + \exp(\mathbf{x}^T \hat{\boldsymbol{\beta}})]^2\}$ with $\hat{\boldsymbol{\beta}}$ an estimate and (\mathbf{x}^T, Y) an independent observation for the p covariates and response.

Tables 5 and 6 and Figure 2 summarize the comparison results of all methods. The conclusions are similar to those in Section 4.2.1. Facilitated by the thresholded parameter space, all methods mimicked very closely the oracle procedure in this nonlinear model for binary data, confirming the theoretical results.

4.2.3 Poisson Regression. We now consider the Poisson regression model (1) with the parameter θ_i for the response Y_i given as in (14). We generated 100 datasets from this model, each of which contains an n -dimensional response vector \mathbf{y} sampled from the Poisson distribution with mean vector $(e^{\theta_1}, \dots, e^{\theta_n})^T$, where $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)^T$ is given in (14). The sample size n and the true regression coefficient vector $\boldsymbol{\beta}$ were

set to be 200 and $\boldsymbol{\beta}_0 = (1, -0.9, 0.8, -1.1, 0.6, 0, \dots, 0)^T$, respectively. The rest of the setting is the same as that in Section 4.2.2. We compared the same concave regularization methods with the oracle procedure, using the same seven prediction and variable selection performance measures as in Section 4.2.1. The prediction error is defined as $E\{Y - \exp(\mathbf{x}^T \hat{\boldsymbol{\beta}})\}^2$ with $\hat{\boldsymbol{\beta}}$ an estimate and (\mathbf{x}^T, Y) an independent observation for the p covariates and response.

Tables 7 and 8 and Figure 3 summarize the comparison results for all methods. As shown in Figure 3, the boxplot for the prediction error of the oracle procedure exhibits some outliers. This is caused by the random design matrix, which may not be well-behaved in some samples, leading to some unstable estimates of the true regression coefficients. The instability comes from the fact that the variance of a Poisson random variable is equal to its mean and thus is generally unbounded if the mean is not bounded. To better compare the performance

Table 3. The means and standard errors (in parentheses) of various performance measures as well as the estimated error standard deviation for all methods in Section 4.2.1 with $(p, r) = (5000, 0.5)$

Measure	Method				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
Setting I					
PE ($\times 0.1$)	2.584 (0.215)	1.958 (0.105)	1.820 (0.062)	2.103 (0.134)	1.715 (0.006)
L_2 -loss ($\times 0.1$)	2.935 (0.343)	1.824 (0.187)	1.555 (0.126)	2.102 (0.243)	1.304 (0.039)
L_1 -loss ($\times 0.1$)	6.750 (0.841)	5.296 (0.523)	3.681 (0.280)	4.618 (0.535)	2.909 (0.089)
L_∞ -loss ($\times 0.01$)	19.28 (2.21)	11.48 (1.25)	10.02 (0.91)	14.31 (1.68)	8.63 (0.29)
FP	0.19 (0.07)	11.33 (1.00)	0.91 (0.17)	0.08 (0.03)	0 (0)
FN	0.41 (0.09)	0.06 (0.03)	0.05 (0.04)	0.21 (0.07)	0 (0)
$\hat{\sigma}$ ($\times 0.1$)	4.394 (0.111)	3.893 (0.061)	3.930 (0.050)	4.169 (0.082)	3.983 (0.031)
Setting II					
PE ($\times 0.1$)	1.655 (0.004)	1.661 (0.004)	1.654 (0.004)	1.652 (0.004)	1.652 (0.004)
L_2 -loss ($\times 0.1$)	0.920 (0.034)	0.951 (0.032)	0.916 (0.031)	0.891 (0.031)	0.894 (0.031)
L_1 -loss ($\times 0.1$)	2.025 (0.079)	2.427 (0.141)	1.996 (0.071)	1.952 (0.071)	1.958 (0.072)
L_∞ -loss ($\times 0.01$)	6.08 (0.22)	6.21 (0.22)	6.19 (0.22)	5.98 (0.22)	6.00 (0.22)
FP	0 (0)	4.82 (1.23)	0 (0)	0 (0)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
$\hat{\sigma}$ ($\times 0.1$)	4.021 (0.020)	3.970 (0.023)	4.012 (0.020)	4.010 (0.020)	4.010 (0.020)
Setting III					
PE ($\times 0.1$)	1.626 (0.003)	1.629 (0.003)	1.626 (0.003)	1.625 (0.003)	1.625 (0.003)
L_2 -loss ($\times 0.1$)	0.676 (0.020)	0.692 (0.020)	0.673 (0.019)	0.661 (0.019)	0.665 (0.019)
L_1 -loss ($\times 0.1$)	1.505 (0.048)	1.713 (0.084)	1.489 (0.045)	1.469 (0.044)	1.473 (0.044)
L_∞ -loss ($\times 0.01$)	4.39 (0.13)	4.43 (0.13)	4.42 (0.13)	4.31 (0.13)	4.36 (0.13)
FP	0 (0)	3.67 (1.11)	0 (0)	0 (0)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
$\hat{\sigma}$ ($\times 0.1$)	4.009 (0.012)	3.993 (0.013)	4.008 (0.012)	4.007 (0.012)	4.006 (0.012)

NOTE: Settings I, II, and III refer to cases of $n = 100, 200,$ and $400,$ respectively.

Table 5. The means and standard errors (in parentheses) of various prediction and variable selection performance measures for all methods in Section 4.2.2

Measure	Method				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
Setting I					
PE ($\times 0.01$)	7.89 (0.03)	7.97 (0.07)	7.86 (0.03)	7.88 (0.04)	7.86 (0.03)
L_2 -loss	0.954 (0.039)	1.033 (0.096)	0.915 (0.052)	0.913 (0.051)	0.897 (0.049)
L_1 -loss	1.927 (0.087)	2.130 (0.271)	1.793 (0.108)	1.788 (0.107)	1.757 (0.103)
L_∞ -loss ($\times 0.1$)	6.354 (0.224)	6.936 (0.509)	6.346 (0.345)	6.348 (0.346)	6.238 (0.333)
FP	0.02 (0.01)	0.09 (0.05)	0 (0)	0.01 (0.01)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
Setting II					
PE ($\times 0.01$)	9.09 (0.07)	9.13 (0.07)	9.00 (0.05)	9.04 (0.06)	8.94 (0.03)
L_2 -loss	1.002 (0.059)	0.998 (0.072)	0.916 (0.059)	0.908 (0.055)	0.855 (0.049)
L_1 -loss	2.044 (0.135)	2.036 (0.168)	1.824 (0.129)	1.802 (0.120)	1.678 (0.103)
L_∞ -loss ($\times 0.1$)	6.549 (0.334)	6.574 (0.408)	6.213 (0.360)	6.154 (0.338)	5.926 (0.314)
FP	0.06 (0.02)	0.16 (0.05)	0.04 (0.02)	0.06 (0.03)	0 (0)
FN	0.01 (0.01)	0 (0)	0 (0)	0 (0)	0 (0)
Setting III					
PE ($\times 0.01$)	7.89 (0.04)	7.96 (0.07)	7.94 (0.07)	7.98 (0.08)	7.85 (0.04)
L_2 -loss	1.060 (0.053)	1.200 (0.089)	1.172 (0.083)	1.175 (0.084)	1.102 (0.079)
L_1 -loss	2.156 (0.123)	2.411 (0.199)	2.337 (0.181)	2.335 (0.182)	2.200 (0.172)
L_∞ -loss ($\times 0.1$)	6.935 (0.301)	8.157 (0.568)	7.997 (0.547)	8.091 (0.560)	7.495 (0.501)
FP	0.03 (0.02)	0.03 (0.02)	0.01 (0.01)	0.01 (0.01)	0 (0)
FN	0 (0)	0.02 (0.01)	0.02 (0.01)	0.03 (0.02)	0 (0)

NOTE: Settings I, II, and III refer to cases of $(p, r) = (1000, 0.25), (1000, 0.5),$ and $(5000, 0.25)$, respectively.

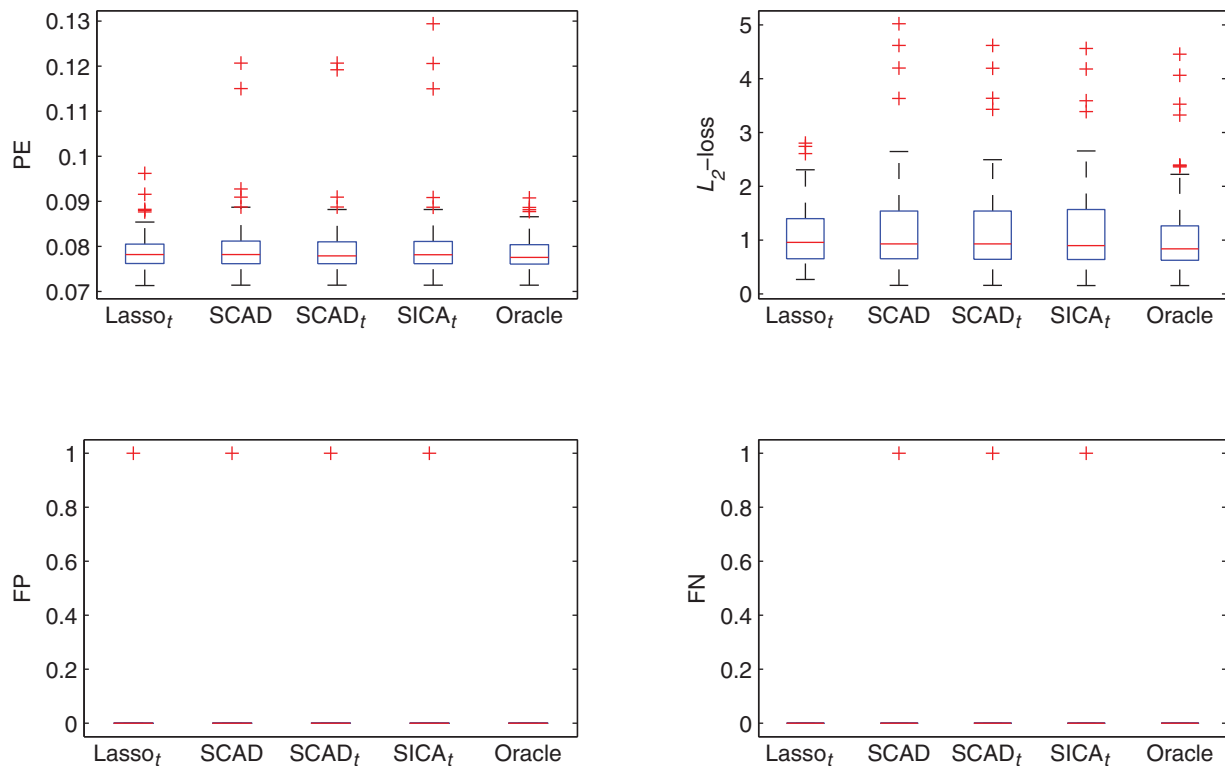


Figure 2. Boxplots of the PE, L_2 -loss, FP, and FN over 100 simulations for all methods in Section 4.2.2, with $p = 5000$. The x-axis represents different methods.

Table 6. Model selection consistency probabilities of all methods in Section 4.2.2

Setting of (p, r)	Model selection consistency probability				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
(1000, 0.25)	0.98	0.95	1	0.99	1
(1000, 0.5)	0.93	0.89	0.96	0.95	1
(5000, 0.25)	0.97	0.95	0.97	0.96	1

Table 8. Model selection consistency probabilities of all methods in Section 4.2.3

Setting of (p, r)	Model selection consistency probability				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
(1000, 0.25)	0.51	0	0.65	0.89	1
(1000, 0.5)	0.17	0	0.68	0.94	1
(5000, 0.25)	0.27	0	0.51	0.72	1

of all methods in such a case, we considered the 5% trimmed means, excluding 5% of values from each tail, and their standard errors of different prediction and variable selection measures. The asymptotic equivalence of concave regularization methods in the thresholded parameter space shown in the theory was also demonstrated in this nonlinear model for count data. But compared to linear models, the finite-sample performance of these methods differs more from that of the oracle procedure, indicating the increased difficulty of model inference for nonlinear models. The improvement of SCAD_t over original SCAD was more profound in this setting.

4.3 Real Data Example

We apply all the methods to the prostate cancer dataset, which was originally studied in Singh et al. (2002) and is available at <http://www.broadinstitute.org/cgi-bin/cancer/datasets.cgi>. This dataset, which was also analyzed in Fan and Fan (2008), consists of 136 patient samples with 77 from the prostate tumor group (labeled as 1) and 59 from the normal group (labeled as 0). For

each patient, we have the gene expression measurements for 12,600 genes.

