

Error estimation and adaptive tuning for unregularized robust M-estimator

Pierre C. Bellec

Department of Statistics

Rutgers University

Piscataway, NJ 08854, USA

PIERRE.BELLECC@RUTGERS.EDU

Takuya Koriyama

Booth School of Business

The University of Chicago

Chicago, IL 60637, USA

TKORIYAM@UCHICAGO.EDU

Editor: Benjamin Guedj

Abstract

We consider unregularized robust M-estimators for linear models under Gaussian design and heavy-tailed noise, in the proportional asymptotics regime where the sample size n and the number of features p are both increasing such that $p/n \rightarrow \gamma \in (0, 1)$. An estimator of the out-of-sample error of a robust M-estimator is analyzed and proved to be consistent for a large family of loss functions that includes the Huber loss. As an application of this result, we propose an adaptive tuning procedure of the scale parameter $\lambda > 0$ of a given loss function ρ : choosing $\hat{\lambda}$ in a given interval I that minimizes the out-of-sample error estimate of the M-estimator constructed with loss $\rho_\lambda(\cdot) = \lambda^2 \rho(\cdot/\lambda)$ leads to the optimal out-of-sample error over I . The proof relies on a smoothing argument: the unregularized M-estimation objective function is perturbed, or smoothed, with a Ridge penalty that vanishes as $n \rightarrow +\infty$, and shows that the unregularized M-estimator of interest inherits properties of its smoothed version.

Keywords: Robust regression, proportional regime, Huber loss, adaptive tuning, high-dimensional statistics.

1. Introduction

Robust statistics originated with the foundational work of Huber (1964), which introduced methods for estimating a location parameter under heavy-tailed noise assumptions. Over time, it has evolved into a crucial area of study, providing tools to address the challenges posed by outliers and non-standard error distributions. The practical applications of robust statistics span diverse fields, including financial modeling (Lambert-Lacroix and Zwald (2011)) and genomic data analysis (Sun et al. (2020)), underscoring its broad utility across scientific and applied domains. For a comprehensive overview of the theoretical foundations and their practical implementations in data analysis, see Loh (2024), Maronna et al. (2019), and references therein.

In this paper, we consider the linear model $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\epsilon}$ where the design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ has i.i.d rows $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}_p, \boldsymbol{\Sigma})$ and the noise $\boldsymbol{\epsilon} \in \mathbb{R}^n$ has an i.i.d. marginal F_ϵ . Our assumption (Assumption 2 below) allows F_ϵ to be heavy-tailed, including distributions

with no finite moments. We focus on the high-dimensional regime where the sample size n and the dimension p are both increasing such that $p/n \rightarrow \gamma \in (0, 1)$. In this setting, we consider the unregularized robust M-estimator

$$\hat{\boldsymbol{\beta}}(\mathbf{y}, \mathbf{X}) \in \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \rho(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}), \quad (1)$$

where $\rho : \mathbb{R} \rightarrow \mathbb{R}$ is a convex and differentiable loss such that its derivative $\psi = \rho' : \mathbb{R} \rightarrow \mathbb{R}$ is bounded and Lipschitz. A well-studied example is the Huber loss, which is defined by

$$\forall x \in \mathbb{R}, \quad \rho(x) = \int_0^{|x|} \min(1, u) du = \begin{cases} x^2/2 & |x| \leq 1 \\ |x| - 1/2 & |x| \geq 1 \end{cases}. \quad (2)$$

Our goal is to select a robust loss achieving a small out-of-sample error in a data-driven manner, i.e., by only looking at the observation (\mathbf{y}, \mathbf{X}) . Here, the out-of-sample error refers to the random quantity

$$R := \|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}}(\mathbf{y}, \mathbf{X}) - \boldsymbol{\beta}^*)\|_2^2 = \mathbb{E} \left[\{\mathbf{x}_0^\top (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\}^2 \mid \mathbf{X}, \mathbf{y} \right], \quad (3)$$

where \mathbf{x}_0 is independent of (\mathbf{y}, \mathbf{X}) and has the same law as any row \mathbf{x}_i of \mathbf{X} . In this paper, we occasionally refer to the random quantity (3) simply as the *Risk*. Since the calculation of the out-of-sample error requires the unobservable quantities $(\boldsymbol{\beta}^*, \boldsymbol{\Sigma})$, we need an observable proxy to the out-of-sample error for such data-driven selection of loss.

