

MIT Open Access Articles

*#1-penalized quantile regression
in high-dimensional sparse models*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Belloni, Alexandre, and Victor Chernozhukov. “# 1 -penalized quantile regression in high-dimensional sparse models.” The Annals of Statistics 39, no. 1 (February 2011): 82-130.

As Published: <http://dx.doi.org/10.1214/10-aos827>

Publisher: Institute of Mathematical Statistics

Persistent URL: <http://hdl.handle.net/1721.1/82630>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



ℓ_1 -PENALIZED QUANTILE REGRESSION IN HIGH-DIMENSIONAL SPARSE MODELS

BY ALEXANDRE BELLONI AND VICTOR CHERNOZHUKOV^{*,†}

Duke University and Massachusetts Institute of Technology

We consider median regression and, more generally, a possibly infinite collection of quantile regressions in high-dimensional sparse models. In these models the number of regressors p is very large, possibly larger than the sample size n , but only at most s regressors have a non-zero impact on each conditional quantile of the response variable, where s grows more slowly than n . Since ordinary quantile regression is not consistent in this case, we consider ℓ_1 -penalized quantile regression (ℓ_1 -QR), which penalizes the ℓ_1 -norm of regression coefficients, as well as the post-penalized QR estimator (post- ℓ_1 -QR), which applies ordinary QR to the model selected by ℓ_1 -QR. First, we show that under general conditions ℓ_1 -QR is consistent at the near-oracle rate $\sqrt{s/n} \sqrt{\log(p \vee n)}$, uniformly in the compact set $\mathcal{U} \subset (0,1)$ of quantile indices. In deriving this result, we propose a partly pivotal, data-driven choice of the penalty level and show that it satisfies the requirements for achieving this rate. Second, we show that under similar conditions post- ℓ_1 -QR is consistent at the near-oracle rate $\sqrt{s/n} \sqrt{\log(p \vee n)}$, uniformly over \mathcal{U} , even if the ℓ_1 -QR-selected models miss some components of the true models, and the rate could be even closer to the oracle rate otherwise. Third, we characterize conditions under which ℓ_1 -QR contains the true model as a submodel, and derive bounds on the dimension of the selected model, uniformly over \mathcal{U} ; we also provide conditions under which hard-thresholding selects the minimal true model, uniformly over \mathcal{U} .

1. Introduction. Quantile regression is an important statistical method for analyzing the impact of regressors on the conditional distribution of a response variable (cf. [27], [24]). It captures the heterogeneous impact of regressors on different parts of the distribution [8], exhibits robustness to outliers [22], has excellent computational properties [34], and has wide applicability [22]. The asymptotic theory for quantile regression has been developed under

^{*}First version: December, 2007, This version: November 2, 2011.

[†]The authors gratefully acknowledge research support from the National Science Foundation.

AMS 2000 subject classifications: Primary 62H12, 62J99; secondary 62J07

Keywords and phrases: median regression, quantile regression, sparse models

both a fixed number of regressors and an increasing number of regressors. The asymptotic theory under a fixed number of regressors is given in [24], [33], [18], [20], [13] and others. The asymptotic theory under an increasing number of regressors is given in [19] and [3, 4], covering the case where the number of regressors p is negligible relative to the sample size n (i.e., $p = o(n)$).

In this paper, we consider quantile regression in high-dimensional sparse models (HDSMs). In such models, the overall number of regressors p is very large, possibly much larger than the sample size n . However, the number of significant regressors for each conditional quantile of interest is at most s , which is smaller than the sample size, that is, $s = o(n)$. HDSMs ([7, 12, 32]) have emerged to deal with many new applications arising in biometrics, signal processing, machine learning, econometrics, and other areas of data analysis where high-dimensional data sets have become widely available.

A number of papers have begun to investigate estimation of HDSMs, focusing primarily on penalized mean regression, with the ℓ_1 -norm acting as a penalty function [7, 12, 26, 32, 39, 41]. [7, 12, 26, 32, 41] demonstrated the fundamental result that ℓ_1 -penalized least squares estimators achieve the rate $\sqrt{s/n}\sqrt{\log p}$, which is very close to the oracle rate $\sqrt{s/n}$ achievable when the true model is known. [39] demonstrated a similar fundamental result on the excess forecasting error loss under both quadratic and non-quadratic loss functions. Thus the estimator can be consistent and can have excellent forecasting performance even under very rapid, nearly exponential, growth of the total number of regressors p . See [7, 9–11, 15, 30, 35] for many other interesting developments and a detailed review of the existing literature.

Our paper’s contribution is to develop a set of results on model selection and rates of convergence for quantile regression within the HDSM framework. Since ordinary quantile regression is inconsistent in HDSMs, we consider quantile regression penalized by the ℓ_1 -norm of parameter coefficients, denoted ℓ_1 -QR. First, we show that under general conditions ℓ_1 -QR estimates of regression coefficients and regression functions are consistent at the near-oracle rate $\sqrt{s/n}\sqrt{\log(p \vee n)}$, uniformly in a compact interval $\mathcal{U} \subset (0, 1)$ of quantile indices.¹ (This result is different from and hence complementary to [39]’s fundamental results on the rates for excess forecasting error loss.) Second, in order to make ℓ_1 -QR practical, we propose a partly pivotal, data-driven choice of the penalty level, and show that this choice leads to the same sharp convergence rate. Third, we show that ℓ_1 -QR correctly selects the true model as a valid submodel when the non-zero coefficients of the true model are well separated from zero. Fourth, we also propose and analyze the post-penalized estimator (post- ℓ_1 -QR), which applies ordi-

¹Under $s \rightarrow \infty$, the oracle rate, uniformly over a proper compact interval \mathcal{U} , is $\sqrt{(s/n)\log n}$, cf. [4]; the oracle rate for a single quantile index is $\sqrt{s/n}$, cf. [19].

nary, unpenalized quantile regression to the model selected by the penalized estimator, and thus aims at reducing the regularization bias of the penalized estimator. We show that under similar conditions post- ℓ_1 -QR can perform as well as ℓ_1 -QR in terms of the rate of convergence, uniformly over \mathcal{U} , even if the ℓ_1 -QR-based model selection misses some components of the true models. This occurs because ℓ_1 -QR-based model selection only misses those components that have relatively small coefficients. Moreover, post- ℓ_1 -QR can perform better than ℓ_1 -QR if the ℓ_1 -QR-based model selection correctly includes all components of the true model as a subset. (Obviously, post- ℓ_1 -QR can perform as well as the oracle if the ℓ_1 -QR perfectly selects the true model, which is, however, unrealistic for many designs of interest.) Fifth, we illustrate the use of ℓ_1 -QR and post- ℓ_1 -QR with a Monte Carlo experiment and an international economic growth example. To the best of our knowledge, all of the above results are new and contribute to the literature on HDSMs. Our results on post-penalized estimators and some proof techniques could also be of interest in other problems. We provide further technical comparisons to the literature in Section 2.

1.1. Notation. In what follows, we implicitly index all parameter values by the sample size n , but we omit the index whenever this does not cause confusion. We use the empirical process notation as defined in [40]. In particular, given a random sample Z_1, \dots, Z_n , let $\mathbb{G}_n(f) = \mathbb{G}_n(f(Z_i)) := n^{-1/2} \sum_{i=1}^n (f(Z_i) - \mathbb{E}[f(Z_i)])$ and $\mathbb{E}_n f = \mathbb{E}_n f(Z_i) := \sum_{i=1}^n f(Z_i)/n$. We use the notation $a \lesssim b$ to denote $a = O(b)$, that is, $a \leq cb$ for some constant $c > 0$ that does not depend on n ; and $a \lesssim_P b$ to denote $a = O_P(b)$. We also use the notation $a \vee b = \max\{a, b\}$ and $a \wedge b = \min\{a, b\}$. We denote the ℓ_2 -norm by $\|\cdot\|$, ℓ_1 -norm by $\|\cdot\|_1$, ℓ_∞ -norm by $\|\cdot\|_\infty$, and the ℓ_0 -“norm” by $\|\cdot\|_0$ (i.e., the number of non-zero components). We denote by $\|\beta\|_{1,n} = \sum_{j=1}^p \hat{\sigma}_j |\beta_j|$ the ℓ_1 -norm weighted by $\hat{\sigma}_j$'s. Finally, given a vector $\delta \in \mathbb{R}^p$, and a set of indices $T \subset \{1, \dots, p\}$, we denote by δ_T the vector in which $\delta_{Tj} = \delta_j$ if $j \in T$, $\delta_{Tj} = 0$ if $j \notin T$.

2. The Estimator, the Penalty Level, and Overview of Rate Results. In this section we formulate the setting and the estimator, and state primitive regularity conditions. We also provide an overview of the main results.

2.1. Basic Setting. The setting of interest corresponds to a parametric quantile regression model, where the dimension p of the underlying model increases with the sample size n . Namely, we consider a response variable y and p -dimensional covariates x such that the u -th conditional quantile function of y given x is given by

$$(2.1) \quad F_{y_i|x_i}^{-1}(u|x_i) = x' \beta(u), \quad \beta(u) \in \mathbb{R}^p, \quad \text{for all } u \in \mathcal{U},$$

where $\mathcal{U} \subset (0, 1)$ is a compact set of quantile indices. Recall that the u -th conditional quantile $F_{y_i|x_i}^{-1}(u|x)$ is the inverse of the conditional distribution function $F_{y_i|x_i}(y|x_i)$ of y_i given x_i . We consider the case where the dimension p of the model is large, possibly much larger than the available sample size n , but the true model $\beta(u)$ has a sparse support

$$T_u = \text{support}(\beta(u)) = \{j \in \{1, \dots, p\} : |\beta_j(u)| > 0\}$$

having only $s_u \leq s \leq n/\log(n \vee p)$ non-zero components for all $u \in \mathcal{U}$.

The population coefficient $\beta(u)$ is known to minimize the criterion function

$$(2.2) \quad Q_u(\beta) = \mathbb{E}[\rho_u(y - x'\beta)],$$

where $\rho_u(t) = (u - 1\{t \leq 0\})t$ is the asymmetric absolute deviation function [24]. Given a random sample $(y_1, x_1), \dots, (y_n, x_n)$, the quantile regression estimator of $\beta(u)$ is defined as a minimizer of the empirical analog of (2.2):

$$(2.3) \quad \hat{Q}_u(\beta) = \mathbb{E}_n [\rho_u(y_i - x_i'\beta)].$$

In high-dimensional settings, particularly when $p \geq n$, ordinary quantile regression is generally inconsistent, which motivates the use of penalization in order to remove all, or at least nearly all, regressors whose population coefficients are zero, thereby possibly restoring consistency. A penalization that has proven quite useful in least squares settings is the ℓ_1 -penalty leading to the Lasso estimator [37].

2.2. Penalized and Post-Penalized Estimators. The ℓ_1 -penalized quantile regression estimator $\hat{\beta}(u)$ is a solution to the following optimization problem:

$$(2.4) \quad \min_{\beta \in \mathbb{R}^p} \hat{Q}_u(\beta) + \frac{\lambda \sqrt{u(1-u)}}{n} \sum_{j=1}^p \hat{\sigma}_j |\beta_j|$$

where $\hat{\sigma}_j^2 = \mathbb{E}_n[x_{ij}^2]$. The criterion function in (2.4) is the sum of the criterion function (2.3) and a penalty function given by a scaled ℓ_1 -norm of the parameter vector. The overall penalty level $\lambda \sqrt{u(1-u)}$ depends on each quantile index u , while λ will depend on the set \mathcal{U} of quantile indices of interest. The ℓ_1 -penalized quantile regression has been considered in [21] under small (fixed) p asymptotics. It is important to note that the penalized quantile regression problem (2.4) is equivalent to a linear programming problem (see Appendix C) with a dual version that is useful for analyzing the sparsity of the solution. When the solution is not unique, we define $\hat{\beta}(u)$ as any optimal basic feasible solution (see, e.g., [6]). Therefore, the problem (2.4) can be solved in polynomial time, avoiding the computational curse of dimensionality. Our goal is to derive the rate of convergence and model selection properties of this estimator.

The post-penalized estimator (post- ℓ_1 -QR) applies ordinary quantile regression to the model \widehat{T}_u selected by the ℓ_1 -penalized quantile regression. Specifically, set

$$\widehat{T}_u = \text{support}(\widehat{\beta}(u)) = \{j \in \{1, \dots, p\} : |\widehat{\beta}_j(u)| > 0\},$$

and define the post-penalized estimator $\widetilde{\beta}(u)$ as

$$(2.5) \quad \widetilde{\beta}(u) \in \arg \min_{\beta \in \mathbb{R}^p : \beta_{\widehat{T}_u^c} = 0} \widehat{Q}_u(\beta),$$

which removes from further estimation the regressors that were not selected. If the model selection works perfectly – that is, $\widehat{T}_u = T_u$ – then this estimator is simply the oracle estimator, whose properties are well known. However, perfect model selection might be unlikely for many designs of interest. Rather, we are interested in the more realistic scenario where the first-step estimator $\widehat{\beta}(u)$ fails to select some components of $\beta(u)$. Our goal is to derive the rate of convergence for the post-penalized estimator and show it can perform well under this scenario.

2.3. The choice of the penalty level λ . In order to describe our choice of the penalty level λ , we introduce the random variable

$$(2.6) \quad \Lambda = n \sup_{u \in \mathcal{U}} \max_{1 \leq j \leq p} \left| \mathbb{E}_n \left[\frac{x_{ij}(u - 1\{u_i \leq u\})}{\widehat{\sigma}_j \sqrt{u(1-u)}} \right] \right|,$$

where u_1, \dots, u_n are i.i.d. uniform $(0, 1)$ random variables, independently distributed from the regressors, x_1, \dots, x_n . The random variable Λ has a known, that is, pivotal, distribution conditional on $X = [x_1, \dots, x_n]'$. We then set

$$(2.7) \quad \lambda = c \cdot \Lambda(1 - \alpha|X),$$

where $\Lambda(1 - \alpha|X) := (1 - \alpha)$ -quantile of Λ conditional on X , and the constant $c > 1$ depends on the design.² Thus the penalty level depends on the pivotal quantity $\Lambda(1 - \alpha|X)$ and the design. Under assumptions D.1-D.4 we can set $c = 2$, similar to [7]’s choice for least squares. Furthermore, we recommend computing $\Lambda(1 - \alpha|X)$ using simulation of Λ .³ Our concrete recommendation for practice is to set $1 - \alpha = 0.9$.

The parameter $1 - \alpha$ is the confidence level in the sense that, as in [7], our (non-asymptotic) bounds on the estimation error will contract at the optimal rate with this probability. We refer the reader to Koenker [23] for an implementation of our choice of penalty level and practical

² c depends only on the constant c_0 appearing in condition D.4; when $c_0 \geq 9$, it suffices to set $c = 2$.

³We also provide analytical bounds on $\Lambda(1 - \alpha|X)$ of the form $C(\alpha, \mathcal{U})\sqrt{n \log p}$ for some numeric constant $C(\alpha, \mathcal{U})$. We recommend simulation because it accounts for correlation among the columns of X in the sample.

suggestions concerning the confidence level. In particular, both here and in Koenker [23], the confidence level $1 - \alpha = 0.9$ gave good performance results in terms of balancing regularization bias with estimation variance. Cross-validation may also be used to choose the confidence level $1 - \alpha$. Finally, we should note that, as in [7], our theoretical bounds allow for any choice of $1 - \alpha$ and are stated as a function of $1 - \alpha$.

The formal rationale behind the choice (2.7) for the penalty level λ is that this choice leads precisely to the optimal rates of convergence for ℓ_1 -QR. (The same or slightly higher choice of λ also guarantees good performance of post- ℓ_1 -QR.) Our general strategy for choosing λ follows [7], who recommend selecting λ so that it dominates a relevant measure of noise in the sample criterion function, specifically the supremum norm of a suitably rescaled gradient of the sample criterion function evaluated at the true parameter value. In our case this general strategy leads precisely to the choice (2.7). Indeed, a (sub)gradient $\hat{S}_u(\beta(u)) = \mathbb{E}_n[(u - 1\{y_i \leq x_i'\beta(u)\})x_i] \in \partial\hat{Q}_u(\beta(u))$ of the quantile regression objective function evaluated at the truth has a pivotal representation, namely $\hat{S}_u(\beta(u)) = \mathbb{E}_n[(u - 1\{u_i \leq u\})x_i]$ for u_1, \dots, u_n i.i.d. uniform $(0, 1)$ conditional on X , and so we can represent Λ as in (2.6), and, thus, choose λ as in (2.7).

2.4. General Regularity Conditions. We consider the following conditions on a sequence of models indexed by n with parameter dimension $p = p_n \rightarrow \infty$. In these conditions all constants can depend on n , but we omit the explicit indexing by n to ease exposition.

D.1. Sampling and Smoothness. Data $(y_i, x_i')', i = 1, \dots, n$, are an i.i.d. sequence of real $(1 + p)$ -vectors, with the conditional u -quantile function given by (2.1) for each $u \in \mathcal{U}$, with the first component of x_i equal to one, and $n \wedge p \geq 3$. For each value x in the support of x_i , the conditional density $f_{y_i|x_i}(y|x)$ is continuously differentiable in y at each $y \in \mathbb{R}$, and $f_{y_i|x_i}(y|x)$ and $\frac{\partial}{\partial y}f_{y_i|x_i}(y|x)$ are bounded in absolute value by constants \bar{f} and \bar{f}' , uniformly in $y \in \mathbb{R}$ and x in the support x_i . Moreover, the conditional density of y_i evaluated at the conditional quantile $x_i'\beta(u)$ is bounded away from zero uniformly in \mathcal{U} , that is $f_{y_i|x_i}(x_i'\beta(u)|x) > \underline{f} > 0$ uniformly in $u \in \mathcal{U}$ and x in the support of x_i .

Condition D.1 imposes only mild smoothness assumptions on the conditional density of the response variable given regressors, and does not impose any normality or homoscedasticity assumptions. The assumption that the conditional density is bounded below at the conditional quantile is standard, but we can replace it by the slightly more general condition $\inf_{u \in \mathcal{U}} \inf_{\delta \neq 0} (\delta' J_u \delta) / (\delta' E[x_i x_i'] \delta) \geq \underline{f} > 0$, on the Jacobian matrices

$$J_u = E[f_{y_i|x_i}(x_i'\beta(u)|x_i)x_i x_i'] \text{ for all } u \in \mathcal{U},$$

throughout the paper.

D.2. Sparsity and Smoothness of $u \mapsto \beta(u)$. Let \mathcal{U} be a compact subset of $(0, 1)$. The coefficients $\beta(u)$ in (2.1) are sparse and smooth with respect to $u \in \mathcal{U}$:

$$\sup_{u \in \mathcal{U}} \|\beta(u)\|_0 \leq s \quad \text{and} \quad \|\beta(u) - \beta(u')\| \leq L|u - u'|, \quad \text{for all } u, u' \in \mathcal{U}$$

where $s \geq 1$, and $\log L \leq C_L \log(p \vee n)$ for some constant C_L .

Condition D.2 imposes sparsity and smoothness on the behavior of the quantile regression coefficients $\beta(u)$ as we vary the quantile index u .

D.3. Well-behaved Covariates. Covariates are normalized such that $\sigma_j^2 = \mathbb{E}[x_{ij}^2] = 1$ for all $j = 1, \dots, p$, and $\hat{\sigma}_j^2 = \mathbb{E}_n[x_{ij}^2]$ obeys $P(\max_{1 \leq j \leq p} |\hat{\sigma}_j - 1| \leq 1/2) \geq 1 - \gamma \rightarrow 1$ as $n \rightarrow \infty$.

Condition D.3 requires that $\hat{\sigma}_j$ does not deviate too much from σ_j and normalizes $\sigma_j^2 = 1$.

In order to state the next assumption, for some $c_0 \geq 0$ and each $u \in \mathcal{U}$, define

$$A_u := \{\delta \in \mathbb{R}^p : \|\delta_{T_u^c}\|_1 \leq c_0 \|\delta_{T_u}\|_1, \|\delta_{T_u^c}\|_0 \leq n\},$$

which will be referred to as the restricted set. Define $\bar{T}_u(\delta, m) \subset \{1, \dots, p\} \setminus T_u$ as the support of the m largest in absolute value components of the vector δ outside of $T_u = \text{support}(\beta(u))$, where $\bar{T}_u(\delta, m)$ is the empty set if $m = 0$.

D.4. Restricted Identifiability and Nonlinearity. For some constants $m \geq 0$ and $c_0 \geq 9$, the matrix $\mathbb{E}[x_i x_i']$ satisfies

$$(RE(c_0, m)) \quad \kappa_m^2 := \inf_{u \in \mathcal{U}} \inf_{\delta \in A_u, \delta \neq 0} \frac{\delta' \mathbb{E}[x_i x_i'] \delta}{\|\delta_{T_u \cup \bar{T}_u(\delta, m)}\|^2} > 0$$

and $\log(\underline{f} \kappa_0^2) \leq C_f \log(n \vee p)$ for some constant C_f . Moreover,

$$(RNI(c_0)) \quad q := \frac{3 \underline{f}^{3/2}}{8 \underline{f}'} \inf_{u \in \mathcal{U}} \inf_{\delta \in A_u, \delta \neq 0} \frac{\mathbb{E}[|x_i' \delta|^2]^{3/2}}{\mathbb{E}[|x_i' \delta|^3]} > 0.$$

The restricted eigenvalue (RE) condition is analogous to the condition in [7] and [12]; see [7] and [12] for different sufficient primitive conditions that yield bounds on κ_m . Also, since κ_m is non-increasing in m , $RE(c_0, m)$ for any $m > 0$ implies $RE(c_0, 0)$. The restricted non-linear impact (RNI) coefficient q appearing in D.4 is a new concept, which controls the quality of minoration of the quantile regression objective function by a quadratic function over the restricted set.

Finally, we state another condition needed to derive results on the post-model selected estimator. In order to state the condition, define the sparse set $\tilde{A}_u(\tilde{m}) = \{\delta \in \mathbb{R}^p : \|\delta_{T_u^c}\|_0 \leq \tilde{m}\}$ for $\tilde{m} \geq 0$ and $u \in \mathcal{U}$.

D.5. Sparse Identifiability and Nonlinearity. The matrix $E[x_i x_i']$ satisfies for some $\tilde{m} \geq 0$:

$$(SE(\tilde{m})) \quad \tilde{\kappa}_{\tilde{m}}^2 := \inf_{u \in \mathcal{U}} \inf_{\delta \in \tilde{A}_u(\tilde{m}), \delta \neq 0} \frac{\delta' E[x_i x_i'] \delta}{\delta' \delta} > 0,$$

and

$$(SNI(\tilde{m})) \quad \tilde{q}_{\tilde{m}} := \frac{3}{8} \frac{f^{3/2}}{f'} \inf_{u \in \mathcal{U}} \inf_{\delta \in \tilde{A}_u(\tilde{m}), \delta \neq 0} \frac{E[|x_i' \delta|^2]^{3/2}}{E[|x_i' \delta|^3]} > 0.$$

We invoke the sparse eigenvalue (SE) condition in order to analyze the post-penalized estimator (2.5). This assumption is similar to the conditions used in [32] and [41] to analyze Lasso. Our form of the SE condition is neither less nor more general than the RE condition. The SNI coefficient $\tilde{q}_{\tilde{m}}$ controls the quality of minoration of the quantile regression objective function by a quadratic function over sparse neighborhoods of the true parameter.

2.5. Examples of Simple Sufficient Conditions. In order to highlight the nature and usefulness of conditions D.1-D.5 it is instructive to state some simple sufficient conditions (note that D.1-D.5 allow for much more general conditions). We relegate the proofs of this section to the Supplementary Material Appendix G for brevity.

DESIGN 1: LOCATION MODEL WITH CORRELATED NORMAL DESIGN. Let us consider estimating a standard location model

$$y = x' \beta^o + \varepsilon,$$

where $\varepsilon \sim N(0, \sigma^2)$, $\sigma > 0$ is fixed, $x = (1, z')'$, with $z \sim N(0, \Sigma)$, where Σ has ones in the diagonal, a minimum eigenvalue bounded away from zero by a constant $\kappa^2 > 0$, and a maximum eigenvalue bounded from above, uniformly in n .

LEMMA 1. *Under Design 1 with $\mathcal{U} = [\xi, 1 - \xi]$, $\xi > 0$, conditions D.1-D.5 are satisfied with*

$$\begin{aligned} \bar{f} &= 1/[\sqrt{2\pi}\sigma], \quad \bar{f}' = \sqrt{e/[2\pi]}/\sigma^2, \quad \underline{f} = 1/\sqrt{2\pi\xi}\sigma, \\ \|\beta(u)\|_0 &\leq \|\beta^o\|_0 + 1, \quad \gamma = 2p \exp(-n/24), \quad L = \sigma/\xi \\ \kappa_m \wedge \tilde{\kappa}_{\tilde{m}} &\geq \kappa, \quad q \wedge \tilde{q}_{\tilde{m}} \geq (3/[32\xi^{3/4}])\sqrt{\sqrt{2\pi}\sigma/e}. \end{aligned}$$

Note that the normality of errors can be easily relaxed by allowing for the disturbance ε to have a smooth density that obeys the conditions stated in D.1. The conditions on the population design matrix can also be replaced by more general primitive conditions specified in Remark 2.1.

DESIGN 2: LOCATION-SCALE MODEL WITH BOUNDED REGRESSORS. Let us consider estimating a standard location-scale model

$$y = x' \beta^o + x' \eta \cdot \varepsilon,$$

where $\varepsilon \sim F$ independent of x , with a continuously differentiable probability density function f . We assume that the population design matrix $E[xx']$ has ones in the diagonal and has eigenvalues uniformly bounded away from zero and from above, $x_1 = 1$, $\max_{1 \leq j \leq p} |x_j| \leq K_B$. Moreover, the vector η is such that $0 < v \leq x' \eta \leq \Upsilon < \infty$ for all values of x .