Following Singh et al. (2002) and Fan and Fan (2008), we randomly split the 136 samples into a training set of 52 samples from the cancer class and 50 samples from the normal class, and a test set of 25 samples from the cancer class and 9 samples from the normal class. For each splitting of the dataset, we fit the logistic regression model to the training data with the regularization methods. We then calculated the classification error using the test data. We repeated the random splitting 50 times, and the means and standard errors of classification errors are summarized in Table 9. We also calculated the median model size by each method: 21 by Lasso_t, 5 by SCAD, 5 by SCAD_t, and 20 by SICA_t. For each method, we computed the percentage of times each gene was selected and listed the most frequently chosen m genes in Table 10, with m equal to the median model size by the method. We see that Lasso_t and SICA_t performed similarly, and SCAD and SCAD_t produced more sparse models than the other two methods.

Table 7. The 5% trimmed means and standard errors (in parentheses) of various prediction and variable selection performance measures for all methods in Section 4.2.3

Measure	Method				
	Lasso _t	SCAD	SCAD _t	SICA _t	Oracle
Setting I					
PE	21.34 (2.23)	13.11 (0.94)	9.00 (0.66)	7.39 (0.50)	6.22 (0.22)
L_2 -loss ($\times 0.01$)	19.62 (1.44)	17.09 (0.71)	11.58 (0.67)	9.05 (0.45)	7.94 (0.31)
L_1 -loss ($\times 0.1$)	4.658 (0.440)	4.714 (0.191)	2.513 (0.173)	1.720 (0.086)	1.513 (0.060)
L_∞ -loss ($\times 0.01$)	12.13 (0.75)	11.39 (0.58)	7.82 (0.43)	6.69 (0.37)	5.69 (0.23)
FP	1.47 (0.28)	11.61 (0.67)	1.20 (0.22)	0.07 (0.03)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
Setting II					
PE	5.934 (0.272)	3.288 (0.065)	2.903 (0.055)	2.754 (0.072)	2.655 (0.030)
L_2 -loss ($\times 0.01$)	38.34 (1.78)	19.12 (0.60)	14.75 (0.69)	12.52 (0.82)	11.54 (0.48)
L_1 -loss ($\times 0.1$)	10.07 (0.651)	5.860 (0.225)	3.254 (0.192)	2.357(0.143)	2.180 (0.097)
L_∞ -loss ($\times 0.01$)	20.96 (0.75)	12.21 (0.49)	9.92 (0.45)	9.15 (0.69)	8.39 (0.35)
FP	2.54 (0.23)	14.84 (0.84)	1.70 (0.34)	0 (0)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0.01 (0.01)	0 (0)
Setting III					
PE	34.86 (3.04)	14.26 (1.00)	10.38 (0.71)	8.28 (0.51)	6.15 (0.20)
L_2 -loss ($\times 0.01$)	31.67 (1.92)	18.42 (0.55)	14.41 (0.76)	11.62 (0.72)	8.35 (0.31)
L_1 -loss ($\times 0.1$)	8.158 (0.617)	5.863 (0.175)	3.600 (0.260)	2.252(0.173)	1.552 (0.062)
L_∞ -loss ($\times 0.01$)	18.29 (0.88)	12.10 (0.48)	9.46 (0.46)	8.48 (0.46)	6.21 (0.23)
FP	2.36 (0.28)	19.12 (0.71)	3.06 (0.49)	0.39 (0.08)	0 (0)
FN	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)

NOTE: Settings I, II, and III refer to cases of $(p, r) = (1000, 0.25), (1000, 0.5),$ and $(5000, 0.25),$ respectively.

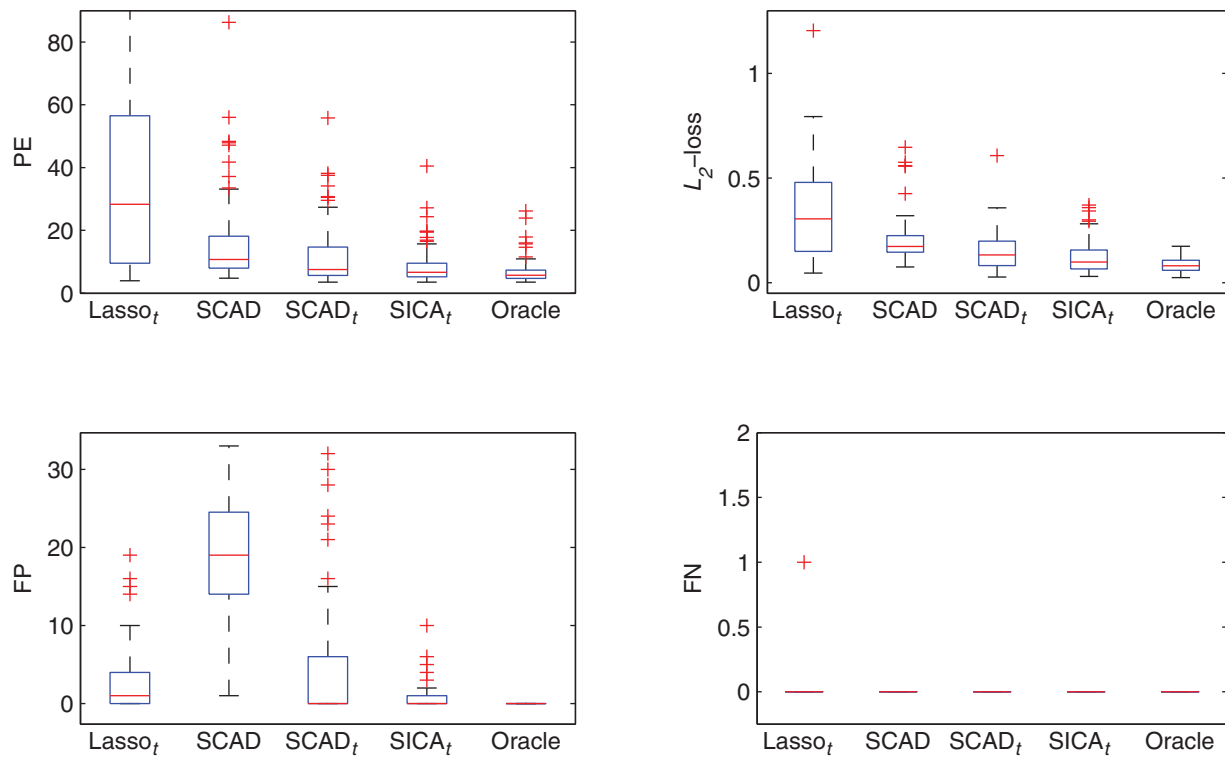


Figure 3. Boxplots of the PE, L_2 -loss, FP, and FN over 100 simulations for all methods in Section 4.2.3, with $p = 5000$. The x -axis represents different methods.

5. DISCUSSIONS

We have studied the asymptotic equivalence of two popular classes of regularization methods with convex penalties and concave penalties, in high-dimensional GLMs. Our framework

Table 9. The means and standard errors of classification errors by different methods over 50 random splittings of the prostate cancer data in Section 4.3

	Lasso _t	SCAD	SCAD _t	SICA _t
Mean	1.42	4.36	3.44	1.30
Standard error	0.20	0.36	0.32	0.18

covers many commonly used regularization methods such as the Lasso and concave ones such as the SCAD, MCP, and SICA. The oracle inequalities as well as the stronger property of the oracle risk inequalities of the global minimizer for the regularization methods have been established to characterize their connections and differences. When the Lasso penalty is considered, our oracle inequalities are consistent with those in Bickel, Ritov, and Tsybakov (2009), with improved sparsity thanks to the introduced thresholded parameter space. The established theoretical results have revealed an interesting phenomenon of phase transition in both linear and nonlinear models, confirmed by our numerical studies. We have also established additional theoretical results to provide insights into the sampling properties of computable solutions.

Table 10. Selection probabilities of most frequently selected genes with number up to median model size by each method across 50 random splittings of the prostate cancer data in Section 4.3

Gene ID	Lasso _t	SCAD	SCAD _t	SICA _t	Gene ID	Lasso _t	SCAD	SCAD _t	SICA _t
1018	—	—	—	0.44	7139	0.38	—	—	0.52
4525	0.96	—	—	0.94	7539	0.94	—	—	0.94
4636	0.42	—	—	0.40	8123	0.44	0.12	—	0.42
5319	0.54	—	—	0.64	9093	0.86	0.12	0.08	0.90
5661	0.68	—	0.12	0.64	9126	—	0.10	—	—
5890	1	0.10	0.28	1	10292	0.36	—	—	0.40
5977	0.58	—	—	0.44	10494	0.80	—	—	0.74
6145	0.28	—	—	0.26	10537	0.82	—	—	0.74
6185	0.94	0.10	0.10	0.98	11215	0.32	—	—	0.32
6390	0.28	—	—	—	11871	1	—	0.24	1
6462	0.36	—	—	—	12547	0.28	—	—	0.28
6512	0.48	—	—	0.46					

To simplify the technical presentation and better illustrate the ideas, we have focused on the setting of GLMs and the Lasso for the convex class of regularization methods. The theoretical results in the article may hold in more general model settings as well. The phase transition phenomenon may also be shown for other convex penalties such as the L_2 -penalty. These problems are beyond the scope of the current article and will be interesting topics for future research.

APPENDIX A. TECHNICAL DETAILS ON CONDITION 1

We show that the two probability bounds in Condition 1 hold for a wide class of error distributions. To this end, note that an application of the Bonferroni inequality gives

$$P(\|n^{-1}\mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty > \lambda/2) \leq \sum_{j=1}^p P(n^{-1}|\tilde{\mathbf{x}}_j^T \boldsymbol{\varepsilon}| > \lambda/2), \quad (\text{A.1})$$

where $(\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_p) = \mathbf{X}$. We consider two cases of error distribution.

Case 1 (Bounded error): Assume that $|\varepsilon_i| \leq a$ for each $1 \leq i \leq n$, with a being some positive constant. Then it follows from Hoeffding's inequality (Hoeffding 1963) that

$$P(n^{-1}|\tilde{\mathbf{x}}_j^T \boldsymbol{\varepsilon}| > \lambda/2) \leq 2 \exp(-\lambda^2 n / (8a^2)), \quad (\text{A.2})$$

since $\|\tilde{\mathbf{x}}_j\|_2 = n^{1/2}$ for each j . This probability bound is of order $O\{p^{-c_0^2/(8a^2)}\}$ when $\lambda = c_0\sqrt{(\log p)/n}$.

Case 2 (Light-tailed error): Assume that there exist positive constants M, v_0 such that

$$E[\exp(M^{-1}|\varepsilon_i|) - 1 - M^{-1}|\varepsilon_i|]M^2 \leq v_0/2 \quad (\text{A.3})$$

holds uniformly for $1 \leq i \leq n$. Then it follows from Bernstein's inequality (Bennett 1962; van der Vaart and Wellner 1996) that

$$P(n^{-1}|\tilde{\mathbf{x}}_j^T \boldsymbol{\varepsilon}| > \lambda/2) \leq 2 \exp\left(-\frac{\lambda^2 n}{8v_0 + 4\|\tilde{\mathbf{x}}_j\|_\infty M \lambda}\right), \quad (\text{A.4})$$

since $\|\tilde{\mathbf{x}}_j\|_2 = n^{1/2}$ for each j . This probability bound is of order $O\{p^{-c_0^2/(8v_0+d)}\}$ with $d = 4c_0M\|\tilde{\mathbf{x}}_j\|_\infty\sqrt{(\log p)/n}$ when $\lambda = c_0\sqrt{(\log p)/n}$, and this bound becomes $O\{p^{-c_0^2/(12v_0)}\}$ if we further assume $\|\tilde{\mathbf{x}}_j\|_\infty \leq (c_0M)^{-1}v_0\sqrt{n/(\log p)}$. This additional assumption means that the maximum absolute element of the design matrix \mathbf{X} is bounded from above by $(c_0M)^{-1}v_0\sqrt{n/(\log p)}$, which is a mild condition. Condition (A.3) was also made in Fan and Lv (2011) for analyzing nonconcave penalized likelihood estimators in GLM and is mild in view of the moment-generating function of distributions in the exponential family.