Classical statistics for M-estimation, for instance confidence intervals for the coefficients of $\boldsymbol{\beta}^*$, require that $p/n \rightarrow 0$, that is, the dimension is negligible compared to the sample size. However, the last decade has identified ample evidence that such classical asymptotic results fail in moderate dimension where dimension and sample size are of the same order (Donoho and Montanari, 2016; El Karoui et al., 2013). Beyond linear models, this phenomenon is also clearly exposed in logistic regression Sur and Candès (2019) with clear departure from classical results as soon as $n \geq 10p$ is violated. From the point of view of classical statistics, a peculiar property of the proportional asymptotic regime where $\lim p/n$ is a positive constant is that *consistency fails*, and that the error $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2$ does not converge to 0 in probability.

1.1 Results at a glance

Our contribution is to propose a consistent estimator of the out-of-sample error for unregularized robust M-estimators. For $\psi = \rho' : \mathbb{R} \rightarrow \mathbb{R}$, let $\psi(\mathbf{r}) = (\psi(r_i))_{i=1}^n \in \mathbb{R}^n$ for any vector $\mathbf{r} \in \mathbb{R}^n$, in other words $\psi : \mathbb{R} \rightarrow \mathbb{R}$ acts componentwise on vectors. Then, under some regularity condition on the loss and noise distribution (see Assumption 1-2), we show that the random quantity

$$\hat{R} = \frac{p \|\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})\|_2^2}{\{\text{tr}[\mathbf{V}]\}^2} \quad \text{with } \mathbf{V} := \frac{\partial \psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})}{\partial \mathbf{y}} = \left(\frac{\partial \psi(y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}})}{\partial y_j} \right)_{i,j} \in \mathbb{R}^{n \times n} \quad (4)$$

is consistent as an estimate of the out-of-sample error, in the sense that

$$\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|_2^2 = \hat{R} + o_P(1). \quad (5)$$

The matrix \mathbf{V} in (4) is the Jacobian of the map $\mathbf{y} \mapsto \psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})$, and its existence will be proved with probability one in Proposition 3. Note that the random quantity \hat{R} in (4) is observable since both of $\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})$ and $\text{tr}[\mathbf{V}]$ can be computed by observation of (\mathbf{y}, \mathbf{X}) only. See Figure 1a for simulation. The formal statement of the consistency result (5) is in Theorem 1 below.

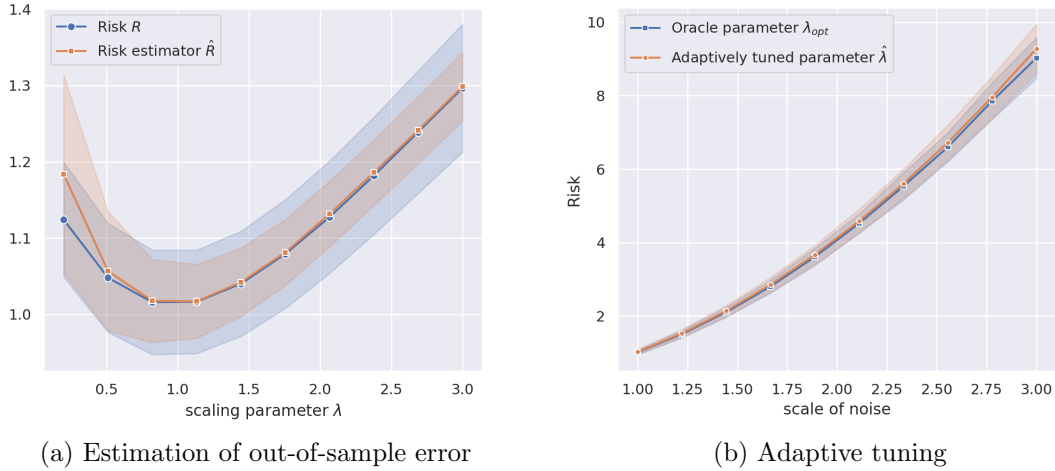


Figure 1: Figure 1a is the plot of the out-of-sample error R and its estimator \hat{R} with the Huber loss for different scaling parameter $\lambda > 0$. Figure 1b is the plot of the oracle out-of-sample error $R(\lambda_{\text{opt}})$ and the out-of-sample error $R(\hat{\lambda})$ with $\hat{\lambda}$ being the minimizer of the estimator \hat{R} among a finite grid I , as the scale of noise changes. See Section 4 for the details.