LEMMA 2. Under Design 2 with $\mathcal{U} = [\xi, 1 - \xi]$, $\xi > 0$, conditions D.1-D.5 are satisfied with

$$\bar{f} \leq \max_{\varepsilon} f(\varepsilon)/v, \quad \bar{f}' \leq \max_{\varepsilon} f'(\varepsilon)/v^2, \quad \underline{f} = \min_{u \in \mathcal{U}} f(F^{-1}(u))/\Upsilon,$$

$$\|\beta(u)\|_0 \leq \|\beta^o\|_0 + \|\eta\|_0 + 1, \quad \gamma = 2p \exp(-n/[8K_B^4]),$$

$$\kappa_m \wedge \tilde{\kappa}_{\tilde{m}} \geq \kappa, \quad L = \|\eta\| \underline{f}$$

$$q \geq \frac{3}{8} \frac{f^{3/2}}{\underline{f}'} \kappa / [10K_B \sqrt{s}], \quad \tilde{q}_{\tilde{m}} \geq \frac{3}{8} \frac{f^{3/2}}{\underline{f}'} \kappa / [K_B \sqrt{s + \tilde{m}}].$$

COMMENT 2.1. (Conditions on $E[x_i x_i']$). The conditions on the population design matrix can also be replaced by more general primitive conditions of the form stated in [7] and [12]. For example, conditions on sparse eigenvalues suffice as shown in [7]. Denote the minimum and maximum eigenvalue of the population design matrix by

$$(2.8) \quad \varphi_{\min}(m) = \min_{\|\delta\|=1, \|\delta\|_0 \leq m} \frac{\delta' E[x_i x_i'] \delta}{\delta' \delta} \quad \text{and} \quad \varphi_{\max}(m) = \max_{\|\delta\|=1, \|\delta\|_0 \leq m} \frac{\delta' E[x_i x_i'] \delta}{\delta' \delta}.$$

Assuming that for some $m \geq s$ we have $m \varphi_{\min}(m + s) \geq c_0^2 s \varphi_{\max}(m)$, then

$$\kappa_m \geq \sqrt{\varphi_{\min}(s + m)} \left(1 - c_0 \sqrt{s \varphi_{\max}(s) / [m \varphi_{\min}(s + m)]} \right) \quad \text{and} \quad \tilde{\kappa}_{\tilde{m}} \geq \varphi_{\min}(s + m).$$

2.6. *Overview of Main Results.* Here we discuss our results under the simple setup of Design 1 and under $1/p \leq \alpha \rightarrow 0$ and $\gamma \rightarrow 0$. These simple assumptions allow us to straightforwardly compare our rate results to those obtained in the literature. We state our more general non-asymptotic results under general conditions in the subsequent sections. Our first main rate result is that ℓ_1 -QR, with our choice (2.7) of parameter λ , satisfies

$$(2.9) \quad \sup_{u \in \mathcal{U}} \|\hat{\beta}(u) - \beta(u)\| \lesssim_P \frac{1}{\underline{f} \kappa_0 \kappa_s} \sqrt{\frac{s \log(n \vee p)}{n}},$$

provided that the upper bound on the number of non-zero components s satisfies

$$(2.10) \quad \frac{\sqrt{s \log(n \vee p)}}{\sqrt{n} \underline{f}^{1/2} \kappa_0 q} \rightarrow 0.$$

Note that κ_0 , κ_s , \underline{f} , and q are bounded away from zero in this example. Therefore, the rate of convergence is $\sqrt{s/n} \cdot \sqrt{\log(n \vee p)}$ uniformly in the set of quantile indices $u \in \mathcal{U}$, which is very close to the oracle rate when p grows polynomially in n . Further, we note that our resulting restriction (2.10) on the dimension s of the true models is very weak; when p is polynomial in n , s can be of almost the same order as n , namely $s = o(n/\log n)$.

Our second main result is that the dimension $\|\hat{\beta}(u)\|_0$ of the model selected by the ℓ_1 -penalized estimator is of the same stochastic order as the dimension s of the true models, namely

$$(2.11) \quad \sup_{u \in \mathcal{U}} \|\hat{\beta}(u)\|_0 \lesssim_P s.$$

Further, if the parameter values of the minimal true model are well separated from zero, then with a high probability the model selected by the ℓ_1 -penalized estimator correctly nests the true minimal model:

$$(2.12) \quad T_u = \text{support}(\beta(u)) \subseteq \hat{T}_u = \text{support}(\hat{\beta}(u)), \text{ for all } u \in \mathcal{U}.$$

Moreover, we provide conditions under which a hard-thresholded version of the estimator selects the correct support.

Our third main result is that the post-penalized estimator, which applies ordinary quantile regression to the selected model, obeys

$$(2.13) \quad \sup_{u \in \mathcal{U}} \|\tilde{\beta}(u) - \beta(u)\| \lesssim_P \frac{1}{\underline{f} \tilde{\kappa}_{\hat{m}}^2} \sqrt{\frac{\hat{m} \log(n \vee p) + s \log n}{n}} + \frac{\sup_{u \in \mathcal{U}} 1\{T_u \not\subseteq \hat{T}_u\}}{\underline{f} \kappa_0 \tilde{\kappa}_{\hat{m}}} \sqrt{\frac{s \log(n \vee p)}{n}},$$

where $\hat{m} = \sup_{u \in \mathcal{U}} \|\hat{\beta}_{T_u^c}(u)\|_0$ is the maximum number of wrong components selected for any quantile index $u \in \mathcal{U}$, provided that the bound on the number of non-zero components s obeys the growth condition (2.10) and

$$(2.14) \quad \frac{\sqrt{\hat{m} \log(n \vee p) + s \log n}}{\sqrt{n} \underline{f}^{1/2} \tilde{\kappa}_{\hat{m}} \tilde{q}_{\hat{m}}} \rightarrow_P 0.$$

(Note that when \mathcal{U} is a singleton, the $s \log n$ factor in (2.13) becomes s .)

We see from (2.13) that post- ℓ_1 -QR can perform well in terms of the rate of convergence even if the selected model \hat{T}_u fails to contain the true model T_u . Indeed, since in this design $\hat{m} \lesssim_P s$, post- ℓ_1 -QR has the rate of convergence $\sqrt{s/n} \cdot \sqrt{\log(n \vee p)}$, which is the same as the rate of convergence of ℓ_1 -QR. The intuition for this result is that the ℓ_1 -QR based model selection can only miss covariates with relatively small coefficients, which then permits post- ℓ_1 -QR to perform as well or even better due to reductions in bias, as confirmed by our computational experiments.

We also see from (2.13) that post- ℓ_1 -QR can perform better than ℓ_1 -QR in terms of the rate of convergence if the number of wrong components selected obeys $\hat{m} = o_P(s)$ and the selected model contains the true model, $\{T_u \subseteq \hat{T}_u\}$ with probability converging to one. In this case post- ℓ_1 -QR has the rate of convergence $\sqrt{(o_P(s)/n) \log(n \vee p) + (s/n) \log n}$, which is faster than the rate of convergence of ℓ_1 -QR. In the extreme case of perfect model selection, that is, when $\hat{m} = 0$, the rate of post- ℓ_1 -QR becomes $\sqrt{(s/n) \log n}$ uniformly in \mathcal{U} . (When \mathcal{U} is a singleton, the $\log n$ factor drops out.) Note that inclusion $\{T_u \subseteq \hat{T}_u\}$ necessarily happens when the coefficients of the true models are well separated from zero, as we stated above. Note also that the condition $\hat{m} = o(s)$ or even $\hat{m} = 0$ could occur under additional conditions on the regressors (such as the mutual coherence conditions that restrict the maximal pairwise correlation of regressors). Finally, we note that our second restriction (2.14) on the dimension s of the true models is very weak in this design; when p is polynomial in n , s can be of almost the same order as n , namely $s = o(n/\log n)$.

To the best of our knowledge, all of the results presented above are new, both for the single ℓ_1 -penalized quantile regression problem as well as for the infinite collection of ℓ_1 -penalized quantile regression problems. These results therefore contribute to the rate results obtained for ℓ_1 -penalized mean regression and related estimators in the fundamental papers of [7, 12, 26, 32, 39, 41]. The results on post- ℓ_1 penalized quantile regression had no analogs in the literature on mean regression, apart from the rather exceptional case of perfect model selection, in which case the post-penalized estimator is simply the oracle. Building on the current work these results have been extended to mean regression in [5]. Our results on the sparsity of ℓ_1 -QR and model selection also contribute to the analogous results for mean regression [32]. Also, our rate results for ℓ_1 -QR are different from, and hence complementary to, the fundamental results in [39] on the excess forecasting loss under possibly non-quadratic loss functions, which also specializes the results to density estimation, mean regression, and logistic regression. In principle we could apply theorems in [39] to the single quantile regression problem to derive the bounds on the excess loss $E[\rho_u(y_i - x_i' \hat{\beta}(u))] - E[\rho_u(y_i - x_i' \beta(u))]$.⁴ However, these bounds would not imply

⁴Of course, such a derivation would entail some difficult work, since we must verify some high-level assumptions made directly on the performance of the oracle and penalized estimators in population and others (cf. [39])'s

our results (2.9), (2.13), (2.11), (2.12), and (2.7), which characterize the rates of estimating coefficients $\beta(u)$ by ℓ_1 -QR and post- ℓ_1 -QR, sparsity and model selection properties, and the data-driven choice of the penalty level.

3. Main Results and Main Proofs. In this section we derive rates of convergence for ℓ_1 -QR and post- ℓ_1 -QR, sparsity bounds, and model selection results.

3.1. Bounds on $\Lambda(1-\alpha|X)$. We start with a characterization of Λ and its $(1-\alpha)$ -quantile, $\Lambda(1-\alpha|X)$, which determines the magnitude of our suggested penalty level λ via equation (2.7).

THEOREM 1 (Bounds on $\Lambda(1-\alpha|X)$). *Let $W_{\mathcal{U}} = \max_{u \in \mathcal{U}} 1/\sqrt{u(1-u)}$. There is a universal constant C_{Λ} such that*

- (i) $P\left(\Lambda \geq k \cdot C_{\Lambda} W_{\mathcal{U}} \sqrt{n \log p} \mid X\right) \leq p^{-k^2+1},$
- (ii) $\Lambda(1-\alpha|X) \leq \sqrt{1 + \log(1/\alpha)/\log p} \cdot C_{\Lambda} W_{\mathcal{U}} \sqrt{n \log p}$ with probability 1.

3.2. Rates of Convergence. In this section we establish the rate of convergence of ℓ_1 -QR. We start with the following preliminary result which shows that if the penalty level exceeds the specified threshold, for each $u \in \mathcal{U}$, the estimator $\widehat{\beta}(u) - \beta(u)$ will belong to the restricted set $A_u := \{\delta \in \mathbb{R}^p : \|\delta_{T_u^c}\|_1 \leq c_0 \|\delta_{T_u}\|_1, \|\delta_{T_u^c}\|_0 \leq n\}$.

LEMMA 3 (Restricted Set). *1. Under D.3, with probability at least $1 - \gamma$ we have for every $\delta \in \mathbb{R}^p$ that*

$$(3.1) \quad \frac{2}{3} \|\delta\|_{1,n} \leq \|\delta\|_1 \leq 2 \|\delta\|_{1,n}.$$

2. Moreover, if for some $\alpha \in (0, 1)$

$$(3.2) \quad \lambda \geq \lambda_0 := \frac{c_0 + 3}{c_0 - 3} \Lambda(1 - \alpha|X),$$

then with probability at least $1 - \alpha - \gamma$, uniformly in $u \in \mathcal{U}$, we have (3.1) and

$$\widehat{\beta}(u) - \beta(u) \in A_u = \{\delta \in \mathbb{R}^p : \|\delta_{T_u^c}\|_1 \leq c_0 \|\delta_{T_u}\|_1, \|\delta_{T_u^c}\|_0 \leq n\}.$$

This result is inspired [7]'s analogous result for least squares.

conditions I.1 and I.2, where I.2 assumes uniform in x_i consistency of the penalized estimator in the population, and does not hold in our main examples, e.g., in Design 1 with normal regressors.)

LEMMA 4 (Identifiability Relations over Restricted Set). *Condition D.4, namely $\text{RE}(c_0, m)$ and $\text{RNI}(c_0)$, implies that for any $\delta \in A_u$ and $u \in \mathcal{U}$,*

$$(3.3) \quad \|(\mathbb{E}[x_i x'_i])^{1/2} \delta\| \leq \|J_u^{1/2} \delta\| / \underline{f}^{1/2},$$

$$(3.4) \quad \|\delta_{T_u}\|_1 \leq \sqrt{s} \|J_u^{1/2} \delta\| / [\underline{f}^{1/2} \kappa_0],$$

$$(3.5) \quad \|\delta\|_1 \leq \sqrt{s}(1 + c_0) \|J_u^{1/2} \delta\| / [\underline{f}^{1/2} \kappa_0],$$

$$(3.6) \quad \|\delta\| \leq \left(1 + c_0 \sqrt{s/m}\right) \|J_u^{1/2} \delta\| / [\underline{f}^{1/2} \kappa_m],$$

$$(3.7) \quad Q_u(\beta(u) + \delta) - Q_u(\beta(u)) \geq (\|J_u^{1/2} \delta\|^2 / 4) \wedge (q \|J_u^{1/2} \delta\|).$$

This second preliminary result derives identifiability relations over A_u . It shows that the coefficients \underline{f} , κ_0 , and κ_m control moduli of continuity between various norms over the restricted set A_u , and the RNI coefficient q controls the quality of minoration of the objective function by a quadratic function over A_u .

Finally, the third preliminary result derives bounds on the empirical error over A_u :

LEMMA 5 (Control of Empirical Error). *Under D.1-4, for any $t > 0$ let*

$$\epsilon(t) := \sup_{u \in \mathcal{U}, \delta \in A_u, \|J_u^{1/2} \delta\| \leq t} \left| \widehat{Q}_u(\beta(u) + \delta) - Q_u(\beta(u) + \delta) - \left(\widehat{Q}_u(\beta(u)) - Q_u(\beta(u)) \right) \right|.$$

Then, there is a universal constant C_E such that for any $A > 1$, with probability at least $1 - 3\gamma - 3p^{-A^2}$

$$\epsilon(t) \leq t \cdot C_E \cdot \frac{(1 + c_0)A}{\underline{f}^{1/2} \kappa_0} \sqrt{\frac{s \log(p \vee [L \underline{f}^{1/2} \kappa_0 / t])}{n}}.$$

In order to prove the lemma we use a combination of chaining arguments and exponential inequalities for contractions [28]. Our use of the contraction principle is inspired by its fundamentally innovative use in [39]; however, the use of the contraction principle alone is not sufficient in our case. Indeed, first we need to make some adjustments to obtain error bounds over the neighborhoods defined by the intrinsic norm $\|J_u^{1/2} \cdot\|$ instead of the $\|\cdot\|_1$ norm; and second, we need to use chaining over $u \in \mathcal{U}$ to get uniformity over \mathcal{U} .

Armed with Lemmas 3-5, we establish the first main result. The result depends on the constants C_Λ , C_E , C_L , and C_f defined in Theorem 1, Lemma 5, D.2, and D.4.

THEOREM 2 (Uniform Bounds on Estimation Error of ℓ_1 -QR). *Assume conditions D.1-4 hold, and let $C > 2C_\Lambda \sqrt{1 + \log(1/\alpha)/\log p} \vee [C_E \sqrt{1 \vee [C_L + C_f + 1/2]}]$. Let λ_0 be defined as in (3.2). Then uniformly in the penalty level λ such that*

$$(3.8) \quad \lambda_0 \leq \lambda \leq C \cdot W_{\mathcal{U}} \sqrt{n \log p},$$

we have that, for any $A > 1$ with probability at least $1 - \alpha - 4\gamma - 3p^{-A^2}$,

$$\begin{aligned} \sup_{u \in \mathcal{U}} \|J_u^{1/2}(\hat{\beta}(u) - \beta(u))\| &\leq 8C \cdot \frac{(1 + c_0)W_{\mathcal{U}}A}{\underline{f}^{1/2}\kappa_0} \cdot \sqrt{\frac{s \log(p \vee n)}{n}}, \\ \sup_{u \in \mathcal{U}} \sqrt{E_x[x'(\hat{\beta}(u) - \beta(u))]^2} &\leq 8C \cdot \frac{(1 + c_0)W_{\mathcal{U}}A}{\underline{f}\kappa_0} \cdot \sqrt{\frac{s \log(p \vee n)}{n}}, \text{ and} \\ \sup_{u \in \mathcal{U}} \|\hat{\beta}(u) - \beta(u)\| &\leq \frac{1 + c_0\sqrt{s/m}}{\kappa_m} \cdot 8C \cdot \frac{(1 + c_0)W_{\mathcal{U}}A}{\underline{f}\kappa_0} \cdot \sqrt{\frac{s \log(p \vee n)}{n}}, \end{aligned}$$

provided s obeys the growth condition

$$(3.9) \quad 2C \cdot (1 + c_0)W_{\mathcal{U}}A \cdot \sqrt{s \log(p \vee n)} < q\underline{f}^{1/2}\kappa_0\sqrt{n}.$$

This result derives the rate of convergence of the ℓ_1 -penalized quantile regression estimator in the intrinsic norm and other norms of interest uniformly in $u \in \mathcal{U}$ as well as uniformly in the penalty level λ in the range specified by (3.8), which includes our recommended choice of λ_0 . We see that the rates of convergence for ℓ_1 -QR generally depend on the number of significant regressors s , the logarithm of the number of regressors p , the strength of identification summarized by κ_0 , κ_m , \underline{f} , and q , and the quantile indices of interest \mathcal{U} (as expected, extreme quantiles can slow down the rates of convergence). These rate results parallel the results of [7] obtained for ℓ_1 -penalized mean regression. Indeed, the role of the parameter \underline{f} is similar to the role of the standard deviation of the disturbance in mean regression. It is worth noting, however, that our results do not rely on normality and homoscedasticity assumptions, and our proofs have to address the non-quadratic nature of the objective function, with parameter q controlling the quality of quadratization. This parameter q enters the results only through the growth restriction (3.9) on s . At this point we refer the reader to Section 2.4 for a further discussion of this result in the context of the correlated normal design. Finally, we note that our proof combines the star-shaped geometry of the restricted set A_u with classical convexity arguments; this insight may be of interest in other problems.

PROOF OF THEOREM 2. We let

$$t := 8C \cdot \frac{(1 + c_0)W_{\mathcal{U}}A}{\underline{f}^{1/2}\kappa_0} \cdot \sqrt{\frac{s \log(p \vee n)}{n}},$$

and consider the following events:

- (i) Ω_1 := the event that (3.1) and $\hat{\beta}(u) - \beta(u) \in A_u$, uniformly in $u \in \mathcal{U}$, hold;
- (ii) Ω_2 := the event that the bound on empirical error $\epsilon(t)$ in Lemma 5 holds;

(iii) $\Omega_3 :=$ the event in which $\Lambda(1 - \alpha|X) \leq \sqrt{1 + \log(1/\alpha)/\log p} \cdot C_\Lambda W_{\mathcal{U}} \sqrt{n \log p}$.

By the choice of λ and Lemma 3, $P(\Omega_1) \geq 1 - \alpha - \gamma$; by Lemma 5 $P(\Omega_2) \geq 1 - 3\gamma - 3p^{-A^2}$; and by Theorem 1 $P(\Omega_3) = 1$, hence $P(\cap_{k=1}^3 \Omega_k) \geq 1 - \alpha - 4\gamma - 3p^{-A^2}$.

Given the event $\cap_{k=1}^3 \Omega_k$, we want to show the event that

$$(3.10) \quad \exists u \in \mathcal{U}, \quad \|J_u^{1/2}(\hat{\beta}(u) - \beta(u))\| > t$$

is impossible, which will prove the first bound. The other two bounds then follow from Lemma 4 and the first bound. First note that the event in (3.10) implies that for some $u \in \mathcal{U}$

$$0 > \min_{\delta \in A_u, \|J_u^{1/2}\delta\| \geq t} \hat{Q}_u(\beta(u) + \delta) - \hat{Q}_u(\beta(u)) + \frac{\lambda \sqrt{u(1-u)}}{n} (\|\beta(u) + \delta\|_{1,n} - \|\beta(u)\|_{1,n}).$$

The key observation is that by convexity of $\hat{Q}_u(\cdot) + \|\cdot\|_{1,n} \lambda \sqrt{u(1-u)}/n$ and by the fact that A_u is a cone, we can replace $\|J_u^{1/2}\delta\| \geq t$ by $\|J_u^{1/2}\delta\| = t$ in the above inequality and still preserve it:

$$0 > \min_{\delta \in A_u, \|J_u^{1/2}\delta\| = t} \hat{Q}_u(\beta(u) + \delta) - \hat{Q}_u(\beta(u)) + \frac{\lambda \sqrt{u(1-u)}}{n} (\|\beta(u) + \delta\|_{1,n} - \|\beta(u)\|_{1,n}).$$

Also, by inequality (3.4) in Lemma 4, for each $\delta \in A_u$

$$\|\beta(u)\|_{1,n} - \|\beta(u) + \delta\|_{1,n} \leq \|\delta_{T_u}\|_{1,n} \leq 2\|\delta_{T_u}\|_1 \leq 2\sqrt{s}\|J_u^{1/2}\delta\|/\underline{f}^{1/2}\kappa_0,$$

which then further implies

$$(3.11) \quad 0 > \min_{\delta \in A_u, \|J_u^{1/2}\delta\| = t} \hat{Q}_u(\beta(u) + \delta) - \hat{Q}_u(\beta(u)) - \frac{\lambda \sqrt{u(1-u)}}{n} \frac{2\sqrt{s}}{\underline{f}^{1/2}\kappa_0} \|J_u^{1/2}\delta\|.$$

Also by Lemma 5, under our choice of $t \geq 1/[\underline{f}^{1/2}\kappa_0\sqrt{n}]$, $\log(L\underline{f}\kappa_0^2) \leq (C_L + C_f) \log(n \vee p)$, and under event Ω_2

$$(3.12) \quad \epsilon(t) \leq t C_E \sqrt{1 \vee [C_L + C_f + 1/2]} \frac{(1 + c_0)A}{\underline{f}^{1/2}\kappa_0} \sqrt{\frac{s \log(p \vee n)}{n}}.$$

Therefore, we obtain from (3.11) and (3.12)

$$\begin{aligned} 0 &\geq \min_{\delta \in A_u, \|J_u^{1/2}\delta\| = t} Q_u(\beta(u) + \delta) - Q_u(\beta(u)) - \frac{\lambda \sqrt{u(1-u)}}{n} \frac{2\sqrt{s}}{\underline{f}^{1/2}\kappa_0} \|J_u^{1/2}\delta\| - \\ &\quad - t C_E \sqrt{1 \vee [C_L + C_f + 1/2]} \frac{(1 + c_0)A}{\underline{f}^{1/2}\kappa_0} \sqrt{\frac{s \log(p \vee n)}{n}}. \end{aligned}$$

Using the identifiability relation (3.7) stated in Lemma 4, we further get

$$0 > \frac{t^2}{4} \wedge (qt) - t \frac{\lambda \sqrt{u(1-u)}}{n} \frac{2\sqrt{s}}{\underline{f}^{1/2} \kappa_0} - t C_E \sqrt{1 \vee [C_L + C_f + 1/2]} \frac{(1+c_0)A}{\underline{f}^{1/2} \kappa_0} \sqrt{\frac{s \log(p \vee n)}{n}}.$$

Using the upper bound on λ under event Ω_3 , we obtain

$$0 > \frac{t^2}{4} \wedge (qt) - t C \frac{2\sqrt{s \log p}}{\sqrt{n}} \frac{W_{\mathcal{U}}}{\underline{f}^{1/2} \kappa_0} - t C_E \sqrt{1 \vee [C_L + C_f + 1/2]} \frac{(1+c_0)A}{\underline{f}^{1/2} \kappa_0} \sqrt{\frac{s \log(p \vee n)}{n}}.$$

Note that qt cannot be smaller than $t^2/4$ under the growth condition (3.9) in the theorem. Thus, using also the lower bound on C given in the theorem, $W_{\mathcal{U}} \geq 1$, and $c_0 \geq 1$, we obtain the relation

$$0 > \frac{t^2}{4} - t \cdot 2C \frac{(1+c_0)W_{\mathcal{U}}A}{\underline{f}^{1/2} \kappa_0} \cdot \sqrt{\frac{s \log(p \vee n)}{n}} = 0,$$

which is impossible. \square

3.3. Sparsity Properties. Next, we derive sparsity properties of the solution to ℓ_1 -penalized quantile regression. Fundamentally, sparsity is linked to the first order optimality conditions of (2.4) and therefore to the (sub)gradient of the criterion function. In the case of least squares, the gradient is a smooth (linear) function of the parameters. In the case of quantile regression, the gradient is a highly non-smooth (piece-wise constant) function. To control the sparsity of $\hat{\beta}(u)$ we rely on empirical process arguments to approximate gradients by smooth functions. In particular, we crucially exploit the fact that the entropy of all m -dimensional submodels of the p -dimensional model is of order $m \log p$, which depends on p only logarithmically.

The statement of the results will depend on the maximal k -sparse eigenvalue of $\mathbb{E}[x_i x_i']$ and $\mathbb{E}_n[x_i x_i']$:

$$(3.13) \quad \varphi_{\max}(k) = \max_{\delta \neq 0, \|\delta\|_0 \leq k} \frac{\mathbb{E}[(x_i' \delta)^2]}{\delta' \delta} \quad \text{and} \quad \phi(k) = \sup_{\delta \neq 0, \|\delta\|_0 \leq k} \frac{\mathbb{E}_n[(x_i' \delta)^2]}{\delta' \delta} \vee \frac{\mathbb{E}[(x_i' \delta)^2]}{\delta' \delta}.$$

In order to establish our main sparsity result, we need two preliminary lemmas.

LEMMA 6 (Empirical Pre-Sparsity). *Let $\hat{s} = \sup_{u \in \mathcal{U}} \|\hat{\beta}(u)\|_0$. Under D.1-4, for any $\lambda > 0$, with probability at least $1 - \gamma$ we have*

$$\hat{s} \leq n \wedge p \wedge [4n^2 \phi(\hat{s}) W_{\mathcal{U}}^2 / \lambda^2].$$

In particular, if $\lambda \geq 2\sqrt{2} W_{\mathcal{U}} \sqrt{n \log(n \vee p) \phi(n / \log(n \vee p))}$ then $\hat{s} \leq n / \log(n \vee p)$.

This lemma establishes an initial bound on the number of non-zero components \widehat{s} as a function of λ and $\phi(\widehat{s})$. Restricting $\lambda \geq 2\sqrt{2}W_{\mathcal{U}}\sqrt{n \log(n \vee p)\phi(n/\log(n \vee p))}$ makes the term $\phi(n/\log(n \vee p))$ appear in subsequent bounds instead of the term $\phi(n)$, which in turn weakens some assumptions. Indeed, not only is the first term smaller than the second, but also there are designs of interest where the second term diverges while the first does not; for instance, in Design 1, if $p \geq 2n$, we have $\phi(n/\log(n \vee p)) \lesssim_P 1$ while $\phi(n) \gtrsim_P \sqrt{\log p}$ by the Supplementary Material Appendix G.

The following lemma establishes a bound on the sparsity as a function of the rate of convergence.