When c_0 is large enough, combining the above two cases with (A.1) leads to the desired probability bound on $P(\mathcal{E}^c)$ in Condition 1. Since $|\alpha_0| = s \leq n$ is assumed implicitly, similar probability bounds hold for the event \mathcal{E}_0^c . Thus we impose Condition 1 instead of making explicit assumptions on the model error distribution and design matrix \mathbf{X} .

APPENDIX B. PROOFS OF MAIN RESULTS

For notational simplicity, we use C to denote a generic positive constant, whose value may change from line to line. Denote by $\mathbf{b}'(\boldsymbol{\theta}) = (b'(\theta_1), \dots, b'(\theta_n))^T$ the n -vector of mean function, $\mathbf{b}''(\boldsymbol{\theta}) = (b''(\theta_1), \dots, b''(\theta_n))^T$ the n -vector of variance function for $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)^T \in \mathbb{R}^n$, and \mathbf{a}_α the subvector of a vector $\mathbf{a} \in \mathbb{R}^p$ formed by components with indices in a given set $\alpha \subset \{1, \dots, p\}$.

B.1 Proof of Proposition 1

Let $k_0 = c_4 n$ be an integer with $c_4 \in (0, 1)$ some constant. For each set $\alpha_1 \subset \{1, \dots, p\}$ with $|\alpha_1| = k_0$, denote by $\boldsymbol{\Sigma}_{\alpha_1, \alpha_1}$ the principal submatrix of $\boldsymbol{\Sigma}$ corresponding to variables in α_1 . We will show that there exist some universal positive constants c_5 and C_1 such that

$$P\{\lambda_{\min}(n^{-1}\mathbf{X}_{\alpha_1}^T \mathbf{X}_{\alpha_1}) < c_5\} \leq \exp(-C_1 n), \quad (\text{B.1})$$

where $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue of a matrix. Note that for any submatrix $n^{-1/2}\mathbf{X}_\alpha$ with $|\alpha| \leq k_0$, its smallest singular value is bounded from below by the smallest singular value of $n^{-1/2}\mathbf{X}_{\alpha_1}$ with $|\alpha_1| = k_0$ and $\alpha_1 \supset \alpha$. It follows that $\lambda_{\min}(n^{-1}\mathbf{X}_\alpha^T \mathbf{X}_\alpha)$ satisfies the same deviation probability bound (B.1). Thus, an application of the Bonferroni inequality with K an integer satisfying $K = 2^{-1}C_1 n / (\log p) \leq k_0$ gives

$$P\left\{\min_{|\alpha| \leq K} \lambda_{\min}(n^{-1}\mathbf{X}_\alpha^T \mathbf{X}_\alpha) < c_5\right\} \leq \sum_{|\alpha| \leq K} \exp(-C_1 n) \leq p^K \exp(-C_1 n) \rightarrow 0.$$

This shows that with asymptotic probability one, $\kappa_c \geq K$ for any $c \leq c_5$.

It remains to prove (B.1). Define $\tilde{\mathbf{X}}_{\alpha_1} = \mathbf{X}_{\alpha_1} \boldsymbol{\Sigma}_{\alpha_1, \alpha_1}^{-1/2}$. Then $\mathbf{X}_{\alpha_1}^T \mathbf{X}_{\alpha_1} = \boldsymbol{\Sigma}_{\alpha_1, \alpha_1}^{1/2} \tilde{\mathbf{X}}_{\alpha_1}^T \tilde{\mathbf{X}}_{\alpha_1} \boldsymbol{\Sigma}_{\alpha_1, \alpha_1}^{1/2}$ and the rows of $\tilde{\mathbf{X}}_{\alpha_1}$ are iid standard Gaussian random vectors. Since $\boldsymbol{\Sigma}$ has smallest eigenvalue bounded from below, we have

$$\begin{aligned} \lambda_{\min}(n^{-1}\mathbf{X}_{\alpha_1}^T \mathbf{X}_{\alpha_1}) &\geq \lambda_{\min}(n^{-1}\tilde{\mathbf{X}}_{\alpha_1}^T \tilde{\mathbf{X}}_{\alpha_1}) \lambda_{\min}(\boldsymbol{\Sigma}_{\alpha_1, \alpha_1}) \\ &\geq C \lambda_{\min}(n^{-1}\tilde{\mathbf{X}}_{\alpha_1}^T \tilde{\mathbf{X}}_{\alpha_1}). \end{aligned}$$

Therefore, we only need to show that $\lambda_{\min}(n^{-1}\tilde{\mathbf{X}}_{\alpha_1}^T \tilde{\mathbf{X}}_{\alpha_1})$ satisfies a similar deviation probability bound as (B.1), which is entailed by the concentration property proved in Fan and Lv (2008) [see their deviation inequality (16)]. This completes the proof.

B.2 Proof of Proposition 2

Since $\hat{\boldsymbol{\beta}} = (\hat{\beta}_1, \dots, \hat{\beta}_p)^T$ is the global minimizer of $Q_n(\boldsymbol{\beta})$, it holds that for each j , $\hat{\beta}_j$ is also the global minimizer of the same objective function along the j th coordinate, that is, $\hat{\beta}_j$ minimizes

$$\tilde{Q}_n(\beta_j) = a_0 - n^{-1} \mathbf{y}^T \tilde{\mathbf{x}}_j \beta_j + \sum_{i=1}^n b(a_i + x_{ij} \beta_j) + \lambda 1_{\{\beta_j \neq 0\}},$$

where $(\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_p) = \mathbf{X}$ with $\tilde{\mathbf{x}}_j = (x_{1j}, \dots, x_{nj})^T$ and a_i 's with $i = 0, 1, \dots, n$ are constants independent of β_j . Note that the first three terms of $\tilde{Q}_n(\beta_j)$ are continuous functions of β_j , while the last term is a step function of β_j . Thus it follows easily that the global minimizer $\hat{\beta}_j$ is either 0 or has magnitude larger than certain positive threshold whose value depends on λ and the continuous part of $\tilde{Q}_n(\beta_j)$, which concludes the proof.

B.3 Lemma 1 and Its Proof

We single out a lemma that is used in the proofs of Theorems 1–3.

Lemma 1. Under Conditions 2–3, we have

$$\|\boldsymbol{\delta}\|_2^2 \leq c^{-1} n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq C(\|n^{-1}\mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty + \lambda) \|\boldsymbol{\delta}\|_1, \quad (\text{B.2})$$

where $\boldsymbol{\delta} = \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0$ is the estimation error for the regularized estimator $\hat{\boldsymbol{\beta}}$ in (6), c is the positive constant in Definition 1, and C is some positive constant.

Proof of Lemma 1. Since $\widehat{\boldsymbol{\beta}}$ is the global minimizer of $Q_n(\boldsymbol{\beta})$ in \mathcal{B}_τ and $\boldsymbol{\beta}_0 \in \mathcal{B}_\tau$, it follows that

$$0 \leq Q_n(\boldsymbol{\beta}_0) - Q_n(\widehat{\boldsymbol{\beta}}) = \frac{1}{n} \{ \mathbf{y}^T \mathbf{X} \boldsymbol{\delta} - \mathbf{1}^T [\mathbf{b}(\mathbf{X}\widehat{\boldsymbol{\beta}}) - \mathbf{b}(\mathbf{X}\boldsymbol{\beta}_0)] \} + \|p_\lambda(\boldsymbol{\beta}_0)\|_1 - \|p_\lambda(\widehat{\boldsymbol{\beta}})\|_1. \quad (\text{B.3})$$

To analyze the nonlinear term $\mathbf{1}^T [\mathbf{b}(\mathbf{X}\widehat{\boldsymbol{\beta}}) - \mathbf{b}(\mathbf{X}\boldsymbol{\beta}_0)]$, we do a second-order Taylor expansion of the function $\mathbf{1}^T [\mathbf{b}(\mathbf{X}\boldsymbol{\beta}) - \mathbf{b}(\mathbf{X}\boldsymbol{\beta}_0)]$ around $\boldsymbol{\beta}_0$ and retain the Lagrange remainder term, which gives

$$\mathbf{1}^T [\mathbf{b}(\mathbf{X}\widehat{\boldsymbol{\beta}}) - \mathbf{b}(\mathbf{X}\boldsymbol{\beta}_0)] = \{\mathbf{b}'(\mathbf{X}\boldsymbol{\beta}_0)\}^T \mathbf{X} \boldsymbol{\delta} + \frac{1}{2} \boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\widetilde{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta}, \quad (\text{B.4})$$

where $\mathbf{H}(\widetilde{\boldsymbol{\beta}}) = \text{diag}\{\mathbf{b}''(\mathbf{X}\widetilde{\boldsymbol{\beta}})\}$ is a diagonal matrix with $\widetilde{\boldsymbol{\beta}} \in \mathbb{R}^p$ lying on the line segment connecting $\boldsymbol{\beta}_0$ and $\widehat{\boldsymbol{\beta}}$. Thus combining inequality (B.3) with representation (B.4) yields

$$0 \leq n^{-1} \left[\boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta} - \frac{1}{2} \boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\widetilde{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta} \right] + \|p_\lambda(\boldsymbol{\beta}_0)\|_1 - \|p_\lambda(\widehat{\boldsymbol{\beta}})\|_1, \quad (\text{B.5})$$

where $\boldsymbol{\varepsilon} = \mathbf{y} - E\mathbf{y} = \mathbf{y} - \mathbf{b}'(\mathbf{X}\boldsymbol{\beta}_0)$ denotes the n -dimensional error vector in the GLM. We observe that the first term on the right-hand side of (B.5) resembles the corresponding one in the case of linear model.

A rearrangement of inequality (B.5) gives

$$(2n)^{-1} \boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\widetilde{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta} \leq n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta} + \|p_\lambda(\boldsymbol{\beta}_0)\|_1 - \|p_\lambda(\widehat{\boldsymbol{\beta}})\|_1. \quad (\text{B.6})$$

It follows from $\boldsymbol{\beta}_0, \widehat{\boldsymbol{\beta}} \in \mathcal{B}_\tau$ that $\|\boldsymbol{\delta}\|_0 \leq \|\boldsymbol{\beta}_0\|_0 + \|\widehat{\boldsymbol{\beta}}\|_0 < \kappa_c$. Thus, by Condition 2 and the robust spark definition, the left-hand side of (B.6) can be bounded as

$$\begin{aligned} n^{-1} \boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\widetilde{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta} &\geq c_2 n^{-1} \|\mathbf{X} \boldsymbol{\delta}\|_2^2 \\ &= c_2 n^{-1} \|\mathbf{X}_{\text{supp}(\boldsymbol{\delta})} \boldsymbol{\delta}_{\text{supp}(\boldsymbol{\delta})}\|_2^2 \geq c c_2 \|\boldsymbol{\delta}\|_2^2. \end{aligned}$$

On the other hand, the first term on the right-hand side of (B.6) can be bounded as

$$|n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}| \leq \|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty \|\boldsymbol{\delta}\|_1.$$

The concavity of the penalty function $p_\lambda(t)$ assumed in Condition 3 entails that $p'_\lambda(t)$ is decreasing in t , which leads to $p'_\lambda(t) \leq p'_\lambda(0+) = c_3 \lambda$ for any $t \geq 0$. Thus, it follows from the mean value theorem and triangular inequality that

$$\| \|p_\lambda(\boldsymbol{\beta}_0)\|_1 - \|p_\lambda(\widehat{\boldsymbol{\beta}})\|_1 \| = \left| \sum_{j=1}^p p'_\lambda(t_j) (|\beta_{0,j}| - |\widehat{\beta}_j|) \right| \leq c_3 \lambda \|\boldsymbol{\delta}\|_1, \quad (\text{B.7})$$

where t_j lies between $|\beta_{0,j}|$ and $|\widehat{\beta}_j|$ for $j = 1, \dots, p$ and $\widehat{\boldsymbol{\beta}} = (\widehat{\beta}_1, \dots, \widehat{\beta}_p)^T$. Combining the above three results with (B.6) completes the proof.

B.4 Proof of Theorem 1

We first show the existence of the global minimizer and then prove the bounds under different losses.