A first application of the above result in statistical inference is the ability to choose, between two competing loss functions ρ and $\tilde{\rho}$, the loss function yielding the M-estimator $\hat{\boldsymbol{\beta}}$ enjoying the narrowest confidence interval for a component β_j^* of $\boldsymbol{\beta}^*$. Lemma 1 in El Karoui et al. (2013) establishes that $\sqrt{p}(\beta_j^* - \hat{\beta}_j)/\|\boldsymbol{\Sigma}^{1/2}(\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}})\|$ is asymptotically normal with variance only depending on the j -th diagonal element of $\boldsymbol{\Sigma}^{-1}$. Thus, the size of the confidence interval is proportional $\|\boldsymbol{\Sigma}^{1/2}(\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}})\|$, and the ability to choose a loss function among $\{\rho, \tilde{\rho}\}$ leading to the smallest error $\|\boldsymbol{\Sigma}^{1/2}(\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}})\|$ through (5) will lead to the smallest confidence interval.

As a second application of the above result, we propose an adaptive tuning procedure of the scale parameter λ in a collection of loss functions $\{\rho_\lambda(\cdot) = \lambda^2 \rho(\cdot/\lambda), \lambda \in I\}$, for some fixed interval $I \subset (0, \infty)$ and some robust loss ρ . Note that this hyperparameter λ controls the sensitivity of the loss ρ_λ to outliers. For each $\lambda > 0$, let $\hat{\boldsymbol{\beta}}(\lambda)$ be the M-estimator (1) with $\rho = \rho_\lambda$ and let $\hat{R}(\lambda)$ be the corresponding estimate of the out-of-sample error (4). Then, we show that there exists a finite subset $J \subset I$ (a finite grid made of regularly spaced parameters) such that

$$\hat{\lambda} \in \arg \min_{\lambda \in J} \hat{R}(\lambda) \quad \text{satisfies} \quad \|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}}(\hat{\lambda}) - \boldsymbol{\beta}^*)\|^2 \leq \|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}}(\lambda_{\text{opt}}) - \boldsymbol{\beta}^*)\|^2 + o_P(1) \quad (6)$$

where λ_{opt} is the oracle parameter achieving the smallest out-of-sample error in probability in the interval I . See Figure 1b for simulation. The formal statement of the result (6) is in Theorem 3 below.

1.2 Related work

Under regularity conditions on the loss ρ and noise distribution F_ϵ , the out-of-sample error of the unregularized M-estimator (1) is known El Karoui et al. (2013); El Karoui (2018, 2013); Donoho and Montanari (2016); Thrampoulidis et al. (2018) to converge in probability to a deterministic value α^2 , which is a solution to the following nonlinear system of equations with positive unknowns (α, κ) :

$$\begin{aligned}\alpha^2\gamma &= \mathbb{E}[\left((\alpha Z + W) - \text{prox}[\kappa\rho](\alpha Z + W)\right)^2] \\ \alpha\gamma &= \mathbb{E}[\left((\alpha Z + W) - \text{prox}[\kappa\rho](\alpha Z + W)\right) \cdot Z]\end{aligned}\tag{7}$$

where $Z \sim \mathcal{N}(0, 1)$, $W \sim F_\epsilon$, and $\text{prox}[f](x) := \arg \min_{u \in \mathbb{R}} (x - u)^2/2 + f(u)$. The convergence in probability of the out-of-sample error of $\hat{\beta}$ to α^2 is granted by Donoho and Montanari (2016); Thrampoulidis et al. (2018) provided that (7) admits a unique solution. The existence of solutions to (7) was established by Donoho and Montanari (2016) for strongly convex losses, and in our recent paper Bellec and Koriyama (2023) for general Lipschitz losses on the side of the phase transition where exact recovery (i.e., $\hat{\beta} = \beta^*$) is not possible. On the other hand, uniqueness is addressed in Thrampoulidis et al. (2018) by strict convexity arguments and in Bellec and Koriyama (2023). The working assumptions of the present paper stated in Section 2 ensure that a solution (α, κ) to (7) exists and is unique, so that the convergence in probability $\|\Sigma^{1/2}(\hat{\beta} - \beta^*)\| \rightarrow^p \alpha$ as $n, p \rightarrow +\infty$ with $p/n \rightarrow \gamma < 1$ holds by the main result of Thrampoulidis et al. (2018).

