# ADAPTIVE ROBUST VARIABLE SELECTION

By Jianqing Fan[1], Yingying Fan[2] and Emre Barut

*Princeton University, University of Southern California and
IBM T. J. Watson Research Center*

Heavy-tailed high-dimensional data are commonly encountered in various scientific fields and pose great challenges to modern statistical analysis. A natural procedure to address this problem is to use penalized quantile regression with weighted $L_1$-penalty, called weighted robust Lasso (WR-Lasso), in which weights are introduced to ameliorate the bias problem induced by the $L_1$-penalty. In the ultra-high dimensional setting, where the dimensionality can grow exponentially with the sample size, we investigate the model selection oracle property and establish the asymptotic normality of the WR-Lasso. We show that only mild conditions on the model error distribution are needed. Our theoretical results also reveal that adaptive choice of the weight vector is essential for the WR-Lasso to enjoy these nice asymptotic properties. To make the WR-Lasso practically feasible, we propose a two-step procedure, called adaptive robust Lasso (AR-Lasso), in which the weight vector in the second step is constructed based on the $L_1$-penalized quantile regression estimate from the first step. This two-step procedure is justified theoretically to possess the oracle property and the asymptotic normality. Numerical studies demonstrate the favorable finite-sample performance of the AR-Lasso.

**1. Introduction.** The advent of modern technology makes it easier to collect massive, large-scale data-sets. A common feature of these data-sets is that the number of covariates greatly exceeds the number of observations, a regime opposite to conventional statistical settings. For example, portfolio allocation with hundreds of stocks in finance involves a covariance matrix of about tens of thousands of parameters, but the sample sizes are often only in the order of hundreds (e.g., daily data over a year period [Fan, Fan and Lv (2008)]). Genome-wide association studies in biology involve hundreds of thousands of single-nucleotide polymorphisms (SNPs), but the available sample size is usually in hundreds, also.

Data-sets with large number of variables but relatively small sample size pose great unprecedented challenges and opportunities for statistical analysis.

Regularization methods have been widely used for high-dimensional variable selection [Tibshirani (1996), Fan and Li (2001), Fan and Peng (2004), Bickel and Li (2006), Candes and Tao (2007), Bickel, Ritov and Tsybakov (2009), Lv and Fan (2009), Meinshausen and Bühlmann (2010), Zhang (2010), Zou (2006)]. Yet, most existing methods such as penalized least-squares or penalized likelihood [Fan and Lv (2011)] are designed for light-tailed distributions. Zhao and Yu (2006) established the irrepresentability conditions for the model selection consistency of the Lasso estimator. Fan and Li (2001) studied the oracle properties of nonconcave penalized likelihood estimators for fixed dimensionality. Lv and Fan (2009) investigated the penalized least-squares estimator with folded-concave penalty functions in the ultra-high dimensional setting and established a nonasymptotic weak oracle property. Fan and Lv (2008) proposed and investigated the sure independence screening method in the setting of light-tailed distributions. The robustness of the aforementioned methods have not yet been thoroughly studied and well understood.

Robust regularization methods such as the least absolute deviation (LAD) regression and quantile regression have been used for variable selection in the case of fixed dimensionality. See, for example, Wang, Li and Jiang (2007), Li and Zhu (2008), Zou and Yuan (2008), Wu and Liu (2009). The penalized composite likelihood method was proposed in Bradic, Fan and Wang (2011) for robust estimation in ultra-high dimensions with focus on the efficiency of the method. They still assumed sub-Gaussian tails. Belloni and Chernozhukov (2011) studied the $L_1$-penalized quantile regression in high-dimensional sparse models where the dimensionality could be larger than the sample size. We refer to their method as robust Lasso (R-Lasso). They showed that the R-Lasso estimate is consistent at the near-oracle rate, and gave conditions under which the selected model includes the true model, and derived bounds on the size of the selected model, uniformly in a compact set of quantile indices. Wang (2013) studied the $L_1$-penalized LAD regression and showed that the estimate achieves near oracle risk performance with a nearly universal penalty parameter and established also a sure screening property for such an estimator. van de Geer and Müller (2012) obtained bounds on the prediction error of a large class of $L_1$-penalized estimators, including quantile regression. Wang, Wu and Li (2012) considered the nonconvex penalized quantile regression in the ultra-high dimensional setting and showed that the oracle estimate belongs to the set of local minima of the nonconvex penalized quantile regression, under mild assumptions on the error distribution.

In this paper, we introduce the penalized quantile regression with the weighted $L_1$-penalty (WR-Lasso) for robust regularization, as in Bradic, Fan and Wang (2011). The weights are introduced to reduce the bias problem induced by the $L_1$-penalty. The flexibility of the choice of the weights provides flexibility in shrinkage estimation of the regression coefficient. WR-Lasso shares a similar spirit

to the folded-concave penalized quantile-regression [Zou and Li (2008), Wang, Wu and Li (2012)], but avoids the nonconvex optimization problem. We establish conditions on the error distribution in order for the WR-Lasso to successfully recover the true underlying sparse model with asymptotic probability one. It turns out that the required condition is much weaker than the sub-Gaussian assumption in Bradic, Fan and Wang (2011). The only conditions we impose is that the density function of error has Lipschitz property in a neighborhood around 0. This includes a large class of heavy-tailed distributions such as the stable distributions, including the Cauchy distribution. It also covers the double exponential distribution whose density function is nondifferentiable at the origin.

Unfortunately, because of the penalized nature of the estimator, WR-Lasso estimate has a bias. In order to reduce the bias, the weights in WR-Lasso need to be chosen adaptively according to the magnitudes of the unknown true regression coefficients, which makes the bias reduction infeasible for practical applications.

To make the bias reduction feasible, we introduce the adaptive robust Lasso (AR-Lasso). The AR-Lasso first runs R-Lasso to obtain an initial estimate, and then computes the weight vector of the weighted $L_1$-penalty according to a decreasing function of the magnitude of the initial estimate. After that, AR-Lasso runs WR-Lasso with the computed weights. We formally establish the model selection oracle property of AR-Lasso in the context of Fan and Li (2001) with no assumptions made on the tail distribution of the model error. In particular, the asymptotic normality of the AR-Lasso is formally established.

This paper is organized as follows. First, we introduce our robust estimators in Section 2. Then, to demonstrate the advantages of our estimator, we show in Section 3 with a simple example that Lasso behaves suboptimally when noise has heavy tails. In Section 4.1, we study the performance of the oracle-assisted regularization estimator. Then in Section 4.2, we show that when the weights are adaptively chosen, WR-Lasso has the model selection oracle property, and performs as well as the oracle-assisted regularization estimate. In Section 4.3, we prove the asymptotic normality of our proposed estimator. The feasible estimator, AR-Lasso, is investigated in Section 5. Section 6 presents the results of the simulation studies. Finally, in Section 7, we present the proofs of the main theorems. Additional proofs, as well as the results of a genome-wide association study, are provided in the supplementary Appendix [Fan, Fan and Barut (2014)].

**2. Adaptive robust Lasso.**   Consider the linear regression model

$$(2.1) \qquad\qquad\qquad \mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon},$$

where $\mathbf{y}$ is an $n$-dimensional response vector, $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)^T = (\tilde{\mathbf{x}}_1, \ldots, \tilde{\mathbf{x}}_p)$ is an $n \times p$ fixed design matrix, $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)^T$ is a $p$-dimensional regression coefficient vector, and $\boldsymbol{\varepsilon} = (\varepsilon_1, \ldots, \varepsilon_n)^T$ is an $n$-dimensional error vector whose components are independently distributed and satisfy $P(\varepsilon_i \leq 0) = \tau$ for some known constant $\tau \in (0, 1)$. Under this model, $\mathbf{x}_i^T \boldsymbol{\beta}$ is the conditional $\tau$th-quantile

of $y_i$ given $\mathbf{x}_i$. We impose no conditions on the heaviness of the tail probability or the homoscedasticity of $\varepsilon_i$. We consider a challenging setting in which $\log p = o(n^b)$ with some constant $b > 0$. To ensure the model identifiability and to enhance the model fitting accuracy and interpretability, the true regression coefficient vector $\boldsymbol{\beta}^*$ is commonly imposed to be sparse with only a small proportion of nonzeros [Tibshirani (1996), Fan and Li (2001)]. Denoting the number of nonzero elements of the true regression coefficients by $s_n$, we allow $s_n$ to slowly diverge with the sample size $n$ and assume that $s_n = o(n)$. To ease the presentation, we suppress the dependence of $s_n$ on $n$ whenever there is no confusion. Without loss of generality, we write $\boldsymbol{\beta}^* = (\boldsymbol{\beta}_1^{*T}, \mathbf{0}^T)^T$, that is, only the first $s$ entries are nonvanishing. The true model is denoted by

$$\mathcal{M}_* = \text{supp}(\boldsymbol{\beta}^*) = \{1, \ldots, s\}$$

and its complement, $\mathcal{M}_*^c = \{s + 1, \ldots, p\}$, represents the set of noise variables.

We consider a fixed design matrix in this paper and denote by $\mathbf{S} = (\mathbf{S}_1, \ldots, \mathbf{S}_n)^T = (\tilde{\mathbf{x}}_1, \ldots, \tilde{\mathbf{x}}_s)$ the submatrix of $\mathbf{X}$ corresponding to the covariates whose coefficients are nonvanishing. These variables will be referred to as the signal covariates and the rest will be called noise covariates. The set of columns that correspond to the noise covariates is denoted by $\mathbf{Q} = (\mathbf{Q}_1, \ldots, \mathbf{Q}_n)^T = (\tilde{\mathbf{x}}_{s+1}, \ldots, \tilde{\mathbf{x}}_p)$. We standardize each column of $\mathbf{X}$ to have $L_2$-norm $\sqrt{n}$.

To recover the true model and estimate $\boldsymbol{\beta}^*$, we consider the following regularization problem:

$$(2.2) \qquad \min_{\boldsymbol{\beta} \in \mathbf{R}^p} \left\{ \sum_{i=1}^{n} \rho_\tau (y_i - \mathbf{x}_i^T \boldsymbol{\beta}) + n\lambda_n \sum_{j=1}^{p} p_{\lambda_n}(|\beta_j|) \right\},$$

where $\rho_\tau(u) = u(\tau - 1\{u \le 0\})$ is the quantile loss function, and $p_{\lambda_n}(\cdot)$ is a nonnegative penalty function on $[0, \infty)$ with a regularization parameter $\lambda_n \ge 0$. The use of quantile loss function in (2.2) is to overcome the difficulty of heavy tails of the error distribution. Since $P(\varepsilon \le 0) = \tau$, (2.2) can be interpreted as the sparse estimation of the conditional $\tau$th quantile. Regarding the choice of $p_{\lambda_n}(\cdot)$, it was demonstrated in Lv and Fan (2009) and Fan and Lv (2011) that folded-concave penalties are more advantageous for variable selection in high dimensions than the convex ones such as the $L_1$-penalty. It is, however, computationally more challenging to minimize the objective function in (2.2) when $p_\lambda(\cdot)$ is folded-concave. Noting that with a good initial estimate $\hat{\boldsymbol{\beta}}^{\text{ini}} = (\hat{\beta}_1^{\text{ini}}, \ldots, \hat{\beta}_p^{\text{ini}})^T$ of the true coefficient vector, we have

$$p_{\lambda_n}(|\beta_j|) \approx p_{\lambda_n}(|\hat{\beta}_j^{\text{ini}}|) + p_{\lambda_n}'(|\hat{\beta}_j^{\text{ini}}|)(|\beta_j| - |\hat{\beta}_j^{\text{ini}}|).$$

Thus, instead of (2.2) we consider the following weighted $L_1$-regularized quantile regression:

$$(2.3) \qquad L_n(\boldsymbol{\beta}) = \sum_{i=1}^{n} \rho_\tau(y_i - \mathbf{x}_i^T \boldsymbol{\beta}) + n\lambda_n \|\mathbf{d} \circ \boldsymbol{\beta}\|_1,$$

where $\mathbf{d} = (d_1, \ldots, d_p)^T$ is the vector of nonnegative weights, and $\circ$ is the Hadamard product, that is, the componentwise product of two vectors. This motivates us to define the weighted robust Lasso (WR-Lasso) estimate as the global minimizer of the convex function $L_n(\boldsymbol{\beta})$ for a given nonstochastic weight vector:

$$(2.4) \qquad \hat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} L_n(\boldsymbol{\beta}).$$

The uniqueness of the global minimizer is easily guaranteed by adding a negligible $L_2$-regularization in implementation. In particular, when $d_j = 1$ for all $j$, the method will be referred to as robust Lasso (R-Lasso).

The adaptive robust Lasso (AR-Lasso) refers specifically to the two-stage procedure in which the stochastic weights $\hat{d}_j = p'_{\lambda_n}(|\hat{\beta}_j^{\text{ini}}|)$ for $j = 1, \ldots, p$ are used in the second step for WR-Lasso and are constructed using a concave penalty $p_{\lambda_n}(\cdot)$ and the initial estimates, $\hat{\beta}_j^{\text{ini}}$, from the first step. In practice, we recommend using R-Lasso as the initial estimate and then using SCAD to compute the weights in AR-Lasso. The asymptotic result of this specific AR-Lasso is summarized in Corollary 1 in Section 5 for the ultra-high dimensional robust regression problem. This is a main contribution of the paper.