LEMMA 7 (Empirical Sparsity). *Assume D.1-4 and let $r = \sup_{u \in \mathcal{U}} \|J_u^{1/2}(\widehat{\beta}(u) - \beta(u))\|$. Then, for any $\varepsilon > 0$, there is a constant $K_\varepsilon \geq \sqrt{2}$ such that with probability at least $1 - \varepsilon - \gamma$*

$$\frac{\sqrt{\widehat{s}}}{W_{\mathcal{U}}} \leq \mu(\widehat{s}) \frac{n}{\lambda}(r \wedge 1) + \sqrt{\widehat{s}} K_\varepsilon \frac{\sqrt{n \log(n \vee p)\phi(\widehat{s})}}{\lambda}, \quad \mu(k) := 2\sqrt{\varphi_{\max}(k)} \left(1 \vee 2\bar{f}/\underline{f}^{1/2}\right).$$

Finally, we combine these results to establish the main sparsity result. In what follows, we define $\bar{\phi}_\varepsilon$ as a constant such that $\phi(n/\log(n \vee p)) \leq \bar{\phi}_\varepsilon$ with probability $1 - \varepsilon$.

THEOREM 3 (Uniform Sparsity Bounds). *Let $\varepsilon > 0$ be any constant, assume D.1-4 hold, and let λ satisfy $\lambda \geq \lambda_0$ and*

$$KW_{\mathcal{U}}\sqrt{n \log(n \vee p)} \leq \lambda \leq K'W_{\mathcal{U}}\sqrt{n \log(n \vee p)}$$

for some constant $K' \geq K \geq 2K_\varepsilon\bar{\phi}_\varepsilon^{1/2}$, for K_ε defined in Lemma 7. Then, for any $A > 1$ with probability at least $1 - \alpha - 2\varepsilon - 4\gamma - p^{-A^2}$

$$\widehat{s} := \sup_{u \in \mathcal{U}} \|\widehat{\beta}(u)\|_0 \leq s \cdot \left[16\mu W_{\mathcal{U}}/\underline{f}^{1/2}\kappa_0\right]^2 [(1 + c_0)AK'/K]^2,$$

where $\mu := \mu(n/\log(n \vee p))$, provided that s obeys the growth condition

$$(3.14) \quad 2K'(1 + c_0)AW_{\mathcal{U}}\sqrt{s \log(n \vee p)} < q\underline{f}^{1/2}\kappa_0\sqrt{n}.$$

The theorem states that by setting the penalty level λ to be possibly higher than our initial recommended choice λ_0 , we can control \widehat{s} , which will be crucial for good performance of the post-penalized estimator. As a corollary, we note that if (a) $\mu \lesssim 1$, (b) $1/(\underline{f}^{1/2}\kappa_0) \lesssim 1$, and (c) $\bar{\phi}_\varepsilon \lesssim 1$ for each $\varepsilon > 0$, then $\widehat{s} \lesssim s$ with a high probability, so the dimension of the selected model is about the same as the dimension of the true model. Conditions (a), (b), and (c) easily hold for the correlated normal design in Design 1. In particular, (c) follows from the concentration

inequalities and from results in classical random matrix theory; see the Supplementary Material Appendix G for proofs. Therefore the possibly higher λ needed to achieve the stated sparsity bound does not slow down the rate of ℓ_1 -QR in this case. The growth condition (3.14) on s is also weak in this case.

PROOF OF THEOREM 3. By the choice of K and Lemma 6, $\hat{s} \leq n/\log(n \vee p)$ with probability $1 - \varepsilon$. With at least the same probability, the choice of λ yields

$$K_\varepsilon \frac{\sqrt{n \log(n \vee p) \phi(\hat{s})}}{\lambda} \leq \frac{K_\varepsilon \bar{\phi}_\varepsilon^{1/2}}{K W_{\mathcal{U}}} \leq \frac{1}{2W_{\mathcal{U}}},$$

so that by virtue of Lemma 7 and by $\mu(\hat{s}) \leq \mu := \mu(n/\log(n \vee p))$,

$$\frac{\sqrt{\hat{s}}}{W_{\mathcal{U}}} \leq \mu \frac{(r \wedge 1)n}{\lambda} + \frac{\sqrt{\hat{s}}}{2W_{\mathcal{U}}} \quad \text{or} \quad \frac{\sqrt{\hat{s}}}{W_{\mathcal{U}}} \leq 2\mu \frac{(r \wedge 1)n}{\lambda},$$

with probability $1 - 2\varepsilon$. Since all conditions of Theorem 2 hold, we obtain the result by plugging in the upper bound on $r = \sup_{u \in \mathcal{U}} \|J_u^{1/2}(\hat{\beta}(u) - \beta(u))\|$ from Theorem 2. \square

3.4. *Model Selection Properties.* Next we turn to the model selection properties of ℓ_1 -QR.

THEOREM 4 (Model Selection Properties of ℓ_1 -QR). *Let $r^o = \sup_{u \in \mathcal{U}} \|\hat{\beta}(u) - \beta(u)\|$. If $\inf_{u \in \mathcal{U}} \min_{j \in T_u} |\beta_j(u)| > r^o$, then*

$$(3.15) \quad T_u := \text{support}(\beta(u)) \subseteq \hat{T}_u := \text{support}(\hat{\beta}(u)) \quad \text{for all } u \in \mathcal{U}.$$

Moreover, the hard-thresholded estimator $\bar{\beta}(u)$, defined for any $\gamma \geq 0$ by

$$(3.16) \quad \bar{\beta}_j(u) = \hat{\beta}_j(u) 1\left\{|\hat{\beta}_j(u)| > \gamma\right\}, \quad u \in \mathcal{U}, \quad j = 1, \dots, p,$$

provided that γ is chosen such that $r^o < \gamma < \inf_{u \in \mathcal{U}} \min_{j \in T_u} |\beta_j(u)| - r^o$, satisfies

$$\text{support}(\bar{\beta}(u)) = T_u \quad \text{for all } u \in \mathcal{U}.$$

These results parallel analogous results in [32] for mean regression. The first result says that if non-zero coefficients are well separated from zero, then the support of ℓ_1 -QR includes the support of the true model. The inclusion of the true support in (3.15) is in general one-sided; the support of the estimator can include some unnecessary components having true coefficients equal to zero. The second result states that if the further conditions are satisfied, additional hard thresholding can eliminate inclusions of such unnecessary components. The value of the hard threshold must explicitly depend on the unknown value $\min_{j \in T_u} |\beta_j(u)|$, characterizing the

separation of non-zero coefficients from zero. The additional conditions stated in this theorem are strong and perfect model selection appears quite unlikely in practice. Certainly it does not work in all real empirical examples we have explored. This motivates our analysis of the post-model-selected estimator under conditions that allow for imperfect model selection, including cases where we miss some non-zero components or have additional unnecessary components.

3.5. The post-penalized estimator. In this section we establish a bound on the rate of convergence of the post-penalized estimator. The proof relies crucially on the identifiability and control of the empirical error over the sparse sets $\tilde{A}_u(\tilde{m}) := \{\delta \in \mathbb{R}^p : \|\delta_{T_u^c}\|_0 \leq \tilde{m}\}$.

LEMMA 8 (Sparse Identifiability and Control of Empirical Error). *1. Suppose D.1 and D.5 hold. Then for all $\delta \in \tilde{A}_u(\tilde{m})$, $u \in \mathcal{U}$, and $\tilde{m} \leq n$, we have that*

$$(3.17) \quad Q_u(\beta(u) + \delta) - Q_u(\beta(u)) \geq \frac{\|J_u^{1/2}\delta\|^2}{4} \wedge \left(\tilde{q}_{\tilde{m}} \|J_u^{1/2}\delta\| \right).$$

2. Suppose D.1-2 and D.5 hold and that $|\cup_{u \in \mathcal{U}} T_u| \leq n$. Then for any $\varepsilon > 0$, there is a constant C_ε such that with probability at least $1 - \varepsilon$ the empirical error

$$\epsilon_u(\delta) := \left| \hat{Q}_u(\beta(u) + \delta) - Q_u(\beta(u) + \delta) - \left(\hat{Q}_u(\beta(u)) - Q_u(\beta(u)) \right) \right|$$

obeys

$$\sup_{u \in \mathcal{U}, \delta \in \tilde{A}_u(\tilde{m}), \delta \neq 0} \frac{\epsilon_u(\delta)}{\|\delta\|} \leq C_\varepsilon \sqrt{\frac{(\tilde{m} \log(n \vee p) + s \log n) \phi(\tilde{m} + s)}{n}} \text{ for all } \tilde{m} \leq n.$$

In order to prove this lemma we exploit the crucial fact that the entropy of all m -dimensional submodels of the p -dimensional model is of order $m \log p$, which depends on p only logarithmically. The following theorem establishes the properties of post-model-selection estimators.

THEOREM 5 (Uniform Bounds on Estimation Error of post- ℓ_1 -QR). *Assume the conditions of Theorem 2 hold, assume that $|\cup_{u \in \mathcal{U}} T_u| \leq n$, and assume D.5 holds with $\hat{m} := \sup_{u \in \mathcal{U}} \|\hat{\beta}_{T_u^c}(u)\|_0$ with probability $1 - \varepsilon$. Then for any $\varepsilon > 0$ there is a constant C_ε such that the bounds*

$$(3.18) \quad \begin{aligned} \sup_{u \in \mathcal{U}} \left\| J_u^{1/2}(\tilde{\beta}(u) - \beta(u)) \right\| &\leq \frac{4C_\varepsilon \sqrt{\phi(\hat{m} + s)}}{\underline{f}^{1/2} \tilde{\kappa}_{\hat{m}}} \cdot \sqrt{\frac{\hat{m} \log(n \vee p) + s \log n}{n}} + \\ &\quad + \sup_{u \in \mathcal{U}} 1\{T_u \not\subseteq \hat{T}_u\} \cdot \frac{4\sqrt{2(1+c_0)A}}{\underline{f}^{1/2} \kappa_0} \cdot C \cdot W_u \sqrt{\frac{s \log(n \vee p)}{n}}, \\ \sup_{u \in \mathcal{U}} \sqrt{\mathbb{E}_x[x'(\tilde{\beta}(u) - \beta(u))]^2} &\leq \sup_{u \in \mathcal{U}} \left\| J_u^{1/2}(\tilde{\beta}(u) - \beta(u)) \right\| / \underline{f}^{1/2}, \\ \sup_{u \in \mathcal{U}} \left\| \tilde{\beta}(u) - \beta(u) \right\| &\leq \sup_{u \in \mathcal{U}} \left\| J_u^{1/2}(\tilde{\beta}(u) - \beta(u)) \right\| / \underline{f}^{1/2} \tilde{\kappa}_{\hat{m}}, \end{aligned}$$

hold with probability at least $1 - \alpha - 3\gamma - 3p^{-A^2} - 2\varepsilon$, provided that s obeys the growth condition

$$\tilde{q}_{\hat{m}} \frac{C_\varepsilon \sqrt{(\hat{m} \log(n \vee p) + s \log n) \phi(\hat{m} + s)}}{\sqrt{n} \underline{f}^{1/2} \tilde{\kappa}_{\hat{m}}} + \sup_{u \in \mathcal{U}} 1\{T_u \not\subseteq \hat{T}_u\} 2A(1+c_0) \cdot C^2 W_{\mathcal{U}}^2 \cdot \frac{s \log(p \vee n)}{n \underline{f} \kappa_0^2} \leq \tilde{q}_{\hat{m}}^2.$$

This theorem describes the performance of post- ℓ_1 -QR. However, an inspection of the proof reveals that it can be applied to any post-model selection estimator. From Theorem 5 we can conclude that in many interesting cases the rates of post- ℓ_1 -QR could be the same or faster than the rate of ℓ_1 -QR. Indeed, first consider the case where the model selection fails to contain the true model, i.e., $\sup_{u \in \mathcal{U}} 1\{T_u \not\subseteq \hat{T}_u\} = 1$ with a non-negligible probability. If (a) $\hat{m} \leq \hat{s} \lesssim_P s$, (b) $\phi(\hat{m} + s) \lesssim_P 1$, and (c) the constants \underline{f} and $\tilde{\kappa}_{\hat{m}}^2$ are of the same order as \underline{f} and $\kappa_0 \kappa_m$, respectively, then the rate of convergence of post- ℓ_1 -QR is the same as the rate of convergence of ℓ_1 -QR. Recall that Theorem 3 provides sufficient conditions needed to achieve (a), which hold in Design 1. Recall also that in Design 1, (b) holds by concentration of measure and classical results in random matrix theory, as shown in the Supplementary Material Appendix G, and (c) holds by the calculations presented in Section 2. This verifies our claim regarding the performance of post- ℓ_1 -QR in the overview, Section 2.4. The intuition for this result is that even though ℓ_1 -QR misses true components, it does not miss very important ones, allowing post- ℓ_1 -QR still to perform well. Second, consider the case where the model selection succeeds in containing the true model, i.e., $\sup_{u \in \mathcal{U}} 1\{T_u \not\subseteq \hat{T}_u\} = 0$ with probability approaching one, and that the number of unnecessary components obeys $\hat{m} = o_P(s)$. In this case the rate of convergence of post- ℓ_1 -QR can be faster than the rate of convergence of ℓ_1 -QR. In the extreme case of perfect model selection, when $\hat{m} = 0$ with a high probability, post- ℓ_1 -QR becomes the oracle estimator with a high probability. We refer the reader to Section 2 for further discussion, and note that this result could be of interest in other problems.

PROOF OF THEOREM 5. Let

$$\hat{\delta}(u) = \hat{\beta}(u) - \beta(u), \quad \tilde{\delta}(u) := \tilde{\beta}(u) - \beta(u), \quad t_u := \|J_u^{1/2} \tilde{\delta}(u)\|,$$

and B_n be a random variable such that $B_n = \sup_{u \in \mathcal{U}} \hat{Q}_u(\hat{\beta}(u)) - \hat{Q}_u(\beta(u))$. By the optimality of $\hat{\beta}(u)$ in (3.16), with probability $1 - \gamma$ we have uniformly in $u \in \mathcal{U}$

$$\begin{aligned} \hat{Q}_u(\hat{\beta}(u)) - \hat{Q}_u(\beta(u)) &\leq \frac{\lambda \sqrt{u(1-u)}}{n} (\|\beta(u)\|_{1,n} - \|\hat{\beta}(u)\|_{1,n}) \\ (3.19) \quad &\leq \frac{\lambda \sqrt{u(1-u)}}{n} \|\hat{\delta}_{T_u}(u)\|_{1,n} \leq \frac{\lambda \sqrt{u(1-u)}}{n} 2 \|\hat{\delta}_{T_u}(u)\|_1, \end{aligned}$$

where the last term in (3.19) is bounded by

$$(3.20) \quad \frac{\lambda \sqrt{u(1-u)}}{n} \frac{2\sqrt{s} \|J_u^{1/2} \widehat{\delta}(u)\|}{\underline{f}^{1/2} \kappa_0} \leq \frac{\lambda \sqrt{u(1-u)}}{n} \frac{2\sqrt{s}}{\underline{f}^{1/2} \kappa_0} \sup_{u \in \mathcal{U}} \|J_u^{1/2} (\widehat{\beta}(u) - \beta(u))\|,$$

using that $\|J_u^{1/2} \widehat{\delta}(u)\| \geq \underline{f}^{1/2} \kappa_0 \|\widehat{\delta}_{T_u}(u)\|$ from $\text{RE}(c_0, 0)$ implied by D.4. Therefore, by Theorem 2 we have

$$B_n \leq \frac{\lambda \sqrt{u(1-u)}}{n} \frac{2\sqrt{s}}{\underline{f}^{1/2} \kappa_0} 8C \cdot \frac{(1+c_0)W_{\mathcal{U}}A}{\underline{f}^{1/2} \kappa_0} \cdot \sqrt{\frac{s \log(p \vee n)}{n}}$$

with probability $1 - \alpha - 3\gamma - 3p^{-A^2}$.

For every $u \in \mathcal{U}$, by optimality of $\widetilde{\beta}(u)$ in (2.5),

$$(3.21) \quad \widehat{Q}_u(\widetilde{\beta}(u)) - \widehat{Q}_u(\beta(u)) \leq 1\{T_u \not\subseteq \widehat{T}_u\} \left(\widehat{Q}_u(\widehat{\beta}(u)) - \widehat{Q}_u(\beta(u)) \right) \leq 1\{T_u \not\subseteq \widehat{T}_u\} B_n.$$

Also, by Lemma 8, with probability at least $1 - \varepsilon$, we have

$$(3.22) \quad \sup_{u \in \mathcal{U}} \frac{\epsilon_u(\widetilde{\delta}(u))}{\|\widetilde{\delta}(u)\|} \leq C_\varepsilon \sqrt{\frac{(\widehat{m} \log(n \vee p) + s \log n) \phi(\widehat{m} + s)}{n}} =: A_{\varepsilon, n}.$$

Recall that $\sup_{u \in \mathcal{U}} \|\widetilde{\delta}_{T_u}(u)\| \leq \widehat{m} \leq n$ so that by D.5 $t_u \geq \underline{f}^{1/2} \widetilde{\kappa}_{\widehat{m}} \|\widetilde{\delta}(u)\|$ for all $u \in \mathcal{U}$ with probability $1 - \varepsilon$. Thus, combining relations (3.21) and (3.22), for every $u \in \mathcal{U}$

$$Q_u(\widetilde{\beta}(u)) - Q_u(\beta(u)) \leq t_u A_{\varepsilon, n} / [\underline{f}^{1/2} \widetilde{\kappa}_{\widehat{m}}] + 1\{T_u \not\subseteq \widehat{T}_u\} B_n$$

with probability at least $1 - 2\varepsilon$. Invoking the sparse identifiability relation (3.17) of Lemma 8, with the same probability, for all $u \in \mathcal{U}$,

$$(t_u^2/4) \wedge (\widetilde{q}_{\widehat{m}} t_u) \leq t_u A_{\varepsilon, n} / [\underline{f}^{1/2} \widetilde{\kappa}_{\widehat{m}}] + 1\{T_u \not\subseteq \widehat{T}_u\} B_n.$$

We then conclude that under the assumed growth condition on s , this inequality implies

$$t_u \leq 4A_{\varepsilon, n} / [\underline{f}^{1/2} \widetilde{\kappa}_{\widehat{m}}] + 1\{T_u \not\subseteq \widehat{T}_u\} \sqrt{4B_n \vee 0}$$

for every $u \in \mathcal{U}$, and the bounds stated in the theorem now follow from the definition of \underline{f} and $\widetilde{\kappa}_{\widehat{m}}$. \square

4. Empirical Performance. In order to access the finite sample practical performance of the proposed estimators, we conducted a Monte Carlo study and an application to international economic growth.

4.1. *Monte Carlo Simulation.* In order to assess the finite sample practical performance of the proposed estimators, we conducted a Monte Carlo study. We will compare the performance of the ℓ_1 -penalized, post- ℓ_1 -penalized, and the ideal oracle quantile regression estimators. Recall that the post-penalized estimator applies canonical quantile regression to the model selected by the penalized estimator. The oracle estimator applies canonical quantile regression to the true model. (Of course, such an estimator is not available outside Monte Carlo experiments.) We focus our attention on the model selection properties of the penalized estimator and biases and empirical risks of these estimators.

We begin by considering the following regression model:

$$y = x'\beta(0.5) + \varepsilon, \quad \beta(0.5) = (1, 1, 1/2, 1/3, 1/4, 1/5, 0, \dots, 0)',$$

where as in Design 1, $x = (1, z')'$ consists of an intercept and covariates $z \sim N(0, \Sigma)$, and the errors ε are independently and identically distributed $\varepsilon \sim N(0, \sigma^2)$. The dimension p of covariates x is 500, and the dimension s of the true model is 6, and the sample size n is 100. We set the regularization parameter λ equal to the 0.9-quantile of the pivotal random variable Λ , following our proposal in Section 2. The regressors are correlated with $\Sigma_{ij} = \rho^{|i-j|}$ and $\rho = 0.5$. We consider two levels of noise, namely $\sigma = 1$ and $\sigma = 0.1$.

We summarize the model selection performance of the penalized estimator in Figures 1 and 2. In the left panels of the figures, we plot the frequencies of the dimensions of the selected model; in the right panels we plot the frequencies of selecting the correct regressors. From the left panels we see that the frequency of selecting a much larger model than the true model is very small in both designs. In the design with a larger noise, as the right panel of Figure 1 shows, the penalized quantile regression never selects the entire true model correctly, always missing the regressors with small coefficients. However, it almost always includes the three regressors with the largest coefficients. (Notably, despite this partial failure of the model selection, post-penalized quantile regression still performs well, as we report below.) On the other hand, we see from the right panel of Figure 2 that in the design with a lower noise level penalized quantile regression rarely misses any component of the true support. These results confirm the theoretical results of Theorem 4, namely, that when the non-zero coefficients are well separated from zero, the penalized estimator should select a model that includes the true model as a subset. Moreover, these results also confirm the theoretical result of Theorem 3, namely, that the dimension of the selected model should be of the same stochastic order as the dimension of the true model. In summary, the model selection performance of the penalized estimator agrees very well with our theoretical results.

We summarize results on estimation performance in Table 1, which records for each estimator $\tilde{\beta}$ the norm of the bias $\|E[\tilde{\beta}] - \beta_0\|$ and also the empirical risk $[E[x'_i(\tilde{\beta} - \beta_0)]^2]^{1/2}$ for recovering

the regression function. Penalized quantile regression has a substantial bias, as we would expect from the definition of the estimator which penalizes large deviations of coefficients from zero. We see that the post-penalized quantile regression drastically improves upon the penalized quantile regression, particularly in terms of reducing the bias, which results in a much lower overall empirical risk. Notably, despite that under the higher noise level the penalized estimator never recovers the true model correctly the post-penalized estimator still performs well. This is because the penalized estimator always manages to select the most important regressors. We also see that the empirical risk of the post-penalized estimator is within a factor of $\sqrt{\log p}$ of the empirical risk of the oracle estimator, as we would expect from our theoretical results. Under the lower noise level, the post-penalized estimator performs almost identically to the ideal oracle estimator. We would expect this since in this case the penalized estimator selects the model especially well, making the post-penalized estimator nearly the oracle. In summary, we find the estimation performance of the penalized and post-penalized estimators to be in close agreement with our theoretical results.

MONTE CARLO RESULTS

Design A ($\sigma = 1$)

	Mean ℓ_0 -norm	Mean ℓ_1 -norm	Bias	Empirical Risk
Penalized QR	3.67	1.28	0.92	1.22
Post-Penalized QR	3.67	2.90	0.27	0.57
Oracle QR	6.00	3.31	0.03	0.33

Design B ($\sigma = 0.1$)

	Mean ℓ_0 -norm	Mean ℓ_1 -norm	Bias	Empirical Risk
Penalized QR	6.09	2.98	0.13	0.19
Post-Penalized QR	6.09	3.28	0.00	0.04
Oracle QR	6.00	3.28	0.00	0.03

TABLE 1

The table displays the average ℓ_0 and ℓ_1 norm of the estimators as well as mean bias and empirical risk. We obtained the results using 100 Monte Carlo repetitions for each design.

4.2. International Economic Growth Example. In this section we apply ℓ_1 -penalized quantile regression to an international economic growth example, using it primarily as a method for model selection. We use the Barro and Lee data consisting of a panel of 138 countries for the period of 1960 to 1985. We consider the national growth rates in gross domestic product (GDP) per capita as a dependent variable y for the periods 1965-75 and 1975-85.⁵ In our analysis, we will consider a model with $p = 60$ covariates, which allows for a total of $n = 90$ complete

⁵The growth rate in GDP over a period from t_1 to t_2 is commonly defined as $\log(GDP_{t_2}/GDP_{t_1}) - 1$.




FIG 1. The figure summarizes the covariate selection results for the design with $\sigma = 1$, based on 100 Monte Carlo repetitions. The left panel plots the histogram for the number of covariates selected out of the possible 500 covariates. The right panel plots the histogram for the number of significant covariates selected; there are in total 6 significant covariates amongst 500 covariates. The sample size for each repetition was $n = 100$.




FIG 2. The figure summarizes the covariate selection results for the design with $\sigma = 0.1$, based on 100 Monte Carlo repetitions. The left panel plots the histogram for the number of covariates selected out of the possible 500 covariates. The right panel plots the histogram for the number of significant covariates selected; there are in total 6 significant covariates amongst 500 covariates. The sample size for each repetition was $n = 100$.

observations. Our goal here is to select a subset of these covariates and briefly compare the resulting models to the standard models used in the empirical growth literature (Barro and Sala-i-Martin [1], Koenker and Machado [25]).

One of the central issues in the empirical growth literature is the estimation of the effect of an initial (lagged) level of GDP per capita on the growth rates of GDP per capita. In particular, a key prediction from the classical Solow-Swan-Ramsey growth model is the hypothesis of convergence, which states that poorer countries should typically grow faster and therefore should tend to catch up with the richer countries. Thus, such a hypothesis states that the effect of the initial level of GDP on the growth rate should be negative. As pointed out in Barro and Sala-i-Martin [2], this hypothesis is rejected using a simple bivariate regression of growth rates on the initial level of GDP. (In our case, median regression yields a positive coefficient of 0.00045.) In order to reconcile the data and the theory, the literature has focused on estimating the effect *conditional* on the pertinent characteristics of countries. Covariates that describe such characteristics can include variables measuring education and science policies, strength of market institutions, trade openness, savings rates and others [2]. The theory then predicts that for countries with similar other characteristics the effect of the initial level of GDP on the growth rate should be negative ([2]).

Given that the number of covariates we can condition on is comparable to the sample size, covariate selection becomes an important issue in this analysis ([29], [36]). In particular, previous findings came under severe criticism for relying on ad hoc procedures for covariate selection. In fact, in some cases, all of the previous findings have been questioned ([29]). Since the number of covariates is high, there is no simple way to resolve the model selection problem using only classical tools. Indeed the number of possible lower-dimensional models is very large, although [29] and [36] attempt to search over several millions of these models. Here we use the Lasso selection device, specifically ℓ_1 -penalized median regression, to resolve this important issue.

Let us now turn to our empirical results. We performed covariate selection using ℓ_1 -penalized median regression, where we initially used our data-driven choice of penalization parameter λ . This initial choice led us to select no covariates, which is consistent with the situations in which the true coefficients are not well-separated from zero. We then proceeded to slowly decrease the penalization parameter in order to allow for some covariates to be selected. We present the model selection results in Table 3. With the first relaxation of the choice of λ , we select the black market exchange rate premium (characterizing trade openness) and a measure of political instability. With a second relaxation of the choice of λ we select an additional set of educational attainment variables, and several others reported in the table. With a third

relaxation of λ we include yet another set of variables also reported in the table. We refer the reader to [1] and [2] for a complete definition and discussion of each of these variables.