Existence of global minimizer: Since the negative log-likelihood function is smooth by Condition 2 and the penalty function is continuous by Condition 3, we see that the objective function $Q_n(\boldsymbol{\beta})$ is continuous. Let R be any subspace of \mathbb{R}^p with dimension less than $\kappa_c/2$, and denote by $L(\boldsymbol{\beta}) = -n^{-1} \{\mathbf{y}^T \mathbf{X} \boldsymbol{\beta} - \mathbf{1}^T \mathbf{b}(\mathbf{X}\boldsymbol{\beta})\}$ the negative log-likelihood function. It follows from Condition 2 and the definition of the robust spark that $L(\boldsymbol{\beta})$ is strictly convex on R and its Hessian matrix has the smallest eigenvalue bounded from below by $c_2 c^2$. Applying the second-order Taylor expansion around $\mathbf{0}$ with the Lagrange remainder term shows that $L(\boldsymbol{\beta})$ is bounded from below by $\widetilde{L}(\boldsymbol{\beta}) = -n^{-1} \mathbf{y}^T \mathbf{X} \boldsymbol{\beta} + b(0) + n^{-1} \{\mathbf{b}'(\mathbf{0})\}^T \mathbf{X} \boldsymbol{\beta} + c_2 c^2 \|\boldsymbol{\beta}\|_2^2$ for any $\boldsymbol{\beta} \in R$, with $L(\mathbf{0}) = \widetilde{L}(\mathbf{0}) = b(0)$. This entails that there exists some

sufficiently large positive number C , which is independent of the subspace R , such that

$$L(\boldsymbol{\beta}) \geq \widetilde{L}(\boldsymbol{\beta}) > \widetilde{L}(\mathbf{0}) = L(\mathbf{0})$$

for any $\boldsymbol{\beta} \in R$ with $\|\boldsymbol{\beta}\|_2 > C$. Thus the global minimizer of $Q_n(\boldsymbol{\beta}) = L(\boldsymbol{\beta}) + \|p_\lambda(\boldsymbol{\beta})\|_1$ on the thresholded parameter space \mathcal{B}_τ must lie in $T = \mathcal{B}_\tau \cap \{\boldsymbol{\beta} \in \mathbb{R}^p : \|\boldsymbol{\beta}\|_2 \leq C\}$. In view of (4), \mathcal{B}_τ is a closed set and thus T is a compact set. Therefore, the existence of the global minimizer of $Q_n(\boldsymbol{\beta})$ over \mathcal{B}_τ is guaranteed by its continuity.

False signs: We use the induction method to prove the result. Let $\boldsymbol{\delta} = (\delta_1, \dots, \delta_p)^T = \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0$. Since $\|\boldsymbol{\delta}\|_0 \leq \|\boldsymbol{\beta}_0\|_0 + \|\widehat{\boldsymbol{\beta}}\|_0 < \kappa_c$, by the Cauchy–Schwarz inequality, we have $\|\boldsymbol{\delta}\|_1 \leq \sqrt{\kappa_c} \|\boldsymbol{\delta}\|_2$. Hence, it follows from Lemma 1 that conditional on the event \mathcal{E} ,

$$\|\boldsymbol{\delta}\|_2^2 \leq C \lambda \|\boldsymbol{\delta}\|_1 \leq C \lambda \sqrt{\kappa_c} \|\boldsymbol{\delta}\|_2. \quad (\text{B.8})$$

Solving for $\|\boldsymbol{\delta}\|_2$ yields

$$\|\boldsymbol{\delta}\|_2 \leq C \lambda \sqrt{\kappa_c}. \quad (\text{B.9})$$

On the other hand, it follows from $\widehat{\boldsymbol{\beta}}, \boldsymbol{\beta}_0 \in \mathcal{B}_\tau$ that $\|\boldsymbol{\delta}\|_2 \geq \{\text{FS}(\widehat{\boldsymbol{\beta}})\}^{1/2} \tau$. This together with (B.9) ensures that

$$\text{FS}(\widehat{\boldsymbol{\beta}}) \leq C(\lambda/\tau)^2 \kappa_c.$$

Thus we have $\|\boldsymbol{\delta}\|_0 \leq \|\boldsymbol{\beta}_0\|_0 + \text{FS}(\widehat{\boldsymbol{\beta}}) \leq s + C(\lambda/\tau)^2 \kappa_c$. Therefore, the upper bound $s + C(\lambda/\tau)^2 \kappa_c$ plays the same role as κ_c in (B.9). Repeating the above derivations with κ_c replaced with $s + C(\lambda/\tau)^2 \kappa_c$ and by induction, we have $\text{FS}(\widehat{\boldsymbol{\beta}}) \leq C s \lambda^2 \tau^{-2} / (1 - C \lambda^2 \tau^{-2})$ conditional on \mathcal{E} , which completes the proof of the result on false signs.

Estimation losses: We first prove the inequalities under the L_1 - and L_2 -norms, and then use Hölder's inequality to prove the general result under the L_q -norm with $q \in (1, 2)$. The result on L_∞ -norm follows immediately from the L_2 -norm result. By default, all arguments are conditioning on \mathcal{E} in Condition 1, which holds with probability at least $1 - O(p^{-c_1})$.

Since $\|\boldsymbol{\delta}\|_0 \leq \|\boldsymbol{\beta}_0\|_0 + \text{FS}(\widehat{\boldsymbol{\beta}})$, by the Cauchy–Schwarz inequality, and the result on $\text{FS}(\widehat{\boldsymbol{\beta}})$ proved above, we have

$$\|\boldsymbol{\delta}\|_1 \leq \|\boldsymbol{\delta}\|_0^{1/2} \|\boldsymbol{\delta}\|_2 \leq C \{s/(1 - C \lambda^2 \tau^{-2})\}^{1/2} \|\boldsymbol{\delta}\|_2.$$

This together with the first inequality in (B.8) entails that

$$\|\boldsymbol{\delta}\|_2 \leq C \lambda \{s/(1 - C \lambda^2 \tau^{-2})\}^{1/2}. \quad (\text{B.10})$$

Combining the above two inequalities, we obtain

$$\|\boldsymbol{\delta}\|_1 \leq C \lambda s / (1 - C \lambda^2 \tau^{-2}). \quad (\text{B.11})$$

Finally, for $q \in (1, 2)$, applying Hölder's inequality and in view of (B.10) and (B.11), we have

$$\begin{aligned} \|\boldsymbol{\delta}\|_q &= \left(\sum_{j=1}^p |\delta_j|^{2-q} |\delta_j|^{2q-2} \right)^{1/q} \leq \|\boldsymbol{\delta}\|_1^{(2-q)/q} \|\boldsymbol{\delta}\|_2^{2(q-1)/q} \\ &\leq C \lambda \{s/(1 - C \lambda^2 \tau^{-2})\}^{1/q}. \end{aligned} \quad (\text{B.12})$$

The oracle inequality on $\|\boldsymbol{\delta}\|_\infty$ follows immediately from $\|\boldsymbol{\delta}\|_\infty \leq \|\boldsymbol{\delta}\|_2$ and (B.10). This completes the proof for the estimation losses.

Prediction loss: The inequality for this loss follows immediately from plugging (B.11) into Lemma 1 and using Condition 1, which concludes the proof.

B.5 Proof of Theorem 2

Define an event $\mathcal{E}_1 = \mathcal{E} \cap \mathcal{E}_0$ with \mathcal{E} and \mathcal{E}_0 defined in (5). We will prove that all results in Theorem 2 hold simultaneously on the event \mathcal{E}_1 .

Then Theorem 2 follows immediately from Condition 1. By default, all arguments in this proof are conditioning on \mathcal{E}_1 .

Sign consistency: Denote by $\alpha = \text{supp}(\widehat{\boldsymbol{\beta}})$ and $\alpha_0 = \text{supp}(\boldsymbol{\beta}_0)$. We use the method of proof by contradiction to show that we must have $\alpha = \alpha_0$. Let $\widehat{\boldsymbol{\beta}}^*$ be the oracle-assisted maximum likelihood estimator. We make use of the following decomposition:

$$Q_n(\widehat{\boldsymbol{\beta}}) - Q_n(\widehat{\boldsymbol{\beta}}^*) = I_1 + I_2,$$

where $I_1 = -n^{-1}\mathbf{y}^T\mathbf{X}(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^*) + n^{-1}\mathbf{1}^T[\mathbf{b}(\mathbf{X}\widehat{\boldsymbol{\beta}}) - \mathbf{b}(\mathbf{X}\widehat{\boldsymbol{\beta}}^*)]$ and $I_2 = \|p_\lambda(\widehat{\boldsymbol{\beta}})\|_1 - \|p_\lambda(\widehat{\boldsymbol{\beta}}^*)\|_1$. We will prove that $\widehat{\boldsymbol{\beta}}^* \in \mathcal{B}_\tau$ and that if $\alpha \neq \alpha_0$, then

$$I_1 \geq cc_2\tau^2/4, \quad (\text{B.13})$$

$$|I_2| \leq o(\tau^2). \quad (\text{B.14})$$

Combining the above results, we have $Q_n(\widehat{\boldsymbol{\beta}}) - Q_n(\widehat{\boldsymbol{\beta}}^*) > 0$ for sufficiently large n , which contradicts with $\widehat{\boldsymbol{\beta}}$ being a global minimizer in \mathcal{B}_τ , and thus we must have $\alpha = \alpha_0$. On the other hand, since $\lambda = o(\tau/\sqrt{s})$, Theorem 1 ensures that for large enough n ,

$$\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_\infty \leq C\sqrt{s}\lambda = o(\tau). \quad (\text{B.15})$$

This together with $\alpha = \alpha_0$ and $\boldsymbol{\beta}_0, \widehat{\boldsymbol{\beta}} \in \mathcal{B}_\tau$ entails that conditioning on \mathcal{E}_1 , $\text{sgn}(\widehat{\boldsymbol{\beta}}) = \text{sgn}(\boldsymbol{\beta}_0)$. Thus, the sign consistency result follows easily from Condition 1.

We first proceed to prove (B.14). By definition, $\text{supp}(\widehat{\boldsymbol{\beta}}^*) = \alpha_0$ and $\widehat{\boldsymbol{\beta}}_{\alpha_0}^* = (\widehat{\beta}_{\alpha_0,1}^*, \dots, \widehat{\beta}_{\alpha_0,s}^*)^T$ minimizes the negative log-likelihood function $Q_n^*(\boldsymbol{\beta}_{\alpha_0}) = -\mathbf{y}^T\mathbf{X}_{\alpha_0}\boldsymbol{\beta}_{\alpha_0} + \mathbf{1}^T\mathbf{b}(\mathbf{X}_{\alpha_0}\boldsymbol{\beta}_{\alpha_0})$ with $\boldsymbol{\beta}_{\alpha_0} \in \mathbb{R}^s$. Thus, $\widehat{\boldsymbol{\beta}}_{\alpha_0}^*$ is a critical point of $Q_n^*(\boldsymbol{\beta}_{\alpha_0})$ and satisfies

$$-\mathbf{X}_{\alpha_0}^T[\mathbf{y} - \mathbf{b}'(\mathbf{X}_{\alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha_0}^*)] = \mathbf{0}. \quad (\text{B.16})$$

Plugging the true model $\mathbf{y} = \mathbf{b}'(\mathbf{X}_{\alpha_0}\boldsymbol{\beta}_{0,\alpha_0}) + \boldsymbol{\varepsilon}$ into (B.16) and applying the mean value theorem componentwise, we have

$$-\mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon} + \mathbf{X}_{\alpha_0}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}}_1, \dots, \widetilde{\boldsymbol{\beta}}_n)\mathbf{X}_{\alpha_0}(\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0}) = \mathbf{0}, \quad (\text{B.17})$$

where $\mathbf{H}(\widetilde{\boldsymbol{\beta}}_1, \dots, \widetilde{\boldsymbol{\beta}}_n) = \text{diag}\{b''(\mathbf{x}_1^T\widetilde{\boldsymbol{\beta}}_1), \dots, b''(\mathbf{x}_n^T\widetilde{\boldsymbol{\beta}}_n)\}$ with each $\widetilde{\boldsymbol{\beta}}_i = (\widetilde{\beta}_{i,1}, \dots, \widetilde{\beta}_{i,p})^T$ lying on the line segment connecting $\boldsymbol{\beta}_0$ and $\widehat{\boldsymbol{\beta}}^*$. The above equation can be rewritten as

$$\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0} = \{\mathbf{X}_{\alpha_0}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}}_1, \dots, \widetilde{\boldsymbol{\beta}}_n)\mathbf{X}_{\alpha_0}\}^{-1}\mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon}. \quad (\text{B.18})$$

Therefore, by Condition 2, we obtain that conditioning on \mathcal{E} ,

$$\|\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0}\|_2 \leq C\|\mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon}\|_2/n \leq C\sqrt{s}\|\mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon}\|_\infty/n \leq C\sqrt{s}\lambda. \quad (\text{B.19})$$

This together with the assumptions $\min_{j \leq s} |\beta_{0,j}| \geq 2\tau$ and $\sqrt{s}\lambda = o(\tau)$ entails that

$$\min_{1 \leq j \leq s} |\widehat{\beta}_j^*| \geq 2\tau - C\sqrt{s}\lambda > \tau \text{ and thus } \widehat{\boldsymbol{\beta}}^* \in \mathcal{B}_\tau.$$

Similarly to (B.7) and by Theorem 1 and (B.19), we can prove

$$|I_2| \leq p'_\lambda(0+)\|\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^*\|_1 \leq c_3\lambda(\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_1 + \|\widehat{\boldsymbol{\beta}}^* - \boldsymbol{\beta}_0\|_1) \leq c_3\lambda(\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_1 + \sqrt{s}\|\widehat{\boldsymbol{\beta}}^* - \boldsymbol{\beta}_0\|_2) \leq Cs\lambda^2 = o(\tau^2). \quad (\text{B.20})$$

This completes the proof of (B.14) and $\widehat{\boldsymbol{\beta}}^* \in \mathcal{B}_\tau$.