**3. Suboptimality of Lasso.** In this section, we use a specific example to illustrate that, in the case of heavy-tailed error distribution, Lasso fails at model selection unless the nonzero coefficients, $\beta_1^*, \ldots, \beta_s^*$, have a very large magnitude. We assume that the errors $\varepsilon_1, \ldots, \varepsilon_n$ have the identical symmetric stable distribution and the characteristic function of $\varepsilon_1$ is given by

$$E\big[\exp(iu\varepsilon_1)\big] = \exp(-|u|^\alpha),$$

where $\alpha \in (0, 2)$. By Nolan (2012), $E|\varepsilon_1|^p$ is finite for $0 < p < \alpha$, and $E|\varepsilon_1|^p = \infty$ for $p \geq \alpha$. Furthermore, as $z \to \infty$,

$$P\big(|\varepsilon_1| \geq z\big) \sim c_\alpha z^{-\alpha},$$

where $c_\alpha = \sin(\frac{\pi\alpha}{2})\Gamma(\alpha)/\pi$ is a constant depending only on $\alpha$, and we use the notation $\sim$ to denote that two terms are equivalent up to some constant. Moreover, for any constant vector $\mathbf{a} = (a_1, \ldots, a_n)^T$, the linear combination $\mathbf{a}^T\boldsymbol{\varepsilon}$ has the following tail behavior:

$$(3.1) \qquad P\big(|\mathbf{a}^T\boldsymbol{\varepsilon}| > z\big) \sim \|\mathbf{a}\|_\alpha^\alpha c_\alpha z^{-\alpha}$$

with $\|\cdot\|_\alpha$ denoting the $L_\alpha$-norm of a vector.

To demonstrate the suboptimality of Lasso, we consider a simple case in which the design matrix satisfies the conditions that $\mathbf{S}^T\mathbf{Q} = \mathbf{0}$, $\frac{1}{n}\mathbf{S}^T\mathbf{S} = \mathbf{I}_s$, the columns of $\mathbf{Q}$ satisfy $|\operatorname{supp}(\tilde{\mathbf{x}}_j)| = m_n = O(n^{1/2})$ and $\operatorname{supp}(\tilde{\mathbf{x}}_k) \cap \operatorname{supp}(\tilde{\mathbf{x}}_j) = \varnothing$ for any $k \neq j$ and $k, j \in \{s+1, \ldots, p\}$. Here, $m_n$ is a positive integer measuring the sparsity level of the columns of $\mathbf{Q}$. We assume that there are only fixed number of true

variables, that is, $s$ is finite, and that $\max_{ij} |x_{ij}| = O(n^{1/4})$. Thus, it is easy to see that $p = O(n^{1/2})$. In addition, we assume further that all nonzero regression coefficients are the same and $\beta_1^* = \cdots = \beta_s^* = \beta_0 > 0$.

We first consider R-Lasso, which is the global minimizer of (2.4). We will later see in Theorem 2 that by choosing the tuning parameter

$$\lambda_n = O\big((\log n)^2 \sqrt{(\log p)/n}\big),$$

R-Lasso can recover the true support $\mathcal{M}_* = \{1, \ldots, s\}$ with probability tending to 1. Moreover, the signs of the true regression coefficients can also be recovered with asymptotic probability one as long as the following condition on signal strength is satisfied:

$$(3.2) \qquad \lambda_n^{-1} \beta_0 \to \infty, \qquad \text{that is, } (\log n)^{-2} \sqrt{n/(\log p)} \beta_0 \to \infty.$$

Now, consider Lasso, which minimizes

$$(3.3) \qquad \widetilde{L}_n(\boldsymbol{\beta}) = \tfrac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + n\lambda_n \|\boldsymbol{\beta}\|_1.$$

We will see that for (3.3) to recover the true model and the correct signs of coefficients, we need a much stronger signal level than that is given in (3.2). By results in optimization theory, the Karush–Kuhn–Tucker (KKT) conditions guaranteeing the necessary and sufficient conditions for $\tilde{\boldsymbol{\beta}}$ with $\mathcal{M} = \mathrm{supp}(\tilde{\boldsymbol{\beta}})$ being a minimizer to (3.3) are

$$\tilde{\boldsymbol{\beta}}_{\mathcal{M}} + n\lambda_n (\mathbf{X}_{\mathcal{M}}^T \mathbf{X}_{\mathcal{M}})^{-1} \mathrm{sgn}(\tilde{\boldsymbol{\beta}}_{\mathcal{M}}) = (\mathbf{X}_{\mathcal{M}}^T \mathbf{X}_{\mathcal{M}})^{-1} \mathbf{X}_{\mathcal{M}}^T \mathbf{y},$$

$$\big\| \mathbf{X}_{\mathcal{M}^c}^T (\mathbf{y} - \mathbf{X}_{\mathcal{M}} \tilde{\boldsymbol{\beta}}_{\mathcal{M}}) \big\|_\infty \le n\lambda_n,$$

where $\mathcal{M}^c$ is the complement of $\mathcal{M}$, $\boldsymbol{\beta}_{\mathcal{M}}$ is the subvector formed by entries of $\boldsymbol{\beta}$ with indices in $\mathcal{M}$, and $\mathbf{X}_{\mathcal{M}}$ and $\mathbf{X}_{\mathcal{M}^c}$ are the submatrices formed by columns of $\mathbf{X}$ with indices in $\mathcal{M}$ and $\mathcal{M}^c$, respectively. It is easy to see from the above two conditions that for Lasso to enjoy the sign consistency, $\mathrm{sgn}(\tilde{\boldsymbol{\beta}}) = \mathrm{sgn}(\boldsymbol{\beta}^*)$ with asymptotic probability one, we must have these two conditions satisfied with $\mathcal{M} = \mathcal{M}^*$ with probability tending to 1. Since we have assumed that $\mathbf{Q}^T \mathbf{S} = \mathbf{0}$ and $n^{-1} \mathbf{S}^T \mathbf{S} = \mathbf{I}$, the above sufficient and necessary conditions can also be written as

$$(3.4) \qquad \tilde{\boldsymbol{\beta}}_{\mathcal{M}^*} + \lambda_n \mathrm{sgn}(\tilde{\boldsymbol{\beta}}_{\mathcal{M}^*}) = \boldsymbol{\beta}_{\mathcal{M}^*}^* + n^{-1} \mathbf{S}^T \boldsymbol{\varepsilon},$$

$$(3.5) \qquad \big\| \mathbf{Q}^T \boldsymbol{\varepsilon} \big\|_\infty \le n\lambda_n.$$

Conditions (3.4) and (3.5) are hard for Lasso to hold simultaneously. The following proposition summarizes the necessary condition, whose proof is given in the supplementary material [Fan, Fan and Barut (2014)].

PROPOSITION 1. *In the above model, with probability at least $1 - e^{-\tilde{c}_0}$, where $\tilde{c}_0$ is some positive constant, Lasso does not have sign consistency, unless the following signal condition is satisfied*

$$(3.6) \qquad n^{(3/4)-(1/\alpha)} \beta_0 \to \infty.$$

Comparing this with (3.2), it is easy to see that even in this simple case, Lasso needs much stronger signal levels than R-Lasso in order to have a sign consistency in the presence of a heavy-tailed distribution.

**4. Model selection oracle property.** In this section, we establish the model selection oracle property of WR-Lasso. The study enables us to see the bias due to penalization, and that an adaptive weighting scheme is needed in order to eliminate such a bias. We need the following condition on the distribution of noise.

CONDITION 1. *There exist universal constants $c_1 > 0$ and $c_2 > 0$ such that for any $u$ satisfying $|u| \le c_1$, $f_i(u)$'s are uniformly bounded away from $0$ and $\infty$ and*

$$\left| F_i(u) - F_i(0) - uf_i(0) \right| \le c_2 u^2,$$

*where $f_i(u)$ and $F_i(u)$ are the density function and distribution function of the error $\varepsilon_i$, respectively.*

Condition 1 implies basically that each $f_i(u)$ is Lipschitz around the origin. Commonly used distributions such as the double-exponential distribution and stable distributions including the Cauchy distribution all satisfy this condition.

Denote by $\mathbf{H} = \text{diag}\{f_1(0), \ldots, f_n(0)\}$. The next condition is on the submatrix of $\mathbf{X}$ that corresponds to signal covariates and the magnitude of the entries of $\mathbf{X}$.

CONDITION 2. *The eigenvalues of $\frac{1}{n}\mathbf{S}^T\mathbf{H}\mathbf{S}$ are bounded from below and above by some positive constants $c_0$ and $1/c_0$, respectively. Furthermore,*

$$\kappa_n \equiv \max_{ij} |x_{ij}| = o(\sqrt{n}s^{-1}).$$

Although Condition 2 is on the fixed design matrix, we note that the above condition on $\kappa_n$ is satisfied with asymptotic probability one when the design matrix is generated from some distributions. For instance, if the entries of $\mathbf{X}$ are independent copies from a subexponential distribution, the bound on $\kappa_n$ is satisfied with asymptotic probability one as long as $s = o(\sqrt{n}/(\log p))$; if the components are generated from sub-Gaussian distribution, then the condition on $\kappa_n$ is satisfied with probability tending to one when $s = o(\sqrt{n/(\log p)})$.

4.1. *Oracle regularized estimator.* To evaluate our newly proposed method, we first study how well one can do with the assistance of the oracle information on the locations of signal covariates. Then we use this to establish the asymptotic property of our estimator without the oracle assistance. Denote by $\hat{\boldsymbol{\beta}}^o = ((\hat{\boldsymbol{\beta}}_1^o)^T, \mathbf{0}^T)^T$ the oracle regularized estimator (ORE) with $\hat{\boldsymbol{\beta}}_1^o \in \mathbf{R}^s$ and $\mathbf{0}$ being the vector of all zeros, which minimizes $L_n(\boldsymbol{\beta})$ over the space $\{\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T \in$

$\mathbf{R}^p : \boldsymbol{\beta}_2 = \mathbf{0} \in \mathbf{R}^{p-s}\}$. The next theorem shows that ORE is consistent, and estimates the correct sign of the true coefficient vector with probability tending to one. We use $\mathbf{d}_0$ to denote the first $s$ elements of $\mathbf{d}$.

THEOREM 1. *Let* $\gamma_n = C_1(\sqrt{s(\log n)/n} + \lambda_n \|\mathbf{d}_0\|_2)$ *with* $C_1 > 0$ *a constant. If Conditions 1 and 2 hold and* $\lambda_n \|\mathbf{d}_0\|_2 \sqrt{s}\kappa_n \to 0$, *then there exists some constant* $c > 0$ *such that*

(4.1) $$P(\|\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*\|_2 \le \gamma_n) \ge 1 - n^{-cs}.$$

*If in addition* $\gamma_n^{-1} \min_{1 \le j \le s} |\beta_j^*| \to \infty$, *then with probability at least* $1 - n^{-cs}$,

$$\mathrm{sgn}(\hat{\boldsymbol{\beta}}_1^o) = \mathrm{sgn}(\boldsymbol{\beta}_1^*),$$

*where the above equation should be understood componentwisely.*

As shown in Theorem 1, the consistency rate of $\hat{\boldsymbol{\beta}}_1^o$ in terms of the vector $L_2$-norm is given by $\gamma_n$. The first component of $\gamma_n$, $C_1\sqrt{s(\log n)/n}$, is the oracle rate within a factor of $\log n$, and the second component $C_1\lambda_n\|\mathbf{d}_0\|_2$ reflects the bias due to penalization. If no prior information is available, one may choose equal weights $\mathbf{d}_0 = (1, 1, \ldots, 1)^T$, which corresponds to R-Lasso. Thus, for R-Lasso, with probability at least $1 - n^{-cs}$, it holds that

(4.2) $$\|\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*\|_2 \le C_1(\sqrt{s(\log n)/n} + \sqrt{s}\lambda_n).$$

4.2. *WR-Lasso.* In this section, we show that even without the oracle information, WR-Lasso enjoys the same asymptotic property as in Theorem 1 when the weight vector is appropriately chosen. Since the regularized estimator $\hat{\boldsymbol{\beta}}$ in (2.4) depends on the full design matrix $\mathbf{X}$, we need to impose the following conditions on the design matrix to control the correlation of columns in $\mathbf{Q}$ and $\mathbf{S}$.

CONDITION 3. *With* $\gamma_n$ *defined in Theorem 1, it holds that*

$$\left\|\frac{1}{n}\mathbf{Q}^T\mathbf{H}\mathbf{S}\right\|_{2,\infty} < \frac{\lambda_n}{2\|\mathbf{d}_1^{-1}\|_\infty \gamma_n},$$

*where* $\|\mathbf{A}\|_{2,\infty} = \sup_{\mathbf{x} \ne 0} \|\mathbf{A}\mathbf{x}\|_\infty / \|\mathbf{x}\|_2$ *for a matrix* $\mathbf{A}$ *and vector* $\mathbf{x}$, *and* $\mathbf{d}_1^{-1} = (d_{s+1}^{-1}, \ldots, d_p^{-1})^T$. *Furthermore,* $\log(p) = o(n^b)$ *for some constant* $b \in (0, 1)$.