We then proceeded to apply ordinary median regression to the selected models and we also report the standard confidence intervals for these estimates. Table 2 shows these results. We should note that the confidence intervals do not take into account that we have selected the models using the data. (In an ongoing companion work, we are working on devising procedures that will account for this.) We find that in all models with additional selected covariates, the median regression coefficients on the initial level of GDP is always negative and the standard confidence intervals do not include zero. Similar conclusions also hold for quantile regressions with quantile indices in the middle range. In summary, we believe that our empirical findings support the hypothesis of convergence from the classical Solow-Swan-Ramsey growth model. Of course, it would be good to find formal inferential methods to fully support this hypothesis. Finally, our findings also agree and thus support the previous findings reported in Barro and Sala-i-Martin [1] and Koenker and Machado [25].

CONFIDENCE INTERVALS AFTER MODEL SELECTION FOR THE INTERNATIONAL GROWTH REGRESSIONS

Penalization Parameter	Real GDP per capita (log)	
$\lambda = 1.077968$	Coefficient	90% Confidence Interval
$\lambda/2$	-0.01691	$[-0.02552, -0.00444]$
$\lambda/3$	-0.04121	$[-0.05485, -0.02976]$
$\lambda/4$	-0.04466	$[-0.06510, -0.03410]$
$\lambda/5$	-0.05148	$[-0.06521, -0.03296]$

TABLE 2

The table above displays the coefficient and a 90% confidence interval associated with each model selected by the corresponding penalty parameter. The selected models are displayed in Table 3.

APPENDIX A: PROOF OF THEOREM 1

PROOF OF THEOREM 1. We note $\Lambda \leq W_{\mathcal{U}} \max_{1 \leq j \leq p} \sup_{u \in \mathcal{U}} n \mathbb{E}_n [(u - 1\{u_i \leq u\})x_{ij}/\hat{\sigma}_j]$. For any $u \in \mathcal{U}$, $j \in \{1, \dots, p\}$ we have by Lemma 1.5 in [28] that $P(|\mathbb{G}_n[(u - 1\{u_i \leq u\})x_{ij}/\hat{\sigma}_j]| \geq \tilde{K}) \leq 2 \exp(-\tilde{K}^2/2)$. Hence by the symmetrization lemma for probabilities, Lemma 2.3.7 in [40], with $\tilde{K} \geq 2\sqrt{\log 2}$ we have

$$\begin{aligned}
 (A.1) \quad P(\Lambda > \tilde{K}\sqrt{n}|X) &\leq 4P\left(\sup_{u \in \mathcal{U}} \max_{1 \leq j \leq p} |\mathbb{G}_n^o[(u - 1\{u_i \leq u\})x_{ij}/\hat{\sigma}_j]| > \tilde{K}/(4W_{\mathcal{U}})|X\right) \\
 &\leq 4p \max_{1 \leq j \leq p} P\left(\sup_{u \in \mathcal{U}} |\mathbb{G}_n^o[(u - 1\{u_i \leq u\})x_{ij}/\hat{\sigma}_j]| > \tilde{K}/(4W_{\mathcal{U}})|X\right),
 \end{aligned}$$

MODEL SELECTION RESULTS FOR THE INTERNATIONAL GROWTH REGRESSIONS

Penalization	
Parameter	Real GDP per capita (log) is included in all models
$\lambda = 1.077968$	Additional Selected Variables
λ	-
$\lambda/2$	Black Market Premium (log) Political Instability
$\lambda/3$	Black Market Premium (log) Political Instability Measure of tariff restriction Infant mortality rate Ratio of real government "consumption" net of defense and education Exchange rate % of "higher school complete" in female population % of "secondary school complete" in male population
$\lambda/4$	Black Market Premium (log) Political Instability Measure of tariff restriction Infant mortality rate Ratio of real government "consumption" net of defense and education Exchange rate % of "higher school complete" in female population % of "secondary school complete" in male population Female gross enrollment ratio for higher education % of "no education" in the male population Population proportion over 65 Average years of secondary schooling in the male population
$\lambda/5$	Black Market Premium (log) Political Instability Measure of tariff restriction Infant mortality rate Ratio of real government "consumption" net of defense and education Exchange rate % of "higher school complete" in female population % of "secondary school complete" in male population Female gross enrollment ratio for higher education % of "no education" in the male population Population proportion over 65 Average years of secondary schooling in the male population Growth rate of population % of "higher school attained" in male population Ratio of nominal government expenditure on defense to nominal GDP Ratio of import to GDP

TABLE 3

For this particular decreasing sequence of penalization parameters we obtained nested models.

where \mathbb{G}_n^o denotes the symmetrized empirical process (see [40]) generated by the Rademacher variables $\varepsilon_i, i = 1, \dots, n$, which are independent of $U = (u_1, \dots, u_n)$ and $X = (x_1, \dots, x_n)$. Let us condition on U and X , and define $\mathcal{F}_j = \{\varepsilon_i x_{ij}(u - 1\{u_i \leq u\})/\hat{\sigma}_j : u \in \mathcal{U}\}$ for $j = 1, \dots, p$. The VC dimension of \mathcal{F}_j is at most 6. Therefore, by Theorem 2.6.7 of [40] for some universal constant $C'_1 \geq 1$ the function class \mathcal{F}_j with envelope function F_j obeys

$$N(\varepsilon \|F_j\|_{\mathbb{P}_{n,2}}, \mathcal{F}_j, L_2(\mathbb{P}_n)) \leq n(\varepsilon, \mathcal{F}_j) = C'_1 \cdot 6 \cdot (16e)^6 (1/\varepsilon)^{10},$$

where $N(\varepsilon, \mathcal{F}, L_2(\mathbb{P}_n))$ denotes the minimal number of balls of radius ε with respect to the $L_2(\mathbb{P}_n)$ norm $\|\cdot\|_{\mathbb{P}_{n,2}}$ needed to cover the class of functions \mathcal{F} ; see [40].

Conditional on the data $U = (u_1, \dots, u_n)$ and $X = (x_1, \dots, x_n)$, the symmetrized empirical process $\{\mathbb{G}_n^o(f), f \in \mathcal{F}_j\}$ is sub-Gaussian with respect to the $L_2(\mathbb{P}_n)$ norm by the Hoeffding inequality; see, e.g., [40]. Since $\|F_j\|_{\mathbb{P}_{n,2}} \leq 1$ and $\rho(\mathcal{F}_j, \mathbb{P}_n) = \sup_{f \in \mathcal{F}_j} \|f\|_{\mathbb{P}_{n,2}}/\|F_j\|_{\mathbb{P}_{n,2}} \leq 1$, we have

$$\|F_j\|_{\mathbb{P}_{n,2}} \int_0^{\rho(\mathcal{F}_j, \mathbb{P}_n)/4} \sqrt{\log n(\varepsilon, \mathcal{F}_j)} d\varepsilon \leq \bar{e} := (1/4) \sqrt{\log(6C'_1(16e)^6)} + (1/4) \sqrt{10 \log 4}.$$

By Lemma 16 with $D = 1$, there is a universal constant c such that for any $K \geq 1$:

$$\begin{aligned} P\left(\sup_{f \in \mathcal{F}_j} |\mathbb{G}_n^o(f)| > Kc\bar{e}|X, U\right) &\leq \int_0^{1/2} \varepsilon^{-1} n(\varepsilon, \mathcal{F}_j)^{-(K^2-1)} d\varepsilon \\ (A.2) \qquad \qquad \qquad &\leq (1/2)[6C'_1(16e)^6]^{-(K^2-1)} \frac{(1/2)^{10(K^2-1)}}{10(K^2-1)}. \end{aligned}$$

By (A.1) and (A.2) for any $k \geq 1$ we have

$$\begin{aligned} P\left(\Lambda \geq k \cdot (4\sqrt{2}c\bar{e})W_U \sqrt{n \log p} | X\right) &\leq 4p \max_{1 \leq j \leq p} \mathbb{E}_U P\left(\sup_{f \in \mathcal{F}_j} |\mathbb{G}_n^o(f)| > k\sqrt{2 \log p} c\bar{e} | X, U\right) \\ &\leq p^{-6k^2+1} \leq p^{-k^2+1} \end{aligned}$$

since $(2k^2 \log p - 1) \geq (\log 2 - 0.5)k^2 \log p$ for $p \geq 2$. Thus, result (i) holds with $C_\Lambda := 4\sqrt{2}c\bar{e}$. Result (ii) follows immediately by choosing $k = \sqrt{1 + \log(1/\alpha)/\log p}$ to make the right side of the display above equal to α . \square

APPENDIX B: PROOFS OF LEMMAS 3-5 (USED IN THEOREM 2)

PROOF OF LEMMA 3. (Restricted Set) Part 1. By condition D.3, with probability $1 - \gamma$, for every $j = 1, \dots, p$ we have $1/2 \leq \hat{\sigma}_j \leq 3/2$, which implies (3.1).

Part 2. Denote the true rankscores by $a_i^*(u) = u - 1\{y_i \leq x_i' \beta(u)\}$ for $i = 1, \dots, n$. Next recall that $\hat{Q}_u(\cdot)$ is a convex function and $\mathbb{E}_n[x_i a_i^*(u)] \in \partial \hat{Q}_u(\beta(u))$. Therefore, we have

$$\hat{Q}_u(\hat{\beta}(u)) \geq \hat{Q}_u(\beta(u)) + \mathbb{E}_n[x_i a_i^*(u)]' (\hat{\beta}(u) - \beta(u)).$$

Let $\widehat{D} = \text{diag}[\widehat{\sigma}_1, \dots, \widehat{\sigma}_p]$ and note that $\lambda\sqrt{u(1-u)}(c_0 - 3)/(c_0 + 3) \geq n\|\widehat{D}^{-1}\mathbb{E}_n[x_i a_i^*(u)]\|_\infty$ with probability at least $1 - \alpha$. By optimality of $\widehat{\beta}(u)$ for the ℓ_1 -penalized problem, we have

$$\begin{aligned} 0 &\leq \widehat{Q}_u(\beta(u)) - \widehat{Q}_u(\widehat{\beta}(u)) + \frac{\lambda\sqrt{u(1-u)}}{n}\|\beta(u)\|_{1,n} - \frac{\lambda\sqrt{u(1-u)}}{n}\|\widehat{\beta}(u)\|_{1,n} \\ &\leq \left| \mathbb{E}_n[x_i a_i^*(u)]' (\widehat{\beta}(u) - \beta(u)) \right| + \frac{\lambda\sqrt{u(1-u)}}{n} \left(\|\beta(u)\|_{1,n} - \|\widehat{\beta}(u)\|_{1,n} \right) \\ &= \left\| \widehat{D}^{-1}\mathbb{E}_n[x_i a_i^*(u)] \right\|_\infty \left\| \widehat{D}(\widehat{\beta}(u) - \beta(u)) \right\|_1 + \frac{\lambda\sqrt{u(1-u)}}{n} \left(\|\beta(u)\|_{1,n} - \|\widehat{\beta}(u)\|_{1,n} \right) \\ &\leq \frac{\lambda\sqrt{u(1-u)}}{n} \sum_{j=1}^p \left(\frac{c_0-3}{c_0+3} \widehat{\sigma}_j \left| \widehat{\beta}_j(u) - \beta_j(u) \right| + \widehat{\sigma}_j |\beta_j(u)| - \widehat{\sigma}_j |\widehat{\beta}_j(u)| \right), \end{aligned}$$

with probability at least $1 - \alpha$. After canceling $\lambda\sqrt{u(1-u)}/n$ we obtain

$$(B.1) \quad \left(1 - \frac{c_0 - 3}{c_0 + 3} \right) \|\widehat{\beta}(u) - \beta(u)\|_{1,n} \leq \sum_{j=1}^p \widehat{\sigma}_j \left(\left| \widehat{\beta}_j(u) - \beta_j(u) \right| + |\beta_j(u)| - |\widehat{\beta}_j(u)| \right).$$

Furthermore, since $\left| \widehat{\beta}_j(u) - \beta_j(u) \right| + |\beta_j(u)| - |\widehat{\beta}_j(u)| = 0$ if $\beta_j(u) = 0$, i.e. $j \in T_u^c$,

$$(B.2) \quad \sum_{j=1}^p \widehat{\sigma}_j \left(\left| \widehat{\beta}_j(u) - \beta_j(u) \right| + |\beta_j(u)| - |\widehat{\beta}_j(u)| \right) \leq 2\|\widehat{\beta}_{T_u}(u) - \beta(u)\|_{1,n}.$$

(B.1) and (B.2) establish that $\|\widehat{\beta}_{T_u^c}(u)\|_{1,n} \leq (c_0/3)\|\widehat{\beta}_{T_u}(u) - \beta(u)\|_{1,n}$ with probability at least $1 - \alpha$. In turn, by Part 1 of this Lemma, $\|\widehat{\beta}_{T_u^c}(u)\|_{1,n} \geq (1/2)\|\widehat{\beta}_{T_u^c}(u)\|_1$ and $\|\widehat{\beta}_{T_u}(u) - \beta(u)\|_{1,n} \leq (3/2)\|\widehat{\beta}_{T_u}(u) - \beta(u)\|_1$, which holds with probability at least $1 - \gamma$. Intersection of these two event holds with probability at least $1 - \alpha - \gamma$. Finally, by Lemma 9, $\|\widehat{\beta}(u)\|_0 \leq n$ with probability 1 uniformly in $u \in \mathcal{U}$. \square

PROOF OF LEMMA 4. (Identification in Population) Part 1. Proof of claims (3.3)-(3.5). By $\text{RE}(c_0, m)$ and by $\delta \in A_u$

$$\|J_u^{1/2}\delta\| \geq \|(\mathbb{E}[x_i x_i'])^{1/2}\delta\| \underline{f}^{1/2} \geq \|\delta_{T_u}\| \underline{f}^{1/2} \kappa_0 \geq \frac{\underline{f}^{1/2} \kappa_0}{\sqrt{s}} \|\delta_{T_u}\|_1 \geq \frac{\underline{f}^{1/2} \kappa_0}{\sqrt{s}(1+c_0)} \|\delta\|_1.$$

Part 2. Proof of claim (3.6). Proceeding similarly to [7], we note that the k th largest in absolute value component of $\delta_{T_u^c}$ is less than $\|\delta_{T_u^c}\|_1/k$. Therefore by $\delta \in A_u$ and $|T_u| \leq s$

$$\|\delta_{(T_u \cup \overline{T}_u(\delta, m))^c}\|^2 \leq \sum_{k \geq m+1} \frac{\|\delta_{T_u^c}\|_1^2}{k^2} \leq \frac{\|\delta_{T_u^c}\|_1^2}{m} \leq c_0^2 \frac{\|\delta_{T_u}\|_1^2}{m} \leq c_0^2 \|\delta_{T_u}\|^2 \frac{s}{m} \leq c_0^2 \|\delta_{T_u \cup \overline{T}_u(\delta, m)}\|^2 \frac{s}{m},$$

so that $\|\delta\| \leq \left(1 + c_0\sqrt{s/m}\right) \|\delta_{T_u \cup \overline{T}_u(\delta, m)}\|$; and the last term is bounded by $\text{RE}(c_0, m)$,

$$\left(1 + c_0\sqrt{s/m}\right) \|(\mathbb{E}[x_i x_i'])^{1/2}\delta\|/\kappa_m \leq \left(1 + c_0\sqrt{s/m}\right) \|J_u^{1/2}\delta\|/[\underline{f}^{1/2}\kappa_m].$$

Part 3. The proof of claim (3.7) proceeds in two steps. Step 1. (Minoration). Define the maximal radius over which the criterion function can be minored by a quadratic function

$$r_{A_u} = \sup_r \left\{ r : Q_u(\beta(u) + \tilde{\delta}) - Q_u(\beta(u)) \geq \frac{1}{4} \|J_u^{1/2} \tilde{\delta}\|^2, \text{ for all } \tilde{\delta} \in A_u, \|J_u^{1/2} \tilde{\delta}\| \leq r \right\}.$$

Step 2 below shows that $r_{A_u} \geq 4q$. By construction of r_{A_u} and the convexity of Q_u ,

$$\begin{aligned} & Q_u(\beta(u) + \delta) - Q_u(\beta(u)) \\ & \geq \frac{\|J_u^{1/2} \delta\|^2}{4} \wedge \left\{ \frac{\|J_u^{1/2} \delta\|}{r_{A_u}} \cdot \inf_{\tilde{\delta} \in A_u, \|J_u^{1/2} \tilde{\delta}\| \geq r_{A_u}} Q_u(\beta(u) + \tilde{\delta}) - Q_u(\beta(u)) \right\} \\ & \geq \frac{\|J_u^{1/2} \delta\|^2}{4} \wedge \left\{ \frac{\|J_u^{1/2} \delta\|}{r_{A_u}} \frac{r_{A_u}^2}{4} \right\} \geq \frac{\|J_u^{1/2} \delta\|^2}{4} \wedge \left\{ q \|J_u^{1/2} \delta\| \right\}, \text{ for any } \delta \in A_u. \end{aligned}$$

Step 2. ($r_{A_u} \geq 4q$) Let $F_{y|x}$ denote the conditional distribution of y given x . From [20], for any two scalars w and v we have that

$$(B.3) \quad \rho_u(w - v) - \rho_u(w) = -v(u - 1\{w \leq 0\}) + \int_0^v (1\{w \leq z\} - 1\{w \leq 0\}) dz.$$

Using (B.3) with $w = y - x'\beta(u)$ and $v = x'\delta$ we conclude $E[-v(u - 1\{w \leq 0\})] = 0$. Using the law of iterated expectations and mean value expansion, we obtain for $\tilde{z}_{x,z} \in [0, z]$

$$\begin{aligned} (B.4) \quad & Q_u(\beta(u) + \delta) - Q_u(\beta(u)) = E \left[\int_0^{x'\delta} F_{y|x}(x'\beta(u) + z) - F_{y|x}(x'\beta(u)) dz \right] \\ & = E \left[\int_0^{x'\delta} z f_{y|x}(x'\beta(u)) + \frac{z^2}{2} f'_{y|x}(x'\beta(u) + \tilde{z}_{x,z}) dz \right] \\ & \geq \frac{1}{2} \|J_u^{1/2} \delta\|^2 - \frac{1}{6} \bar{f}' E[|x'\delta|^3] \geq \frac{1}{4} \|J_u^{1/2} \delta\|^2 + \frac{1}{4} \underline{f} E[|x'\delta|^2] - \frac{1}{6} \bar{f}' E[|x'\delta|^3]. \end{aligned}$$

Note that for $\delta \in A_u$, if $\|J_u^{1/2} \delta\| \leq 4q \leq (3/2) \cdot (\underline{f}^{3/2}/\bar{f}') \cdot \inf_{\delta \in A_u, \delta \neq 0} E[|x'\delta|^2]^{3/2} / E[|x'\delta|^3]$, it follows that $(1/6) \bar{f}' E[|x'\delta|^3] \leq (1/4) \underline{f} E[|x'\delta|^2]$. This and (B.4) imply $r_{A_u} \geq 4q$. \square

PROOF OF LEMMA 5. (Control of Empirical Error) We divide the proof in four steps.

Step 1. (Main Argument) Let

$$\mathcal{A}(t) := \epsilon(t) \sqrt{n} = \sup_{u \in \mathcal{U}, \|J_u^{1/2} \delta\| \leq t, \delta \in A_u} |\mathbb{G}_n[\rho_u(y_i - x'_i(\beta(u) + \delta)) - \rho_u(y_i - x'_i\beta(u))]|$$

Let Ω_1 be the event in which $\max_{1 \leq j \leq p} |\hat{\sigma}_j - 1| \leq 1/2$, where $P(\Omega_1) \geq 1 - \gamma$.

In order to apply the symmetrization lemma, Lemma 2.3.7 in [40], to bound the tail probability of $\mathcal{A}(t)$ first note that for any fixed $\delta \in A_u$, $u \in \mathcal{U}$ we have

$$\text{var}(\mathbb{G}_n[\rho_u(y_i - x'_i(\beta(u) + \delta)) - \rho_u(y_i - x'_i\beta(u))]) \leq E[(x'_i\delta)^2] \leq t^2/\underline{f}$$

Then application of the symmetrization lemma for probabilities, Lemma 2.3.7 in [40], yields

$$(B.5) \quad P(\mathcal{A}(t) \geq M) \leq \frac{2P(\mathcal{A}^o(t) \geq M/4)}{1 - t^2/(\underline{f}M^2)} \leq \frac{2P(\mathcal{A}^o(t) \geq M/4|\Omega_1) + 2P(\Omega_1^c)}{1 - t^2/(\underline{f}M^2)},$$

where $\mathcal{A}^o(t)$ is the symmetrized version of $\mathcal{A}(t)$, constructed by replacing the empirical process \mathbb{G}_n with its symmetrized version \mathbb{G}_n^o , and $P(\Omega_1^c) \leq \gamma$. We set $M > M_1 := t(3/\underline{f})^{1/2}$, which makes the denominator on right side of (B.5) greater than $2/3$. Further, Step 3 below shows that $P(\mathcal{A}^o(t) \geq M/4|\Omega_1) \leq p^{-A^2}$ for

$$M/4 \geq M_2 := t \cdot A \cdot 18\sqrt{2} \cdot \Gamma \cdot \sqrt{2 \log p + \log(2 + 4\sqrt{2}L\underline{f}^{1/2}\kappa_0/t)}, \quad \Gamma = \sqrt{s}(1 + c_0)/[\underline{f}^{1/2}\kappa_0].$$

We conclude that with probability at least $1 - 3\gamma - 3p^{-A^2}$, $\mathcal{A}(t) \leq M_1 \vee (4M_2)$.

Therefore, there is a universal constant C_E such that with probability at least $1 - 3\gamma - 3p^{-A^2}$,

$$\mathcal{A}(t) \leq t \cdot C_E \cdot \frac{(1 + c_0)A}{\underline{f}^{1/2}\kappa_0} \sqrt{s \log(p \vee [L\underline{f}^{1/2}\kappa_0/t])}$$

and the result follows.

Step 2. (Bound on $P(\mathcal{A}^o(t) \geq K|\Omega_1)$). We begin by noting that Lemma 3 and 4 imply that $\|\delta\|_{1,n} \leq \frac{3}{2}\sqrt{s}(1 + c_0)\|J_u^{1/2}\delta\|/[\underline{f}^{1/2}\kappa_0]$ so that for all $u \in \mathcal{U}$

$$(B.6) \quad \{\delta \in A_u : \|J_u^{1/2}\delta\| \leq t\} \subseteq \{\delta \in \mathbb{R}^p : \|\delta\|_{1,n} \leq 2t\Gamma\}, \quad \Gamma := \sqrt{s}(1 + c_0)/[\underline{f}^{1/2}\kappa_0].$$

Further, we let $\mathcal{U}_k = \{\hat{u}_1, \dots, \hat{u}_k\}$ be an ε -net of quantile indices in \mathcal{U} with

$$(B.7) \quad \varepsilon \leq t\Gamma/(2\sqrt{2}sL) \text{ and } k \leq 1/\varepsilon.$$

By $\rho_u(y_i - x'_i(\beta(u) + \delta)) - \rho_u(y_i - x'_i\beta(u)) = ux'_i\delta + w_i(x'_i\delta, u)$, for $w_i(b, u) := (y_i - x'_i\beta(u) - b)_- - (y_i - x'_i\beta(u))_-$, and by (B.6) we have that $\mathcal{A}^o(t) \leq \mathcal{B}^o(t) + \mathcal{C}^o(t)$, where

$$\mathcal{B}^o(t) := \sup_{u \in \mathcal{U}, \|\delta\|_{1,n} \leq 2t\Gamma} |\mathbb{G}_n^o[x'_i\delta]| \text{ and } \mathcal{C}^o(t) := \sup_{u \in \mathcal{U}, \|\delta\|_{1,n} \leq 2t\Gamma} |\mathbb{G}_n^o[w_i(\delta, u)]|.$$

Then we compute the bounds

$$\begin{aligned} P[\mathcal{B}^o(t) > K|\Omega_1] &\leq \min_{\lambda \geq 0} e^{-\lambda K} \mathbb{E}[e^{\lambda \mathcal{B}^o(t)}|\Omega_1] \text{ by Markov} \\ &\leq \min_{\lambda \geq 0} e^{-\lambda K} 2p \exp((2\lambda t\Gamma)^2/2) \text{ by Step 3} \\ &\leq 2p \exp(-K^2/(2\sqrt{2}t\Gamma)^2) \text{ by setting } \lambda = K/(2t\Gamma)^2, \\ P[\mathcal{C}^o(t) > K|\Omega_1] &\leq \min_{\lambda \geq 0} e^{-\lambda K} \mathbb{E}[e^{\lambda \mathcal{C}^o(t)}|\Omega_1, X] \text{ by Markov} \\ &\leq \min_{\lambda \geq 0} \exp(-\lambda K) 2(p/\varepsilon) \exp((16\lambda t\Gamma)^2/2) \text{ by Step 4} \\ &\leq \varepsilon^{-1} 2p \exp(-K^2/(16\sqrt{2}t\Gamma)^2) \text{ by setting } \lambda = K/(16t\Gamma)^2, \end{aligned}$$

so that

$$\begin{aligned} P[\mathcal{A}^o(t) > 2\sqrt{2}K + 16\sqrt{2}K|\Omega_1] &\leq P[\mathcal{B}^o(t) > 2\sqrt{2}K|\Omega_1] + P[\mathcal{C}^o(t) > 16\sqrt{2}K|\Omega_1] \\ &\leq 2p(1 + \varepsilon^{-1}) \exp(-K^2/(t\Gamma)^2). \end{aligned}$$

Setting $K = A \cdot t \cdot \Gamma \cdot \sqrt{\log\{2p^2(1 + \varepsilon^{-1})\}}$, for $A \geq 1$, we get $P[\mathcal{A}^o(t) \geq 18\sqrt{2}K|\Omega_1] \leq p^{-A^2}$.

Step 3. (Bound on $E[e^{\lambda\mathcal{B}^o(t)}|\Omega_1]$) We bound

$$\begin{aligned} E[e^{\lambda\mathcal{B}^o(t)}|\Omega_1] &\leq E[\exp(2\lambda t\Gamma \max_{j \leq p} |\mathbb{G}_n^o(x_{ij})/\hat{\sigma}_j|)|\Omega_1] \\ &\leq 2p \max_{j \leq p} E[\exp(2\lambda t\Gamma \mathbb{G}_n^o(x_{ij})/\hat{\sigma}_j)|\Omega_1] \leq 2p \exp((2\lambda t\Gamma)^2/2), \end{aligned}$$

where the first inequality follows from $|\mathbb{G}_n^o[x'_i\delta]| \leq 2\|\delta\|_{1,n} \max_{1 \leq j \leq p} |\mathbb{G}_n^o(x_{ij})/\hat{\sigma}_j|$ holding under event Ω_1 , the penultimate inequality follows from the simple bound

$$E[\max_{j \leq p} e^{z_j}] \leq p \max_{j \leq p} E[e^{z_j}] \leq p \max_{j \leq p} E[e^{z_j} + e^{-z_j}] \leq 2p \max_{j \leq p} E[e^{z_j}]$$

holding for symmetric random variables z_j , and the last inequality follows from the law of iterated expectations and from $E[\exp(2\lambda t\Gamma \mathbb{G}_n^o(x_{ij})/\hat{\sigma}_j)|\Omega_1, X] \leq \exp((2\lambda t\Gamma)^2/2)$ holding by the Hoeffding inequality (more precisely, by the intermediate step in the proof of the Hoeffding inequality, see, e.g., p. 100 in [40]). Here $E[\cdot|\Omega_1, X]$ denotes the expectation over the symmetrizing Rademacher variables entering the definition of the symmetrized process \mathbb{G}_n^o .