It remains to prove (B.13). Applying the second-order Taylor expansion around $\widehat{\boldsymbol{\beta}}^*$ with the Lagrange remainder term, I_1 can be decomposed as

$$I_1 = -\frac{1}{n}[\mathbf{y} - \mathbf{b}'(\mathbf{X}\widehat{\boldsymbol{\beta}}^*)]^T\mathbf{X}(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^*) + \frac{1}{2n}(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^*)^T\mathbf{X}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}}^*)\mathbf{X}(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^*) \equiv I_{1,1} + I_{1,2}, \quad (\text{B.21})$$

where $\mathbf{H}(\widetilde{\boldsymbol{\beta}}^*) = \text{diag}\{b''(\mathbf{x}^T\widetilde{\boldsymbol{\beta}}^*)\}$ with $\widetilde{\boldsymbol{\beta}}^*$ lying on the line segment connecting $\widehat{\boldsymbol{\beta}}^*$ and $\widehat{\boldsymbol{\beta}}$. It follows from $\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\beta}}^* \in \mathcal{B}_\tau$ that $\|\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^*\|_0 < \kappa_c$.

Thus, by Condition 2, the robust spark definition, and $\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\beta}}^* \in \mathcal{B}_\tau$, we have

$$I_{1,2} \geq \frac{1}{2}c_2c\tau\|\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^*\|_2^2 \geq \frac{1}{2}c_2c\tau(\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1 + \|\widehat{\boldsymbol{\beta}}_{\alpha_0 \setminus \alpha}^*\|_1). \quad (\text{B.22})$$

We now consider the term $I_{1,1}$ in (B.21). By (B.16), we have

$$I_{1,1} = -n^{-1}[\mathbf{y} - \mathbf{b}'(\mathbf{X}_{\alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha_0}^*)]^T\mathbf{X}_{\alpha_0}(\widehat{\boldsymbol{\beta}}_{\alpha_0} - \widehat{\boldsymbol{\beta}}_{\alpha_0}^*) - n^{-1}[\mathbf{y} - \mathbf{b}'(\mathbf{X}\widehat{\boldsymbol{\beta}}^*)]^T\mathbf{X}_{\alpha \setminus \alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0} = -n^{-1}[\mathbf{y} - \mathbf{b}'(\mathbf{X}\widehat{\boldsymbol{\beta}}^*)]^T\mathbf{X}_{\alpha \setminus \alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}. \quad (\text{B.23})$$

Plugging $\mathbf{y} = \mathbf{b}'(\mathbf{X}\boldsymbol{\beta}_0) + \boldsymbol{\varepsilon}$ into (B.23) and by the mean value theorem, we have

$$I_{1,1} = -n^{-1}\boldsymbol{\varepsilon}^T\mathbf{X}_{\alpha \setminus \alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0} + n^{-1}(\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0})^T\mathbf{X}_{\alpha_0}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}})\mathbf{X}_{\alpha \setminus \alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}, \quad (\text{B.24})$$

where $\mathbf{H}(\widetilde{\boldsymbol{\beta}}) = \text{diag}\{b''(\mathbf{x}^T\widetilde{\boldsymbol{\beta}})\}$ with $\widetilde{\boldsymbol{\beta}}$ lying on the line segment connecting $\boldsymbol{\beta}_0$ and $\widehat{\boldsymbol{\beta}}^*$. Conditioning on \mathcal{E} , the first term of (B.24) can be bounded as

$$n^{-1}|\boldsymbol{\varepsilon}^T\mathbf{X}_{\alpha \setminus \alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}| \leq \|n^{-1}\mathbf{X}_{\alpha \setminus \alpha_0}^T\boldsymbol{\varepsilon}\|_\infty\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1 \leq \lambda\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1. \quad (\text{B.25})$$

Next we study the second term of (B.24). We will make use of (B.18). By the Cauchy-Schwarz inequality, Condition 2, and the robust spark definition, we have

$$\begin{aligned} \|\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0}\|_\infty &\leq \|\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0}\|_2 \\ &\leq \|\{n^{-1}\mathbf{X}_{\alpha_0}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}}_1, \dots, \widetilde{\boldsymbol{\beta}}_n)\mathbf{X}_{\alpha_0}\}^{-1}\|_2\|n^{-1}\mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon}\|_2 \\ &\leq C\sqrt{s}\|n^{-1}\mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon}\|_\infty \leq C\sqrt{s(\log n)/n}. \end{aligned} \quad (\text{B.26})$$

Recall that $\widetilde{\boldsymbol{\beta}}$ defined in (B.24) lies on the line segment connecting $\boldsymbol{\beta}_0$ and $\widehat{\boldsymbol{\beta}}^*$. Thus, by (B.19), we have $\|\widetilde{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_2 \leq \|\widehat{\boldsymbol{\beta}}^* - \boldsymbol{\beta}_0\|_2 \leq C\sqrt{s}\lambda$, which ensures that $\widetilde{\boldsymbol{\beta}} \in \mathcal{B}_1^*$ with \mathcal{B}_1^* defined in (8). Since Theorem 1 ensures that $|\alpha \setminus \alpha_0| \leq \text{FS}(\widehat{\boldsymbol{\beta}}) \leq s$ for large enough n , it follows from inequality (B.26) that the second term of (B.24) can be bounded as

$$\begin{aligned} n^{-1}|\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0})^T\mathbf{X}_{\alpha_0}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}})\mathbf{X}_{\alpha \setminus \alpha_0}\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}| \\ \leq \|n^{-1}\mathbf{X}_{\alpha \setminus \alpha_0}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}})\mathbf{X}_{\alpha_0}\|_\infty\|\widehat{\boldsymbol{\beta}}_{\alpha_0}^* - \boldsymbol{\beta}_{0,\alpha_0}\|_\infty\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1 \\ \leq C\gamma_n\sqrt{s(\log n)/n}\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1. \end{aligned} \quad (\text{B.27})$$

Combining (B.25) and (B.27) and in view of (B.24), we obtain that

$$|I_{1,1}| \leq [\lambda + C\gamma_n\sqrt{s(\log n)/n}]\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1.$$

This together with (B.21), (B.22), $\widehat{\boldsymbol{\beta}}^*, \widehat{\boldsymbol{\beta}} \in \mathcal{B}_\tau$, and the assumption $\tau \gg \max\{\lambda, \gamma_n\sqrt{s(\log n)/n}\}$ ensures that if $\alpha \neq \alpha_0$, then for large enough n ,

$$I_1 \geq \frac{1}{2}c_2c\tau\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1 + \left[\frac{1}{2}c_2c\tau - \lambda - C\gamma_n\sqrt{s(\log n)/n}\right]\|\widehat{\boldsymbol{\beta}}_{\alpha \setminus \alpha_0}\|_1 \geq cc_2\tau^2/4,$$

which proves (B.13) and completes the proof of sign consistency.

Estimation losses: We first prove for the L_∞ -estimation loss. By (B.15) and the sign consistency proved above, we have $\text{supp}(\widehat{\boldsymbol{\beta}}) = \alpha_0$ and $\min_{1 \leq j \leq s} |\widehat{\beta}_j| \geq \min_{1 \leq j \leq s} |\beta_{0,j}| - o(\tau) > \tau$ for large enough n . Thus, $\widehat{\boldsymbol{\beta}}$ is an interior point of \mathcal{B}_τ . Since $\widehat{\boldsymbol{\beta}}$ is the global minimizer, it follows that $\widehat{\boldsymbol{\beta}}_{\alpha_0}$ is a critical point of $Q_n(\boldsymbol{\beta}_{\alpha_0}, \mathbf{0})$ and satisfies

$$-\mathbf{X}_{\alpha_0}^T\mathbf{y} + \mathbf{X}_{\alpha_0}^T\mathbf{b}'(\mathbf{X}\widehat{\boldsymbol{\beta}}) + n\bar{p}_\lambda(\widehat{\boldsymbol{\beta}}_{\alpha_0}) = \mathbf{0}, \quad (\text{B.28})$$

where $\bar{p}_\lambda(\widehat{\boldsymbol{\beta}}_{\alpha_0})$ is an s -dimensional vector with components $p'_\lambda(|\widehat{\beta}_j|)\text{sgn}(\widehat{\beta}_j)$ for $j \in \alpha_0$. Similarly to (B.17), plugging $\mathbf{y} = \mathbf{b}'(\mathbf{X}\boldsymbol{\beta}_0) + \boldsymbol{\varepsilon}$ into (B.28) and applying the mean value theorem componentwise, we have

$$n\bar{p}_\lambda(\widehat{\boldsymbol{\beta}}_{\alpha_0}) = -\mathbf{X}_{\alpha_0}^T\mathbf{H}(\widetilde{\boldsymbol{\beta}}_1, \dots, \widetilde{\boldsymbol{\beta}}_n)\mathbf{X}_{\alpha_0}(\widehat{\boldsymbol{\beta}}_{\alpha_0} - \boldsymbol{\beta}_{0,\alpha_0}) + \mathbf{X}_{\alpha_0}^T\boldsymbol{\varepsilon}, \quad (\text{B.29})$$

where $\mathbf{H}(\tilde{\boldsymbol{\beta}}_1, \dots, \tilde{\boldsymbol{\beta}}_n)$ is defined similarly as in (B.17) with each $\tilde{\boldsymbol{\beta}}_i$ lying on the line segment connecting $\boldsymbol{\beta}_0$ and $\hat{\boldsymbol{\beta}}$. Thus, for each $1 \leq i \leq n$, $\text{supp}(\tilde{\boldsymbol{\beta}}_i) = \alpha_0$, and by Theorem 1, $\|\tilde{\boldsymbol{\beta}}_i - \boldsymbol{\beta}_0\|_2 \leq \|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_2 \leq 2C\lambda\sqrt{s}$ for large enough n , which ensures that $\tilde{\boldsymbol{\beta}}_i \in \mathcal{B}_1^*$ with \mathcal{B}_1^* defined in (8). Therefore, (B.29) can be rewritten as

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{\alpha_0} - \boldsymbol{\beta}_{0,\alpha_0} &= [\mathbf{X}_{\alpha_0}^T \mathbf{H}(\tilde{\boldsymbol{\beta}}_1, \dots, \tilde{\boldsymbol{\beta}}_n) \mathbf{X}_{\alpha_0}]^{-1} \mathbf{X}_{\alpha_0}^T \boldsymbol{\varepsilon} \\ &\quad - n [\mathbf{X}_{\alpha_0}^T \mathbf{H}(\tilde{\boldsymbol{\beta}}_1, \dots, \tilde{\boldsymbol{\beta}}_n) \mathbf{X}_{\alpha_0}]^{-1} \bar{p}_\lambda(\hat{\boldsymbol{\beta}}_{\alpha_0}) \\ &\equiv I_1(\alpha_0) + I_2(\alpha_0). \end{aligned} \quad (\text{B.30})$$

We first study $I_1(\alpha_0)$. Since $\tilde{\boldsymbol{\beta}}_i \in \mathcal{B}_1^*$, by the Cauchy–Schwarz inequality,

$$\begin{aligned} \|I_1(\alpha_0)\|_\infty &\leq \left\| [n^{-1} \mathbf{X}_{\alpha_0}^T \mathbf{H}(\tilde{\boldsymbol{\beta}}_1, \dots, \tilde{\boldsymbol{\beta}}_n) \mathbf{X}_{\alpha_0}]^{-1} \right\|_\infty \|n^{-1} \mathbf{X}_{\alpha_0}^T \boldsymbol{\varepsilon}\|_\infty \\ &\leq C\gamma_n^* \sqrt{(\log n)/n}. \end{aligned} \quad (\text{B.31})$$

Next we study $I_2(\alpha_0)$. Similarly, since $p_\lambda(t)$ is a concave penalty, $\text{supp}(\hat{\boldsymbol{\beta}}) = \alpha_0$, and $\hat{\boldsymbol{\beta}} \in \mathcal{B}_\tau$, we can prove that

$$\begin{aligned} \|I_2(\alpha_0)\|_\infty &\leq \left\| [n^{-1} \mathbf{X}_{\alpha_0}^T \mathbf{H}(\tilde{\boldsymbol{\beta}}_1, \dots, \tilde{\boldsymbol{\beta}}_n) \mathbf{X}_{\alpha_0}]^{-1} \right\|_\infty \|\bar{p}_\lambda(\hat{\boldsymbol{\beta}}_{\alpha_0})\|_\infty \\ &\leq \gamma_n^* p'_\lambda(\tau). \end{aligned} \quad (\text{B.32})$$

Therefore, if $p'_\lambda(\tau) = O(\sqrt{(\log n)/n})$, then conditioning on \mathcal{E}_1 ,

$$\|\hat{\boldsymbol{\beta}}_{\alpha_0} - \boldsymbol{\beta}_{0,\alpha_0}\|_\infty \leq \|I_1(\alpha_0)\|_\infty + \|I_2(\alpha_0)\|_\infty \leq C\gamma_n^* \sqrt{(\log n)/n},$$

which completes the proof of the oracle inequality under the L_∞ -estimation loss.