To understand the implications of Condition 3, we consider the case of $f_1(0) = \cdots = f_n(0) \equiv f(0)$. In the special case of $\mathbf{Q}^T\mathbf{S} = \mathbf{0}$, Condition 3 is satisfied automatically. In the case of equal correlation, that is, $n^{-1}\mathbf{X}^T\mathbf{X}$ having off-diagonal elements all equal to $\rho$, the above Condition 3 reduces to

$$|\rho| < \frac{\lambda_n}{4f(0)\|\mathbf{d}_1^{-1}\|_\infty \sqrt{s}\gamma_n}.$$

This puts an upper bound on the correlation coefficient $\rho$ for such a dense matrix.

It is well known that for Gaussian errors, the optimal choice of regularization parameter $\lambda_n$ has the order $\sqrt{(\log p)/n}$ [Bickel, Ritov and Tsybakov (2009)]. The distribution of the model noise with heavy tails demands a larger choice of $\lambda_n$ to filter the noise for R-Lasso. When $\lambda_n \geq \sqrt{(\log n)/n}$, $\gamma_n$ given in (4.2) is in the order of $C_1 \lambda_n \sqrt{s}$. In this case, Condition 3 reduces to

$$(4.3) \qquad \|n^{-1}\mathbf{Q}^T\mathbf{H}\mathbf{S}\|_{2,\infty} < O(s^{-1/2}).$$

For WR-Lasso, if the weights are chosen such that $\|\mathbf{d}_0\|_2 = O(\sqrt{s(\log n)/n}/\lambda_n)$ and $\|\mathbf{d}_1\|_\infty = O(1)$, then $\gamma_n$ is in the order of $C_1\sqrt{s(\log n)/n}$, and correspondingly, Condition 3 becomes

$$\|n^{-1}\mathbf{Q}^T\mathbf{H}\mathbf{S}\|_{2,\infty} < O\big(\lambda_n\sqrt{n/(s(\log n))}\big).$$

This is a more relaxed condition than (4.3), since with heavy-tailed errors, the optimal $\lambda_n$ should be larger than $\sqrt{(\log p)/n}$. In other words, WR-Lasso not only reduces the bias of the estimate, but also allows for stronger correlations among the signal and noise covariates. However, the above choice of weights depends on unknown locations of signals. A data-driven choice will be given in Section 5, in which the resulting AR-Lasso estimator will be studied.

The following theorem shows the model selection oracle property of the WR-Lasso estimator.

THEOREM 2. *Suppose Conditions 1–3 hold. In addition, assume that* $\min_{j \geq s+1} d_j > c_3$ *with some constant* $c_3 > 0$,

$$(4.4) \quad \gamma_n s^{3/2}\kappa_n^2(\log_2 n)^2 = o(n\lambda_n^2), \qquad \lambda_n\|\mathbf{d}_0\|_2\kappa_n\max\{\sqrt{s}, \|\mathbf{d}_0\|_2\} \to 0$$

*and* $\lambda_n > 2\sqrt{(1+c)(\log p)/n}$, *where* $\kappa_n$ *is defined in Condition 2,* $\gamma_n$ *is defined in Theorem 1, and* $c$ *is some positive constant. Then, with probability at least* $1 - O(n^{-cs})$, *there exists a global minimizer* $\hat{\boldsymbol{\beta}} = ((\hat{\boldsymbol{\beta}}_1^o)^T, \hat{\boldsymbol{\beta}}_2^T)^T$ *of* $L_n(\boldsymbol{\beta})$ *which satisfies*

(1) $\hat{\boldsymbol{\beta}}_2 = 0$;
(2) $\|\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*\|_2 \leq \gamma_n$.

Theorem 2 shows that the WR-Lasso estimator enjoys the same property as ORE with probability tending to one. However, we impose nonadaptive assumptions on the weight vector $\mathbf{d} = (\mathbf{d}_0^T, \mathbf{d}_1^T)^T$. For noise covariates, we assume $\min_{j>s} d_j > c_3$, which implies that each coordinate needs to be penalized. For the signal covariates, we impose (4.4), which requires $\|\mathbf{d}_0\|_2$ to be small.

When studying the nonconvex penalized quantile regression, Wang, Wu and Li (2012) assumed that $\kappa_n$ is bounded and the density functions of $\varepsilon_i$'s are uniformly bounded away from 0 and $\infty$ in a small neighborhood of 0. Their assumption on

the error distribution is weaker than our Condition 1. We remark that the difference is because we have weaker conditions on $\kappa_n$ and the penalty function [see Condition 2 and (4.4)]. In fact, our Condition 1 can be weakened to the same condition as that in Wang, Wu and Li (2012) at the cost of imposing stronger assumptions on $\kappa_n$ and the weight vector $\mathbf{d}$.

Belloni and Chernozhukov (2011) and Wang (2013) imposed the restricted eigenvalue assumption of the design matrix and studied the $L_1$-penalized quantile regression and LAD regression, respectively. We impose different conditions on the design matrix and allow flexible shrinkage by choosing $\mathbf{d}$. In addition, our Theorem 2 provides a stronger result than consistency; we establish model selection oracle property of the estimator.

4.3. *Asymptotic normality.* We now present the asymptotic normality of our estimator. Define $\mathbf{V}_n = (\mathbf{S}^T \mathbf{H} \mathbf{S})^{-1/2}$ and $\mathbf{Z}_n = (\mathbf{Z}_{n1}, \ldots, \mathbf{Z}_{nn})^T = \mathbf{S} \mathbf{V}_n$ with $\mathbf{Z}_{nj} \in \mathbf{R}^s$ for $j = 1, \ldots, n$.

THEOREM 3. *Assume the conditions of Theorem 2 hold, the first and second order derivatives $f_i'(u)$ and $f_i''(u)$ are uniformly bounded in a small neighborhood of 0 for all $i = 1, \ldots, n$, and that $\|\mathbf{d}_0\|_2 = O(\sqrt{s/n}/\lambda_n)$, $\max_i \|\mathbf{H}^{1/2} \mathbf{Z}_{ni}\|_2 = o(s^{-7/2}(\log s)^{-1})$, and $\sqrt{n/s} \min_{1 \leq j \leq s} |\beta_j^*| \to \infty$. Then, with probability tending to 1 there exists a global minimizer $\hat{\boldsymbol{\beta}} = ((\hat{\boldsymbol{\beta}}_1^o)^T, \hat{\boldsymbol{\beta}}_2^T)^T$ of $L_n(\boldsymbol{\beta})$ such that $\hat{\boldsymbol{\beta}}_2 = 0$. Moreover,*

$$\mathbf{c}^T (\mathbf{Z}_n^T \mathbf{Z}_n)^{-1/2} \mathbf{V}_n^{-1} \left[ (\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*) + \frac{n\lambda_n}{2} \mathbf{V}_n^2 \tilde{\mathbf{d}}_0 \right] \xrightarrow{\mathscr{D}} N(0, \tau(1 - \tau)),$$

*where $\mathbf{c}$ is an arbitrary $s$-dimensional vector satisfying $\mathbf{c}^T \mathbf{c} = 1$, and $\tilde{\mathbf{d}}_0$ is an $s$-dimensional vector with the $j$th element $d_j \operatorname{sgn}(\beta_j^*)$.*

The proof of Theorem 3 is an extension of the proof on the asymptotic normality theorem for the LAD estimator in Pollard (1991), in which the theorem is proved for fixed dimensionality. The idea is to approximate $L_n(\boldsymbol{\beta}_1, \mathbf{0})$ in (2.4) by a sequence of quadratic functions, whose minimizers converge to normal distribution. Since $L_n(\boldsymbol{\beta}_1, \mathbf{0})$ and the quadratic approximation are close, their minimizers are also close, which results in the asymptotic normality in Theorem 3.

Theorem 3 assumes that $\max_i \|\mathbf{H}^{1/2} \mathbf{Z}_{ni}\|_2 = o(s^{-7/2}(\log s)^{-1})$. Since by definition $\sum_{i=1}^n \|\mathbf{H}^{1/2} \mathbf{Z}_{ni}\|_2^2 = s$, it is seen that the condition implies $s = o(n^{1/8})$. This assumption is made to guarantee that the quadratic approximation is close enough to $L_n(\boldsymbol{\beta}_1, \mathbf{0})$. When $s$ is finite, the condition becomes $\max_i \|\mathbf{Z}_{ni}\|_2 = o(1)$, as in Pollard (1991). Another important assumption is $\lambda_n \sqrt{n} \|\mathbf{d}_0\|_2 = O(\sqrt{s})$, which is imposed to make sure that the bias $2^{-1} n \lambda_n \mathbf{c}^T \mathbf{V}_n \tilde{\mathbf{d}}_0$ caused by the penalty term does not diverge. For instance, using R-Lasso will create a nondiminishing bias, and thus cannot be guaranteed to have asymptotic normality.

Note that we do not assume a parametric form of the error distribution. Thus, our oracle estimator is in fact a semiparametric estimator with the error density as the nuisance parameter. Heuristically speaking, Theorem 3 shows that the asymptotic variance of $\sqrt{n}(\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*)$ is $n\tau(1 - \tau)\mathbf{V}_n\mathbf{Z}_n^T\mathbf{Z}_n\mathbf{V}_n$. Since $\mathbf{V}_n = (\mathbf{S}^T\mathbf{H}\mathbf{S})^{-1/2}$ and $\mathbf{Z}_n = \mathbf{S}\mathbf{V}_n$, if the model errors $\varepsilon_i$ are i.i.d. with density function $f_\varepsilon(\cdot)$, then this asymptotic variance reduces to $\tau(1 - \tau)(n^{-1}f_\varepsilon^2(0)\mathbf{S}^T\mathbf{S})^{-1}$. In the random design case where the true covariate vectors $\{\mathbf{S}_i\}_{i=1}^n$ are i.i.d. observations, $n^{-1}\mathbf{S}^T\mathbf{S}$ converges to $E[\mathbf{S}_1^T\mathbf{S}_1]$ as $n \to \infty$, and the asymptotic variance reduces to $\tau(1 - \tau)(f_\varepsilon^2(0)E[\mathbf{S}_1^T\mathbf{S}_1])^{-1}$. This is the semiparametric efficiency bound derived by Newey and Powell (1990) for random designs. In fact, if we assume that $(\mathbf{x}_i, y_i)$ are i.i.d., then the conditions of Theorem 3 can hold with asymptotic probability one. Using similar arguments, it can be formally shown that $\sqrt{n}(\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*)$ is asymptotically normal with covariance matrix equal to the aforementioned semiparametric efficiency bound. Hence, our oracle estimator is semiparametric efficient.

## 5. Properties of the adaptive robust Lasso.

In previous sections, we have seen that the choice of the weight vector $\mathbf{d}$ plays a pivotal role for the WR-Lasso estimate to enjoy the model selection oracle property and asymptotic normality. In fact, conditions in Theorem 2 require that $\min_{j \geq s+1} d_j > c_3$ and that $\|\mathbf{d}_0\|_2$ does not diverge too fast. Theorem 3 imposes an even more stringent condition, $\|\mathbf{d}_0\|_2 = O(\sqrt{s/n}/\lambda_n)$, on the weight vector $\mathbf{d}_0$. For R-Lasso, $\|\mathbf{d}_0\|_2 = \sqrt{s}$ and these conditions become very restrictive. For example, the condition in Theorem 3 becomes $\lambda_n = O(n^{-1/2})$, which is too low for a thresholding level even for Gaussian errors. Hence, an adaptive choice of weights is needed to ensure that those conditions are satisfied. To this end, we propose a two-step procedure.

In the first step, we use R-Lasso, which gives the estimate $\hat{\boldsymbol{\beta}}^{\text{ini}}$. As has been shown in Belloni and Chernozhukov (2011) and Wang (2013), R-Lasso is consistent at a near-oracle rate $\sqrt{s(\log p)/n}$ and selects the true model $\mathcal{M}^*$ as a submodel [in other words, R-Lasso has the sure screening property using the terminology of Fan and Lv (2008)] with asymptotic probability one, namely,

$$\text{supp}(\hat{\boldsymbol{\beta}}^{\text{ini}}) \supseteq \text{supp}(\boldsymbol{\beta}^*) \quad \text{and} \quad \|\hat{\boldsymbol{\beta}}_1^{\text{ini}} - \boldsymbol{\beta}_1^*\|_2 = O(\sqrt{s(\log p)/n}).$$

We remark that our Theorem 2 also ensures the consistency of R-Lasso. Compared to Belloni and Chernozhukov (2011), Theorem 2 presents stronger results but also needs more restrictive conditions for R-Lasso. As will be shown in latter theorems, only the consistency of R-Lasso is needed in the study of AR-Lasso, so we quote the results and conditions on R-Lasso in Belloni and Chernozhukov (2011) with the mind of imposing weaker conditions.