Step 4. (Bound on $E[e^{\lambda\mathcal{C}^o(t)}|\Omega_1]$) We bound

$$\begin{aligned} \mathcal{C}^o(t) &\leq \sup_{u \in \mathcal{U}, |u - \hat{u}| \leq \varepsilon, \hat{u} \in \mathcal{U}_k} \sup_{\|\delta\|_{1,n} \leq 2t\Gamma} |\mathbb{G}_n^o[w_i(x'_i(\delta + \beta(u) - \beta(\hat{u})), \hat{u})]| \\ &\quad + \sup_{u \in \mathcal{U}, |u - \hat{u}| \leq \varepsilon, \hat{u} \in \mathcal{U}_k} |\mathbb{G}_n[w_i(x'_i(\beta(u) - \beta(\hat{u})), \hat{u})]| \\ &\leq 2 \sup_{\hat{u} \in \mathcal{U}_k, \|\delta\|_{1,n} \leq 4t\Gamma} |\mathbb{G}_n^o[w_i(x'_i\delta, \hat{u})]| =: \mathcal{D}^o(t), \end{aligned}$$

where the first inequality is elementary, and the second inequality follows from the inequality

$$\sup_{|u - \hat{u}| \leq \varepsilon} \|\beta(u) - \beta(\hat{u})\|_{1,n} \leq \sqrt{2s}L(2 \max_{1 \leq j \leq p} \sigma_j)\varepsilon \leq \sqrt{2s}L(2 \cdot 3/2)\varepsilon \leq 2t\Gamma,$$

holding by our choice (B.7) of ε and by event Ω_1 .

Next we bound $E[e^{\lambda\mathcal{D}^o(t)}|\Omega_1]$

$$\begin{aligned} E[e^{\lambda\mathcal{D}^o(t)}|\Omega_1] &\leq (1/\varepsilon) \max_{\hat{u} \in \mathcal{U}_k} E[\exp(2\lambda \sup_{\|\delta\|_{1,n} \leq 4t\Gamma} |\mathbb{G}_n^o[w_i(x'_i\delta, \hat{u})])|\Omega_1] \\ &\leq (1/\varepsilon) \max_{\hat{u} \in \mathcal{U}_k} E[\exp(4\lambda \sup_{\|\delta\|_{1,n} \leq 4t\Gamma} |\mathbb{G}_n^o[x'_i\delta])|\Omega_1] \\ &\leq 2(p/\varepsilon) \max_{j \leq p} E[\exp(16\lambda t\Gamma \mathbb{G}_n^o(x_{ij})/\hat{\sigma}_j)|\Omega_1] \leq 2(p/\varepsilon) \exp((16\lambda t\Gamma)^2/2), \end{aligned}$$

where the first inequality follows from the definition of w_i and by $k \leq 1/\varepsilon$, the second inequality follows from the exponential moment inequality for contractions (Theorem 4.12 of Ledoux and Talagrand [28]) and from the contractive property $|w_i(a, \hat{u}) - w_i(b, \hat{u})| \leq |a - b|$, and the last two inequalities follow exactly as in Step 3. \square

APPENDIX C: PROOF OF LEMMAS 6-7 (USED IN THEOREM 3)

In order to characterize the sparsity properties of $\hat{\beta}(u)$, we will exploit the fact that (2.4) can be written as the following linear programming problem:

$$(C.1) \quad \min_{\xi^+, \xi^-, \beta^+, \beta^- \in \mathbb{R}_+^{2n+2p}} \quad \mathbb{E}_n [u\xi_i^+ + (1-u)\xi_i^-] + \frac{\lambda\sqrt{u(1-u)}}{n} \sum_{j=1}^p \hat{\sigma}_j(\beta_j^+ + \beta_j^-)$$

$$\xi_i^+ - \xi_i^- = y_i - x_i'(\beta^+ - \beta^-), \quad i = 1, \dots, n.$$

Our theoretical analysis of the sparsity of $\hat{\beta}(u)$ relies on the dual of (C.1):

$$(C.2) \quad \max_{a \in \mathbb{R}^n} \quad \mathbb{E}_n [y_i a_i]$$

$$|\mathbb{E}_n [x_{ij} a_i]| \leq \lambda\sqrt{u(1-u)}\hat{\sigma}_j/n, \quad j = 1, \dots, p,$$

$$(u-1) \leq a_i \leq u, \quad i = 1, \dots, n.$$

The dual program maximizes the correlation between the response variable and the rank scores subject to the condition requiring the rank scores to be approximately uncorrelated with the regressors. The optimal solution $\hat{a}(u)$ to (C.2) plays a key role in determining the sparsity of $\hat{\beta}(u)$.

LEMMA 9 (Signs and Interpolation Property). (1) For any $j \in \{1, \dots, p\}$

$$(C.3) \quad \begin{aligned} \hat{\beta}_j(u) > 0 & \quad \text{iff} \quad \mathbb{E}_n [x_{ij}\hat{a}_i(u)] = \lambda\sqrt{u(1-u)}\hat{\sigma}_j/n, \\ \hat{\beta}_j(u) < 0 & \quad \text{iff} \quad \mathbb{E}_n [x_{ij}\hat{a}_i(u)] = -\lambda\sqrt{u(1-u)}\hat{\sigma}_j/n, \end{aligned}$$

(2) $\|\hat{\beta}(u)\|_0 \leq n \wedge p$ uniformly over $u \in \mathcal{U}$. (3) If y_1, \dots, y_n are absolutely continuous conditional on x_1, \dots, x_n , then the number of interpolated data points, $I_u = |\{i : y_i = x_i'\hat{\beta}(u)\}|$, is equal to $\|\hat{\beta}(u)\|_0$ with probability one uniformly over $u \in \mathcal{U}$.

PROOF OF LEMMA 9. Step 1. Part (1) follows from the complementary slackness condition for linear programming problems, see Theorem 4.5 of [6].

Step 2. To show part (2) consider any $u \in \mathcal{U}$. Trivially we have $\|\hat{\beta}(u)\|_0 \leq p$. Let $Y = (y_1, \dots, y_n)'$, $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_p)'$, X be the $n \times p$ matrix with rows x_i' , $i = 1, \dots, n$, $c_u = (ue', (1-u)e', \lambda\sqrt{u(1-u)}\hat{\sigma}', \lambda\sqrt{u(1-u)}\hat{\sigma}')'$, and $A = [I \quad -I \quad X \quad -X]$, where $e = (1, 1, \dots, 1)'$ denotes

an n -vectors of ones, and I denotes the $n \times n$ identity matrix. For $w = (\xi^+, \xi^-, \beta^+, \beta^-)$, the primal problem (C.1) can be written as $\min_w \{c'_u w : Aw = Y, w \geq 0\}$. Matrix A has rank n , since it has linearly independent rows. By Theorem 2.4 of [6] there is at least one optimal basic solution $\hat{w}(u) = (\hat{\xi}^+(u), \hat{\xi}^-(u), \hat{\beta}^+(u), \hat{\beta}^-(u))$, and all basic solutions have at most n non-zero components. Since $\hat{\beta}(u) = \hat{\beta}^+(u) - \hat{\beta}^-(u)$, $\hat{\beta}(u)$ has at most n non-zero components.

Let I_u denote the number of interpolated points in (2.4) at the quantile index u . We have that $n - I_u$ components of $\hat{\xi}^+(u)$ and $\hat{\xi}^-(u)$ are non-zero. Therefore, $\|\hat{\beta}(u)\|_0 + (n - I_u) \leq n$, which leads to $\|\hat{\beta}(u)\|_0 \leq I_u$. By step 3 below this holds with equality with probability 1 uniformly over $u \in \mathcal{U}$, thus establishing part (3).

Step 3. Consider the dual problem $\max_a \{Y'a : A'a \leq c_u\}$ for all $u \in \mathcal{U}$. Conditional on X the feasible region of this problem is the polytope $R_u = \{a : A'a \leq c_u\}$. Since $c_u > 0$, R_u is non-empty for all $u \in \mathcal{U}$. Moreover, the form of A' implies that $R_u \subset [-1, 1]^n$ so R_u is bounded. Therefore, if the solution of the dual is not unique for some $u \in \mathcal{U}$ there exist vertices a^1, a^2 connected by an edge of R_u such that $Y'(a^1 - a^2) = 0$. Note that the matrix A' is the same for all $u \in \mathcal{U}$ so that the direction $\frac{a^1 - a^2}{\|a^1 - a^2\|}$ of the edge linking a^1 and a^2 is generated by a finite number of intersections of hyperplanes associated with the rows of A' . Thus, the event $Y'(a^1 - a^2) = 0$ is a zero probability event uniformly in $u \in \mathcal{U}$ since Y is absolutely continuous conditional on X and the number of different edge directions is finite. Therefore the dual problem has a unique solution with probability one uniformly in $u \in \mathcal{U}$. If the dual basic solution is unique, we have that the primal basic solution is non-degenerate, that is, the number of non-zero variables equals n , see [6]. Therefore, with probability one $\|\hat{\beta}(u)\|_0 + (n - I_u) = n$, or $\|\hat{\beta}(u)\|_0 = I_u$ for all $u \in \mathcal{U}$. \square

PROOF OF LEMMA 6. (Empirical Pre-Sparsity) That $\hat{s} \leq n \wedge p$ follows from Lemma 9. We proceed to show the last bound.

Let $\hat{a}(u)$ be the solution of the dual problem (C.2), $\hat{T}_u = \text{support}(\hat{\beta}(u))$, and $\hat{s}_u = \|\hat{\beta}(u)\|_0 = |\hat{T}_u|$. For any $j \in \hat{T}_u$, from (C.3) we have $(X'\hat{a}(u))_j = \text{sign}(\hat{\beta}_j(u))\lambda\hat{\sigma}_j\sqrt{u(1-u)}$ and, for $j \notin \hat{T}_u$ we have $\text{sign}(\hat{\beta}_j(u)) = 0$. Therefore, by the Cauchy-Schwarz inequality, and by D.3, with probability $1 - \gamma$ we have

$$\begin{aligned} \hat{s}_u \lambda &= \text{sign}(\hat{\beta}(u))' \text{sign}(\hat{\beta}(u)) \lambda \leq \text{sign}(\hat{\beta}(u))' (X'\hat{a}(u)) / \min_{j=1, \dots, p} \hat{\sigma}_j \sqrt{u(1-u)} \\ &\leq 2 \|X \text{sign}(\hat{\beta}(u))\| \|\hat{a}(u)\| / \sqrt{u(1-u)} \leq 2 \sqrt{n \phi(\hat{s}_u)} \|\text{sign}(\hat{\beta}(u))\| \|\hat{a}(u)\| / \sqrt{u(1-u)}, \end{aligned}$$

where we used that $\|\text{sign}(\hat{\beta}(u))\|_0 = \hat{s}_u$ and $\min_{1 \leq j \leq p} \hat{\sigma}_j \geq 1/2$ with probability $1 - \gamma$. Since $\|\hat{a}(u)\| \leq \sqrt{n}$, and $\|\text{sign}(\hat{\beta}(u))\| = \sqrt{\hat{s}_u}$ we have $\hat{s}_u \lambda \leq 2n \sqrt{\hat{s}_u \phi(\hat{s}_u)} W_{\mathcal{U}}$. Taking the supremum over $u \in \mathcal{U}$ on both sides yields the first result.

To establish the second result, note that $\hat{s} \leq \bar{m} = \max \{m : m \leq n \wedge p \wedge 4n^2 \phi(m) W_{\mathcal{U}}^2 / \lambda^2\}$.

Suppose that $\bar{m} > m_0 = n/\log(n \vee p)$, so that $\bar{m} = m_0 \ell$ for some $\ell > 1$, since $\bar{m} \leq n$ is finite. By definition, \bar{m} satisfies $\bar{m} \leq 4n^2 \phi(\bar{m}) W_{\mathcal{U}}^2 / \lambda^2$. Insert the lower bound on λ , m_0 , and $\bar{m} = m_0 \ell$ in this inequality, and using Lemma 23 we obtain:

$$\bar{m} = m_0 \ell \leq \frac{4n^2 W_{\mathcal{U}}^2}{8W_{\mathcal{U}}^2 n \log(n \vee p)} \frac{\phi(m_0 \ell)}{\phi(m_0)} \leq \frac{n}{2 \log(n \vee p)} \lceil \ell \rceil < \frac{n}{\log(n \vee p)} \ell = m_0 \ell,$$

which is a contradiction. \square

PROOF OF LEMMA 7. (Empirical Sparsity) It is convenient to define:

1. the true rank scores, $a_i^*(u) = u - 1\{y_i \leq x_i' \beta(u)\}$ for $i = 1, \dots, n$;
2. the estimated rank scores, $a_i(u) = u - 1\{y_i \leq x_i' \hat{\beta}(u)\}$ for $i = 1, \dots, n$;
3. the dual optimal rank scores, $\hat{a}(u)$, that solve the dual program (C.2).

Let \hat{T}_u denote the support of $\hat{\beta}(u)$, and $\hat{s}_u = \|\hat{\beta}(u)\|_0$. Let $\tilde{x}_{i\hat{T}_u} = (x_{ij}/\hat{\sigma}_j, j \in \hat{T}_u)'$, and $\hat{\beta}_{\hat{T}_u}(u) = (\hat{\beta}_j(u), j \in \hat{T}_u)'$. From the complementary slackness characterizations (C.3)

$$(C.4) \quad \sqrt{\hat{s}_u} = \|\text{sign}(\hat{\beta}_{\hat{T}_u}(u))\| = \left\| \frac{n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} \hat{a}_i(u)]}{\lambda \sqrt{u(1-u)}} \right\|.$$

Therefore we can bound the number \hat{s}_u of non-zero components of $\hat{\beta}(u)$ provided we can bound the empirical expectation in (C.4). This is achieved in the next step by combining the maximal inequalities and assumptions on the design matrix.

Using the triangle inequality in (C.4), write

$$\lambda \sqrt{\hat{s}} \leq \sup_{u \in \mathcal{U}} \left\{ \frac{\left\| n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} (\hat{a}_i(u) - a_i(u))] \right\| + \left\| n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} (a_i(u) - a_i^*(u))] \right\| + \left\| n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} a_i^*(u)] \right\|}{\sqrt{u(1-u)}} \right\}.$$

This leads to the inequality

$$\begin{aligned} \lambda \sqrt{\hat{s}} &\leq \frac{W_{\mathcal{U}}}{\min_{j=1, \dots, p} \hat{\sigma}_j} \left(\sup_{u \in \mathcal{U}} \left\| n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} (\hat{a}_i(u) - a_i(u))] \right\| + \sup_{u \in \mathcal{U}} \left\| n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} (a_i(u) - a_i^*(u))] \right\| \right) \\ &\quad + \sup_{u \in \mathcal{U}} \left\| n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} a_i^*(u) / \sqrt{u(1-u)}] \right\|. \end{aligned}$$

Then we bound each of the three components in this display.

(a) To bound the first term, we observe that $\hat{a}_i(u) \neq a_i(u)$ only if $y_i = x_i' \hat{\beta}(u)$. By Lemma 9 the penalized quantile regression fit can interpolate at most $\hat{s}_u \leq \hat{s}$ points with probability one uniformly over $u \in \mathcal{U}$. This implies that $\mathbb{E}_n [|\hat{a}_i(u) - a_i(u)|^2] \leq \hat{s}/n$. Therefore,

$$\begin{aligned} \sup_{u \in \mathcal{U}} \left\| n \mathbb{E}_n [\tilde{x}_{i\hat{T}_u} (\hat{a}_i(u) - a_i(u))] \right\| &\leq n \sup_{\|\alpha\|_0 \leq \hat{s}, \|\alpha\| \leq 1} \sup_{u \in \mathcal{U}} \mathbb{E}_n [|\alpha' x_i| |\hat{a}_i(u) - a_i(u)|] \\ &\leq n \sup_{\|\alpha\|_0 \leq \hat{s}, \|\alpha\| \leq 1} \sqrt{\mathbb{E}_n [|\alpha' x_i|^2]} \sup_{u \in \mathcal{U}} \sqrt{\mathbb{E}_n [|\hat{a}_i(u) - a_i(u)|^2]} \leq \sqrt{n \phi(\hat{s}) \hat{s}}. \end{aligned}$$

(b) To bound the second term, note that

$$\begin{aligned} & \sup_{u \in \mathcal{U}} \left\| n \mathbb{E}_n \left[x_{i\widehat{T}_u} (a_i(u) - a_i^*(u)) \right] \right\| \\ & \leq \sup_{u \in \mathcal{U}} \left\| \sqrt{n} \mathbb{G}_n \left(x_{i\widehat{T}_u} (a_i(u) - a_i^*(u)) \right) \right\| + \sup_{u \in \mathcal{U}} \left\| n \mathbb{E} \left[x_{i\widehat{T}_u} (a_i(u) - a_i^*(u)) \right] \right\| \\ & \leq \sqrt{n} \epsilon_1(r, \widehat{s}) + \sqrt{n} \epsilon_2(r, \widehat{s}). \end{aligned}$$

where for $\psi_i(\beta, u) = (1\{y_i \leq x'_i \beta\} - u)x_i$,

$$\begin{aligned} (C.5) \quad \epsilon_1(r, m) &:= \sup_{u \in \mathcal{U}, \beta \in R_u(r, m), \alpha \in \mathbb{S}(\beta)} |\mathbb{G}_n(\alpha' \psi_i(\beta, u)) - \mathbb{G}_n(\alpha' \psi_i(\beta(u), u))|, \\ \epsilon_2(r, m) &:= \sup_{u \in \mathcal{U}, \beta \in R_u(r, m), \alpha \in \mathbb{S}(\beta)} \sqrt{n} |\mathbb{E}[\alpha' \psi_i(\beta, u)] - \mathbb{E}[\alpha' \psi_i(\beta(u), u)]|, \text{ and} \end{aligned}$$

$$\begin{aligned} (C.6) \quad R_u(r, m) &:= \{ \beta \in \mathbb{R}^p : \beta - \beta(u) \in A_u : \|\beta\|_0 \leq m, \|J_u^{1/2}(\beta - \beta(u))\| \leq r \}, \\ \mathbb{S}(\beta) &:= \{ \alpha \in \mathbb{R}^p : \|\alpha\| \leq 1, \text{support}(\alpha) \subseteq \text{support}(\beta) \}. \end{aligned}$$

By Lemma 12 there is a constant $A_{\varepsilon/2}^1$ such that $\sqrt{n} \epsilon_1(r, \widehat{s}) \leq A_{\varepsilon/2}^1 \sqrt{n \widehat{s} \log(n \vee p)} \sqrt{\phi(\widehat{s})}$ with probability $1 - \varepsilon/2$. By Lemma 10 we have $\sqrt{n} \epsilon_2(r, \widehat{s}) \leq n(\mu(\widehat{s})/2)(r \wedge 1)$.

(c) To bound the last term, by Theorem 1 there exists a constant $A_{\varepsilon/2}^0$ such that with probability $1 - \varepsilon/2$

$$\sup_{u \in \mathcal{U}} \left\| n \mathbb{E}_n \left[\tilde{x}_{i\widehat{T}_u} a_i^*(u) / \sqrt{u(1-u)} \right] \right\| \leq \sqrt{\widehat{s}} \Lambda \leq \sqrt{\widehat{s}} A_{\varepsilon/2}^0 W_{\mathcal{U}} \sqrt{n \log p},$$

where we used that $a_i^*(u) = u - 1\{u_i \leq u\}$, $i = 1, \dots, n$, for u_1, \dots, u_n i.i.d. uniform $(0, 1)$.

Combining bounds in (a)-(c), using that $\min_{j=1, \dots, p} \widehat{\sigma}_j \geq 1/2$ by condition D.3 with probability $1 - \gamma$, we have

$$\frac{\sqrt{\widehat{s}}}{W_{\mathcal{U}}} \leq \mu(\widehat{s}) \frac{n}{\lambda} (r \wedge 1) + \sqrt{\widehat{s}} K_{\varepsilon} \frac{\sqrt{n \log(n \vee p) \phi(\widehat{s})}}{\lambda},$$

with probability at least $1 - \varepsilon - \gamma$, for $K_{\varepsilon} = 2(1 + A_{\varepsilon/2}^0 + A_{\varepsilon/2}^1)$. □

Next we control the linearization error ϵ_2 defined in (C.5).

LEMMA 10 (Controlling linearization error ϵ_2). *Under D.1-2*

$$\epsilon_2(r, m) \leq \sqrt{n} \sqrt{\varphi_{\max}(m)} \left\{ 1 \wedge \left(2[\bar{f}/\underline{f}^{1/2}]r \right) \right\} \text{ for all } r > 0 \text{ and } m \leq n.$$

PROOF. By definition

$$\epsilon_2(r, m) = \sup_{u \in \mathcal{U}, \beta \in R_u(r, m), \alpha \in \mathbb{S}(\beta)} \sqrt{n} |\mathbb{E}[(\alpha' x_i) (1\{y_i \leq x'_i \beta\} - 1\{y_i < x'_i \beta(u)\})]|.$$

By Cauchy-Schwarz, and using that $\varphi_{\max}(m) = \sup_{\|\alpha\| \leq 1, \|\alpha\|_0 \leq m} \mathbb{E}[|\alpha' x_i|^2]$

$$\epsilon_2(r, m) \leq \sqrt{n} \sqrt{\varphi_{\max}(m)} \sup_{u \in \mathcal{U}, \beta \in R_u(r, m)} \sqrt{\mathbb{E}[(1\{y_i \leq x'_i \beta\} - 1\{y_i < x'_i \beta(u)\})^2]}.$$

Then, since for any $\beta \in R_u(r, m)$, $u \in \mathcal{U}$,

$$\begin{aligned} \mathbb{E}[(1\{y_i \leq x'_i \beta\} - 1\{y_i < x'_i \beta(u)\})^2] &\leq \mathbb{E}[1\{|y_i - x'_i \beta(u)| \leq |x'_i(\beta - \beta(u))|\}] \\ &\leq \mathbb{E}[(2\bar{f}|x'_i(\beta - \beta(u))|) \wedge 1] \leq \left\{ 2\bar{f} (\mathbb{E}[|x'_i(\beta - \beta(u))|^2])^{1/2} \right\} \wedge 1 \end{aligned}$$

and $(\mathbb{E}[|x'_i(\beta - \beta(u))|^2])^{1/2} \leq \|J_u^{1/2}(\beta - \beta(u))\|/\underline{f}^{1/2}$ by Lemma 4, the result follows. \square

Next we proceed to control the empirical error ϵ_1 defined in (C.5). We shall need the following preliminary result on the uniform L_2 covering numbers ([40]) of a relevant function class.

LEMMA 11. (1) Consider a fixed subset $T \subset \{1, 2, \dots, p\}$, $|T| = m$. The class of functions

$$\mathcal{F}_T = \{\alpha'(\psi_i(\beta, u) - \psi_i(\beta(u), u)) : u \in \mathcal{U}, \alpha \in \mathbb{S}(\beta), \text{support}(\beta) \subseteq T\}$$

has a VC index bounded by cm for some universal constant c . (2) There are universal constants C and c such that for any $m \leq n$ the function class

$$\mathcal{F}_m = \{\alpha'(\psi_i(\beta, u) - \psi_i(\beta(u), u)) : u \in \mathcal{U}, \beta \in \mathbb{R}^p, \|\beta\|_0 \leq m, \alpha \in \mathbb{S}(\beta)\}$$

has the the uniform covering numbers bounded as

$$\sup_Q N(\epsilon \|F_m\|_{Q,2}, \mathcal{F}_m, L_2(Q)) \leq C \left(\frac{16e}{\epsilon} \right)^{2(cm-1)} \left(\frac{ep}{m} \right)^m, \quad \epsilon > 0.$$

PROOF. The proof involves standard combinatorial arguments and is relegated to the Supplementary Material Appendix G. \square

LEMMA 12 (Controlling empirical error ϵ_1). Under D.1-2 there exists a universal constant A such that with probability $1 - \delta$

$$\epsilon_1(r, m) \leq A\delta^{-1/2} \sqrt{m \log(n \vee p)} \sqrt{\phi(m)} \quad \text{uniformly for all } r > 0 \text{ and } m \leq n.$$

PROOF. By definition $\epsilon_1(r, m) \leq \sup_{f \in \mathcal{F}_m} |\mathbb{G}_n(f)|$. From Lemma 11 the uniform covering number of \mathcal{F}_m is bounded by $C(16e/\epsilon)^{2(cm-1)} (ep/m)^m$. Using Lemma 19 with $N = n$ and $\theta_m = p$ we have that uniformly in $m \leq n$, with probability at least $1 - \delta$

$$(C.7) \quad \sup_{f \in \mathcal{F}_m} |\mathbb{G}_n(f)| \leq A\delta^{-1/2} \sqrt{m \log(n \vee p)} \max \left\{ \sup_{f \in \mathcal{F}_m} \mathbb{E}[f^2]^{1/2}, \sup_{f \in \mathcal{F}_m} \mathbb{E}_n[f^2]^{1/2} \right\}$$

By $|\alpha'(\psi_i(\beta, u) - \psi_i(\beta(u), u))| \leq |\alpha'x_i|$ and definition of $\phi(m)$

$$(C.8) \quad \mathbb{E}_n[f^2] \leq \mathbb{E}_n[|\alpha'x_i|^2] \leq \phi(m) \quad \text{and} \quad \mathbb{E}[f^2] \leq \mathbb{E}[|\alpha'x_i|^2] \leq \phi(m).$$

Combining (C.8) with (C.7) we obtain the result. \square

(c) The next lemma provides a bound on maximum k -sparse eigenvalues, which we used in some of the derivations presented earlier.