We now study the L_q -estimation loss with $q \in [1, 2]$. Similarly as in Theorem 1, we first prove results under the L_1 - and L_2 -norms, and then use Hölder's inequality to prove the general results. Since $\text{supp}(\hat{\boldsymbol{\beta}}) = \alpha_0$, $p'_\lambda(\tau) = O(\sqrt{(\log n)/n})$, and $p'_\lambda(t)$ is decreasing in $t \in (0, \infty)$, inequality (B.7) in the proof of Lemma 1 can be bounded as

$$\begin{aligned} \left| \|p_\lambda(\boldsymbol{\beta}_0)\|_1 - \|p_\lambda(\hat{\boldsymbol{\beta}})\|_1 \right| &= \left| \sum_{j \in \alpha_0} p'_\lambda(t_j) (|\beta_{0,j}| - |\hat{\beta}_j|) \right| \\ &\leq p'_\lambda(\tau) \|\boldsymbol{\delta}\|_1 \leq C\sqrt{(\log n)/n} \|\boldsymbol{\delta}\|_1, \end{aligned}$$

where the second step is because t_j is between $|\beta_{0,j}|$ and $|\hat{\beta}_j|$ and thus $t_j \geq \tau$ for each $j \in \alpha_0$. Using similar proof as in Lemma 1 and the above inequality, we obtain that conditioning on \mathcal{E}_1 ,

$$\|\boldsymbol{\delta}\|_2^2 \leq c^{-1} n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq C\sqrt{(\log n)/n} \|\boldsymbol{\delta}\|_1. \quad (\text{B.33})$$

Since the sign consistency proved above ensures $\|\boldsymbol{\delta}\|_1 \leq \sqrt{s} \|\boldsymbol{\delta}\|_2$, it follows from (B.33) that

$$\begin{aligned} \|\boldsymbol{\delta}\|_2 &\leq \sqrt{s} \|\boldsymbol{\delta}\|_2 / \|\boldsymbol{\delta}\|_1 \leq C\sqrt{s(\log n)/n} \quad \text{and} \\ \|\boldsymbol{\delta}\|_1 &\leq \sqrt{s} \|\boldsymbol{\delta}\|_2 \leq Cs\sqrt{(\log n)/n}. \end{aligned} \quad (\text{B.34})$$

The oracle inequalities under the L_q -estimation loss with $q \in (1, 2)$ follow immediately from Hölder's inequality and (B.34), as in (B.12). Thus, the results on estimation losses are proved.

Prediction loss: Since $E\mathbf{Y} = \mathbf{b}'(\mathbf{X}\boldsymbol{\beta}_0)$, it follows from the second-order Taylor expansion around $\boldsymbol{\beta}_0$ with the Lagrange remainder term that

$$D(\hat{\boldsymbol{\beta}}) = \frac{1}{2} \boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\tilde{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta}, \quad (\text{B.35})$$

where $\mathbf{H}(\tilde{\boldsymbol{\beta}}) = \text{diag}(\mathbf{b}''(\mathbf{X}\tilde{\boldsymbol{\beta}}))$ with $\tilde{\boldsymbol{\beta}}$ lying on the line segment connecting $\boldsymbol{\beta}_0$ and $\hat{\boldsymbol{\beta}}$. Since $\text{supp}(\hat{\boldsymbol{\beta}}) = \alpha_0$, it follows from Condition 2, (B.33), and (B.34) that

$$\boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\tilde{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta} \leq c_2^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq C\sqrt{n(\log n)} \|\boldsymbol{\delta}\|_1 \leq Cs(\log n).$$

Thus, combining the above inequality with (B.35) completes the proof.

B.6 Proof of Theorem 3

Define $\mathcal{E}_1 = \mathcal{E} \cap \mathcal{E}_0$ as in the proof of Theorem 2. Then all results in Theorem 2 hold simultaneously on the event \mathcal{E}_1 , which satisfies $P(\mathcal{E}_1^c) = O(n^{-c_1})$ by Condition 1. Denote by $\boldsymbol{\delta} = \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0$.

Estimation risks: Similarly as in Theorem 2, we first prove the results under the L_2 - and L_1 -losses, and then use Hölder's inequality to prove the general result under the L_q -loss with $q \in (1, 2)$. We first show that

$$E \|\boldsymbol{\delta}\|_2^2 \leq Cs(\log n)/n. \quad (\text{B.36})$$

The key is to prove the following three inequalities:

$$E \{ \|\boldsymbol{\delta}\|_2^2 1_{\mathcal{E}_1} \} \leq Cs(\log n)/n, \quad (\text{B.37})$$

$$E \{ \|\boldsymbol{\delta}\|_2^2 1_{\mathcal{E}^c} \} = O(p^{-c_1/2} \kappa_c), \quad (\text{B.38})$$

$$E \{ \|\boldsymbol{\delta}\|_2^2 1_{\mathcal{E}_0^c \cap \mathcal{E}} \} = O(\lambda s n^{-c_1}). \quad (\text{B.39})$$

Since $\mathcal{E}_1^c = \mathcal{E}^c \cup (\mathcal{E}_0^c \cap \mathcal{E})$, c_1 can be chosen arbitrarily large, and $\kappa_c \leq n + 1$, inequality (B.36) follows immediately by combining (B.37)–(B.39).

We first proceed to prove (B.37). By (B.34) in the proof of Theorem 2, (B.37) can be proved as follows:

$$E[\|\boldsymbol{\delta}\|_2 1_{\mathcal{E}_1}] \leq E[C\sqrt{s(\log n)/n} 1_{\mathcal{E}_1}] \leq C\sqrt{s(\log n)/n}.$$

Next we prove (B.38). By the Cauchy–Schwarz inequality and Condition 1, we have

$$\begin{aligned} E[\|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty^2 1_{\mathcal{E}^c}] &\leq n^{-2} \{E[\|\mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty^4] P(\mathcal{E}^c)\}^{1/2} \\ &\leq O(p^{-c_1/2}) n^{-2} \left\{ E \left[\max_{1 \leq j \leq p} \|\tilde{\mathbf{x}}_j\|_2^4 \|\boldsymbol{\varepsilon}\|_2^4 \right] \right\}^{1/2} \\ &= O(p^{-c_1/2}), \end{aligned} \quad (\text{B.40})$$

where the last step is because of $\|\tilde{\mathbf{x}}_j\|_2 = \sqrt{n}$ and the assumption $\max_{1 \leq i \leq n} E\varepsilon_i^4 \leq C$. Similarly, we can prove that

$$\begin{aligned} E[\|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty 1_{\mathcal{E}^c}] &= O(p^{-c_1/2}) \quad \text{and} \quad E[\|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty^2 1_{\mathcal{E}_0^c}] \\ &= O(n^{-c_1/2}). \end{aligned} \quad (\text{B.41})$$

Since $\|\boldsymbol{\delta}\|_0 \leq \|\boldsymbol{\beta}_0\|_0 + \|\hat{\boldsymbol{\beta}}\|_0 < \kappa_c$, it follows that $\|\boldsymbol{\delta}\|_1 \leq \sqrt{\kappa_c} \|\boldsymbol{\delta}\|_2$. This together with (B.2) in Lemma 1 yields

$$\|\boldsymbol{\delta}\|_2 1_{\mathcal{E}^c} \leq C\sqrt{\kappa_c} (\|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty + \lambda) 1_{\mathcal{E}^c}. \quad (\text{B.42})$$

Thus, by Condition 1 and (B.40), inequality (B.38) is proved as follows:

$$\begin{aligned} E[\|\boldsymbol{\delta}\|_2^2 1_{\mathcal{E}^c}] &\leq C\kappa_c E[\|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty^2 1_{\mathcal{E}^c}] + C\kappa_c \lambda^2 P(\mathcal{E}^c) \\ &= O(p^{-c_1/2} \kappa_c). \end{aligned} \quad (\text{B.43})$$

Finally we prove (B.39). To this end, we first prove

$$E[\|\boldsymbol{\delta}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] = O(\lambda s n^{-c_1}). \quad (\text{B.44})$$

Then by Lemma 1 and the definition of \mathcal{E} , (B.39) can be proved as follows:

$$E[\|\boldsymbol{\delta}\|_2^2 1_{\mathcal{E}_0^c \cap \mathcal{E}}] \leq C\lambda E[\|\boldsymbol{\delta}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] = O(\lambda^2 s n^{-c_1}). \quad (\text{B.45})$$

It remains to prove (B.44). We first study $E[\|\boldsymbol{\delta}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}]$ by decomposing it into two terms:

$$E[\|\boldsymbol{\delta}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] = E[\|\boldsymbol{\delta}_{\alpha_0^c}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] + E[\|\boldsymbol{\delta}_{\alpha_0}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}]. \quad (\text{B.46})$$

We now consider the first term on the right-hand side of (B.46). Since $\hat{\boldsymbol{\beta}} \in \mathcal{B}_\tau$, it follows that $\|\boldsymbol{\delta}\|_2^2 = \|\boldsymbol{\delta}_{\alpha_0}\|_2^2 + \|\boldsymbol{\delta}_{\alpha_0^c}\|_2^2 \geq s^{-1} \|\boldsymbol{\delta}_{\alpha_0}\|_1^2 + \tau \|\boldsymbol{\delta}_{\alpha_0^c}\|_1$. Thus, by Lemma 1, we have conditioning on $\mathcal{E}_0^c \cap \mathcal{E}$,

$$s^{-1} \|\boldsymbol{\delta}_{\alpha_0}\|_1^2 + \tau \|\boldsymbol{\delta}_{\alpha_0^c}\|_1 \leq \|\boldsymbol{\delta}\|_2^2 \leq C\lambda \|\boldsymbol{\delta}\|_1 = C\lambda (\|\boldsymbol{\delta}_{\alpha_0}\|_1 + \|\boldsymbol{\delta}_{\alpha_0^c}\|_1).$$

A rearrangement of the above inequality yields

$$(\|\boldsymbol{\delta}_{\alpha_0}\|_1 - C\lambda)^2 \leq s(C\lambda - \tau) \|\boldsymbol{\delta}_{\alpha_0^c}\|_1 + Cs^2\lambda^2. \quad (\text{B.47})$$

Since the left-hand side of (B.47) is always nonnegative and $\lambda = o(\tau/\sqrt{s})$, we have $\|\delta_{\alpha_0^c}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}} \leq C\tau^{-1}s\lambda^2 1_{\mathcal{E}_0^c \cap \mathcal{E}}$. Thus, it follows from Condition 1 and $\lambda = o(\tau/\sqrt{s})$ that

$$E[\|\delta_{\alpha_0^c}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] \leq C\tau^{-1}s\lambda^2 P(\mathcal{E}_0^c \cap \mathcal{E}) = o(s\lambda n^{-c_1}). \quad (\text{B.48})$$

Since $\lambda = o(\tau/\sqrt{s})$, the first term on the right-hand side of (B.47) is negative for sufficiently large n . Therefore, it follows from (B.47) that conditioning on $\mathcal{E}_0^c \cap \mathcal{E}$, $\|\delta_{\alpha_0}\|_1 - Cs\lambda \leq Cs\lambda$. Hence, we obtain that

$$E[\|\delta_{\alpha_0}\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] \leq Cs\lambda E[1_{\mathcal{E}_0^c \cap \mathcal{E}}] = O(s\lambda n^{-c_1}).$$

This together with (B.48) proves (B.44), which completes the proof of (B.39). Consequently, (B.36) follows and the result under the L_2 -loss is proved.