In the second step, we set $\hat{\mathbf{d}} = (\hat{d}_1, \ldots, \hat{d}_p)^T$ with $\hat{d}_j = p'_{\lambda_n}(|\hat{\beta}_j^{\text{ini}}|)$ where $p_{\lambda_n}(|\cdot|)$ is a folded-concave penalty function, and then solve the regularization problem (2.4) with a newly computed weight vector. Thus, vector $\hat{\mathbf{d}}_0$ is expected

to be close to the vector $(p'_{\lambda_n}(|\beta_1^*|), \ldots, p'_{\lambda_n}(|\beta_s^*|))^T$ under $L_2$-norm. If a folded-concave penalty such as SCAD is used, then $p'_{\lambda_n}(|\beta_j^*|)$ will be close, or even equal, to zero for $1 \le j \le s$, and thus the magnitude of $\|\hat{\mathbf{d}}_0\|_2$ is negligible.

Now, we formally establish the asymptotic properties of AR-Lasso. We first present a more general result and then highlight our recommended procedure, which uses R-Lasso as the initial estimate and then uses SCAD to compute the stochastic weights, in Corollary 1. Denote by $\mathbf{d}^* = (d_1^*, \ldots, d_p^*)$ with $d_j^* = p'_{\lambda_n}(|\beta_j^*|)$. Using the weight vector $\hat{\mathbf{d}}$, AR-Lasso minimizes the following objective function:

$$(5.1) \qquad \widehat{L}_n(\boldsymbol{\beta}) = \sum_{i=1}^n \rho_\tau (y_i - \mathbf{x}_i^T \boldsymbol{\beta}) + n\lambda_n \|\hat{\mathbf{d}} \circ \boldsymbol{\beta}\|_1.$$

We also need the following conditions to show the model selection oracle property of the two-step procedure.

CONDITION 4. *With asymptotic probability one, the initial estimate satisfies* $\|\hat{\boldsymbol{\beta}}^{\mathrm{ini}} - \boldsymbol{\beta}^*\|_2 \le C_2\sqrt{s(\log p)/n}$ *with some constant $C_2 > 0$.*

As discussed above, if R-Lasso is used to obtain the initial estimate, it satisfies the above condition. Our second condition is on the penalty function.

CONDITION 5. $p'_{\lambda_n}(t)$ *is nonincreasing in $t \in (0, \infty)$ and is Lipschitz with constant $c_5$, that is,*

$$\left| p'_{\lambda_n}(|\beta_1|) - p'_{\lambda_n}(|\beta_2|) \right| \le c_5 |\beta_1 - \beta_2|$$

*for any $\beta_1, \beta_2 \in \mathbf{R}$. Moreover, $p'_{\lambda_n}(C_2\sqrt{s(\log p)/n}) > \frac{1}{2}p'_{\lambda_n}(0+)$ for large enough $n$, where $C_2$ is defined in Condition 4.*

For the SCAD [Fan and Li (2001)] penalty, $p'_{\lambda_n}(\beta)$ is given by

$$(5.2) \qquad p'_{\lambda_n}(\beta) = 1\{\beta \le \lambda_n\} + \frac{(a\lambda_n - \beta)_+}{(a-1)\lambda_n}1\{\beta > \lambda_n\}$$

for a given constant $a > 2$, and it can be easily verified that Condition 5 holds if $\lambda_n > 2(a+1)^{-1}C_2\sqrt{s(\log p)/n}$.

THEOREM 4. *Assume conditions of Theorem 2 hold with $\mathbf{d} = \mathbf{d}^*$ and $\gamma_n = a_n$, where*

$$a_n = C_3\left(\sqrt{s(\log n)/n} + \lambda_n\left(\|\mathbf{d}_0^*\|_2 + C_2c_5\sqrt{s(\log p)/n}\right)\right)$$

*with some constant $C_3 > 0$ and $\lambda_n s\kappa_n\sqrt{(\log p)/n} \to 0$. Then, under Conditions 4 and 5, with probability tending to one, there exists a global minimizer $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}_1^T, \hat{\boldsymbol{\beta}}_2^T)^T$ of (5.1) such that $\hat{\boldsymbol{\beta}}_2 = \mathbf{0}$ and $\|\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*\|_2 \le a_n$.*

The results in Theorem 4 are analogous to those in Theorem 2. The extra term $\lambda_n \sqrt{s(\log p)/n}$ in the convergence rate $a_n$, compared to the convergence rate $\gamma_n$ in Theorem 2, is caused by the bias of the initial estimate $\hat{\boldsymbol{\beta}}^{\text{ini}}$. Since the regularization parameter $\lambda_n$ goes to zero, the bias of AR-Lasso is much smaller than that of the initial estimator $\hat{\boldsymbol{\beta}}^{\text{ini}}$. Moreover, the AR-Lasso $\hat{\boldsymbol{\beta}}$ possesses the model selection oracle property.

Now we present the asymptotic normality of the AR-Lasso estimate.

CONDITION 6. *The smallest signal satisfies* $\min_{1 \le j \le s} |\beta_j^*| > 2C_2\sqrt{(s \log p)/n}$. *Moreover, it holds that* $p_{\lambda_n}''(|\beta|) = o(s^{-1}\lambda_n^{-1}(n \log p)^{-1/2})$ *for any* $|\beta| > 2^{-1}\min_{1 \le j \le s} |\beta_j^*|$.

The above condition on the penalty function is satisfied when the SCAD penalty is used and $\min_{1 \le j \le s} |\beta_j^*| \ge 2a\lambda_n$ where $a$ is the parameter in the SCAD penalty (5.2).

THEOREM 5. *Assume conditions of Theorem 3 hold with* $\mathbf{d} = \mathbf{d}^*$ *and* $\gamma_n = a_n$, *where* $a_n$ *is defined in Theorem 4. Then, under Conditions 4–6, with asymptotic probability one, there exists a global minimizer* $\hat{\boldsymbol{\beta}}$ *of* (5.1) *having the same asymptotic properties as those in Theorem 3.*

With the SCAD penalty, conditions in Theorems 4 and 5 can be simplified and AR-Lasso still enjoys the same asymptotic properties, as presented in the following corollary.

COROLLARY 1. *Assume* $\lambda_n = O(\sqrt{s(\log p)(\log \log n)/n})$, $\log p = o(\sqrt{n})$, $\min_{1 \le j \le s} |\beta_j^*| \ge 2a\lambda_n$ *with* $a$ *the parameter in the SCAD penalty and* $\kappa_n = o(n^{1/4}s^{-1/2}(\log n)^{-3/2}(\log p)^{1/2})$. *Further assume that* $\|n^{-1}\mathbf{Q}^T\mathbf{H}\mathbf{S}\|_{2,\infty} < C_4\sqrt{(\log p)(\log \log n)/\log n}$ *with* $C_4$ *some positive constant. Then, under Conditions 1 and 2, with asymptotic probability one, there exists a global minimizer* $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}_1^T, \hat{\boldsymbol{\beta}}_2^T)^T$ *of* $\widehat{L}_n(\boldsymbol{\beta})$ *such that*

$$\|\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*\|_2 \le O(\sqrt{s(\log n)/n}), \qquad \text{sgn}(\hat{\boldsymbol{\beta}}_1) = \text{sgn}(\boldsymbol{\beta}_1^*) \quad and \quad \hat{\boldsymbol{\beta}}_2 = \mathbf{0}.$$

*If in addition,* $\max_i \|\mathbf{H}^{1/2}\mathbf{Z}_{ni}\|_2 = o(s^{-7/2}(\log s)^{-1})$, *then we also have*

$$\mathbf{c}^T(\mathbf{Z}_n^T\mathbf{Z}_n)^{-1/2}\mathbf{V}_n^{-1}(\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*) \xrightarrow{\mathscr{D}} N(0, \tau(1-\tau)),$$

*where* $\mathbf{c}$ *is an arbitrary* $s$-*dimensional vector satisfying* $\mathbf{c}^T\mathbf{c} = 1$.

Corollary 1 provides sufficient conditions for ensuring the variable selection sign consistency of AR-Lasso. These conditions require that R-Lasso in the initial step has the sure screening property. We remark that in implementation,

AR-Lasso is able to select the variables missed by R-Lasso, as demonstrated in our numerical studies in the next section. The theoretical comparison of the variable selection results of R-Lasso and AR-Lasso would be an interesting topic for future study. One set of $(p, n, s, \kappa_n)$ satisfying conditions in Corollary 1 is $\log p = O(n^{b_1})$, $s = o(n^{(1-b_1)/2})$ and $\kappa_n = o(n^{b_1/4}(\log n)^{-3/2})$ with $b_1 \in (0, 1/2)$ some constant. Corollary 1 gives one specific choice of $\lambda_n$, not necessarily the smallest $\lambda_n$, which makes our procedure work. In fact, the condition on $\lambda_n$ can be weakened to $\lambda_n > 2(a + 1)^{-1}\|\hat{\boldsymbol{\beta}}_1^{\text{ini}} - \boldsymbol{\beta}_1\|_\infty$. Currently, we use the $L_2$-norm $\|\hat{\boldsymbol{\beta}}_1^{\text{ini}} - \boldsymbol{\beta}_1\|_2$ to bound this $L_\infty$-norm, which is too crude. If one can establish $\|\hat{\boldsymbol{\beta}}_1^{\text{ini}} - \boldsymbol{\beta}_1\|_\infty = O_p(\sqrt{n^{-1}\log p})$ for an initial estimator $\hat{\boldsymbol{\beta}}_1^{\text{ini}}$, then the choice of $\lambda_n$ can be as small as $O(\sqrt{n^{-1}\log p})$, the same order as that used in Wang (2013). On the other hand, since we are using AR-Lasso, the choice of $\lambda_n$ is not as sensitive as R-Lasso.

**6. Numerical studies.** In this section, we evaluate the finite sample property of our proposed estimator with synthetic data. Please see the supplementary material [Fan, Fan and Barut (2014)] for a real life data-set analysis, where we provide results of an eQTL study on the *CHRNA6* gene.

To assess the performance of the proposed estimator and compare it with other methods, we simulated data from the high-dimensional linear regression model

$$y_i = \mathbf{x}_i^T \boldsymbol{\beta}_0 + \varepsilon_i, \qquad \mathbf{x} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_{\mathbf{x}}),$$

where the data had $n = 100$ observations and the number of parameters was chosen as $p = 400$. We fixed the true regression coefficient vector as

$$\boldsymbol{\beta}_0 = \{2, 0, 1.5, 0, 0.80, 0, 0, 1, 0, 1.75, 0, 0, 0.75, 0, 0, 0.3, 0, \ldots, 0\}.$$

For the distribution of the noise, $\varepsilon$, we considered six symmetric distributions: normal with variance 2 ($\mathcal{N}(0, 2)$), a scale mixture of Normals for which $\sigma_i^2 = 1$ with probability 0.9 and $\sigma_i^2 = 25$ otherwise (MN$_1$), a different scale mixture model where $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$ and $\sigma_i \sim \text{Unif}(1, 5)$ (MN$_2$), Laplace, Student's $t$ with degrees of freedom 4 with doubled variance ($\sqrt{2} \times t_4$) and Cauchy. We take $\tau = 0.5$, corresponding to $L_1$-regression, throughout the simulation. Correlation of the covariates, $\boldsymbol{\Sigma}_{\mathbf{x}}$ were either chosen to be identity (i.e., $\boldsymbol{\Sigma}_{\mathbf{x}} = I_p$) or they were generated from an AR(1) model with correlation 0.5, that is $\boldsymbol{\Sigma}_{\mathbf{x}_{(i,j)}} = 0.5^{|i-j|}$.

We implemented five methods for each setting:

1. $L_2$-*Oracle*, which is the least squares estimator based on the signal covariates.
2. *Lasso*, the penalized least-squares estimator with $L_1$-penalty as in Tibshirani (1996).
3. *SCAD*, the penalized least-squares estimator with SCAD penalty as in Fan and Li (2001).

4. *R-Lasso*, the robust Lasso defined as the minimizer of (2.4) with $\mathbf{d} = \mathbf{1}$.
5. *AR-Lasso*, which is the adaptive robust Lasso whose adaptive weights on the penalty function were computed based on the SCAD penalty using the R-Lasso estimate as an initial value.

The tuning parameter, $\lambda_n$, was chosen optimally based on 100 validation data-sets. For each of these data-sets, we ran a grid search to find the best $\lambda_n$ (with the lowest $L_2$ error for $\boldsymbol{\beta}$) for the particular setting. This optimal $\lambda_n$ was recorded for each of the 100 validation data-sets. The median of these 100 optimal $\lambda_n$ were used in the simulation studies. We preferred this procedure over cross-validation because of the instability of the $L_2$ loss under heavy tails.

The following four performance measures were calculated:

1. $L_2$ loss, which is defined as $\|\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}}\|_2$.
2. $L_1$ loss, which is defined as $\|\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}}\|_1$.
3. Number of noise covariates that are included in the model, that is the number of false positives (FP).
4. Number of signal covariates that are not included, that is, the number of false negatives (FN).