LEMMA 13. *Let M be a semi-definite positive matrix and $\phi_M(k) = \sup\{\alpha'M\alpha : \alpha \in \mathbb{R}^p, \|\alpha\| = 1, \|\alpha\|_0 \leq k\}$. For any integers k and ℓk with $\ell \geq 1$, we have $\phi_M(\ell k) \leq \lceil \ell \rceil \phi_M(k)$.*

PROOF. Let $\bar{\alpha}$ achieve $\phi_M(\ell k)$. Moreover let $\sum_{i=1}^{\lceil \ell \rceil} \alpha_i = \bar{\alpha}$ such that $\sum_{i=1}^{\lceil \ell \rceil} \|\alpha_i\|_0 = \|\bar{\alpha}\|_0$. We can choose α_i 's such that $\|\alpha_i\|_0 \leq k$ since $\lceil \ell \rceil k \geq \ell k$. Since M is positive semi-definite, for any i, j w $\alpha_i'M\alpha_i + \alpha_j'M\alpha_j \geq 2|\alpha_i'M\alpha_j|$. Therefore

$$\begin{aligned} \phi_M(\ell k) &= \bar{\alpha}'M\bar{\alpha} = \sum_{i=1}^{\lceil \ell \rceil} \alpha_i'M\alpha_i + \sum_{i=1}^{\lceil \ell \rceil} \sum_{j \neq i} \alpha_i'M\alpha_j \leq \sum_{i=1}^{\lceil \ell \rceil} \{\alpha_i'M\alpha_i + (\lceil \ell \rceil - 1)\alpha_i'M\alpha_i\} \\ &\leq \lceil \ell \rceil \sum_{i=1}^{\lceil \ell \rceil} \|\alpha_i\|^2 \phi_M(\|\alpha_i\|_0) \leq \lceil \ell \rceil \max_{i=1, \dots, \lceil \ell \rceil} \phi_M(\|\alpha_i\|_0) \leq \lceil \ell \rceil \phi_M(k) \end{aligned}$$

where we used that $\sum_{i=1}^{\lceil \ell \rceil} \|\alpha_i\|^2 = 1$. \square

APPENDIX D: PROOF OF THEOREM 4

PROOF OF THEOREM 4. By assumption $\sup_{u \in \mathcal{U}} \|\hat{\beta}(u) - \beta(u)\|_\infty \leq \sup_{u \in \mathcal{U}} \|\hat{\beta}(u) - \beta(u)\| \leq r^o < \inf_{u \in \mathcal{U}} \min_{j \in T_u} |\beta_j(u)|$, which immediately implies the inclusion event (3.15), since the converse of this event implies $\|\hat{\beta}(u) - \beta(u)\|_\infty \geq \inf_{u \in \mathcal{U}} \min_{j \in T_u} |\beta_j(u)|$.

Consider the hard-thresholded estimator next. To establish the inclusion, we note that $\inf_{u \in \mathcal{U}} \min_{j \in T_u} |\hat{\beta}_j(u)| \geq \inf_{u \in \mathcal{U}} \min_{j \in T_u} \{|\beta_j(u)| - |\beta_j(u) - \hat{\beta}_j(u)|\} > \inf_{u \in \mathcal{U}} \min_{j \in T_u} |\beta_j(u)| - r^o > \gamma$, by assumption on γ . Therefore $\inf_{u \in \mathcal{U}} \min_{j \in T_u} |\hat{\beta}_j(u)| > \gamma$ and $\text{support}(\beta(u)) \subseteq \text{support}(\hat{\beta}(u))$ for all $u \in \mathcal{U}$. To establish the opposite inclusion, consider $e_n = \sup_{u \in \mathcal{U}} \max_{j \notin T_u} |\hat{\beta}_j(u)|$. By definition of r^o , $e_n \leq r^o$ and therefore $e_n < \gamma$ by the assumption on γ . By the hard-threshold rule, all components smaller than γ are excluded from the support of $\bar{\beta}(u)$ which yields $\text{support}(\bar{\beta}(u)) \subseteq \text{support}(\beta(u))$. \square

APPENDIX E: PROOF OF LEMMA 8 (USED IN THEOREM 5)

PROOF OF LEMMA 8. (Sparse Identifiability and Control of Empirical Error) The proof of claim (3.17) of this lemma follows identically the proof of claim (3.7) of Lemma 4, given in Appendix B, after replacing A_u with \tilde{A}_u . Next we bound the empirical error

$$(E.1) \quad \sup_{u \in \mathcal{U}, \delta \in \tilde{A}_u(\tilde{m}), \delta \neq 0} \frac{|\epsilon_u(\delta)|}{\|\delta\|} \leq \sup_{u \in \mathcal{U}, \delta \in \tilde{A}_u(\tilde{m}), \delta \neq 0} \frac{1}{\|\delta\| \sqrt{n}} \left| \int_0^1 \delta' \mathbb{G}_n(\psi_i(\beta(u) + \gamma\delta, u)) d\gamma \right|$$

$$\leq \frac{1}{\sqrt{n}} \epsilon_3(\tilde{m})$$

where $\epsilon_3(\tilde{m}) := \sup_{f \in \tilde{\mathcal{F}}_{\tilde{m}}} |\mathbb{G}_n(f)|$ and the class of functions $\tilde{\mathcal{F}}_{\tilde{m}}$ is defined in Lemma 14. The result follows from the bound on $\epsilon_3(\tilde{m})$ holding uniformly in $\tilde{m} \leq n$ given in Lemma 15. \square

Next we control the empirical error ϵ_3 defined in (E.1) for $\tilde{\mathcal{F}}_{\tilde{m}}$ defined below. We first bound uniform covering numbers of $\tilde{\mathcal{F}}_{\tilde{m}}$.

LEMMA 14. Consider a fixed subset $T \subset \{1, 2, \dots, p\}$, $T_u = \text{support}(\beta(u))$ such that $|T \setminus T_u| \leq \tilde{m}$ and $|T_u| \leq s$ for some $u \in \mathcal{U}$. The class of functions

$$\mathcal{F}_{T,u} = \{ \alpha' x_i (1\{y_i \leq x_i' \beta\} - u) : \alpha \in \mathbb{S}(\beta), \text{support}(\beta) \subseteq T \}$$

has a VC index bounded by $c(\tilde{m} + s) + 2$. The class of functions

$$\tilde{\mathcal{F}}_{\tilde{m}} = \{ \mathcal{F}_{T,u} : u \in \mathcal{U}, T \subset \{1, 2, \dots, p\}, |T \setminus T_u| \leq \tilde{m} \},$$

obeys, for some universal constants C and c and each $\epsilon > 0$,

$$\sup_Q N(\epsilon \|\tilde{F}_{\tilde{m}}\|_{Q,2}, \tilde{\mathcal{F}}_{\tilde{m}}, L_2(Q)) \leq C (32e/\epsilon)^{4(c(\tilde{m}+s)+2)} p^{2\tilde{m}} |\cup_{u \in \mathcal{U}} T_u|^{2s}.$$

PROOF. The proof involves standard combinatorial arguments and is relegated to the Supplementary Material Appendix G. \square

LEMMA 15 (Controlling empirical error ϵ_3). Suppose that D.1 holds and $|\cup_{u \in \mathcal{U}} T_u| \leq n$. There exists a universal constant A such that with probability at least $1 - \delta$,

$$\epsilon_3(\tilde{m}) := \sup_{f \in \tilde{\mathcal{F}}_{\tilde{m}}} |\mathbb{G}_n(f)| \leq A \delta^{-1/2} \sqrt{(\tilde{m} \log(n \vee p) + s \log n) \phi(\tilde{m} + s)} \text{ for all } \tilde{m} \leq n.$$

PROOF. Lemma 14 bounds the uniform covering number of $\tilde{\mathcal{F}}_{\tilde{m}}$. Using Lemma 19 with $N \leq 2n$, $m = \tilde{m} + s$ and $\theta_m = p^{2([m-s]/m)} \cdot n^{2(s/m)} = p^{2(\tilde{m}/[\tilde{m}+s])} \cdot n^{2(s/[\tilde{m}+s])}$, we conclude that uniformly in $0 \leq \tilde{m} \leq n$

$$(E.2) \quad \sup_{f \in \tilde{\mathcal{F}}_{\tilde{m}}} |\mathbb{G}_n(f)| \leq A \delta^{-1/2} \sqrt{(\tilde{m} + s) \log(n \vee \theta_m)} \cdot \max \left\{ \sup_{f \in \tilde{\mathcal{F}}_{\tilde{m}}} \mathbb{E}[f^2]^{1/2}, \sup_{f \in \tilde{\mathcal{F}}_{\tilde{m}}} \mathbb{E}_n[f^2]^{1/2} \right\}$$

$$\leq A' \delta^{-1/2} \sqrt{\tilde{m} \log(n \vee p) + s \log n} \cdot \max \left\{ \sup_{f \in \tilde{\mathcal{F}}_{\tilde{m}}} \mathbb{E}[f^2]^{1/2}, \sup_{f \in \tilde{\mathcal{F}}_{\tilde{m}}} \mathbb{E}_n[f^2]^{1/2} \right\}$$

with probability at least $1 - \delta$. The result follows, since for any $f \in \tilde{\mathcal{F}}_{\tilde{m}}$, the corresponding vector α obeys $\|\alpha\|_0 \leq \tilde{m} + s$, so that $\mathbb{E}_n[f^2] \leq \mathbb{E}_n[|\alpha' x_i|^2] \leq \phi(\tilde{m} + s)$ and $\mathbb{E}[f^2] \leq \mathbb{E}[|\alpha' x_i|^2] \leq \phi(\tilde{m} + s)$ by definition of $\phi(\tilde{m} + s)$. \square

APPENDIX F: MAXIMAL INEQUALITIES FOR A COLLECTION OF EMPIRICAL PROCESSES

The main results here are Lemma 16 and Lemma 19, used in the proofs of Theorem 1 and Theorems 3 and 5, respectively. Lemma 19 gives a maximal inequality that controls the empirical process uniformly over a collection of classes of functions using class-dependent bounds. We need this lemma because the standard maximal inequalities applied to the union of function classes yield a single class-independent bound that is too large for our purposes. We prove Lemma 19 by first stating Lemma 16, giving a bound on tail probabilities of a separable sub-Gaussian process, stated in terms of uniform covering numbers. Here we want to explicitly trace the impact of covering numbers on the tail probability, since these covering numbers grow rapidly under increasing parameter dimension and thus help to tighten the probability bound. Using the symmetrization approach, we then obtain Lemma 18, giving a bound on tail probabilities of a general separable empirical process, also stated in terms of uniform covering numbers. Finally, given a growth rate on the covering numbers, we obtain Lemma 19.

LEMMA 16 (Exponential Inequality for Sub-Gaussian Process). *Consider any linear zero-mean separable process $\{\mathbb{G}(f) : f \in \mathcal{F}\}$, whose index set \mathcal{F} includes zero, is equipped with a $L_2(P)$ norm, and has envelope F . Suppose further that the process is sub-Gaussian, namely for each $g \in \mathcal{F} - \mathcal{F}$: $\mathbb{P}\{|\mathbb{G}(g)| > \eta\} \leq 2 \exp\left(-\frac{1}{2}\eta^2/D^2\|g\|_{P,2}^2\right)$ for any $\eta > 0$, with D a positive constant; and suppose that we have the following upper bound on the $L_2(P)$ covering numbers for \mathcal{F} :*

$$N(\epsilon\|F\|_{P,2}, \mathcal{F}, L_2(P)) \leq n(\epsilon, \mathcal{F}, P) \text{ for each } \epsilon > 0,$$

where $n(\epsilon, \mathcal{F}, P)$ is increasing in $1/\epsilon$, and $\epsilon\sqrt{\log n(\epsilon, \mathcal{F}, P)} \rightarrow 0$ as $1/\epsilon \rightarrow \infty$ and is decreasing in $1/\epsilon$. Then for $K > D$, for some universal constant $c < 30$, $\rho(\mathcal{F}, P) := \sup_{f \in \mathcal{F}} \|f\|_{P,2}/\|F\|_{P,2}$,

$$\mathbb{P}\left\{\frac{\sup_{f \in \mathcal{F}} |\mathbb{G}(f)|}{\|F\|_{P,2} \int_0^{\rho(\mathcal{F}, P)/4} \sqrt{\log n(x, \mathcal{F}, P)} dx} > cK\right\} \leq \int_0^{\rho(\mathcal{F}, P)/2} \epsilon^{-1} n(\epsilon, \mathcal{F}, P)^{-\{(K/D)^2 - 1\}} d\epsilon.$$

The result of Lemma 16 is in spirit of the Talagrand tail inequality for Gaussian processes. Our result is less sharp than Talagrand's result in the Gaussian case (by a log factor), but it applies to more general sub-Gaussian processes.

In order to prove a bound on tail probabilities of a general separable empirical process, we need to go through a symmetrization argument. Since we use a data-dependent threshold, we

need an appropriate extension of the classical symmetrization lemma to allow for this. Let us call a threshold function $x : \mathbb{R}^n \mapsto \mathbb{R}$ k -sub-exchangeable if, for any $v, w \in \mathbb{R}^n$ and any vectors \tilde{v}, \tilde{w} created by the pairwise exchange of the components in v with components in w , we have that $x(\tilde{v}) \vee x(\tilde{w}) \geq [x(v) \vee x(w)]/k$. Several functions satisfy this property, in particular $x(v) = \|v\|$ with $k = \sqrt{2}$ and constant functions with $k = 1$. The following result generalizes the standard symmetrization lemma for probabilities (Lemma 2.3.7 of [40]) to the case of a random threshold x that is sub-exchangeable.

LEMMA 17 (Symmetrization with Data-dependent Thresholds). *Consider arbitrary independent stochastic processes Z_1, \dots, Z_n and arbitrary functions $\mu_1, \dots, \mu_n : \mathcal{F} \mapsto \mathbb{R}$. Let $x(Z) = x(Z_1, \dots, Z_n)$ be a k -sub-exchangeable random variable and for any $\tau \in (0, 1)$ let q_τ denote the τ quantile of $x(Z)$, $\bar{p}_\tau := P(x(Z) \leq q_\tau) \geq \tau$, and $p_\tau := P(x(Z) < q_\tau) \leq \tau$. Then*

$$P\left(\left\|\sum_{i=1}^n Z_i\right\|_{\mathcal{F}} > x_0 \vee x(Z)\right) \leq \frac{4}{\bar{p}_\tau} P\left(\left\|\sum_{i=1}^n \varepsilon_i (Z_i - \mu_i)\right\|_{\mathcal{F}} > \frac{x_0 \vee x(Z)}{4k}\right) + p_\tau$$

where x_0 is a constant such that $\inf_{f \in \mathcal{F}} P(|\sum_{i=1}^n Z_i(f)| \leq \frac{x_0}{2}) \geq 1 - \frac{\bar{p}_\tau}{2}$.

Note that we can recover the classical symmetrization lemma for fixed thresholds by setting $k = 1$, $\bar{p}_\tau = 1$, and $p_\tau = 0$.

LEMMA 18 (Exponential inequality for separable empirical process). *Consider a separable empirical process $\mathbb{G}_n(f) = n^{-1/2} \sum_{i=1}^n \{f(Z_i) - \mathbb{E}[f(Z_i)]\}$ and the empirical measure \mathbb{P}_n for Z_1, \dots, Z_n , an underlying i.i.d. data sequence. Let $K > 1$ and $\tau \in (0, 1)$ be constants, and $e_n(\mathcal{F}, \mathbb{P}_n) = e_n(\mathcal{F}, Z_1, \dots, Z_n)$ be a k -sub-exchangeable random variable, such that*

$$\|F\|_{\mathbb{P}_n, 2} \int_0^{\rho(\mathcal{F}, \mathbb{P}_n)/4} \sqrt{\log n(\epsilon, \mathcal{F}, \mathbb{P}_n)} d\epsilon \leq e_n(\mathcal{F}, \mathbb{P}_n) \text{ and } \sup_{f \in \mathcal{F}} \text{var}_{\mathbb{P}} f \leq \frac{\tau}{2} (4kcK e_n(\mathcal{F}, \mathbb{P}_n))^2$$

for the same constant $c > 0$ as in Lemma 16, then

$$\mathbb{P}\left\{\sup_{f \in \mathcal{F}} |\mathbb{G}_n(f)| \geq 4kcK e_n(\mathcal{F}, \mathbb{P}_n)\right\} \leq \frac{4}{\tau} \mathbb{E}_{\mathbb{P}}\left(\left[\int_0^{\rho(\mathcal{F}, \mathbb{P}_n)/2} \epsilon^{-1} n(\epsilon, \mathcal{F}, \mathbb{P}_n)^{-\{K^2-1\}} d\epsilon\right] \wedge 1\right) + \tau.$$

Finally, our main result in this section is as follows.

LEMMA 19 (Maximal Inequality for a Collection of Empirical Processes). *Consider a collection of separable empirical processes $\mathbb{G}_n(f) = n^{-1/2} \sum_{i=1}^n \{f(Z_i) - \mathbb{E}[f(Z_i)]\}$, where Z_1, \dots, Z_n is an underlying i.i.d. data sequence, defined over function classes $\mathcal{F}_m, m = 1, \dots, N$ with envelopes $F_m = \sup_{f \in \mathcal{F}_m} |f(x)|, m = 1, \dots, N$, and with upper bounds on the uniform covering*

numbers of \mathcal{F}_m given for all m by

$$n(\epsilon, \mathcal{F}_m, \mathbb{P}_n) = (N \vee n \vee \theta_m)^m (\omega/\epsilon)^{vm}, \quad 0 < \epsilon < 1,$$

with some constants $\omega > 1$, $v > 1$, and $\theta_m \geq \theta_0$. For a constant $C := (1 + \sqrt{2v})/4$ set

$$e_n(\mathcal{F}_m, \mathbb{P}_n) = C \sqrt{m \log(N \vee n \vee \theta_m \vee \omega)} \max \left\{ \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}, 2}, \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n, 2} \right\}.$$

Then, for any $\delta \in (0, 1/6)$, and any constant $K \geq \sqrt{2/\delta}$ we have

$$\sup_{f \in \mathcal{F}_m} |\mathbb{G}_n(f)| \leq 4\sqrt{2}cK e_n(\mathcal{F}_m, \mathbb{P}_n), \quad \text{for all } m \leq N,$$

with probability at least $1 - \delta$, provided that $N \vee n \vee \theta_0 \geq 3$; the constant c is the same as in Lemma 16.

PROOF OF LEMMA 16. The strategy of the proof is similar to the proof of Lemma 19.34 in [38], page 286 given for the expectation of a supremum of a process; here we instead bound tail probabilities and also compute all constants explicitly.

Step 1. There exists a sequence of nested partitions of \mathcal{F} , $\{(\mathcal{F}_{qi}, i = 1, \dots, N_q), q = q_0, q_0 + 1, \dots\}$ where the q -th partition consists of sets of $L_2(P)$ radius at most $\|F\|_{P, 2} 2^{-q}$, and q_0 is the largest positive integer such that $2^{-q_0} \leq \rho(\mathcal{F}, P)/4$ so that $q_0 \geq 2$. The existence of such a partition follows from a standard argument, e.g. [38], page 286.

Let f_{qi} be an arbitrary point of \mathcal{F}_{qi} . Set $\pi_q(f) = f_{qi}$ if $f \in \mathcal{F}_{qi}$. By separability of the process, we can replace \mathcal{F} by $\cup_{q,i} f_{qi}$, since the supremum norm of the process can be computed by taking this set only. In this case, we can decompose $f - \pi_{q_0}(f) = \sum_{q=q_0+1}^{\infty} (\pi_q(f) - \pi_{q-1}(f))$. Hence by linearity $\mathbb{G}(f) - \mathbb{G}(\pi_{q_0}(f)) = \sum_{q=q_0+1}^{\infty} \mathbb{G}(\pi_q(f) - \pi_{q-1}(f))$, so that

$$\begin{aligned} \mathbb{P} \left\{ \sup_{f \in \mathcal{F}} |\mathbb{G}(f)| > \sum_{q=q_0}^{\infty} \eta_q \right\} &\leq \sum_{q=q_0+1}^{\infty} \mathbb{P} \left\{ \max_f |\mathbb{G}(\pi_q(f) - \pi_{q-1}(f))| > \eta_q \right\} \\ &+ \mathbb{P} \left\{ \max_f |\mathbb{G}(\pi_{q_0}(f))| > \eta_{q_0} \right\}, \end{aligned}$$

for constants η_q chosen below.

Step 2. By construction of the partition sets $\|\pi_q(f) - \pi_{q-1}(f)\|_{P, 2} \leq 2\|F\|_{P, 2} 2^{-(q-1)} \leq 4\|F\|_{P, 2} 2^{-q}$, for $q \geq q_0 + 1$. Setting $\eta_q = 8K\|F\|_{P, 2} 2^{-q} \sqrt{\log N_q}$, using sub-Gaussianity, setting $K > D$, using that $2 \log N_q \geq \log N_q N_{q-1} \geq \log n_q$, using that $q \mapsto \log n_q$ is increasing in q ,

and $2^{-q_0} \leq \rho(\mathcal{F}, P)/4$, we obtain

$$\begin{aligned}
& \sum_{q=q_0+1}^{\infty} \mathbb{P} \left\{ \max_f |\mathbb{G}(\pi_q(f) - \pi_{q-1}(f))| > \eta_q \right\} \leq \sum_{q=q_0+1}^{\infty} N_q N_{q-1} 2 \exp(-\eta_q^2 / (4D \|F\|_{P,2} 2^{-q})^2) \\
& \leq \sum_{q=q_0+1}^{\infty} N_q N_{q-1} 2 \exp(-(K/D)^2 2 \log N_q) \leq \sum_{q=q_0+1}^{\infty} 2 \exp(-\{(K/D)^2 - 1\} \log n_q) \\
& \leq \int_{q_0}^{\infty} 2 \exp(-\{(K/D)^2 - 1\} \log n_q) dq = \int_0^{\rho(\mathcal{F}, P)/4} (x \ln 2)^{-1} 2n(x, \mathcal{F}, P)^{-\{(K/D)^2 - 1\}} dx.
\end{aligned}$$

By Jensen's inequality we have $\sqrt{\log N_q} \leq a_q := \sum_{j=q_0}^q \sqrt{\log n_j}$, so that we obtain $\sum_{q=q_0+1}^{\infty} \eta_q \leq 8 \sum_{q=q_0+1}^{\infty} K \|F\|_{P,2} 2^{-q} a_q$. Letting $b_q = 2 \cdot 2^{-q}$, noting $a_{q+1} - a_q = \sqrt{\log n_{q+1}}$ and $b_{q+1} - b_q = -2^{-q}$, we get using summation by parts

$$\begin{aligned}
& \sum_{q=q_0+1}^{\infty} 2^{-q} a_q = - \sum_{q=q_0+1}^{\infty} (b_{q+1} - b_q) a_q = -a_q b_q|_{q_0+1}^{\infty} + \sum_{q=q_0+1}^{\infty} b_{q+1} (a_{q+1} - a_q) \\
& = 2 \cdot 2^{-(q_0+1)} \sqrt{\log n_{q_0+1}} + \sum_{q=q_0+1}^{\infty} 2 \cdot 2^{-(q+1)} \sqrt{\log n_{q+1}} = 2 \sum_{q=q_0+1}^{\infty} 2^{-q} \sqrt{\log n_q},
\end{aligned}$$

where we use the assumption that $2^{-q} \sqrt{\log n_q} \rightarrow 0$ as $q \rightarrow \infty$, so that $-a_q b_q|_{q_0+1}^{\infty} = 2 \cdot 2^{-(q_0+1)} \sqrt{\log n_{q_0+1}}$. Using that $2^{-q} \sqrt{\log n_q}$ is decreasing in q by assumption,

$$2 \sum_{q=q_0+1}^{\infty} 2^{-q} \sqrt{\log n_q} \leq 2 \int_{q_0}^{\infty} 2^{-q} \sqrt{\log n(2^{-q}, \mathcal{F}, P)} dq.$$

Using a change of variables and that $2^{-q_0} \leq \rho(\mathcal{F}, P)/4$, we finally conclude that

$$\sum_{q=q_0+1}^{\infty} \eta_q \leq K \|F\|_{P,2} \frac{16}{\log 2} \int_0^{\rho(\mathcal{F}, P)/4} \sqrt{\log n(x, \mathcal{F}, P)} dx.$$

Step 3. Letting $\eta_{q_0} = K \|F\|_{P,2} \rho(\mathcal{F}, P) \sqrt{2 \log N_{q_0}}$, recalling that $N_{q_0} = n_{q_0}$, using that $\|\pi_{q_0}(f)\|_{P,2} \leq \|F\|_{P,2}$ and sub-Gaussianity, we conclude

$$\begin{aligned}
& \mathbb{P} \left\{ \max_f |\mathbb{G}(\pi_{q_0}(f))| > \eta_{q_0} \right\} \leq n_{q_0} 2 \exp(-(K/D)^2 \log n_{q_0}) \leq 2 \exp(-\{(K/D)^2 - 1\} \log n_{q_0}) \\
& \leq \int_{q_0-1}^{q_0} 2 \exp(-\{(K/D)^2 - 1\} \log n_q) dq = \int_{\rho(\mathcal{F}, P)/4}^{\rho(\mathcal{F}, P)/2} (x \ln 2)^{-1} 2n(x, \mathcal{F}, P)^{-\{(K/D)^2 - 1\}} dx.
\end{aligned}$$

Also, since $n_{q_0} = n(2^{-q_0}, \mathcal{F}, P)$, $2^{-q_0} \leq \rho(\mathcal{F}, P)/4$, and $n(x, \mathcal{F}, P)$ is increasing in $1/x$, we obtain $\eta_{q_0} \leq 4\sqrt{2} K \|F\|_{P,2} \int_0^{\rho(\mathcal{F}, P)/4} \sqrt{\log n(x, \mathcal{F}, P)} dx$.

Step 4. Finally, adding the bounds on tail probabilities from Steps 2 and 3 we obtain the tail bound stated in the main text. Further, adding bounds on η_q from Steps 2 and 3, and using $c = 16/\log 2 + 4\sqrt{2} < 30$, we obtain $\sum_{q=q_0}^{\infty} \eta_q \leq cK\|F\|_{P,2} \int_0^{\rho(\mathcal{F},P)/4} \sqrt{\log n(x, \mathcal{F}, P)} dx$.