We now consider $E\|\delta\|_1$ under the L_1 -estimation loss by using the following decomposition

$$E\|\delta\|_1 = E[\|\delta\|_1 1_{\mathcal{E}_1}] + E[\|\delta\|_1 1_{\mathcal{E} \cap \mathcal{E}_0^c}] + E[\|\delta\|_1 1_{\mathcal{E}^c}]. \quad (\text{B.49})$$

First, by Theorem 2, the first term on the right-hand side of (B.49) can be bounded as

$$E[\|\delta\|_1 1_{\mathcal{E}_1}] \leq s\sqrt{(\log n)/n} P(\mathcal{E}_1) \leq Cs\sqrt{(\log n)/n}. \quad (\text{B.50})$$

The second term of (B.49) has already been considered in (B.44). Therefore, we only need to study the third term. Since $\|\delta\|_1 \leq \sqrt{\kappa_c}\|\delta\|_2$, by (B.42) and (B.41), we can bound the third term as

$$E[\|\delta\|_1 1_{\mathcal{E}^c}] \leq \sqrt{\kappa_c} E[\|\delta\|_2 1_{\mathcal{E}^c}] \leq C\kappa_c E[\|n^{-1}\mathbf{X}^T \boldsymbol{\varepsilon}\|_{\infty} 1_{\mathcal{E}^c}] + C\kappa_c \lambda P(\mathcal{E}^c) = O(p^{-c_1/2} \kappa_c). \quad (\text{B.51})$$

Since c_1 can be chosen arbitrarily large and $\kappa_c \leq n+1$, the above inequality together with (B.50), (B.44), and (B.49) leads to

$$E\|\delta\|_1 \leq Cs(\log n)/n. \quad (\text{B.52})$$

Thus, the risk result under the L_1 -estimation loss is proved.

Now, applying Hölder's inequality and by (B.36) and (B.52), we can prove the risk inequalities under the L_q -estimation loss with $q \in (1, 2)$, as in (B.12).

Finally we consider the L_{∞} -estimation loss. By (B.36) and Condition 1,

$$E[\|\delta\|_{\infty} 1_{\mathcal{E}_1^c}] \leq E[\|\delta\|_2 1_{\mathcal{E}_1^c}] \leq \{E[\|\delta\|_2^2] P(\mathcal{E}_1^c)\}^{1/2} = O\{s^{1/2} n^{-(c_1+1)/2} \sqrt{\log n}\}.$$

Moreover, by Theorem 2, we have $\|\delta\|_{\infty} 1_{\mathcal{E}_1} \leq C\gamma_n^* \sqrt{(\log n)/n}$. Since c_1 can be chosen arbitrarily large, it follows that

$$E\|\delta\|_{\infty} = E[\|\delta\|_{\infty} 1_{\mathcal{E}_1}] + E[\|\delta\|_{\infty} 1_{\mathcal{E}_1^c}] \leq C\gamma_n^* \sqrt{(\log n)/n},$$

which completes the proof for estimation risks.

Prediction risk: By (B.35) and Condition 2, we have

$$E\{D(\widehat{\boldsymbol{\beta}})\} = E[2^{-1} \boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\widehat{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta}] \leq (2c_2)^{-1} (I_1 + I_2 + I_3), \quad (\text{B.53})$$

where $I_1 = E[\|\mathbf{X} \boldsymbol{\delta}\|_2^2 1_{\mathcal{E}_1}]$, $I_2 = E[\|\mathbf{X} \boldsymbol{\delta}\|_2^2 1_{\mathcal{E}_0^c \cap \mathcal{E}}]$, and $I_3 = E[\|\mathbf{X} \boldsymbol{\delta}\|_2^2 1_{\mathcal{E}^c}]$. We first consider $I_1 = E[\|\mathbf{X} \boldsymbol{\delta}\|_2^2 1_{\mathcal{E}_1}]$. By the second inequality in (B.33) and (B.52),

$$I_1 \leq C\sqrt{(\log n)n} E[\|\delta\|_1 1_{\mathcal{E}_1}] \leq C\sqrt{(\log n)n} E[\|\delta\|_1] \leq Cs(\log n). \quad (\text{B.54})$$

Next, we study the term I_2 . By Lemma 1, the definition of \mathcal{E} , and (B.44), we have

$$I_2 \leq Cn\lambda E[\|\delta\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] = O(s\lambda^2 n^{1-c_1}). \quad (\text{B.55})$$

Now we consider the last term I_3 . It follows from the proof of Lemma 1 that

$$I_3 \leq CE[\|\boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}\| 1_{\mathcal{E}^c}] + Cn\lambda E[\|\delta\|_1 1_{\mathcal{E}^c}] \equiv I_{3,1} + I_{3,2}. \quad (\text{B.56})$$

Since $\|\delta\|_1 \leq \sqrt{\kappa_c}\|\delta\|_2$, by (B.38) and (B.40), we can bound $I_{3,1}$ as

$$\begin{aligned} I_{3,1} &= E[\|\boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}\| 1_{\mathcal{E}^c}] \leq CE[\|\delta\|_1 \|\mathbf{X} \boldsymbol{\varepsilon}\|_{\infty} 1_{\mathcal{E}^c}] \\ &\leq C \{E[\|\delta\|_1^2 1_{\mathcal{E}^c}]\}^{1/2} \{E[\|\mathbf{X} \boldsymbol{\varepsilon}\|_{\infty}^2 1_{\mathcal{E}^c}]\}^{1/2} \\ &\leq C\sqrt{\kappa_c} \{E[\|\delta\|_2^2 1_{\mathcal{E}^c}]\}^{1/2} \{E[\|\mathbf{X}^T \boldsymbol{\varepsilon}\|_{\infty}^2 1_{\mathcal{E}^c}]\}^{1/2} \\ &= O(np^{-c_1/2} \kappa_c). \end{aligned} \quad (\text{B.57})$$

By (B.51), we have $I_{3,2} = O(\lambda np^{-c_1/2} \kappa_c)$. This together with (B.56) and (B.57) entails

$$I_3 = O(np^{-c_1/2} \kappa_c). \quad (\text{B.58})$$

Combining (B.58) with (B.53)–(B.55) and noting that c_1 can be chosen arbitrarily large, we finish the proof for prediction risk.

Sign risk: Since $\widehat{\boldsymbol{\beta}} \in \mathcal{B}_{\tau}$ and $p_{\lambda}(t)$ is increasing in $t \in [0, \infty)$, we have $\|p_{\lambda}(\widehat{\boldsymbol{\beta}})\|_1 = \sum_{j=1}^p p_{\lambda}(|\widehat{\beta}_j|) \geq \|\widehat{\boldsymbol{\beta}}\|_0 p_{\lambda}(\tau)$ and thus $\|\widehat{\boldsymbol{\beta}}\|_0 \leq [p_{\lambda}(\tau)]^{-1} \|p_{\lambda}(\widehat{\boldsymbol{\beta}})\|_1$. This together with (B.6) and Condition 2 gives

$$\begin{aligned} \text{FS}(\widehat{\boldsymbol{\beta}}) &\leq \|\widehat{\boldsymbol{\beta}}\|_0 + s \\ &\leq s + [p_{\lambda}(\tau)]^{-1} \left[\|p_{\lambda}(\boldsymbol{\beta}_0)\|_1 + n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta} - \frac{1}{2n} \boldsymbol{\delta}^T \mathbf{X}^T \mathbf{H}(\widehat{\boldsymbol{\beta}}) \mathbf{X} \boldsymbol{\delta} \right] \\ &\leq s + [p_{\lambda}(\tau)]^{-1} [\|p_{\lambda}(\boldsymbol{\beta}_0)\|_1 + n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}]. \end{aligned} \quad (\text{B.59})$$

Since $|n^{-1} \boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}| \leq \|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_{\infty} \|\delta\|_1 \leq \lambda \|\delta\|_1$ on the event \mathcal{E} , by (B.57) and (B.44), we have

$$\begin{aligned} E[n^{-1} |\boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}| 1_{\mathcal{E}_1^c}] &= E[n^{-1} |\boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}| 1_{\mathcal{E}^c}] + E[n^{-1} |\boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}| 1_{\mathcal{E}_0^c \cap \mathcal{E}}] \\ &\leq O(p^{-c_1/2} \kappa_c) + \lambda E[\|\delta\|_1 1_{\mathcal{E}_0^c \cap \mathcal{E}}] \\ &= O(p^{-c_1/2} \kappa_c) + O(s\lambda^2 n^{-c_1}). \end{aligned} \quad (\text{B.60})$$

Thus, combining (B.59) with (B.60) and noting $\|p_{\lambda}(\boldsymbol{\beta}_0)\|_1 \geq sp_{\lambda}(\tau)$, we obtain

$$\begin{aligned} E[\text{FS}(\widehat{\boldsymbol{\beta}}) 1_{\mathcal{E}_1^c}] &\leq P(\mathcal{E}_1^c) \{s + [p_{\lambda}(\tau)]^{-1} \|p_{\lambda}(\boldsymbol{\beta}_0)\|_1\} + [p_{\lambda}(\tau)]^{-1} E[n^{-1} |\boldsymbol{\varepsilon}^T \mathbf{X} \boldsymbol{\delta}| 1_{\mathcal{E}_1^c}] \\ &= [p_{\lambda}(\tau)]^{-1} [\|p_{\lambda}(\boldsymbol{\beta}_0)\|_1 O(n^{-c_1}) + O(p^{-c_1/2} \kappa_c) + O(s\lambda^2 n^{-c_1})]. \end{aligned}$$

On the other hand, Theorem 2 shows that $\text{FS}(\widehat{\boldsymbol{\beta}}) = 0$ on the event \mathcal{E}_1 . Thus, we have $E[\text{FS}(\widehat{\boldsymbol{\beta}}) 1_{\mathcal{E}_1}] = 0$, which leads to $E[\text{FS}(\widehat{\boldsymbol{\beta}})] = E[\text{FS}(\widehat{\boldsymbol{\beta}}) 1_{\mathcal{E}_1^c}]$. This concludes the proof.