For each setting, we present the average of the performance measure based on 100 simulations. The results are depicted in Tables 1 and 2. A boxplot of the $L_2$

TABLE 1
*Simulation results with independent covariates*

| | | $L_2$ Oracle | Lasso | SCAD | R-Lasso | AR-Lasso |
|---|---|---|---|---|---|---|
| $\mathcal{N}(0, 2)$ | $L_2$ loss | 0.833 | 4.114 | 3.412 | 5.342 | 2.662 |
| | $L_1$ loss | 0.380 | 1.047 | 0.819 | 1.169 | 0.785 |
| | FP, FN | – | 27.00, 0.49 | 29.60, 0.51 | 36.81, 0.62 | 17.27, 0.70 |
| MN$_1$ | $L_2$ loss | 0.977 | 5.232 | 4.736 | 4.525 | 2.039 |
| | $L_1$ loss | 0.446 | 1.304 | 1.113 | 1.028 | 0.598 |
| | FP, FN | – | 26.80, 0.73 | 29.29, 0.68 | 34.26, 0.51 | 16.76, 0.51 |
| MN$_2$ | $L_2$ loss | 1.886 | 7.563 | 7.583 | 8.121 | 5.647 |
| | $L_1$ loss | 0.861 | 2.085 | 2.007 | 2.083 | 1.845 |
| | FP, FN | – | 20.39, 2.28 | 23.25, 2.19 | 24.64, 2.29 | 11.97, 2.57 |
| Laplace | $L_2$ loss | 0.795 | 4.056 | 3.395 | 4.610 | 2.025 |
| | $L_1$ loss | 0.366 | 1.016 | 0.799 | 1.039 | 0.573 |
| | FP, FN | – | 26.87, 0.62 | 29.98, 0.49 | 34.76, 0.48 | 18.81, 0.40 |
| $\sqrt{2} \times t_4$ | $L_2$ loss | 1.087 | 5.303 | 5.859 | 6.185 | 3.266 |
| | $L_1$ loss | 0.502 | 1.378 | 1.256 | 1.403 | 0.951 |
| | FP, FN | – | 24.61, 0.85 | 36.95, 0.76 | 33.84, 0.84 | 18.53, 0.82 |
| Cauchy | $L_2$ loss | 37.451 | 211.699 | 266.088 | 6.647 | 3.587 |
| | $L_1$ loss | 17.136 | 30.052 | 40.041 | 1.646 | 1.081 |
| | FP, FN | – | 27.39, 5.78 | 34.32, 5.94 | 27.33, 1.41 | 17.28, 1.10 |

TABLE 2
*Simulation results with correlated covariates*

| | | $L_2$ Oracle | Lasso | SCAD | R-Lasso | AR-Lasso |
|---|---|---|---|---|---|---|
| $\mathcal{N}(0,2)$ | $L_2$ loss | 0.836 | 3.440 | 3.003 | 4.185 | 2.580 |
| | $L_1$ loss | 0.375 | 0.943 | 0.803 | 1.079 | 0.806 |
| | FP, FN | – | 20.62, 0.59 | 23.13, 0.56 | 22.72, 0.77 | 14.49, 0.74 |
| $MN_1$ | $L_2$ loss | 1.081 | 4.415 | 3.589 | 3.652 | 1.829 |
| | $L_1$ loss | 0.495 | 1.211 | 1.055 | 0.901 | 0.593 |
| | FP, FN | – | 18.66, 0.77 | 15.71, 0.75 | 26.65, 0.60 | 13.29, 0.51 |
| $MN_2$ | $L_2$ loss | 1.858 | 6.427 | 6.249 | 6.882 | 4.890 |
| | $L_1$ loss | 0.844 | 1.899 | 1.876 | 1.916 | 1.785 |
| | FP, FN | – | 15.16, 2.08 | 14.77, 1.96 | 18.22, 1.91 | 7.86, 2.71 |
| Laplace | $L_2$ loss | 0.803 | 3.341 | 2.909 | 3.606 | 1.785 |
| | $L_1$ loss | 0.371 | 0.931 | 0.781 | 0.927 | 0.573 |
| | FP, FN | – | 19.32, 0.62 | 21.60, 0.38 | 24.44, 0.46 | 12.90, 0.55 |
| $\sqrt{2} \times t_4$ | $L_2$ loss | 1.122 | 4.474 | 4.259 | 4.980 | 2.855 |
| | $L_1$ loss | 0.518 | 1.222 | 1.201 | 1.299 | 0.946 |
| | FP, FN | – | 20.00, 0.76 | 18.49, 0.91 | 23.56, 0.79 | 13.40, 1.05 |
| Cauchy | $L_2$ loss | 31.095 | 217.395 | 243.141 | 5.388 | 3.286 |
| | $L_1$ loss | 13.978 | 31.361 | 36.624 | 1.461 | 1.074 |
| | FP, FN | – | 25.59, 5.48 | 32.01, 5.43 | 20.80, 1.16 | 12.45, 1.17 |

losses under different noise settings is also given in Figure 1 (the $L_2$ loss boxplot for the independent covariate setting is similar and omitted). For the results in Tables 1 and 2, one should compare the performance between Lasso and R-Lasso and that between SCAD and AR-Lasso. This comparison reflects the effectiveness of $L_1$-regression in dealing with heavy-tail distributions. Furthermore, comparing Lasso with SCAD, and R-Lasso with AR-Lasso, shows the effectiveness of using adaptive weights in the penalty function.

Our simulation results reveal the following facts. The quantile based estimators were more robust in dealing with the outliers. For example, for the first mixture model ($MN_1$) and Cauchy, R-Lasso outperformed Lasso, and AR-Lasso outperformed SCAD in all of the four metrics, and significantly so when the error distribution is the Cauchy distribution. On the other hand, for the light-tail distributions such as the normal distribution, the efficiency loss was limited. When the tails get heavier, for instance, for the Laplace distribution, quantile based methods started to outperform the least-squares based approaches, more so when the tails got heavier.

The effectiveness of weights in AR-Lasso is self-evident. SCAD outperformed Lasso and AR-Lasso outperformed R-Lasso in almost all of the settings. Furthermore, for all of the error settings AR-Lasso had significantly lower $L_2$ and $L_1$ loss as well as a smaller model size compared to other estimators.

(a) Normal(0,2)                    (b) $MN_1$

(c) $MN_2$                    (d) Laplace

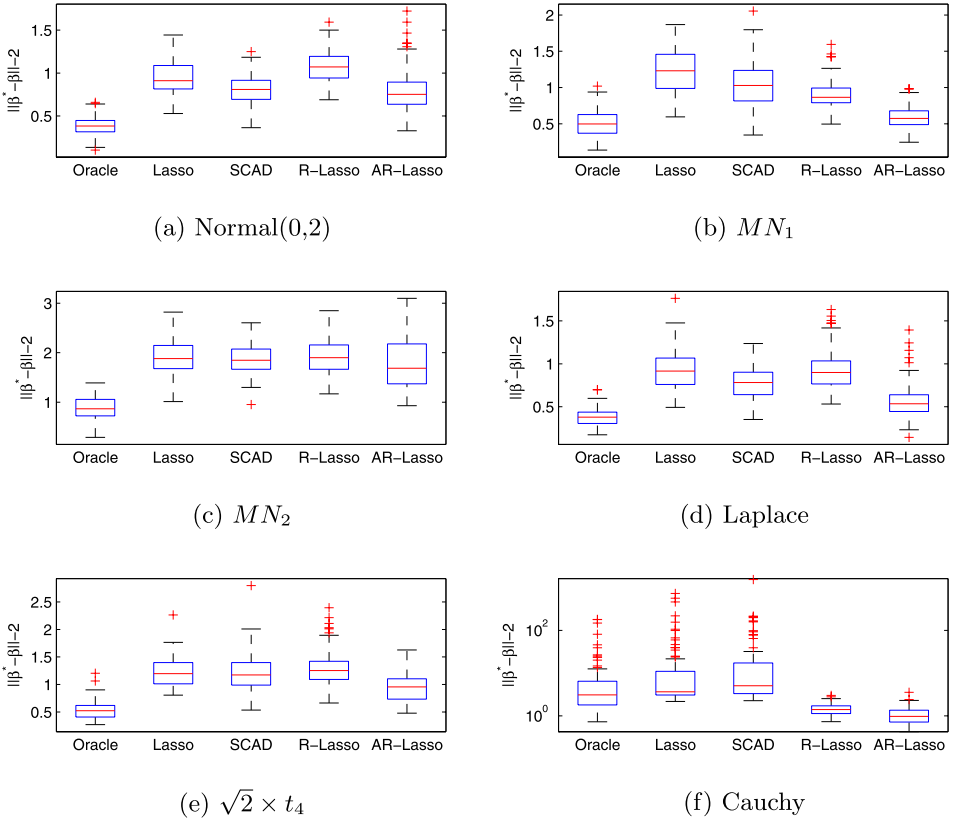(e) $\sqrt{2} \times t_4$                    (f) Cauchy

FIG. 1.    *Boxplots for $L_2$ loss with correlated covariates.*

It is seen that when the noise does not have heavy tails, that is for the normal and the Laplace distribution, all the estimators are comparable in terms of $L_1$ loss. As expected, estimators that minimize squared loss worked better than R-Lasso and AR-Lasso estimators under Gaussian noise, but their performances deteriorated as the tails got heavier. In addition, in the two heteroscedastic settings, AR-Lasso had the best performance among others.

For Cauchy noise, least squares methods could only recover 1 or 2 of the true variables on average. On the other hand, $L_1$-estimators (R-Lasso and AR-Lasso) had very few false negatives, and as evident from $L_2$ loss values, these estimators only missed variables with smaller magnitudes.

In addition, AR-Lasso consistently selected a smaller set of variables than R-Lasso. For instance, for the setting with independent covariates, under the Laplace distribution, R-Lasso and AR-Lasso had on average 34.76 and 18.81 false positives, respectively. Also note that AR-Lasso consistently outperformed R-Lasso: it estimated $\boldsymbol{\beta}^*$ (lower $L_1$ and $L_2$ losses), and the support of $\boldsymbol{\beta}^*$ (lower averages for the number of false positives) more efficiently.

**7. Proofs.** In this section, we prove Theorems 1, 2 and 4 and provide the lemmas used in these proofs. The proofs of Theorems 3 and 5 and Proposition 1 are given in the supplementary Appendix [Fan, Fan and Barut (2014)].

We use techniques from empirical process theory to prove the theoretical results. Let $v_n(\boldsymbol{\beta}) = \sum_{i=1}^{n} \rho_\tau(y_i - \mathbf{x}_i^T \boldsymbol{\beta})$. Then $L_n(\boldsymbol{\beta}) = v_n(\boldsymbol{\beta}) + n\lambda_n \sum_{j=1}^{p} d_j |\beta_j|$. For a given deterministic $M > 0$, define the set

$$\mathcal{B}_0(M) = \{\boldsymbol{\beta} \in \mathbf{R}^p : \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2 \leq M, \operatorname{supp}(\boldsymbol{\beta}) \subseteq \operatorname{supp}(\boldsymbol{\beta}^*)\}.$$

Then define the function

$$(7.1) \qquad Z_n(M) = \sup_{\boldsymbol{\beta} \in \mathcal{B}_0(M)} \frac{1}{n} |(v_n(\boldsymbol{\beta}) - v_n(\boldsymbol{\beta}^*)) - E(v_n(\boldsymbol{\beta}) - v_n(\boldsymbol{\beta}^*))|.$$

Lemma 1 in Section 7.4 gives the rate of convergence for $Z_n(M)$.

7.1. *Proof of Theorem* 1. We first show that for any $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \mathbf{0}^T)^T \in \mathcal{B}_0(M)$ with $M = o(\kappa_n^{-1} s^{-1/2})$,

$$(7.2) \qquad E[v_n(\boldsymbol{\beta}) - v_n(\boldsymbol{\beta}^*)] \geq \tfrac{1}{2} c_0 c n \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2^2$$

for sufficiently large $n$, where $c$ is the lower bound for $f_i(\cdot)$ in the neighborhood of 0. The intuition follows from the fact that $\boldsymbol{\beta}^*$ is the minimizer of the function $Ev_n(\boldsymbol{\beta})$, and hence in Taylor's expansion of $E[v_n(\boldsymbol{\beta}) - v_n(\boldsymbol{\beta}^*)]$ around $\boldsymbol{\beta}^*$, the first-order derivative is zero at the point $\boldsymbol{\beta} = \boldsymbol{\beta}^*$. The left-hand side of (7.2) will be controlled by $Z_n(M)$. This yields the $L_2$-rate of convergence in Theorem 1.

To prove (7.2), we set $a_i = |\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)|$. Then, for $\boldsymbol{\beta} \in \mathcal{B}_0(M)$,

$$|a_i| \leq \|\mathbf{S}_i\|_2 \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2 \leq \sqrt{s}\kappa_n M \to 0.$$

Thus, if $\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) > 0$, by $E1\{\varepsilon_i \leq 0\} = \tau$, Fubini's theorem, mean value theorem and Condition 1 it is easy to derive that

$$(7.3) \qquad \begin{aligned} &E[\rho_\tau(\varepsilon_i - a_i) - \rho_\tau(\varepsilon_i)] \\ &\quad = E[a_i(1\{\varepsilon_i \leq a_i\} - \tau) - \varepsilon_i 1\{0 \leq \varepsilon_i \leq a_i\}] \\ &\quad = E\left[\int_0^{a_i} 1\{0 \leq \varepsilon_i \leq s\} ds\right] \\ &\quad = \int_0^{a_i} (F_i(s) - F_i(0)) ds = \frac{1}{2} f_i(0) a_i^2 + o(1) a_i^2, \end{aligned}$$

where the $o(1)$ is uniformly over all $i = 1, \dots, n$. When $\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) < 0$, the same result can be obtained. Furthermore, by Condition 2,

$$\sum_{i=1}^{n} f_i(0) a_i^2 = (\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)^T \mathbf{S}^T \mathbf{H} \mathbf{S}(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) \geq c_0 n \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2^2.$$

This together with (7.3) and the definition of $v_n(\boldsymbol{\beta})$ proves (7.2).