□

PROOF OF LEMMA 17. The proof proceeds analogously to the proof of Lemma 2.3.7 (page 112) in [40] with the necessary adjustments. Letting q_τ be the τ quantile of $x(Z)$ we have

$$P \left\{ \left\| \sum_{i=1}^n Z_i \right\|_{\mathcal{F}} > x_0 \vee x(Z) \right\} \leq P \left\{ x(Z) \geq q_\tau, \left\| \sum_{i=1}^n Z_i \right\|_{\mathcal{F}} > x_0 \vee x(Z) \right\} + P\{x(Z) < q_\tau\}.$$

Next we bound the first term of the expression above. Let $Y = (Y_1, \dots, Y_n)$ be an independent copy of $Z = (Z_1, \dots, Z_n)$, suitably defined on a product space. Fix a realization of Z such that $x(Z) \geq q_\tau$ and $\|\sum_{i=1}^n Z_i\|_{\mathcal{F}} > x_0 \vee x(Z)$. Therefore $\exists f_Z \in \mathcal{F}$ such that $|\sum_{i=1}^n Z_i(f_Z)| > x_0 \vee x(Z)$. Conditional on such a Z and using the triangular inequality we have that

$$\begin{aligned} P_Y \left\{ x(Y) \leq q_\tau, |\sum_{i=1}^n Y_i(f_Z)| \leq \frac{x_0}{2} \right\} &\leq P_Y \left\{ |\sum_{i=1}^n (Y_i - Z_i)(f_Z)| > \frac{x_0 \vee x(Z) \vee x(Y)}{2} \right\} \\ &\leq P_Y \left\{ \left\| \sum_{i=1}^n (Y_i - Z_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(Z) \vee x(Y)}{2} \right\}. \end{aligned}$$

By definition of x_0 we have $\inf_{f \in \mathcal{F}} P \left\{ |\sum_{i=1}^n Y_i(f)| \leq \frac{x_0}{2} \right\} \geq 1 - \bar{p}_\tau/2$. Since $P_Y \{x(Y) \leq q_\tau\} = \bar{p}_\tau$, by Bonferroni inequality we have that the left hand side is bounded from below by $\bar{p}_\tau - \bar{p}_\tau/2 = \bar{p}_\tau/2$. Therefore, over the set $\{Z : x(Z) \geq q_\tau, \|\sum_{i=1}^n Z_i\|_{\mathcal{F}} > x_0 \vee x(Z)\}$ we have $\frac{\bar{p}_\tau}{2} \leq P_Y \left\{ \left\| \sum_{i=1}^n (Y_i - Z_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(Z) \vee x(Y)}{2} \right\}$. Integrating over Z we obtain

$$\frac{\bar{p}_\tau}{2} P \left\{ x(Z) \geq q_\tau, \left\| \sum_{i=1}^n Z_i \right\|_{\mathcal{F}} > x_0 \vee x(Z) \right\} \leq P_Z P_Y \left\{ \left\| \sum_{i=1}^n (Y_i - Z_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(Z) \vee x(Y)}{2} \right\}.$$

Let $\varepsilon_1, \dots, \varepsilon_n$ be an independent sequence of Rademacher random variables. Given $\varepsilon_1, \dots, \varepsilon_n$, set $(\tilde{Y}_i = Y_i, \tilde{Z}_i = Z_i)$ if $\varepsilon_i = 1$ and $(\tilde{Y}_i = Z_i, \tilde{Z}_i = Y_i)$ if $\varepsilon_i = -1$. That is, we create vectors \tilde{Y} and \tilde{Z} by pairwise exchanging their components; by construction, conditional on each $\varepsilon_1, \dots, \varepsilon_n$, (\tilde{Y}, \tilde{Z}) has the same distribution as (Y, Z) . Therefore,

$$P_Z P_Y \left\{ \left\| \sum_{i=1}^n (Y_i - Z_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(Z) \vee x(Y)}{2} \right\} = E_\varepsilon P_Z P_Y \left\{ \left\| \sum_{i=1}^n (\tilde{Y}_i - \tilde{Z}_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(\tilde{Z}) \vee x(\tilde{Y})}{2} \right\}.$$

By $x(\cdot)$ being k -sub-exchangeable, and since $\varepsilon_i(Y_i - Z_i) = (\tilde{Y}_i - \tilde{Z}_i)$, we have that

$$E_\varepsilon P_Z P_Y \left\{ \left\| \sum_{i=1}^n (\tilde{Y}_i - \tilde{Z}_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(\tilde{Z}) \vee x(\tilde{Y})}{2} \right\} \leq E_\varepsilon P_Z P_Y \left\{ \left\| \sum_{i=1}^n \varepsilon_i (Y_i - Z_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(Z) \vee x(Y)}{2k} \right\}.$$

By the triangular inequality and removing $x(Y)$ or $x(Z)$, the latter is bounded by

$$P \left\{ \left\| \sum_{i=1}^n \varepsilon_i (Y_i - \mu_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(Y)}{4k} \right\} + P \left\{ \left\| \sum_{i=1}^n \varepsilon_i (Z_i - \mu_i) \right\|_{\mathcal{F}} > \frac{x_0 \vee x(Z)}{4k} \right\}.$$

□

PROOF OF LEMMA 18. Let $\mathbb{G}_n^o(f) = n^{-1/2} \sum_{i=1}^n \{\varepsilon_i f(Z_i)\}$ be the symmetrized empirical process, where $\varepsilon_1, \dots, \varepsilon_n$ are i.i.d. Rademacher random variables, i.e., $P(\varepsilon_i = 1) = P(\varepsilon_i = -1) = 1/2$, which are independent of Z_1, \dots, Z_n . By the Chebyshev's inequality and the assumption on $e_n(\mathcal{F}, \mathbb{P}_n)$, namely $\sup_{f \in \mathcal{F}} \text{var}_{\mathbb{P}} f \leq (\tau/2)(4kcKe_n(\mathcal{F}, \mathbb{P}_n))^2$, we have for the constant τ fixed in the statement of the lemma

$$P(|\mathbb{G}_n(f)| > 4kcKe_n(\mathcal{F}, \mathbb{P}_n)) \leq \tau/2.$$

Therefore, by the symmetrization Lemma 17 we obtain

$$\mathbb{P} \left\{ \sup_{f \in \mathcal{F}} |\mathbb{G}_n(f)| > 4kcKe_n(\mathcal{F}, \mathbb{P}_n) \right\} \leq \frac{4}{\tau} \mathbb{P} \left\{ \sup_{f \in \mathcal{F}} |\mathbb{G}_n^o(f)| > cKe_n(\mathcal{F}, \mathbb{P}_n) \right\} + \tau.$$

We then condition on the values of Z_1, \dots, Z_n , denoting the conditional probability measure as \mathbb{P}_ε . Conditional on Z_1, \dots, Z_n , by the Hoeffding inequality the symmetrized process \mathbb{G}_n^o is sub-Gaussian for the $L_2(\mathbb{P}_n)$ norm, namely, for $g \in \mathcal{F} - \mathcal{F}$, $\mathbb{P}_\varepsilon\{\mathbb{G}_n^o(g) > x\} \leq 2 \exp(-x^2/[2\|g\|_{\mathbb{P}_n,2}^2])$. Hence by Lemma 16 with $D = 1$, we can bound

$$\mathbb{P}_\varepsilon \left\{ \sup_{f \in \mathcal{F}} |\mathbb{G}_n^o(f)| \geq cKe_n(\mathcal{F}, \mathbb{P}_n) \right\} \leq \left[\int_0^{\rho(\mathcal{F}, \mathbb{P}_n)/2} \epsilon^{-1} n(\epsilon, \mathcal{F}, P)^{-\{K^2-1\}} d\epsilon \right] \wedge 1.$$

The result follows from taking the expectation over Z_1, \dots, Z_n . □

PROOF OF LEMMA 19. Step 1. (Main Step) In this step we prove the main result. First, we observe that the bound $\epsilon \mapsto n(\epsilon, \mathcal{F}_m, \mathbb{P}_n)$ satisfies the monotonicity hypotheses of Lemma 18 uniformly in $m \leq N$.

Second, recall $e_n(\mathcal{F}_m, \mathbb{P}_n) := C \sqrt{m \log(N \vee n \vee \theta_m \vee \omega)} \max\{\sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P},2}, \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n,2}\}$ for $C = (1 + \sqrt{2v})/4$. Note that $\sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n,2}$ is $\sqrt{2}$ -sub-exchangeable and $\rho(\mathcal{F}_m, \mathbb{P}_n) := \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n,2} / \|F_m\|_{\mathbb{P}_n,2} \geq 1/\sqrt{n}$ by Step 2 below. Thus, uniformly in $m \leq N$:

$$\begin{aligned} & \|F_m\|_{\mathbb{P}_n,2} \int_0^{\rho(\mathcal{F}_m, \mathbb{P}_n)/4} \sqrt{\log n(\epsilon, \mathcal{F}_m, \mathbb{P}_n)} d\epsilon \\ & \leq \|F_m\|_{\mathbb{P}_n,2} \int_0^{\rho(\mathcal{F}_m, \mathbb{P}_n)/4} \sqrt{m \log(N \vee n \vee \theta_m) + vm \log(\omega/\epsilon)} d\epsilon \\ & \leq (1/4) \sqrt{m \log(N \vee n \vee \theta_m)} \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n,2} + \|F_m\|_{\mathbb{P}_n,2} \int_0^{\rho(\mathcal{F}_m, \mathbb{P}_n)/4} \sqrt{vm \log(\omega/\epsilon)} d\epsilon \\ & \leq \sqrt{m \log(N \vee n \vee \theta_m \vee \omega)} \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n,2} (1 + \sqrt{2v})/4 \leq e_n(\mathcal{F}_m, \mathbb{P}_n), \end{aligned}$$

which follows by $\int_0^\rho \sqrt{\log(\omega/\epsilon)} d\epsilon \leq (\int_0^\rho 1 d\epsilon)^{1/2} (\int_0^\rho \log(\omega/\epsilon) d\epsilon)^{1/2} \leq \rho \sqrt{2 \log(n \vee \omega)}$, for $1/\sqrt{n} \leq \rho \leq 1$.

Third, for any $K \geq \sqrt{2/\delta} > 1$ we have $(K^2 - 1) \geq 1/\delta$, and let $\tau_m = \delta/(4m \log(N \vee n \vee \theta_0))$. Recall that $4\sqrt{2}cC > 4$ where $4 < c < 30$ is defined in Lemma 16. Note that for any $m \leq N$ and $f \in \mathcal{F}_m$, we have by Chebyshev's inequality and since $e_n(\mathcal{F}_m, \mathbb{P}_n) \geq C\sqrt{m \log(N \vee n \vee \theta_m \vee \omega)} \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P},2}$

$$P(|\mathbb{G}_n(f)| > 4\sqrt{2}cK e_n(\mathcal{F}_m, \mathbb{P}_n)) \leq \frac{\delta/2}{(4\sqrt{2}cC)^2 m \log(N \vee n \vee \theta_0)} \leq \tau_m/2.$$

By Lemma 18 with our choice of τ_m , $m \leq N$, $\omega > 1$, $v > 1$, and $\rho(\mathcal{F}_m, \mathbb{P}_n) \leq 1$,

$$\begin{aligned} & \mathbb{P} \left\{ \sup_{f \in \mathcal{F}_m} |\mathbb{G}_n(f)| > 4\sqrt{2}cK e_n(\mathcal{F}_m, \mathbb{P}_n), \exists m \leq N \right\} \\ & \leq \sum_{m=1}^N \mathbb{P} \left\{ \sup_{f \in \mathcal{F}_m} |\mathbb{G}_n(f)| > 4\sqrt{2}cK e_n(\mathcal{F}_m, \mathbb{P}_n) \right\} \\ & \leq \sum_{m=1}^N \left[\frac{4(N \vee n \vee \theta_m)^{-m/\delta}}{\tau_m} \int_0^{1/2} (\omega/\epsilon)^{(-vm/\delta)+1} d\epsilon + \tau_m \right] \\ & \leq 4 \sum_{m=1}^N \frac{(N \vee n \vee \theta_m)^{-m/\delta}}{\tau_m} \frac{1}{vm/\delta} + \sum_{m=1}^N \tau_m \\ & < 16 \frac{(N \vee n \vee \theta_0)^{-1/\delta}}{1 - (N \vee n \vee \theta_0)^{-1/\delta}} \log(N \vee n \vee \theta_0) + \frac{\delta}{4} \frac{(1 + \log N)}{\log(N \vee n \vee \theta_0)} \leq \delta, \end{aligned}$$

where the last inequality follows by $N \vee n \vee \theta_0 \geq 3$ and $\delta \in (0, 1/6)$.

Step 2. (Auxiliary calculations.) To establish that $\sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n,2}$ is $\sqrt{2}$ -sub-exchangeable, let \tilde{Z} and \tilde{Y} be created by exchanging any components in Z with corresponding components in Y . Then

$$\begin{aligned} & \sqrt{2}(\sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(\tilde{Z}),2} \vee \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(\tilde{Y}),2}) \geq (\sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(\tilde{Z}),2}^2 + \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(\tilde{Y}),2}^2)^{1/2} \\ & \geq (\sup_{f \in \mathcal{F}_m} \mathbb{E}_n[f(\tilde{Z}_i)^2] + \mathbb{E}_n[f(\tilde{Y}_i)^2])^{1/2} = (\sup_{f \in \mathcal{F}_m} \mathbb{E}_n[f(Z_i)^2] + \mathbb{E}_n[f(Y_i)^2])^{1/2} \\ & \geq (\sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(Z),2}^2 \vee \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(Y),2}^2)^{1/2} = \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(Z),2} \vee \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n(Y),2}. \end{aligned}$$

Next we show that $\rho(\mathcal{F}_m, \mathbb{P}_n) := \sup_{f \in \mathcal{F}_m} \|f\|_{\mathbb{P}_n,2}/\|F_m\|_{\mathbb{P}_n,2} \geq 1/\sqrt{n}$ for $m \leq N$. The latter bound follows from $\mathbb{E}_n[F_m^2] = \mathbb{E}_n[\sup_{f \in \mathcal{F}_m} |f(Z_i)|^2] \leq \sup_{i \leq n} \sup_{f \in \mathcal{F}_m} |f(Z_i)|^2$, and from $\sup_{f \in \mathcal{F}_m} \mathbb{E}_n[|f(Z_i)|^2] \geq \sup_{f \in \mathcal{F}_m} \sup_{i \leq n} |f(Z_i)|^2/n$. \square

APPENDIX G: SUPPLEMENTARY MATERIAL

Recall the notation for the unit sphere $\mathbb{S}^{n-1} = \{\alpha \in \mathbb{R}^n : \|\alpha\| = 1\}$ and the k -sparse unit sphere $\mathbb{S}_p^k = \{\alpha \in \mathbb{R}^p : \|\alpha\| = 1, \|\alpha\|_0 \leq k\}$. Define also the sparse sphere associated with a given vector β as $\mathbb{S}(\beta) = \{\alpha \in \mathbb{R}^p : \|\alpha\| \leq 1, \text{support}(\alpha) \subseteq \text{support}(\beta)\}$.

G.1. Proofs for Examples of Simple Sufficient Conditions. In this section we provide the proofs for Lemmas 1 and 2 which show that two designs of interest imply conditions D.1-5 under mild assumptions. We restate the designs for the reader's convenience.

DESIGN 1: LOCATION MODEL WITH CORRELATED NORMAL DESIGN. Let us consider estimating a standard location model

$$y = x' \beta^o + \varepsilon,$$

where $\varepsilon \sim N(0, \sigma^2)$, $\sigma > 0$ is fixed, $x = (1, z')'$, with $z \sim N(0, \Sigma)$, where Σ has ones in the diagonal, minimum eigenvalue bounded away from zero by constant κ^2 and maximum eigenvalue bounded from above, uniformly in n .

LEMMA 20. *Under Design 1 with $\mathcal{U} = [\xi, 1 - \xi]$, $\xi > 0$, conditions D.1-D.5 are satisfied with*

$$\begin{aligned} \bar{f} &= 1/[\sqrt{2\pi}\sigma], \quad \bar{f}' = \sqrt{e/[2\pi]}/\sigma^2, \quad \underline{f} = 1/\sqrt{2\pi\xi}\sigma, \\ \|\beta(u)\|_0 &\leq \|\beta^o\|_0 + 1, \quad \gamma = 2p \exp(-n/24), \quad L = \sigma/\xi \\ \kappa_m \wedge \tilde{\kappa}_{\tilde{m}} &\geq \kappa, \quad q \wedge \tilde{q}_{\tilde{m}} \geq (3/[32\xi^{3/4}])\sqrt{\sqrt{2\pi}\sigma/e}. \end{aligned}$$

PROOF. This model implies a linear quantile model with coefficients $\beta_1(u) = \beta_1^o + \sigma\Phi^{-1}(u)$ and $\beta_j(u) = \beta_j^o$ for $j = 2, \dots, p$. Let

$$\bar{f}' = \sup_z \phi'(z/\sigma)/\sigma^2, \quad \bar{f} = \sup_z \phi(z/\sigma)/\sigma, \quad \underline{f} = \min_{u \in \mathcal{U}} \phi(\Phi^{-1}(u))/\sigma,$$

so that D.1 holds with the constants \bar{f} and \bar{f}' . D.2 holds, since $\|\beta(u)\|_0 \leq \|\beta^o\|_0 + 1$ and $u \mapsto \beta(u)$ is Lipschitz over \mathcal{U} with the constant $L = \sup_{u \in \mathcal{U}} \sigma/\phi(\Phi^{-1}(u))$, which trivially obeys $\log L \lesssim \log(n \vee p)$. D.4 also holds, in particular by Chernoff's tail bound

$$P \left\{ \max_{1 \leq j \leq p} |\hat{\sigma}_j - 1| \leq 1/2 \right\} \geq 1 - \gamma = 1 - 2p \exp(-n/24),$$

where $1 - \gamma$ approaches 1 if $n/\log p \rightarrow \infty$. Furthermore, the smallest eigenvalue of the population design matrix Σ is at least $(1 - |\rho|)/(2 + 2|\rho|)$ and the maximum eigenvalue is at most $(1 + |\rho|)/(1 - |\rho|)$. Thus, D.4 and D.5 hold with

$$\kappa_m \wedge \tilde{\kappa}_{\tilde{m}} \geq \kappa,$$

for all $m, \tilde{m} \geq 0$. If the covariates x have a log-concave density, then

$$q \geq 3\underline{f}^{3/2}/(8K_\ell \bar{f}') \text{ for a universal constant } K_\ell.$$

In the case of normal variables you can take $K_\ell = 4/\sqrt{2\pi}$. The bound follows from $\mathbb{E}[|x'_i \delta|^3] \leq K_\ell \mathbb{E}[|x'_i \delta|^2]^{3/2}$ holding for log-concave x for some universal constant K_ℓ by Theorem 5.22 of [31]. The bound for $\tilde{q}_{\tilde{m}}$ is the same. \square

DESIGN 2: LOCATION-SCALE SHIFT WITH BOUNDED REGRESSORS. Let us consider estimating a standard location model

$$y = x' \beta^o + \varepsilon \cdot x' \eta,$$

where $\varepsilon \sim F$ independent of x , with probability density function f . We assume that the population design matrix $\mathbb{E}[xx']$ has ones in the diagonal and has eigenvalues uniformly bounded away from zero and from above, $x_1 = 1$, $\max_{1 \leq j \leq p} |x_j| \leq K_B$. Moreover, the vector η is such that $0 < v \leq x' \eta \leq \Upsilon < \infty$ for all values of x .

LEMMA 21. *Under Design 2 with $\mathcal{U} = [\xi, 1 - \xi]$, $\xi > 0$, conditions D.1-D.5 are satisfied with*

$$\begin{aligned} \bar{f} &\leq \max_{\varepsilon} f(\varepsilon)/v, \quad \bar{f}' \leq \max_{\varepsilon} f'(\varepsilon)/v^2, \quad \underline{f} = \min_{u \in \mathcal{U}} f(F^{-1}(u))/\Upsilon, \\ \|\beta(u)\|_0 &\leq \|\beta^o\|_0 + \|\eta\|_0 + 1, \quad \gamma = 2p \exp(-n/[8K_B^4]), \\ \kappa_m \wedge \tilde{\kappa}_{\tilde{m}} &\geq \kappa, \quad L = \|\eta\| \underline{f} \\ q &\geq \frac{3}{8} \frac{\underline{f}^{3/2}}{\bar{f}'} \kappa / [10K_B \sqrt{s}], \quad \tilde{q}_{\tilde{m}} \geq \frac{3}{8} \frac{\underline{f}^{3/2}}{\bar{f}'} \kappa / [K_B \sqrt{s + \tilde{m}}], \end{aligned}$$

PROOF. This model implies a linear quantile model with coefficients $\beta(u) = \beta_1^o + F^{-1}(u)\eta$. We have

$$\bar{f}' = \max_y f'(y)/v^2, \quad \bar{f} = \max_y f(y)/v, \quad \underline{f} \geq \min_{u \in \mathcal{U}} f(F^{-1}(u))/\Upsilon,$$

so that D.1 holds with the constants \bar{f} and \bar{f}' . D.2 holds, since $\|\beta(u)\|_0 \leq \|\beta^o\|_0 + \|\eta\|_0 + 1$ and $u \mapsto \beta(u)$ is Lipschitz over \mathcal{U} with the constant $L = \|\eta\| \max_{u \in \mathcal{U}} \Upsilon / f(F^{-1}(u))$ uniformly in n , which obeys $\log L \lesssim \log(n \vee p)$. Next recall that $x_{ij}^2 \leq K_B^2$. Then, by Hoeffding inequality we have

$$P(|\mathbb{E}_n[x_{ij}^2] - 1| \geq 1/2) \leq 2 \exp(-n/[8K_B^4]).$$

Applying a union bound D.3 holds with $\gamma = 2p \exp(-n/[8K_B^4])$ which approaches 0 if $n/\log p \rightarrow \infty$. Furthermore, the smallest eigenvalue of the population design matrix is bounded away from zero. Thus, D.4 and D.5 hold with $c_0 = 9$ (in fact with any $c_0 > 0$) and

$$\kappa_m \wedge \tilde{\kappa}_{\tilde{m}} \geq \sqrt{\text{mineig}(\mathbb{E}[xx'])},$$

for all $m, \tilde{m} \geq 0$.

Finally, the restricted nonlinear impact coefficient satisfies $q \geq 3\underline{f}^{3/2}\kappa_0/(8\bar{f}'K_B(1+c_0)\sqrt{s})$. Indeed, the latter bound follows from $\mathbb{E}[|x'_i\delta|^3] \leq \mathbb{E}[|x'_i\delta|^2]K_B\|\delta\|_1 \leq \mathbb{E}[|x'_i\delta|^2]K_B(1+c_0)\sqrt{s}\|\delta_{T_u}\| \leq \mathbb{E}[|x'_i\delta|^2]^{3/2}K_B(1+c_0)\sqrt{s}/\kappa_0$ holding since $\delta \in A_u$ so that $\|\delta\|_1 \leq (1+c_0)\|\delta_{T_u}\|_1 \leq \sqrt{s}(1+c_0)\|\delta_{T_u}\|$. Similarly, one can show $\tilde{q}_{\tilde{m}} \geq (3/8)\underline{f}^{3/2}\tilde{\kappa}_{\tilde{m}}/(\bar{f}'K_B\sqrt{\tilde{m}+s})$. \square

G.2. Lemma 9: Proof of the VC index bound.

LEMMA 22. *Consider a fixed subset $T \subset \{1, 2, \dots, p\}$, $|T| = m$. The class of functions*

$$\mathcal{F}_T = \{\alpha'(\psi_i(\beta, u) - \psi_i(\beta(u), u)) : u \in \mathcal{U}, \alpha \in \mathbb{S}(\beta), \text{support}(\beta) \subseteq T\}$$

has a VC index bounded by cm for some universal constant c .

PROOF. Consider the classes of functions $\mathcal{W} := \{x'\alpha : \text{support}(\alpha) \subseteq T\}$ and $\mathcal{V} := \{1\{y \leq x'\beta\} : \text{support}(\beta) \subseteq T\}$ (for convenience let $Z = (y, x)$), $\mathcal{Z} := \{1\{y \leq x'\beta(u)\} : u \in \mathcal{U}\}$. Since T is fixed and has cardinality m , the VC indices of \mathcal{W} and \mathcal{V} are bounded by $m+2$; see, for example, van der Vaart and Wellner [40] Lemma 2.6.15. On the other hand, since $u \mapsto x'\beta(u)$ is monotone, the VC index of \mathcal{Z} is 1. Next consider $f \in \mathcal{F}_T$ which can be written in the form $f(Z) := g(Z)(1\{h(Z) \leq 0\} - 1\{p(Z) \leq 0\})$ where $g \in \mathcal{W}$, $1\{h \leq 0\} \in \mathcal{V}$, and $1\{p \leq 0\} \in \mathcal{Z}$. The VC index of \mathcal{F}_T is by definition equal to the VC index of the class of sets $\{(Z, t) : f(Z) \leq t\}, f \in \mathcal{F}_T, t \in \mathbb{R}$. We have that

$$\begin{aligned} \{(Z, t) : f(Z) \leq t\} &= \{(Z, t) : g(Z)(1\{h(Z) \leq 0\} - 1\{p(Z) \leq 0\}) \leq t\} \\ &= \{(Z, t) : h(Z) > 0, p(Z) > 0, t \geq 0\} \cup \\ &\cup \{(Z, t) : h(Z) \leq 0, p(Z) \leq 0, t \geq 0\} \cup \\ &\cup \{(Z, t) : h(Z) \leq 0, p(Z) > 0, g(Z) \leq t\} \cup \\ &\cup \{(Z, t) : h(Z) > 0, p(Z) \leq 0, g(Z) \geq t\}. \end{aligned}$$

Thus each set $\{(Z, t) : f(Z) \leq t\}$ is created by taking finite unions, intersections, and complements of the basic sets $\{Z : h(Z) > 0\}$, $\{Z : p(Z) \leq 0\}$, $\{t \geq 0\}$, $\{(Z, t) : g(Z) \geq t\}$, and $\{(Z, t) : g(Z) \leq t\}$. These basic sets form VC classes, each having VC index of order m . Therefore, the VC index of a class of sets $\{(Z, t) : f(Z) \leq t\}, f \in \mathcal{F}_T, t \in \mathbb{R}$ is of the same order as the sum of the VC indices of the set classes formed by the basic VC sets; see van der Vaart and Wellner [40] Lemma 2.6.17. \square

G.3. Gaussian Sparse Eigenvalue. It will be convenient to recall the following result.

LEMMA 23. *Let M be a semi-definite positive matrix and $\phi_M(k) = \sup\{\alpha'M\alpha : \alpha \in \mathbb{S}_p^k\}$. For any integers k and ℓk with $\ell \geq 1$, we have $\phi_M(\ell k) \leq \lceil \ell \rceil \phi_M(k)$.*

The following lemmas characterize the behavior of the maximal sparse eigenvalue for the case of correlated Gaussian regressors. We start by establishing an upper bound on $\phi_{\mathbb{E}_n[x_i x'_i]}(k)$ that holds with high probability.