B.7 Proof of Theorem 4

To simplify the technical presentation, we first consider the case of linear model. Then the penalized negative log-likelihood minimization problem in (6) becomes the penalized least-squares problem with $Q_n(\boldsymbol{\beta}) = (2n)^{-1} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \|p_{\lambda}(\boldsymbol{\beta})\|_1$. Note that in the case of linear model, $\boldsymbol{\mu}(\boldsymbol{\theta}) = \boldsymbol{\theta}$ and thus

$$\|n^{-1} \mathbf{X}_{\alpha}^T [\boldsymbol{\mu}(\mathbf{X}\boldsymbol{\beta}) - \boldsymbol{\mu}(\mathbf{X}\boldsymbol{\beta}_0)]\|_2 = \|n^{-1} \mathbf{X}_{\alpha}^T \mathbf{X}(\boldsymbol{\beta} - \boldsymbol{\beta}_0)\|_2 \geq c_4 \|\boldsymbol{\beta} - \boldsymbol{\beta}_0\|_2$$

holds for any $\boldsymbol{\beta} \in \mathcal{B}_{\tau}$, with $c_4 = c^2$ and $\alpha = \text{supp}(\boldsymbol{\beta}) \cup \text{supp}(\boldsymbol{\beta}_0)$. Denote by $\boldsymbol{\delta} = (\delta_1, \dots, \delta_p)^T = \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0$ with $\widehat{\boldsymbol{\beta}} = (\widehat{\beta}_1, \dots, \widehat{\beta}_p)^T$. Let $\alpha_0 = \text{supp}(\boldsymbol{\beta}_0)$ and $\alpha_1 = \text{supp}(\widehat{\boldsymbol{\beta}})$. Clearly, $\text{supp}(\boldsymbol{\delta}) \subset \alpha = \alpha_0 \cup \alpha_1$. It follows from $\boldsymbol{\beta}_0, \widehat{\boldsymbol{\beta}} \in \mathcal{B}_{\tau}$ that $|\alpha_0| < \kappa_c/2$, $|\alpha_1| < \kappa_c/2$, and $|\alpha| \leq |\alpha_0| + |\alpha_1| < \kappa_c$. Thus by the definition of the robust spark κ_c , we have $\lambda_{\min}(n^{-1} \mathbf{X}_{\alpha}^T \mathbf{X}_{\alpha}) \geq c^2$, which leads to

$$\|\delta\|_2 = \|\delta_{\alpha}\|_2 \leq c^{-2} \|n^{-1} \mathbf{X}_{\alpha}^T \mathbf{X}_{\alpha} \delta_{\alpha}\|_2 = c^{-2} \|n^{-1} \mathbf{X}_{\alpha}^T \mathbf{X} \boldsymbol{\delta}\|_2. \quad (\text{B.61})$$

Since $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\varepsilon}$ in linear model, we have $\mathbf{X}\boldsymbol{\delta} = \mathbf{X}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) = -(\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}}) + \boldsymbol{\varepsilon}$ and thus

$$n^{-1} \mathbf{X}_{\alpha}^T \mathbf{X} \boldsymbol{\delta} = n^{-1} \mathbf{X}_{\alpha}^T [-(\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}}) + \boldsymbol{\varepsilon}] = -n^{-1} \mathbf{X}_{\alpha}^T (\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}}) + n^{-1} \mathbf{X}_{\alpha}^T \boldsymbol{\varepsilon}.$$

This representation together with (B.61) yields

$$\begin{aligned} \|\delta\|_2 &\leq c^{-2} \|n^{-1} \mathbf{X}_\alpha^T \mathbf{X} \delta\|_2 \\ &\leq c^{-2} \|n^{-1} \mathbf{X}_\alpha^T (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}})\|_2 + c^{-2} \|n^{-1} \mathbf{X}_\alpha^T \boldsymbol{\varepsilon}\|_2. \end{aligned} \quad (\text{B.62})$$

Such an inequality provides an effective way to bound the size of the set α .

By Condition 1, the event $\mathcal{E} = \{\|n^{-1} \mathbf{X}^T \boldsymbol{\varepsilon}\|_\infty \leq \lambda/2\}$ has a large probability. We condition on this event hereafter. Then it holds that

$$\|n^{-1} \mathbf{X}_\alpha^T \boldsymbol{\varepsilon}\|_2 \leq |\alpha|^{1/2} \|n^{-1} \mathbf{X}_\alpha^T \boldsymbol{\varepsilon}\|_\infty \leq |\alpha|^{1/2} \lambda/2. \quad (\text{B.63})$$

Since $\eta_n = \|n^{-1} \mathbf{X}^T [\mathbf{y} - \boldsymbol{\mu}(\mathbf{X} \hat{\boldsymbol{\beta}})]\|_\infty = \|n^{-1} \mathbf{X}^T (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}})\|_\infty$, we have

$$\|n^{-1} \mathbf{X}_\alpha^T (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}})\|_2 \leq |\alpha|^{1/2} \|n^{-1} \mathbf{X}_\alpha^T (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}})\|_\infty \leq |\alpha|^{1/2} \eta_n. \quad (\text{B.64})$$

Let $k = |\alpha_1 \setminus \alpha_0|$. Clearly, $|\alpha| = |\alpha_0| + |\alpha_1 \setminus \alpha_0| = s + k$. Note that for each $j \in \alpha_1 \setminus \alpha_0$, we have $\delta_j = \hat{\beta}_j - \beta_{0,j} = \hat{\beta}_j$ and thus $|\delta_j| \geq \tau$, which entails that

$$\|\delta\|_2 \geq k^{1/2} \tau. \quad (\text{B.65})$$

Combining inequalities (B.62)–(B.65) along with $|\alpha| = s + k$ gives

$$k^{1/2} \tau \leq c^{-2} (s + k)^{1/2} (\eta_n + 2^{-1} \lambda),$$

which ensures that

$$k \leq \frac{c^{-4} (\eta_n + 2^{-2} \lambda)^2 / \tau^2}{1 - c^{-4} (\eta_n + 2^{-2} \lambda)^2 / \tau^2} s. \quad (\text{B.66})$$

Since $\eta_n + \lambda = o(\tau)$, it follows from the bound in (B.66) that $k \leq s$ for large enough n . Thus, applying similar arguments as above results in

$$\|\delta\|_2 \leq c^{-2} (2s)^{1/2} (\eta_n + 2^{-1} \lambda). \quad (\text{B.67})$$

Since $\min_{1 \leq j \leq s} |\beta_{0,j}| > c_5 s^{1/2} (\eta_n + \lambda)$ with c_5 some sufficiently large positive constant, inequality (B.67) entails that for large enough n , $\hat{\beta}_j \neq 0$ for each $j \in \alpha_0$. This shows that $\text{supp}(\hat{\boldsymbol{\beta}}) \supset \alpha_0 = \text{supp}(\boldsymbol{\beta}_0)$. Note that by assumption, $\hat{\boldsymbol{\beta}}$ is the global minimizer of problem (6) when constrained on the subspace given by its support. Observe that all arguments in the proofs of Theorems 1–3 on the global minimizer equally apply to the computable solution $\hat{\boldsymbol{\beta}}$ as long as $\text{supp}(\hat{\boldsymbol{\beta}}) \supset \text{supp}(\boldsymbol{\beta}_0)$. Therefore, $\hat{\boldsymbol{\beta}}$ enjoys the same asymptotic properties as for any global minimizer in Theorems 1–3 under the same conditions therein.

For the case of nonlinear model, by assumption, we have

$$\|n^{-1} \mathbf{X}_\alpha^T [\boldsymbol{\mu}(\mathbf{X} \hat{\boldsymbol{\beta}}) - \boldsymbol{\mu}(\mathbf{X} \boldsymbol{\beta}_0)]\|_2 \geq c_4 \|\delta\|_2,$$

which together with $\boldsymbol{\varepsilon} = \mathbf{y} - \boldsymbol{\mu}(\mathbf{X} \boldsymbol{\beta}_0)$ leads to

$$\|\delta\|_2 \leq c_4^{-1} \|n^{-1} \mathbf{X}_\alpha^T [\mathbf{y} - \boldsymbol{\mu}(\mathbf{X} \hat{\boldsymbol{\beta}})]\|_2 + c_4^{-1} \|n^{-1} \mathbf{X}_\alpha^T \boldsymbol{\varepsilon}\|_2. \quad (\text{B.68})$$

Observe that inequality (B.68) is of similar form as (B.62). Thus an application of similar arguments as above completes the proof.

[Received May 2012. Revised March 2013.]

REFERENCES

Antoniadis, A., and Fan, J. (2001), “Regularization of Wavelets Approximations” (with discussion), *Journal of the American Statistical Association*, 96, 939–967. [1045,1049]

Barron, A., Birge, L., and Massart, P. (1999), “Risk Bounds for Model Selection via Penalization,” *Probability Theory Related Fields*, 113, 301–413. [1046]

Bennett, G. (1962), “Probability Inequalities for the Sum of Independent Random Variables,” *Journal of the American Statistical Association*, 57, 33–45. [1056]

Bickel, P. J., Ritov, Y., and Tsybakov, A. (2009), “Simultaneous Analysis of Lasso and Dantzig Selector,” *The Annals of Statistics*, 37, 1705–1732. [1044,1045,1046,1047,1055]

Bunea, F., Tsybakov, A., and Wegkamp, M. H. (2007), “Sparsity Oracle Inequalities for the LASSO,” *Electronic Journal of Statistics*, 1, 169–194. [1044]

Candes, E., and Tao, T. (2007), “The Dantzig Selector: Statistical Estimation When p Is Much Larger Than n ” (with discussion), *The Annals of Statistics*, 35, 2313–2404. [1045]

Donoho, D., and Elad, M. (2003), “Optimally Sparse Representation in General (Nonorthogonal) Dictionaries via ℓ^1 Minimization,” *Proceedings of the National Academy of Sciences of the United States of America*, 100, 2197–2202. [1046]

Donoho, D. L., Elad, M., and Temlyakov, V. (2006), “Stable Recovery of Sparse Overcomplete Representations in the Presence of Noise,” *IEEE Transactions on Information Theory*, 52, 6–18. [1044]

Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), “Least Angle Regression” (with discussion), *The Annals of Statistics*, 32, 407–499. [1044]

Fan, J., and Fan, Y. (2008), “High-Dimensional Classification Using Features Annealed Independence Rules,” *The Annals of Statistics*, 36, 2605–2637. [1054]

Fan, J., and Li, R. (2001), “Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties,” *Journal of the American Statistical Association*, 96, 1348–1360. [1044,1049]

Fan, J., and Lv, J. (2008), “Sure Independence Screening for Ultrahigh Dimensional Feature Space” (with discussion), *Journal of the Royal Statistical Society, Series B*, 70, 849–911. [1056]

——— (2011), “Nonconcave Penalized Likelihood With NP-Dimensionality,” *IEEE Transactions on Information Theory*, 57, 5467–5484. [1044,1047,1049,1056]

Fan, J., and Peng, H. (2004), “Nonconcave Penalized Likelihood With Diverging Number of Parameters,” *The Annals of Statistics*, 32, 928–961. [1044]

Fan, Y., and Lv, J. (2012), “Asymptotic Properties of Global Optimum for Combined L_1 and Concave Regularization,” unpublished manuscript. [1046]

Frank, I. E., and Friedman, J. H. (1993), “A Statistical View of Some Chemometrics Regression Tools” (with discussion), *Technometrics*, 35, 109–148. [1048]

Friedman, J., Hastie, T., Höfling, H., and Tibshirani, R. (2007), “Pathwise Coordinate Optimization,” *The Annals of Applied Statistics*, 1, 302–332. [1044]

Hoeffding, W. (1963), “Probability Inequalities for Sums of Bounded Random Variables,” *Journal of the American Statistical Association*, 58, 13–30. [1056]

Lin, W., and Lv, J. (2013), “High-Dimensional Sparse Additive Hazards Regression,” *Journal of the American Statistical Association*, 108, 247–264. [1049]

Lv, J., and Fan, Y. (2009), “A Unified Approach to Model Selection and Sparse Recovery Using Regularized Least Squares,” *The Annals of Statistics*, 37, 3498–3528. [1044,1046,1047]

Singh, D., Febbo, P. G., Ross, K., Jackson, D. G., Manola, J., Ladd, C., Tamayo, P., Renshaw, A. A., D’Amico, A. V., Richie, J. P., Lander, E. S., Loda, M., Kantoff, P. W., Golub, T. R., and Sellers, W. R. (2002), “Gene Expression Correlates of Clinical Prostate Cancer Behavior,” *Cancer Cell*, 1, 203–209. [1054]

Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [1044]

van de Geer, S. (2008), “High-Dimensional Generalized Linear Models and the LASSO,” *The Annals of Statistics*, 36, 614–645. [1044]

van der Vaart, A. W., and Wellner, J. A. (1996), *Weak Convergence and Empirical Processes*, New York: Springer. [1056]

Wu, T. T., and Lange, K. (2008), “Coordinate Descent Algorithms for Lasso Penalized Regression,” *The Annals of Applied Statistics*, 2, 224–244. [1044]

Zhang, C.-H. (2010), “Nearly Unbiased Variable Selection Under Minimax Concave Penalty,” *The Annals of Statistics*, 38, 894–942. [1044]

Zhang, C.-H., and Huang, J. (2008), “The Sparsity and Bias of the LASSO Selection in High-Dimensional Linear Regression,” *The Annals of Statistics*, 36, 1567–1594. [1044]

Zhao, P., and Yu, B. (2006), “On Model Selection Consistency of Lasso,” *Journal of Machine Learning Research*, 7, 2541–2563. [1044,1047]

Zheng, Z., Fan, Y., and Lv, J. (2014), “High-dimensional thresholded regression and shrinkage effect,” *Journal of the Royal Statistical Society, Series B*, 76, forthcoming. [1046]

Zou, H. (2006), “The Adaptive Lasso and Its Oracle Properties,” *Journal of the American Statistical Association*, 101, 1418–1429. [1044,1048]

Zou, H., and Li, R. (2008), “One-Step Sparse Estimates in Nonconcave Penalized Likelihood Models” (with discussion), *The Annals of Statistics*, 36, 1509–1566. [1049]