The inequality (7.2) holds for any $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \mathbf{0}^T)^T \in \mathcal{B}_0(M)$, yet $\hat{\boldsymbol{\beta}}^o = ((\hat{\boldsymbol{\beta}}_1^o)^T, \mathbf{0}^T)^T$ may not be in the set. Thus, we let $\tilde{\boldsymbol{\beta}} = (\tilde{\boldsymbol{\beta}}_1^T, \mathbf{0}^T)^T$, where

$$\tilde{\boldsymbol{\beta}}_1 = u\hat{\boldsymbol{\beta}}_1^o + (1-u)\boldsymbol{\beta}_1^* \qquad \text{with } u = M/(M + \|\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*\|_2),$$

which falls in the set $\mathcal{B}_0(M)$. Then, by the convexity and the definition of $\hat{\boldsymbol{\beta}}_1^o$,

$$L_n(\tilde{\boldsymbol{\beta}}) \le uL_n(\hat{\boldsymbol{\beta}}_1^o, \mathbf{0}) + (1-u)L_n(\boldsymbol{\beta}_1^*, \mathbf{0}) \le L_n(\boldsymbol{\beta}_1^*, \mathbf{0}) = L_n(\boldsymbol{\beta}^*).$$

Using this and the triangle inequality, we have

(7.4)
$$\begin{aligned}
E[v_n(\tilde{\boldsymbol{\beta}}) - v_n(\boldsymbol{\beta}^*)] \\
&= \{v_n(\boldsymbol{\beta}^*) - Ev_n(\boldsymbol{\beta}^*)\} - \{v_n(\tilde{\boldsymbol{\beta}}) - Ev_n(\tilde{\boldsymbol{\beta}})\} \\
&\quad + L_n(\tilde{\boldsymbol{\beta}}) - L_n(\boldsymbol{\beta}^*) + n\lambda_n\|\mathbf{d}_0 \circ \boldsymbol{\beta}_1^*\|_1 - n\lambda_n\|\mathbf{d}_0 \circ \tilde{\boldsymbol{\beta}}_1\|_1 \\
&\le nZ_n(M) + n\lambda_n\|\mathbf{d}_0 \circ (\boldsymbol{\beta}_1^* - \tilde{\boldsymbol{\beta}}_1)\|_1.
\end{aligned}$$

By the Cauchy–Schwarz inequality, the very last term is bounded by $n\lambda_n\|\mathbf{d}_0\|_2\|\tilde{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*\|_2 \le n\lambda_n\|\mathbf{d}_0\|_2 M$.

Define the event $\mathcal{E}_n = \{Z_n(M) \le 2Mn^{-1/2}\sqrt{s\log n}\}$. Then by Lemma 1,

(7.5)                    $$P(\mathcal{E}_n) \ge 1 - \exp(-c_0 s(\log n)/8).$$

On the event $\mathcal{E}_n$, by (7.4), we have

$$E[v_n(\tilde{\boldsymbol{\beta}}) - v_n(\boldsymbol{\beta}^*)] \le 2M\sqrt{sn(\log n)} + n\lambda_n\|\mathbf{d}_0\|_2 M.$$

Taking $M = 2\sqrt{s/n} + \lambda_n\|\mathbf{d}_0\|_2$. By Condition 2 and the assumption $\lambda_n\|\mathbf{d}_0\|_2\sqrt{s}\kappa_n \to 0$, it is easy to check that $M = o(\kappa_n^{-1}s^{-1/2})$. Combining these two results with (7.2), we obtain that on the event $\mathcal{E}_n$,

$$\tfrac{1}{2}c_0 n\|\tilde{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*\|_2^2 \le (2\sqrt{sn(\log n)} + n\lambda_n\|\mathbf{d}_0\|_2)(2\sqrt{s/n} + \lambda_n\|\mathbf{d}_0\|_2),$$

which entails that

$$\|\boldsymbol{\beta}_1^* - \tilde{\boldsymbol{\beta}}_1\|_2 \le O(\lambda_n\|\mathbf{d}_0\|_2 + \sqrt{s(\log n)/n}).$$

Note that $\|\boldsymbol{\beta}_1^* - \tilde{\boldsymbol{\beta}}\|_2 \le M/2$ implies $\|\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*\|_2 \le M$. Thus, on the event $\mathcal{E}_n$,

$$\|\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*\|_2 \le O(\lambda_n\|\mathbf{d}_0\|_2 + \sqrt{s(\log n)/n}).$$

The second result follows trivially.

7.2. *Proof of Theorem* 2. Since $\hat{\boldsymbol{\beta}}_1^o$ defined in Theorem 1 is a minimizer of $L_n(\boldsymbol{\beta}_1, \mathbf{0})$, it satisfies the KKT conditions. To prove that $\hat{\boldsymbol{\beta}} = ((\hat{\boldsymbol{\beta}}_1^o)^T, \mathbf{0}^T)^T \in \mathbf{R}^p$ is a global minimizer of $L_n(\boldsymbol{\beta})$ in the original $\mathbf{R}^p$ space, we only need to check the following condition:

$$\text{(7.6)} \qquad \left\| \mathbf{d}_1^{-1} \circ \mathbf{Q}^T \rho_\tau'(\mathbf{y} - \mathbf{S}\hat{\boldsymbol{\beta}}_1^o) \right\|_\infty < n\lambda_n,$$

where $\rho_\tau'(\mathbf{u}) = (\rho_\tau'(u_1), \ldots, \rho_\tau'(u_n))^T$ for any $n$-vector $\mathbf{u} = (u_1, \ldots, u_n)^T$ with $\rho_\tau'(u_i) = \tau - 1\{u_i \le 0\}$. Here, $\mathbf{d}_1^{-1}$ denotes the vector $(d_{s+1}^{-1}, \ldots, d_p^{-1})^T$. Then the KKT conditions and the convexity of $L_n(\boldsymbol{\beta})$ together ensure that $\hat{\boldsymbol{\beta}}$ is a global minimizer of $L(\boldsymbol{\beta})$.

Define events

$$A_1 = \left\{ \|\hat{\boldsymbol{\beta}}_1^o - \boldsymbol{\beta}_1^*\|_2 \le \gamma_n \right\}, \qquad A_2 = \left\{ \sup_{\boldsymbol{\beta} \in \mathcal{N}} \left\| \mathbf{d}_1^{-1} \circ \mathbf{Q}^T \rho_\tau'(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1) \right\|_\infty < n\lambda_n \right\},$$

where $\gamma_n$ is defined in Theorem 1 and

$$\mathcal{N} = \left\{ \boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T \in R^p : \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2 \le \gamma_n, \boldsymbol{\beta}_2 = \mathbf{0} \in \mathbf{R}^{p-s} \right\}.$$

Then by Theorem 1 and Lemma 2 in Section 7.4, $P(A_1 \cap A_2) \ge 1 - o(n^{-cs})$. Since $\hat{\boldsymbol{\beta}} \in \mathcal{N}$ on the event $A_1$, the inequality (7.6) holds on the event $A_1 \cap A_2$. This completes the proof of Theorem 2.

7.3. *Proof of Theorem* 4. The idea of the proof follows those used in the proof of Theorems 1 and 2. We first consider the minimizer of $\widehat{L}_n(\boldsymbol{\beta})$ in the subspace $\{\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T \in \mathbf{R}^p : \boldsymbol{\beta}_2 = \mathbf{0}\}$. Let $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \mathbf{0})^T$, where $\boldsymbol{\beta}_1 = \boldsymbol{\beta}_1^* + \tilde{a}_n \mathbf{v}_1 \in \mathbf{R}^s$ with $\tilde{a}_n = \sqrt{s(\log n)/n} + \lambda_n(\|\mathbf{d}_0^*\|_2 + C_2 c_5 \sqrt{s(\log p)/n})$, $\|\mathbf{v}_1\|_2 = C$, and $C > 0$ is some large enough constant. By the assumptions in the theorem, we have $\tilde{a}_n = o(\kappa_n^{-1} s^{-1/2})$. Note that

$$\text{(7.7)} \qquad \widehat{L}_n(\boldsymbol{\beta}_1^* + \tilde{a}_n \mathbf{v}_1, \mathbf{0}) - \widehat{L}_n(\boldsymbol{\beta}_1^*, \mathbf{0}) = I_1(\mathbf{v}_1) + I_2(\mathbf{v}_1),$$

where $I_1(\mathbf{v}_1) = \|\rho_\tau(\mathbf{y} - \mathbf{S}(\boldsymbol{\beta}_1^* + \tilde{a}_n \mathbf{v}_1))\|_1 - \|\rho_\tau(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1^*)\|_1$ and $I_2(\mathbf{v}_1) = n\lambda_n(\|\hat{\mathbf{d}}_0 \circ (\boldsymbol{\beta}_1^* + \tilde{a}_n \mathbf{v}_1)\|_1 - \|\hat{\mathbf{d}}_0 \circ \boldsymbol{\beta}_1^*\|_1)$ with $\|\rho_\tau(\mathbf{u})\|_1 = \sum_{i=1}^n \rho_\tau(u_i)$ for any vector $\mathbf{u} = (u_1, \ldots, u_n)^T$. By the results in the proof of Theorem 1, $E[I_1(\mathbf{v}_1)] \ge 2^{-1} c_0 n \|\tilde{a}_n \mathbf{v}_1\|_2^2$, and moreover, with probability at least $1 - n^{-cs}$,

$$\left| I_1(\mathbf{v}_1) - E[I_1(\mathbf{v}_1)] \right| \le n Z_n(C a_n) \le 2\tilde{a}_n \sqrt{s(\log n)n} \|\mathbf{v}_1\|_2.$$

Thus, by the triangle inequality,

$$\text{(7.8)} \qquad I_1(\mathbf{v}_1) \ge 2^{-1} c_0 \tilde{a}_n^2 n \|\mathbf{v}_1\|_2^2 - 2\tilde{a}_n \sqrt{s(\log n)n} \|\mathbf{v}_1\|_2.$$

The second term on the right-hand side of (7.7) can be bounded as

$$\text{(7.9)} \qquad \left| I_2(\mathbf{v}_1) \right| \le n\lambda_n \|\hat{\mathbf{d}} \circ (\tilde{a}_n \mathbf{v}_1)\|_1 \le n a_n \lambda_n \|\hat{\mathbf{d}}_0\|_2 \|\mathbf{v}_1\|_2.$$

By triangle inequality and Conditions 4 and 5, it holds that

$$
(7.10) \quad
\begin{aligned}
\|\hat{\mathbf{d}}_0\|_2 &\le \|\hat{\mathbf{d}}_0 - \mathbf{d}_0^*\|_2 + \|\mathbf{d}_0^*\|_2 \\
&\le c_5 \|\hat{\boldsymbol{\beta}}_1^{\mathrm{ini}} - \boldsymbol{\beta}_1^*\|_2 + \|\mathbf{d}_0^*\|_2 \le C_2 c_5 \sqrt{s(\log p)/n} + \|\mathbf{d}_0^*\|_2.
\end{aligned}
$$

Thus, combining (7.7)–(7.10) yields

$$
\begin{aligned}
\widehat{L}_n(\boldsymbol{\beta}^* + \tilde{a}_n \mathbf{v}_1) - \widehat{L}_n(\boldsymbol{\beta}^*) \ge{}& 2^{-1} c_0 n a_n^2 \|\mathbf{v}_1\|_2^2 - 2\tilde{a}_n \sqrt{s(\log n)n} \|\mathbf{v}_1\|_2 \\
&- n a_n \lambda_n \big(\|\mathbf{d}_0^*\|_2 + C_2 c_5 \sqrt{s(\log p)/n}\big)\|\mathbf{v}_1\|_2.
\end{aligned}
$$

Making $\|\mathbf{v}_1\|_2 = C$ large enough, we obtain that with probability tending to one, $\widehat{L}_n(\boldsymbol{\beta}^* + \tilde{a}_n \mathbf{v}) - \widehat{L}_n(\boldsymbol{\beta}^*) > 0$. Then it follows immediately that with asymptotic probability one, there exists a minimizer $\hat{\boldsymbol{\beta}}_1$ of $\widehat{L}_n(\boldsymbol{\beta}_1, \mathbf{0})$ such that $\|\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*\|_2 \le C_3 \tilde{a}_n \equiv a_n$ with some constant $C_3 > 0$.