Bean et al. (2013) studied the construction of the optimal loss that minimizes the solution α of (7) when F_ϵ is log-concave and known. However, constructing the optimal loss in this way requires the knowledge of F_ϵ to minimize the corresponding α in (7). Since F_ϵ is typically unknown, this construction cannot be implemented in practice.

In this paper, we focus on the underparametrized regime $p/n \rightarrow \gamma$ with $\gamma \in (0, 1)$. In the overparametrized regime with $\gamma > 1$, the out-of-sample error of the unregularized M-estimator (1) explodes (cf. (Thrampoulidis et al., 2018, Remark 5.1.1)). This curse of dimensionality can be overcome by using a penalty function g and computing the corresponding penalized M-estimator $\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} n^{-1} \sum_{i=1}^n \rho(y_i - \mathbf{x}_i^\top \beta) + g(\beta)$, for which the out-of-sample error will be finite under suitable assumptions on the penalty g Bayati and Montanari (2011); Thrampoulidis et al. (2018); Loureiro et al. (2021). If $\gamma < 1$, however, some advantages of the unregularized M-estimator(1) include its invariance with respect to (Σ, β^*) (El Karoui et al., 2013, Lemma 1) and that confidence intervals for components β_j^* of β^* using the asymptotic normality of $\hat{\beta}_j$ do not require knowledge of Σ ; on the other hand, the de-debiasing correction $\mathbf{e}_j^\top \Sigma^{-1/2} \mathbf{X}^\top \psi(\mathbf{y} - \mathbf{X}\hat{\beta})$ necessary for asymptotic normality of a regularized M-estimator with penalty $g(\cdot)$ as above does require knowledge (or some estimate) of Σ (Bellec et al., 2022).

Most similar works to (5) are Bellec (2023); Bellec and Shen (2022); Bellec (2025); they show that for general pairs of loss ρ and penalty g , the out-of-sample error of the penalized

M-estimator $\hat{\boldsymbol{\beta}} \in \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} n^{-1} \sum_{i=1}^n \rho(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}) + g(\boldsymbol{\beta})$ enjoys the approximation

$$\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|_2^2 \approx \text{tr}[\mathbf{V}]^{-2} (\|\hat{\boldsymbol{\psi}}\|_2^2 (2\hat{\text{df}} - p) + \|\boldsymbol{\Sigma}^{-1/2} \mathbf{X}^\top \hat{\boldsymbol{\psi}}\|_2^2), \quad (8)$$

where $\hat{\boldsymbol{\psi}} := \boldsymbol{\psi}(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \in \mathbb{R}^n$, $\hat{\text{df}} := \text{tr}[(\partial/\partial \mathbf{y})(\mathbf{X}\hat{\boldsymbol{\beta}})]$, and \mathbf{V} as in (4). Furthermore, these derivatives have closed forms for a certain choice of (ρ, g) ; for instance when ρ is the Huber loss and g is the L_1 penalty, we have $\hat{\text{df}} = |\hat{S}|$ and $\text{tr}[\mathbf{V}] = |\hat{I}| - |\hat{S}|$, where $\hat{S} = \{j \in [p] : \hat{\beta}_j \neq 0\}$ is the active set and $\hat{I} = \{i \in [n] : |y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}}| \leq 1\}$ is the set of inliers (Bellec, 2023, Propositions 2.2 and 2.3).

With the KKT conditions of (1) giving $\mathbf{X}^\top \hat{\boldsymbol{\psi}} = \mathbf{0}_p$, we may regard (4) as a special case of (8) by setting $\hat{\text{df}} = p$. However, the proof techniques in the aforementioned papers Bellec (2023); Bellec and Shen (2022); Bellec (2025) rely either on the strong convexity of the penalty g (excluding the $g = 0$ case we study here), or on a sufficiently sparse structure in $\boldsymbol{\beta}^*$ in which case the L1 penalty also allows for approximations of the form (8) Celentano et al. (2023); Bellec (2023).