LEMMA 24. *Consider $x_i = \Sigma^{1/2} z_i$, where $z_i \sim N(0, I_p)$, $p \geq n$, and $\sup_{\alpha \in \mathbb{S}_p^k} \alpha' \Sigma \alpha \leq \sigma^2(k)$. Let $\phi_{\mathbb{E}_n[x_i x'_i]}(k)$ be the maximal k -sparse eigenvalue of $\mathbb{E}_n[x_i x'_i]$, for $k \leq n$. Then with probability converging to one, uniformly in $k \leq n$,*

$$\sqrt{\phi_{\mathbb{E}_n[x_i x'_i]}(k)} \lesssim \sigma(k) \left(1 + \sqrt{k/n} \sqrt{\log p}\right).$$

PROOF. By Lemma 23 it suffices to establish the result for $k \leq n/2$. Let Z be the $n \times p$ matrix collecting vectors z'_i , $i = 1, \dots, n$ as rows. Consider the Gaussian process $\mathcal{G}_k : (\alpha, \tilde{\alpha}) \mapsto \tilde{\alpha}' Z \alpha / \sqrt{n}$, where $(\alpha, \tilde{\alpha}) \in \mathbb{S}_p^k \times \mathbb{S}^{n-1}$. Note that

$$\|\mathcal{G}_k\| = \sup_{(\alpha, \tilde{\alpha}) \in \mathbb{S}_p^k \times \mathbb{S}^{n-1}} |\tilde{\alpha}' Z \alpha / \sqrt{n}| = \sup_{\alpha \in \mathbb{S}_p^k} \sqrt{\alpha' \mathbb{E}_n[z_i z'_i] \alpha} = \sqrt{\phi_{\mathbb{E}_n[z_i z'_i]}(k)}.$$

Using Borell's concentration inequality for the Gaussian process (see van der Vaart and Wellner [40] Lemma A.2.1) we have that $P(\|\mathcal{G}_k\| - \text{median}\|\mathcal{G}_k\| > r) \leq e^{-nr^2/2}$. Also, by classical results on the behavior of the maximal eigenvalues of the Gaussian covariance matrices (see German [17]), as $n \rightarrow \infty$, for any $k/n \rightarrow \gamma \in [0, 1]$, we have that $\lim_{k,n}(\text{median}\|\mathcal{G}_k\| - 1 - \sqrt{k/n}) = 0$. Since k/n lies within $[0, 1]$, any sequence k_n/n has convergent subsequence with limit in $[0, 1]$. Therefore, we can conclude that, as $n \rightarrow \infty$, $\limsup_{k,n}(\text{median}\|\mathcal{G}_{k_n}\| - 1 - \sqrt{k_n/n}) \leq 0$. This further implies $\limsup_n \sup_{k \leq n}(\text{median}\|\mathcal{G}_k\| - 1 - \sqrt{k/n}) \leq 0$. Therefore, for any $r_0 > 0$ there exists n_0 large enough such that for all $n \geq n_0$ and all $k \leq n$, $P(\|\mathcal{G}_k\| > 1 + \sqrt{k/n} + r_0) \leq e^{-nr^2/2}$. There are at most $\binom{p}{k}$ subvectors of z_i containing k elements, so we conclude that for $n \geq n_0$,

$$P\left(\sup_{\alpha \in \mathbb{S}_p^k} \sqrt{\alpha' \mathbb{E}_n[z_i z'_i] \alpha} > 1 + \sqrt{k/n} + r_k + r_0\right) \leq \binom{p}{k} e^{-nr_k^2/2}.$$

Summing over $k \leq n$ we obtain

$$\sum_{k=1}^n P\left(\sup_{\alpha \in \mathbb{S}_p^k} \sqrt{\alpha' \mathbb{E}_n[z_i z'_i] \alpha} > 1 + \sqrt{k/n} + r_k + r_0\right) \leq \sum_{k=1}^n \binom{p}{k} e^{-nr_k^2/2}.$$

Setting $r_k = \sqrt{ck/n \log p}$ for $c > 1$ and using that $\binom{p}{k} \leq p^k$, we bound the right side by $\sum_{k=1}^n e^{(1-c)k \log p} \rightarrow 0$ as $n \rightarrow \infty$. We conclude that with probability converging to one, uniformly for all k : $\sup_{\alpha \in \mathbb{S}_p^k} \sqrt{\alpha' \mathbb{E}_n[z_i z'_i] \alpha} \lesssim 1 + \sqrt{k/n} \sqrt{\log p}$. Furthermore, since $\sup_{\alpha \in \mathbb{S}_p^k} \alpha' \Sigma \alpha \leq \sigma^2(k)$, we conclude that with probability converging to one, uniformly for all k :

$$\sup_{\alpha \in \mathbb{S}_p^k} \sqrt{\alpha' \mathbb{E}_n[x_i x'_i] \alpha} \lesssim \sigma(k) (1 + \sqrt{k/n} \sqrt{\log p}).$$

□

Next, relying on Sudakov's minoration, we show a lower bound on the expectation of the maximum k -sparse eigenvalue. We do not use the lower bound in the analysis, but the result shows that the upper bound is sharp in terms of the rate dependence on k, p , and n .

LEMMA 25. Consider $x_i = \Sigma^{1/2} z_i$, where $z_i \sim N(0, I_p)$, and $\inf_{\alpha \in \mathbb{S}_p^k} \alpha' \Sigma \alpha \geq \underline{\sigma}^2(k)$. Let $\phi_{\mathbb{E}_n[x_i x_i']}(k)$ be the maximal k -sparse eigenvalue of $\mathbb{E}_n[x_i x_i']$, for $k \leq n < p$. Then for any even k we have that:

$$(1) \quad \mathbb{E} \left[\sqrt{\phi_{\mathbb{E}_n[x_i x_i']}(k)} \right] \geq \frac{\underline{\sigma}(2k)}{3\sqrt{n}} \sqrt{(k/2) \log(p-k)} \text{ and}$$

$$(2) \quad \sqrt{\phi_{\mathbb{E}_n[x_i x_i']}(k)} \gtrsim_P \frac{\underline{\sigma}(2k)}{3\sqrt{n}} \sqrt{(k/2) \log(p-k)}.$$

PROOF. Let X be the $n \times p$ matrix collecting vectors x'_i , $i = 1, \dots, n$ as rows. Consider the Gaussian process $(\alpha, \tilde{\alpha}) \mapsto \tilde{\alpha}' X \alpha / \sqrt{n}$, where $(\alpha, \tilde{\alpha}) \in \mathbb{S}_p^k \times \mathbb{S}^{n-1}$. Note that $\sqrt{\phi_{\mathbb{E}_n[x_i x_i']}(k)}$ is the supremum of this Gaussian process

$$(G.1) \quad \sup_{(\alpha, \tilde{\alpha}) \in \mathbb{S}_p^k \times \mathbb{S}^{n-1}} |\tilde{\alpha}' X \alpha / \sqrt{n}| = \sup_{\alpha \in \mathbb{S}_p^k} \sqrt{\alpha' \mathbb{E}_n[x_i x_i'] \alpha} = \sqrt{\phi_{\mathbb{E}_n[x_i x_i']}(k)}.$$

Hence we proceed in three steps: In Step 1, we consider the uncorrelated case and prove the lower bound (1) on the expectation of the supremum using Sudakov's minoration, using a lower bound on a relevant packing number. In Step 2, we derive the lower bound on the packing number. In Step 3, we generalize Step 1 to the correlated case. In Step 4, we prove the lower bound (2) on the supremum itself using Borell's concentration inequality.

Step 1. In this step we consider the case where $\Sigma = I$ and show the result using Sudakov's minoration. By fixing $\tilde{\alpha} = (1, \dots, 1)' / \sqrt{n} \in \mathbb{S}^{n-1}$, we have $\sqrt{\phi_{\mathbb{E}_n[x_i x_i']}(k)} \geq \sup_{\alpha \in \mathbb{S}_p^k} \mathbb{E}_n[x'_i \alpha] = \sup_{\alpha \in \mathbb{S}_p^k} Z_\alpha$, where $\alpha \mapsto Z_\alpha := \mathbb{E}_n[x'_i \alpha]$ is a Gaussian process on \mathbb{S}_p^k . We will bound $E[\sup_{\alpha \in \mathbb{S}_p^k} Z_\alpha]$ from below using Sudakov's minoration.

We consider the standard deviation metric on \mathbb{S}_p^k induced by Z : for any $t, s \in \mathbb{S}_p^k$,

$$d(s, t) = \sqrt{\sigma^2(Z_t - Z_s)} = \sqrt{\mathbb{E}[(Z_t - Z_s)^2]} = \sqrt{\mathbb{E}[\mathbb{E}_n[(x'_i(t-s))^2]]} = \|t - s\| / \sqrt{n}.$$

Consider the packing number $D(\epsilon, \mathbb{S}_p^k, d)$, the largest number of disjoint closed balls of radius ϵ with respect to the metric d that can be packed into \mathbb{S}_p^k , see [14]. We will bound the packing number from below for $\epsilon = \frac{1}{\sqrt{n}}$. In order to do this we restrict attention to the collection \mathcal{T} of

elements $t = (t_1, \dots, t_p) \in \mathbb{S}_p^k$ such that $t_i = 1/\sqrt{k}$ for exactly k components and $t_i = 0$ in the remaining $p - k$ components. There are $|T| = \binom{p}{k}$ of such elements. Consider any $s, t \in \mathcal{T}$ such that the support of s agrees with the support of t in at most $k/2$ elements. In this case

$$(G.2) \quad \|s - t\|^2 = \sum_{j=1}^p |t_j - s_j|^2 \geq \sum_{\substack{j \in \text{support}(t) \\ \setminus \text{support}(s)}} \frac{1}{k} + \sum_{\substack{j \in \text{support}(s) \\ \setminus \text{support}(t)}} \frac{1}{k} \geq 2 \frac{k}{2} \frac{1}{k} = 1.$$

Let \mathcal{P} be the set of the maximal cardinality, consisting of elements in \mathcal{T} such that $|\text{support}(t) \setminus \text{support}(s)| \geq k/2$ for every $s, t \in \mathcal{P}$. By the inequality (G.2) we have that $D(1/\sqrt{n}, \mathbb{S}_p^k, d) \geq |\mathcal{P}|$. Furthermore, by Step 2 given below we have that $|\mathcal{P}| \geq (p - k)^{k/2}$.

Using Sudakov's minoration ([16], Theorem 4.1.4), we conclude that

$$\mathbb{E} \left[\sup_{t \in \mathbb{S}_p^k} Z_t \right] \geq \sup_{\epsilon > 0} \frac{\epsilon}{3} \sqrt{\log D(\epsilon, \mathbb{S}_p^k, d)} \geq \sqrt{\log D(1/\sqrt{n}, \mathbb{S}_p^k, d)} \geq \frac{1}{3} \sqrt{k \log(p - k)/(2n)},$$

proving the claim of the lemma for the case $\Sigma = I$.

Step 2. In this step we show that $|\mathcal{P}| \geq (p - k)^{k/2}$.

It is convenient to identify every element $t \in \mathcal{T}$ with the set $\text{support}(t)$, where $\text{support}(t) = (j \in (1, \dots, p) : t_j = 1/\sqrt{k})$, which has cardinality k . For any $t \in \mathcal{T}$ let $\mathcal{N}(t) = (s \in \mathcal{T} : |\text{support}(t) \setminus \text{support}(s)| \leq k/2)$. By construction we have that $\max_{t \in \mathcal{T}} |\mathcal{N}(t)| |\mathcal{P}| \geq |\mathcal{T}|$. Since as shown below $\max_{t \in \mathcal{T}} |\mathcal{N}(t)| \leq K := \binom{k}{k/2} \binom{p-k/2}{k/2}$ for every t , we conclude that $|\mathcal{P}| \geq |\mathcal{T}|/K = \binom{p}{k}/K \geq (p - k)^{k/2}$.

It remains only to show that $|\mathcal{N}(t)| \leq \binom{k}{k/2} \binom{p-k/2}{k/2}$. Consider an arbitrary $t \in \mathcal{T}$. Fix any $k/2$ components of $\text{support}(t)$, and generate elements $s \in \mathcal{N}(t)$ by switching any of the remaining $k/2$ components in $\text{support}(t)$ to any of the possible $p - k/2$ values. This gives us at most $\binom{p-k/2}{k/2}$ such elements $s \in \mathcal{N}(t)$. Next let us repeat this procedure for all other combinations of initial $k/2$ components of $\text{support}(t)$, where the number of such combinations is bounded by $\binom{k}{k/2}$. In this way we generate every element $s \in \mathcal{N}(t)$. From the construction we conclude that $|\mathcal{N}(t)| \leq \binom{k}{k/2} \binom{p-k/2}{k/2}$.

Step 3. The case where $\Sigma \neq I$ follows similarly noting that the new metric, $d(s, t) = \sqrt{\sigma^2(Z_t - Z_s)} = \sqrt{\mathbb{E}[(Z_t - Z_s)^2]}$, satisfies

$$d(s, t) \geq \underline{\sigma}(2k) \|s - t\|/\sqrt{n} \quad \text{since} \quad \|s - t\|_0 \leq 2k.$$

Step 4. Using Borell's concentration inequality (see van der Vaart and Wellner [40] Lemma A.2.1) for the supremum of the Gaussian process defined in (G.1), we have $P(|\sqrt{\phi_{\mathbb{E}_n[x_i x_i']}(k)} - \mathbb{E}[\sqrt{\phi_{\mathbb{E}_n[x_i x_i']}(k)}] > r) \leq 2e^{-nr^2/2}$, which proves the second claim of the lemma. \square

Next we combine the previous lemmas to control the empirical sparse eigenvalues of the following example.

EXAMPLE 1 (Correlated Normal Design). Let us consider estimating the median ($u = 1/2$) of the following regression model

$$y = x' \beta_0 + \varepsilon,$$

where the covariate $x_1 = 1$ is the intercept and the covariates $x_{-1} \sim N(0, \Sigma)$, where $\Sigma_{ij} = \rho^{|i-j|}$ and $-1 < \rho < 1$ is fixed.

LEMMA 26. For $k \leq n$, under the design of Example 1 with $p \geq 2n$, we have

$$\phi_{\mathbb{E}_n[x_i x_i']}(k) \lesssim_P \frac{1+|\rho|}{1-|\rho|} \left(1 + \sqrt{\frac{k \log p}{n}}\right) \quad \text{and} \quad \phi_{\mathbb{E}_n[x_i x_i']}(k) \gtrsim_P \frac{1-|\rho|}{1+|\rho|} \left(1 + \sqrt{\frac{k \log p}{n}}\right).$$

PROOF. Consider first the design in Example 1 with $\rho = 0$. Let $x_{i,-1}$ denote the i th observation without the first component. Write

$$\mathbb{E}_n[x_i x_i'] = \begin{bmatrix} 1 & \mathbb{E}_n[x_{i,-1}'] \\ \mathbb{E}_n[x_{i,-1}] & 0 \end{bmatrix} + \mathbb{E}_n \begin{bmatrix} 0 & 0 \\ 0 & \mathbb{E}_n[x_{i,-1} x_{i,-1}'] \end{bmatrix} = M + N.$$

We first bound $\phi_N(k)$. Letting $N_{-1,-1} = \mathbb{E}_n[x_{i,-1} x_{i,-1}']$ we have $\phi_N(k) = \phi_{N_{-1,-1}}(k)$. Lemma 24 implies that $\phi_N(k) \lesssim_P 1 + \sqrt{k/n} \sqrt{\log p}$. Lemma 25 bounds $\phi_N(k)$ from below because $\phi_{N_{-1,-1}}(k) \gtrsim_P \sqrt{(k/2n) \log(p-k)}$.

We then bound $\phi_M(k)$. Since $M_{11} = 1$, we have $\phi_M(1) \geq 1$. To produce an upper bound let $w = (a, b')'$ achieve $\phi_M(k)$ where $a \in \mathbb{R}$, $b \in \mathbb{R}^{p-1}$. By definition we have $\|w\| = 1$, $\|w\|_0 \leq k$. Note that $|a| \leq 1$, $\|b\| = \sqrt{1 - |a|^2} \leq 1$, $\|b\|_1 \leq \sqrt{k} \|b\|$. Therefore

$$\begin{aligned} \phi_M(k) = w' M w &= a^2 + 2ab' \mathbb{E}_n[x_{i,-1}] \leq 1 + 2b' \mathbb{E}_n[x_{i,-1}] \\ &\leq 1 + 2\|b\|_1 \|\mathbb{E}_n[x_{i,-1}]\|_\infty \leq 1 + 2\sqrt{k} \|b\| \|\mathbb{E}_n[x_{i,-1}]\|_\infty. \end{aligned}$$

Next we bound $\|\mathbb{E}_n[x_{i,-1}]\|_\infty = \max_{j=2,\dots,p} |\mathbb{E}_n[x_{ij}]|$. Since $\mathbb{E}_n[x_{ij}] \sim N(0, 1/n)$ for $j = 2, \dots, p$, we have $\|\mathbb{E}_n[x_{i,-1}]\|_\infty \lesssim_P \sqrt{(1/n) \log p}$. Therefore we have $\phi_M(k) \lesssim_P 1 + 2\sqrt{k/n} \sqrt{\log p}$.

Finally, we bound $\phi_{\mathbb{E}_n[x_i x_i']}(k)$. Note that $\phi_{\mathbb{E}_n[x_i x_i']}(k) = \sup_{\alpha \in \mathbb{S}_p^k} \alpha' (M + N) \alpha = \sup_{\alpha \in \mathbb{S}_p^k} \alpha' M \alpha + \alpha' N \alpha \leq \phi_M(k) + \phi_N(k)$. On the other hand, $\phi_{\mathbb{E}_n[x_i x_i']}(k) \geq 1 \vee \phi_{N_{-1,-1}}(k)$ since the covariates contain an intercept. The result follows by using the bounds derived above.

The proof for an arbitrary ρ in Example 1 is similar. Since $-1 < \rho < 1$ is fixed, the bounds on the eigenvalues of the population design matrix Σ are given by $\sigma^2(k) = \sup_{\alpha \in \mathbb{S}_p^k} \alpha' \Sigma \alpha \leq (1+|\rho|)/(1-|\rho|)$ and $\underline{\sigma}^2(k) = \inf_{\alpha \in \mathbb{S}_p^k} \alpha' \Sigma \alpha \geq \frac{1}{2}(1-|\rho|)/(1+|\rho|)$. So we can apply Lemmas 24 and 25. To bound $\phi_M(k)$ we use comparison theorems, e.g. Corollary 3.12 of [28], which allows for the same bound as for the uncorrelated design to hold. \square

ACKNOWLEDGEMENTS

We would like to thank Arun Chandrasekhar, Denis Chetverikov, Moshe Cohen, Brigham Fradsen, Joonhwan Lee, Ye Luo, and Pierre-Andre Maugis for thorough proof-reading of several versions of this paper and their detailed comments that helped us considerably improve the paper. We also would like to thank Don Andrews, Alexandre Tsybakov, the editor Susan Murphy, the Associate Editor, and three anonymous referees for their comments that also helped us considerably improve the paper. We would also like to thank the participants of seminars in Brown University, CEMMAP Quantile Regression conference at UCL, Columbia University, Cowles Foundation Lecture at the Econometric Society Summer Meeting, Harvard-MIT, Latin American Meeting 2008 of the Econometric Society, Winter 2007 North American Meeting of the Econometric Society, London Business School, PUC-Rio, the Stats in the Chateau, the Triangle Econometrics Conference, and University College London.

REFERENCES

- [1] R. J. BARRO AND J.-W. LEE (1994). Data set for a panel of 139 countries, Discussion paper, NBER, <http://www.nber.org/pub/barro.lee>.
- [2] R. J. BARRO AND X. SALA-I-MARTIN (1995). *Economic Growth*. McGraw-Hill, New York.
- [3] A. BELLONI AND V. CHERNOZHUKOV (2008). On the Computational Complexity of MCMC-based Estimators in Large Samples, *The Annals of Statistics*, Volume 37, Number 4, 2011–2055.
- [4] A. BELLONI, V. CHERNOZHUKOV AND I. FERNANDEZ-VAL (2011). Conditional Quantile Processes based on Series or Many Regressors, arXiv:1105.6154.
- [5] A. BELLONI AND V. CHERNOZHUKOV (2009). Least Squares After Model Selection in High-dimensional Sparse Models, arXiv:1001.0188, forthcoming at Bernoulli.
- [6] D. BERTSIMAS AND J. TSITSIKLIS (1997). *Introduction to Linear Optimization*, Athena Scientific.
- [7] P. J. BICKEL, Y. RITOV AND A. B. TSYBAKOV (2009). Simultaneous analysis of Lasso and Dantzig selector, *Ann. Statist.* Volume 37, Number 4, 1705–1732.
- [8] M. BUCHINSKY (1994). Changes in the U.S. Wage Structure 1963-1987: Application of Quantile Regression *Econometrica*, Vol. 62, No. 2 (Mar.), pp. 405–458.
- [9] F. BUNEA, A. B. TSYBAKOV, AND M. H. WEGKAMP (2006). Aggregation and Sparsity via ℓ_1 Penalized Least Squares, in *Proceedings of 19th Annual Conference on Learning Theory (COLT 2006)* (G. Lugosi and H. U. Simon, eds.), *Lecture Notes in Artificial Intelligence* 4005 379-391. Springer, Berlin.
- [10] F. BUNEA, A. B. TSYBAKOV, AND M. H. WEGKAMP (2007). Aggregation for Gaussian regression, *The Annals of Statistics*, Vol. 35, No. 4, 1674-1697.
- [11] F. BUNEA, A. TSYBAKOV, AND M. H. WEGKAMP (2007). Sparsity oracle inequalities for the Lasso, *Electronic Journal of Statistics*, Vol. 1, 169-194.
- [12] E. CANDÈS AND T. TAO (2007). The Dantzig selector: statistical estimation when p is much larger than n . *Ann. Statist.* Volume 35, Number 6, 2313–2351.
- [13] V. CHERNOZHUKOV (2005). Extremal quantile regression. *Ann. Statist.* 33, no. 2, 806–839.
- [14] R. DUDLEY (2000). *Uniform Central Limit Theorems*, Cambridge Studies in advanced mathematics.
- [15] J. FAN AND J. LV (2008). Sure Independence Screening for Ultra-High Dimensional Feature Space, *Journal of the Royal Statistical Society Series B*, 70, 849–911.
- [16] X. FERNIQUE (1997). Fonctions aléatoires gaussiennes, vecteurs aléatoires gaussiens. *Publications du Centre de Recherches Mathématiques*, Montréal., *Ann. Probab.* Volume 8, Number 2, 252-261.
- [17] S. GERMAN (1980). A Limit Theorem for the Norm of Random Matrices, *Ann. Probab.* Volume 8, Number 2, 252-261.

- [18] C. GUTENBRUNNER AND J. JUREČKOVÁ (1992). Regression Rank Scores and Regression Quantiles The Annals of Statistics, Vol. 20, No. 1 (Mar.), pp. 305–330.
- [19] X. HE AND Q.-M. SHAO (2000). On Parameters of Increasing Dimenions, Journal of Multivariate Analysis, 73, 120–135.
- [20] K. KNIGHT (1998). Limiting distributions for L_1 regression estimators under general conditions, Annals of Statistics, 26, no. 2, 755–770.
- [21] K. KNIGHT AND FU, W. J. (2000). Asymptotics for Lasso-type estimators. Ann. Statist. 28 1356–1378.
- [22] R. KOENKER (2005). Quantile regression, Econometric Society Monographs, Cambridge University Press.
- [23] R. KOENKER (2010). Additive models for quantile regression: model selection and confidence bandaids, working paper, <http://www.econ.uiuc.edu/~roger/research/bandaids/bandaids.pdf>.
- [24] R. KOENKER AND G. BASSET (1978). Regression Quantiles, Econometrica, Vol. 46, No. 1, January, 33–50.
- [25] R. KOENKER AND J. MACHADO (1999). Goodness of fit and related inference process for quantile regression Journal of the American Statistical Association, 94, 1296–1310.
- [26] V. KOLTCHINSKII (2009). Sparsity in penalized empirical risk minimization, Ann. Inst. H. Poincar Probab. Statist. Volume 45, Number 1, 7–57.
- [27] P.-S. LAPLACE (1818). *Théorie analytique des probabilités*. Éditions Jacques Gabay (1995), Paris.
- [28] M. LEDOUX AND M. TALAGRAND (1991). Probability in Banach Spaces (Isoperimetry and processes). Ergebnisse der Mathematik und ihrer Grenzgebiete, Springer-Verlag.
- [29] R. LEVINE AND D. RENELT (1992). A Sensitivity Analysis of Cross-Country Growth Regressions, The American Economic Review, Vol. 82, No. 4, pp. 942–963.
- [30] K. LOUNICI, M. PONTIL, A. B. TSYBAKOV, AND S. VAN DE GEER (2009). Taking Advantage of Sparsity in Multi-Task Learning, arXiv:0903.1468v1 [stat.ML].
- [31] L. LOVÁSZ AND S. VEMPALA(2007). The geometry of logconcave functions and sampling algorithms, Random Structures and Algorithms, Volume 30 Issue 3, pages 307–358.
- [32] N. MEINSHAUSEN AND B. YU (2009). Lasso-type recovery of sparse representations for high-dimensional data, The Annals of Statistics, Vol. 37, No. 1, 246–270.
- [33] S. PORTNOY (1991). Asymptotic behavior of regression quantiles in nonstationary, dependent cases. J. Multivariate Anal. 38, no. 1, 100–113.
- [34] S. PORTNOY AND R. KOENKER (1997). The Gaussian hare and the Laplacian tortoise: computability of squared-error versus absolute-error estimators. Statist. Sci. Volume 12, Number 4, 279–300.
- [35] M. ROSENBAUM AND A. B. TSYBAKOV (2008). Sparse recovery under matrix uncertainty, arXiv:0812.2818v1 [math.ST].
- [36] X. SALA-I-MARTIN(1997). I Just Ran Two Million Regressions, The American Economic Review, Vol. 87, No. 2, pp. 178–183.
- [37] R. TIBSHIRANI (1996). Regression shrinkage and selection via the Lasso. J. Roy. Statist. Soc. Ser. B, 58, 267–288.
- [38] A. W. VAN DER VAART (1998). Asymptotic Statistics, Cambridge Series in Statistical and Probabilistic Mathematics.
- [39] S. A. VAN DE GEER (2008). High-dimensional generalized linear models and the Lasso, Annals of Statistics, Vol. 36, No. 2, 614–645.
- [40] A. W. VAN DER VAART AND J. A. WELLNER (1996). Weak Convergence and Empirical Processes, Springer Series in Statistics.
- [41] C.-H. ZHANG AND J. HUANG (2008). The sparsity and bias of the Lasso selection in high-dimensional linear regression. Ann. Statist. Volume 36, Number 4, 1567–1594.