It remains to prove that with asymptotic probability one,

$$
(7.11) \quad \big\|\hat{\mathbf{d}}_1^{-1} \circ \mathbf{Q}^T \rho_\tau'(\mathbf{y} - \mathbf{S}\hat{\boldsymbol{\beta}}_1)\big\|_\infty < n\lambda_n.
$$

Then by KKT conditions, $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}_1^T, \mathbf{0}^T)^T$ is a global minimizer of $\widehat{L}_n(\boldsymbol{\beta})$.

Now we proceed to prove (7.11). Since $\beta_j^* = 0$ for all $j = s + 1, \ldots, p$, we have that $d_j^* = p_{\lambda_n}'(0+)$. Furthermore, by Condition 4, it holds that $|\hat{\beta}_j^{\mathrm{ini}}| \le C_2 \sqrt{s(\log p)/n}$ with asymptotic probability one. Then, it follows that

$$
\min_{j > s} p_{\lambda_n}'\big(|\hat{\beta}_j^{\mathrm{ini}}|\big) \ge p_{\lambda_n}'\big(C_2 \sqrt{s(\log p)/n}\big).
$$

Therefore, by Condition 5 we conclude that

$$
(7.12) \quad \big\|(\hat{\mathbf{d}}_1)^{-1}\big\|_\infty = \Big(\min_{j > s} p_{\lambda_n}'\big(|\hat{\beta}_j^{\mathrm{ini}}|\big)\Big)^{-1} < 2/p_{\lambda_n}'(0+) = 2\big\|(\mathbf{d}_1^*)^{-1}\big\|_\infty.
$$

From the conditions of Theorem 2 with $\gamma_n = a_n$, it follows from Lemma 2 [inequality (7.20)] that, with probability at least $1 - o(p^{-c})$,

$$
(7.13) \quad \sup_{\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2 \le C_3 a_n} \big\|\mathbf{Q}^T \rho_\tau'(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1)\big\|_\infty < \frac{n\lambda_n}{2\|(\mathbf{d}_1^*)^{-1}\|_\infty}(1 + o(1)).
$$

Combining (7.12)–(7.13) and by the triangle inequality, it holds that with asymptotic probability one,

$$
\sup_{\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2 \le C_3 a_n} \big\|(\hat{\mathbf{d}}_1)^{-1} \circ \mathbf{Q}^T \rho_\tau'(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1)\big\|_\infty < n\lambda_n.
$$

Since the minimizer $\hat{\boldsymbol{\beta}}_1$ satisfies $\|\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*\|_2 < C_3 a_n$ with asymptotic probability one, the above inequality ensures that (7.11) holds with probability tending to one. This completes the proof.

7.4. *Lemmas.* This subsection contains lemmas used in proofs of Theorems 1, 2 and 4.

LEMMA 1. *Under Condition 2, for any $t > 0$, we have*

$$(7.14) \qquad P\big(Z_n(M) \geq 4M\sqrt{s/n} + t\big) \leq \exp\big(-nc_0 t^2/(8M^2)\big).$$

PROOF. Define $\rho(s, y) = (y - s)(\tau - 1\{y - s \leq 0\})$. Then $v_n(\boldsymbol{\beta})$ in (7.1) can be rewritten as $v_n(\boldsymbol{\beta}) = \sum_{i=1}^n \rho(\mathbf{x}_i^T \boldsymbol{\beta}, y_i)$. Note that the following Lipschitz condition holds for $\rho(\cdot, y_i)$:

$$(7.15) \qquad \big|\rho(s_1, y_i) - \rho(s_2, y_i)\big| \leq \max\{\tau, 1 - \tau\}|s_1 - s_2| \leq |s_1 - s_2|.$$

Let $W_1, \ldots, W_n$ be a Rademacher sequence, independent of model errors $\varepsilon_1, \ldots, \varepsilon_n$. The Lipschitz inequality (7.15) combined with the symmetrization theorem and concentration inequality [see, e.g., Theorems 14.3 and 14.4 in Bühlmann and van de Geer (2011)] yields that

$$(7.16) \qquad \begin{aligned} E\big[Z_n(M)\big] &\leq 2E \sup_{\boldsymbol{\beta} \in \mathcal{B}_0(M)} \left| \frac{1}{n} \sum_{i=1}^n W_i\big(\rho(\mathbf{x}_i^T \boldsymbol{\beta}, y_i) - \rho(\mathbf{x}_i^T \boldsymbol{\beta}^*, y_i)\big) \right| \\ &\leq 4E \sup_{\boldsymbol{\beta} \in \mathcal{B}_0(M)} \left| \frac{1}{n} \sum_{i=1}^n W_i\big(\mathbf{x}_i^T \boldsymbol{\beta} - \mathbf{x}_i^T \boldsymbol{\beta}^*\big) \right|. \end{aligned}$$

On the other hand, by the Cauchy–Schwarz inequality

$$\left| \sum_{i=1}^n W_i\big(\mathbf{x}_i^T \boldsymbol{\beta} - \mathbf{x}_i^T \boldsymbol{\beta}^*\big) \right| = \left| \sum_{j=1}^s \left( \sum_{i=1}^n W_i x_{ij} \right)(\beta_j - \beta_j^*) \right|$$

$$\leq \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2 \left\{ \sum_{j=1}^s \left| \sum_{i=1}^n W_i x_{ij} \right|^2 \right\}^{1/2}.$$

By Jensen's inequality and concavity of the square root function, $E(X^{1/2}) \leq (EX)^{1/2}$ for any nonnegative random variable $X$. Thus, these two inequalities ensure that the very right-hand side of (7.16) can be further bounded by

$$(7.17) \qquad \begin{aligned} &\sup_{\boldsymbol{\beta} \in \mathcal{B}_0(M)} \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2 E \left\{ \sum_{j=1}^s \left| \frac{1}{n} \sum_{i=1}^n W_i x_{ij} \right|^2 \right\}^{1/2} \\ &\leq M \left\{ \sum_{j=1}^s E \left| \frac{1}{n} \sum_{i=1}^n W_i x_{ij} \right|^2 \right\}^{1/2} = M\sqrt{s/n}. \end{aligned}$$

Therefore, it follows from (7.16) and (7.17) that

$$(7.18) \qquad E\big[Z_n(M)\big] \leq 4M\sqrt{s/n}.$$

Next, since $n^{-1} \mathbf{S}^T \mathbf{S}$ has bounded eigenvalues, for any $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \mathbf{0}^T) \in \mathcal{B}_0(M)$,

$$\frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i^T (\boldsymbol{\beta} - \boldsymbol{\beta}^*))^2 = \frac{1}{n} (\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)^T \mathbf{S}^T \mathbf{S} (\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) \leq c_0^{-1} \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2^2 \leq c_0^{-1} M^2.$$

Combining this with the Lipschitz inequality (7.15), (7.18) and applying Massart's concentration theorem [see Theorem 14.2 in Bühlmann and van de Geer (2011)] yields that for any $t > 0$,

$$P(Z_n(M) \geq 4M\sqrt{s/n} + t) \leq \exp(-nc_0 t^2/(8M^2)).$$

This proves the lemma.  $\square$

LEMMA 2.   *Consider a ball in $R^s$ around $\boldsymbol{\beta}^*$: $\mathcal{N} = \{\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T \in R^p : \boldsymbol{\beta}_2 = 0, \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2 \leq \gamma_n\}$ with some sequence $\gamma_n \to 0$. Assume that $\min_{j>s} d_j > c_3$, $\sqrt{1 + \gamma_n s^{3/2} \kappa_n^2} \log_2 n = o(\sqrt{n}\lambda_n)$, $n^{1/2}\lambda_n (\log p)^{-1/2} \to \infty$, and $\kappa_n \gamma_n^2 = o(\lambda_n)$. Then under Conditions 1–3, there exists some constant $c > 0$ such that*

$$P\Big(\sup_{\boldsymbol{\beta} \in \mathcal{N}} \|\mathbf{d}_1^{-1} \circ \mathbf{Q}^T \rho_\tau'(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1)\|_\infty \geq n\lambda_n\Big) \leq o(p^{-c}),$$

*where $\rho_\tau'(u) = \tau - 1\{u \leq 0\}$.*

PROOF.    For a fixed $j \in \{s+1, \ldots, p\}$ and $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T \in \mathcal{N}$, define

$$\gamma_{\boldsymbol{\beta},j}(\mathbf{x}_i, y_i) = x_{ij}[\rho_\tau'(y_i - \mathbf{x}_i^T \boldsymbol{\beta}) - \rho_\tau'(\varepsilon_i) - E[\rho_\tau'(y_i - \mathbf{x}_i^T \boldsymbol{\beta}) - \rho_\tau'(\varepsilon_i)]],$$

where $\mathbf{x}_i^T = (x_{i1}, \ldots, x_{ip})$ is the $i$th row of the design matrix. The key for the proof is to use the following decomposition:

$$\sup_{\boldsymbol{\beta} \in \mathcal{N}} \Big\| \frac{1}{n} \mathbf{Q}^T \rho_\tau'(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1) \Big\|_\infty$$

(7.19)
$$\leq \sup_{\boldsymbol{\beta} \in \mathcal{N}} \Big\| \frac{1}{n} \mathbf{Q}^T E[\rho_\tau'(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1) - \rho_\tau'(\boldsymbol{\varepsilon})] \Big\|_\infty$$

$$+ \Big\| \frac{1}{n} \mathbf{Q}^T \rho_\tau'(\boldsymbol{\varepsilon}) \Big\|_\infty + \max_{j>s} \sup_{\boldsymbol{\beta} \in \mathcal{N}} \frac{1}{n} \sum_{i=1}^n |\gamma_{\boldsymbol{\beta},j}(\mathbf{x}_i, y_i)|.$$

We will prove that with probability at least $1 - o(p^{-c})$,

(7.20)    $I_1 \equiv \sup\limits_{\boldsymbol{\beta} \in \mathcal{N}} \Big\| \dfrac{1}{n} \mathbf{Q}^T E[\rho_\tau'(\mathbf{y} - \mathbf{S}\boldsymbol{\beta}_1) - \rho_\tau'(\boldsymbol{\varepsilon})] \Big\|_\infty < \dfrac{\lambda_n}{2\|\mathbf{d}_1^{-1}\|_\infty} + o(\lambda_n),$

(7.21)    $I_2 \equiv n^{-1} \|\mathbf{Q}^T \rho_\tau'(\boldsymbol{\varepsilon})\|_\infty = o(\lambda_n),$

(7.22)    $I_3 \equiv \max\limits_{j>s} \sup\limits_{\boldsymbol{\beta} \in \mathcal{N}} \Big| \dfrac{1}{n} \sum\limits_{i=1}^n \gamma_{\boldsymbol{\beta},j}(\mathbf{x}_i, y_i) \Big| = o_p(\lambda_n).$

347 ADAPTIVE ROBUST VARIABLE SELECTION 347

Combining (7.19)–(7.22) with the assumption $\min_{j>s} d_j > c_3$ completes the proof of the lemma.

Now we proceed to prove (7.20). Note that $I_1$ can be rewritten as

$$(7.23) \qquad I_1 = \max_{j>s} \sup_{\boldsymbol{\beta} \in \mathcal{N}} \left| \frac{1}{n} \sum_{i=1}^{n} x_{ij} E[\rho_\tau'(\varepsilon_i) - \rho_\tau'(y_i - \mathbf{x}_i^T \boldsymbol{\beta})] \right|.$$

By Condition 1,

$$
\begin{aligned}
&E[\rho_\tau'(\varepsilon_i) - \rho_\tau'(y_i - \mathbf{x}_i^T \boldsymbol{\beta})] \\
&\qquad = F_i(\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)) - F_i(0) = f_i(0)\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) + \widetilde{I}_i,
\end{aligned}
$$

where $F(t)$ is the cumulative distribution function of $\varepsilon_i$, and $\widetilde{I}_i = F_i(\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)) - F_i(0) - f_i(0)\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)$. Thus, for any $j > s$,

$$\sum_{i=1}^{n} x_{ij} E[\rho_\tau'(\varepsilon_i) - \rho_\tau'(y_i - \mathbf{x}_i^T \boldsymbol{\beta})] = \sum_{i=1}^{n} (f_i(0)x_{ij}\mathbf{S}_i^T)(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) + \sum_{i=1}^{n} x_{ij} \widetilde{I}_i.$$

This together with (7.23) and Cauchy–Schwarz inequality entails that

$$(7.24) \qquad I_1 \leq \left\| \frac{1}{n} \mathbf{Q}^T \mathbf{H} \mathbf{S}(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) \right\|_\infty + \max_{j>s} \left| \frac{1}{n} \sum_{i=1}^{n} x_{ij} \widetilde{I}_i \right|,$$

where $\mathbf{H} = \operatorname{diag}\{f_1(0), \ldots, f_n(0)\}$. We consider the two terms on the right-hand side of (7.24) one by one. By Condition 3, the first term can be bounded as

$$(7.25) \qquad \left\| \frac{1}{n} \mathbf{Q}^T \mathbf{H} \mathbf{S}(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*) \right\|_\infty \leq \left\| \frac{1}{n} \mathbf{Q}^T \mathbf{H} \mathbf{S} \right\|_{2,\infty} \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2 < \frac{\lambda_n}{2\|\mathbf{d}_1^{-1}\|_\infty}.$$

By Condition 1, $|\widetilde{I}_i| \leq c(\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*))^2$. This together with Condition 2 ensures that the second term of (7.24) can be bounded as

$$
\begin{aligned}
\max_{j>s} \left| \frac{1}{n} \sum_{i=1}^{n} x_{ij} \widetilde{I}_i \right| &\leq \frac{\kappa_n}{n} \sum_{i=1}^{n} |\widetilde{I}_1| \\
&\leq C \frac{\kappa_n}{n} \sum_{i=1}^{n} (\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*))^2 \\
&\leq C\kappa_n \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*\|_2^2.
\end{aligned}
$$

Since $\boldsymbol{\beta} \in \mathcal{N}$, it follows from the assumption $\lambda_n^{-1}\kappa_n\gamma_n^2 = o(1)$ that

$$\max_{j>s} \left| \frac{1}{n} \sum_{i=1}^{n} x_{ij} \widetilde{I}_i \right| \leq C\kappa_n\gamma_n^2 = o(\lambda_n).$$

Plugging the above inequality and (7.25) into (7.24) completes the proof of (7.20).