For the unregularized case (1) studied here, the results from the papers El Karoui et al. (2013); El Karoui (2018, 2013); Donoho and Montanari (2016); Thrampoulidis et al. (2018) are not sufficient to establish the approximation (5), say for the Huber loss, on the one hand because El Karoui et al. (2013); El Karoui (2018, 2013); Donoho and Montanari (2016) rely on strong convexity on either the loss or the penalty, but more importantly because these results do not study the trace of the Jacobian $\text{tr}[\mathbf{V}]$ in (4) and how this quantity relates to the solution (α, κ) of the nonlinear system (7). For instance, the Convex Gaussian Min-max Theorem (CGMT) of Thrampoulidis et al. (2018) provides the limit in probability of $\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|^2$, and reversing the argument (replacing the design matrix by its transpose) provides the limit in probability of $\|\boldsymbol{\psi}\|^2/n$ in (1), which equals $\alpha^2 \gamma / \kappa^2$. However, characterizing a limit in probability for the quantity $\text{tr}[\mathbf{V}]$ has so far remained out of reach of the CGMT. The trace of the Jacobian $\text{tr}[\mathbf{V}]$ appears in the leave-one-out analysis of El Karoui et al. (2013); El Karoui (2013, 2018), however these works study the regularized estimate $\hat{\boldsymbol{\beta}} \in \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} n^{-1} \sum_{i=1}^n \rho(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}) + \mu \|\boldsymbol{\beta}\|^2/2$ with an additive Ridge penalty $\mu \|\boldsymbol{\beta}\|^2/2$. Recently, Bellec (2025) proposed an alternative regularization technique to study unregularized estimates of the form (1), but this crucially requires the loss ρ to be twice continuously differentiable and $\rho''(x) > 0$ for all $x \in \mathbb{R}$, which rules out the Huber loss (2), one of the most common robust loss functions and a major application of the present paper.

1.3 Precise analysis of a perturbed M-estimator

To overcome these difficulties and prove (5) for the Huber loss and its variants, we use the “*Ridge-smoothing*” technique, which is used in previous works El Karoui (2013); Celentano and Montanari (2022); Loureiro et al. (2022). By rotational and translation invariance, assume without loss of generality that $\boldsymbol{\beta}^* = \mathbf{0}_p$ and $\boldsymbol{\Sigma} = \mathbf{I}_p$. Then, (5) is reduced to

$$\|\hat{\boldsymbol{\beta}}\|_2^2 = p \|\boldsymbol{\psi}\|_2^2 \cdot (\text{tr}[(\partial/\partial \boldsymbol{\epsilon})\boldsymbol{\psi}])^{-2} + o_P(1) \quad (9)$$

where $\boldsymbol{\psi} : \boldsymbol{\epsilon} \in \mathbb{R}^n \mapsto \boldsymbol{\psi}(\boldsymbol{\epsilon} - \mathbf{X}\hat{\boldsymbol{\beta}}) \in \mathbb{R}^n$ is a vector field and $(\partial/\partial \boldsymbol{\epsilon})\boldsymbol{\psi}$ is the Jacobian matrix. To show (9), we instead consider the ridge-regularized M-estimator $\hat{\boldsymbol{\beta}}_{\text{ridge}}$ with a diminishing

regularization parameter n^{-c} for some constant $c > 0$

$$\hat{\boldsymbol{\beta}}_{\text{ridge}} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \rho(\epsilon_i - \mathbf{x}_i^\top \boldsymbol{\beta}) + \frac{n^{-c}}{2} \|\boldsymbol{\beta}\|_2^2, \quad (10)$$

and define the vector field $\boldsymbol{\psi}_{\text{ridge}} : \boldsymbol{\epsilon} \mapsto \boldsymbol{\psi}(\boldsymbol{\epsilon} - \mathbf{X} \hat{\boldsymbol{\beta}}_{\text{ridge}})$. We prove the target (9) by the following two steps:

- (I) Prove (9) for the regularized $\hat{\boldsymbol{\beta}}_{\text{ridge}}$, i.e., $\|\hat{\boldsymbol{\beta}}_{\text{ridge}}\|_2^2 \approx p \|\boldsymbol{\psi}_{\text{ridge}}\|_2^2 \cdot (\text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}_{\text{ridge}}])^{-2}$.
- (II) Prove that each quantity in (9) is approximately the same for $\hat{\boldsymbol{\beta}}_{\text{ridge}}$ and for $\hat{\boldsymbol{\beta}}$, as the Ridge penalty coefficient in (10) converges to 0 polynomially in n :

$$\|\hat{\boldsymbol{\beta}}_{\text{ridge}}\|_2^2 \approx \|\hat{\boldsymbol{\beta}}\|_2^2, \quad \frac{1}{n} \|\boldsymbol{\psi}_{\text{ridge}}\|_2^2 \approx \frac{1}{n} \|\boldsymbol{\psi}\|_2^2, \quad \frac{1}{n} \text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}_{\text{ridge}}] \approx \frac{1}{n} \text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}].$$