Next, we prove (7.21). By Hoeffding's inequality, if $\lambda_n > 2\sqrt{(1+c)(\log p)/n}$ with $c$ is some positive constant, then

$$P\left(\left\|\mathbf{Q}^T \rho'_\tau(\boldsymbol{\varepsilon})\right\|_\infty \geq n\lambda_n\right)$$

$$\leq \sum_{j=s+1}^{p} 2\exp\left(-\frac{n^2\lambda_n^2}{4\sum_{i=1}^n x_{ij}^2}\right) = 2\exp(\log(p-s) - n\lambda_n^2/4) \leq O(p^{-c}).$$

Thus, with probability at least $1 - O(p^{-c})$, (7.21) holds.

We now apply Corollary 14.4 in Bühlmann and van de Geer (2011) to prove (7.22). To this end, we need to check conditions of the corollary. For each fixed $j$, define the functional space $\Gamma_j = \{\gamma_{\boldsymbol{\beta},j} : \boldsymbol{\beta} \in \mathcal{N}\}$. First note that $E[\gamma_{\boldsymbol{\beta},j}(\mathbf{x}_i, y_i)] = 0$ for any $\gamma_{\boldsymbol{\beta},j} \in \Gamma_j$. Second, since the $\rho'_\tau$ function is bounded, we have

$$\frac{1}{n}\sum_{i=1}^n \gamma_{\boldsymbol{\beta},j}^2(\mathbf{x}_i, y_i)$$

$$= \frac{1}{n}\sum_{i=1}^n x_{ij}^2 \left(\rho'_\tau(y_i - \mathbf{x}_i^T\boldsymbol{\beta}) - \rho'_\tau(\varepsilon_i) - E(\rho'_\tau(y_i - \mathbf{x}_i^T\boldsymbol{\beta}) - \rho_\tau(\varepsilon_i))\right)^2 \leq 4.$$

Thus, $\|\gamma_{\boldsymbol{\beta},j}\|_n \equiv (n^{-1}\sum_{i=1}^n \gamma_{\boldsymbol{\beta},j}^2(\mathbf{x}_i, y_i))^{1/2} \leq 2$.

Third, we will calculate the covering number of the functional space $\Gamma_j$, $N(\cdot, \Gamma_j, \|\cdot\|_2)$. For any $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T \in \mathcal{N}$ and $\tilde{\boldsymbol{\beta}} = (\tilde{\boldsymbol{\beta}}_1^T, \tilde{\boldsymbol{\beta}}_2^T)^T \in \mathcal{N}$, by Condition 1 and the mean value theorem,

$$E\left[\rho'_\tau(y_i - \mathbf{x}_i^T\boldsymbol{\beta}) - \rho'_\tau(\varepsilon_i)\right] - E\left[\rho'_\tau(y_i - \mathbf{x}_i^T\tilde{\boldsymbol{\beta}}) - \rho'_\tau(\varepsilon_i)\right]$$

$$(7.26) \qquad = F_i\left(\mathbf{S}_i^T(\tilde{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*)\right) - F_i\left(\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)\right)$$

$$= f_i(a_{1i})\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \tilde{\boldsymbol{\beta}}_1),$$

where $F(t)$ is the cumulative distribution function of $\varepsilon_i$, and $a_{1i}$ lies on the segment connecting $\mathbf{S}_i^T(\tilde{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1^*)$ and $\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_1^*)$. Let $\kappa_n = \max_{ij}|x_{ij}|$. Since $f_i(u)$'s are uniformly bounded, by (7.26),

$$\left|x_{ij}E\left[\rho'_\tau(y_i - \mathbf{x}_i^T\boldsymbol{\beta}) - \rho'_\tau(\varepsilon_i)\right] - x_{ij}E\left[\rho'_\tau(y_i - \mathbf{x}_i^T\tilde{\boldsymbol{\beta}}) - \rho'_\tau(\varepsilon_i)\right]\right|$$

$$(7.27) \qquad \leq C\left|x_{ij}\mathbf{S}_i^T(\boldsymbol{\beta}_1 - \tilde{\boldsymbol{\beta}}_1)\right|$$

$$\leq C\|x_{ij}\mathbf{S}_i\|_2\|\boldsymbol{\beta}_1 - \tilde{\boldsymbol{\beta}}_1\|_2 \leq C\sqrt{s}\kappa_n^2\|\boldsymbol{\beta}_1 - \tilde{\boldsymbol{\beta}}_1\|_2,$$

where $C > 0$ is some generic constant. It is known [see, e.g., Lemma 14.27 in Bühlmann and van de Geer (2011)] that the ball $\mathcal{N}$ in $\mathbf{R}^s$ can be covered by $(1 + 4\gamma_n/\delta)^s$ balls with radius $\delta$. Since $\rho'_\tau(y_i - \mathbf{x}_i^T\boldsymbol{\beta}) - \rho'_\tau(\varepsilon_i)$ can only take 3 different values $\{-1, 0, 1\}$, it follows from (7.27) that the covering number of $\Gamma_j$ is $N(2^{2-k}, \Gamma_j, \|\cdot\|_2) = 3(1 + C^{-1}2^k\gamma_n s^{1/2}\kappa_n^2)^s$. Thus, by calculus, for any

$0 \le k \le (\log_2 n)/2$,

$$\log\big(1 + N\big(2^{2-k}, \Gamma, \|\cdot\|_2\big)\big)$$
$$\le \log(6) + s\log\big(1 + C^{-1}2^k\gamma_n s^{1/2}\kappa_n^2\big)$$
$$\le \log(6) + C^{-1}2^k\gamma_n s^{3/2}\kappa_n^2 \le 4\big(1 + C^{-1}\gamma_n s^{3/2}\kappa_n^2\big)2^{2k}.$$

Hence, conditions of Corollary 14.4 in Bühlmann and van de Geer (2011) are checked and we obtain that for any $t > 0$,

$$P\left(\sup_{\boldsymbol{\beta}\in\mathcal{N}}\left|\frac{1}{n}\sum_{i=1}^{n}\gamma_{\boldsymbol{\beta},j}(\mathbf{x}_i, y_i)\right| \ge \frac{8}{\sqrt{n}}\Big(3\sqrt{1 + C^{-1}\gamma_n s^{3/2}\kappa_n^2}\log_2 n + 4 + 4t\Big)\right)$$
$$\le 4\exp\left(-\frac{nt^2}{8}\right).$$

Taking $t = \sqrt{C(\log p)/n}$ with $C > 0$ large enough constant, we obtain that

$$P\left(\max_{j>s}\sup_{\boldsymbol{\beta}\in\mathcal{N}}\left|\frac{1}{n}\sum_{i=1}^{n}\gamma_{\boldsymbol{\beta},j}(\mathbf{x}_i, y_i)\right| \ge \frac{24}{\sqrt{n}}\sqrt{1 + C^{-1}\gamma_n s^{3/2}\kappa_n^2}\log_2 n\right)$$
$$\le 4(p-s)\exp\left(-\frac{C\log p}{8}\right) \to 0.$$

Thus, if $\sqrt{1 + \gamma_n s^{3/2}\kappa_n^2}\log_2 n = o(\sqrt{n}\lambda_n)$, then with probability at least $1 - o(p^{-c})$, (7.22) holds. This completes the proof of the lemma. □

## SUPPLEMENTARY MATERIAL

**Supplementary material for: Adaptive robust variable selection** (DOI: 10.1214/13-AOS1191SUPP; .pdf). Due to space constraints, the proofs of Theorems 3 and 5 and the results of the real life data-set study are relegated to the supplement [Fan, Fan and Barut (2014)].

## REFERENCES

BELLONI, A. and CHERNOZHUKOV, V. (2011). $\ell_1$-penalized quantile regression in high-dimensional sparse models. *Ann. Statist.* **39** 82–130. MR2797841

BICKEL, P. J. and LI, B. (2006). Regularization in statistics. *TEST* **15** 271–344. With comments and a rejoinder by the authors. MR2273731

BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. MR2533469

BRADIC, J., FAN, J. and WANG, W. (2011). Penalized composite quasi-likelihood for ultrahigh dimensional variable selection. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **73** 325–349. MR2815779

BÜHLMANN, P. and VAN DE GEER, S. (2011). *Statistics for High-Dimensional Data*: *Methods*, *Theory and Applications*. Springer, Heidelberg. MR2807761

CANDES, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when $p$ is much larger than $n$. *Ann. Statist.* **35** 2313–2351. MR2382644

FAN, J., FAN, Y. and BARUT, E. (2014). Supplement to "Adaptive robust variable selection." DOI:10.1214/13-AOS1191SUPP.

FAN, J., FAN, Y. and LV, J. (2008). High dimensional covariance matrix estimation using a factor model. *J. Econometrics* **147** 186–197. MR2472991

FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. MR1946581

FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 849–911. MR2530322

FAN, J. and LV, J. (2011). Nonconcave penalized likelihood with $np$-dimensionality. *IEEE Trans. Inform. Theory* **57** 5467–5484.

FAN, J. and PENG, H. (2004). Nonconcave penalized likelihood with a diverging number of parameters. *Ann. Statist.* **32** 928–961. MR2065194

LI, Y. and ZHU, J. (2008). $L_1$-norm quantile regression. *J. Comput. Graph. Statist.* **17** 163–185. MR2424800

LV, J. and FAN, Y. (2009). A unified approach to model selection and sparse recovery using regularized least squares. *Ann. Statist.* **37** 3498–3528. MR2549567

MEINSHAUSEN, N. and BÜHLMANN, P. (2010). Stability selection. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **72** 417–473. MR2758523

NEWEY, W. K. and POWELL, J. L. (1990). Efficient estimation of linear and type I censored regression models under conditional quantile restrictions. *Econometric Theory* **6** 295–317. MR1085576

NOLAN, J. P. (2012). *Stable Distributions—Models for Heavy-Tailed Data*. Birkhauser, Cambridge. (In progress, Chapter 1 online at academic2.american.edu/~jpnolan).

POLLARD, D. (1991). Asymptotics for least absolute deviation regression estimators. *Econometric Theory* **7** 186–199. MR1128411

TIBSHIRANI, R. (1996). Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **58** 267–288. MR1379242

VAN DE GEER, S. and MÜLLER, P. (2012). Quasi-likelihood and/or robust estimation in high dimensions. *Statist. Sci.* **27** 469–480. MR3025129

WANG, L. (2013). $L_1$ penalized LAD estimator for high dimensional linear regression. *J. Multivariate Anal.* **120** 135–151. MR3072722

WANG, H., LI, G. and JIANG, G. (2007). Robust regression shrinkage and consistent variable selection through the LAD-Lasso. *J. Bus. Econom. Statist.* **25** 347–355. MR2380753

WANG, L., WU, Y. and LI, R. (2012). Quantile regression for analyzing heterogeneity in ultra-high dimension. *J. Amer. Statist. Assoc.* **107** 214–222. MR2949353

WU, Y. and LIU, Y. (2009). Variable selection in quantile regression. *Statist. Sinica* **19** 801–817. MR2514189

ZHANG, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38** 894–942. MR2604701

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

ZOU, H. (2006). The adaptive Lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. MR2279469

ZOU, H. and LI, R. (2008). One-step sparse estimates in nonconcave penalized likelihood models. *Ann. Statist.* **36** 1509–1533. MR2435443

ZOU, H. and YUAN, M. (2008). Composite quantile regression and the oracle model selection theory. *Ann. Statist.* **36** 1108–1126. MR2418651

J. FAN
DEPARTMENT OF OPERATIONS RESEARCH
  AND FINANCIAL ENGINEERING
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: jqfan@princeton.edu

Y. FAN
DATA SCIENCES AND OPERATIONS DEPARTMENT
MARSHALL SCHOOL OF BUSINESS
UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CALIFORNIA 90089
USA
E-MAIL: fanyingy@marshall.usc.edu

E. BARUT
IBM T. J. WATSON RESEARCH CENTER
YORKTOWN HEIGHTS, NEW YORK 10598
USA
E-MAIL: abarut@us.ibm.com