We prove (I) by a chi-square type moment inequality given in (Bellec, 2023, Section 7). The more subtle and novel part of the proof is to derive the three approximations in (II), in particular the challenging approximation

$$\text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}_{\text{ridge}}] \approx \text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}]. \quad (11)$$

Indeed, the closeness of the two vector fields $(\boldsymbol{\psi}, \boldsymbol{\psi}_{\text{ridge}})$ in the Euclidean norm does not necessarily imply the closeness of their divergence. To show (11), we leverage the assumption that the noise distribution is sufficiently smooth: concretely, Assumption 2 below grants that the noise distribution is a convolution of two probability distributions (F, \tilde{F}) , that is,

$$\boldsymbol{\epsilon} = \mathbf{z} + \boldsymbol{\delta}, \quad \mathbf{z} \perp\!\!\!\perp \boldsymbol{\delta}, \quad (z_i)_{i=1}^n \sim F, \quad (\delta_i)_{i=1}^n \sim \tilde{F}, \quad (12)$$

where F has density $z \mapsto \exp(-\phi(z))$ such that $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is twice continuously differentiable with bounded second derivative (on the other hand, \tilde{F} is unrestricted and may have arbitrarily fat tails). Then, using a variant of the second order Stein's formula (Bellec and Zhang, 2021, Section 2.4) extended to the distribution with density $z \mapsto \exp(-\phi(z))$ (see Theorem 6 below), we will argue in Lemma 5 below that $\mathbb{E}[(\text{tr}[(\partial/\partial \mathbf{z})(\boldsymbol{\psi}_{\text{ridge}} - \boldsymbol{\psi})] - \phi'(\mathbf{z})^\top (\boldsymbol{\psi}_{\text{ridge}} - \boldsymbol{\psi}))^2]$ is relatively small. As a consequence, using the chain rule for the Jacobian, we have

$$\text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}_{\text{ridge}}] - \text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}] = \text{tr}[(\partial/\partial \mathbf{z})(\boldsymbol{\psi}_{\text{ridge}} - \boldsymbol{\psi})] \approx \phi'(\mathbf{z})^\top (\boldsymbol{\psi}_{\text{ridge}} - \boldsymbol{\psi}). \quad (13)$$

Roughly speaking, Assumption 2 below grants that the noise distribution is sufficiently smooth, and thanks to this assumption, closeness of the vector fields in the form $\|\boldsymbol{\psi} - \boldsymbol{\psi}_{\text{ridge}}\|_2^2/n = o_P(1)$ carries over to the two divergences, and implies

$$\text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}]/n = \text{tr}[(\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}_{\text{ridge}}]/n + o_P(1). \quad (14)$$

It is not clear at this point if (14) holds without this smoothness assumption on the noise distribution.

Adaptive tuning of the scale parameter λ was investigated before in several papers, including Loh (2021); Wang et al. (2021) and references therein. These works do not assume a proportional asymptotics regime and do not aim for exact multiplicative constants: the resulting risk is only guaranteed to be less than $C \times$ (optimal risk) for a constant $C > 1$. On the other hand, our goal in the present paper is to achieve the optimal risk with multiplicative constant 1.

1.4 Organization

In Section 2, we derive several results on the consistency of an estimator of the out-of-sample error. In Section 3, we discuss the adaptive tuning of scale parameters. Section 4 is devoted to numerical simulations, and Section 5 gives an outline of the proof. The rigorous proofs are provided in appendix.

1.5 Notation

For any vector \mathbf{u} , we denote by $\|\mathbf{u}\|$ the Euclidean norm $\sqrt{\sum_i u_i^2}$. For any matrix \mathbf{X} , let $\|\mathbf{X}\|_{op}$ be the operator norm, i.e., the maximum singular value of \mathbf{X} . Given a vector field $\mathbf{f} : \mathbf{z} \in \mathbb{R}^n \mapsto \mathbf{f}(\mathbf{z}) \in \mathbb{R}^m$, let $(\partial/\partial\mathbf{z})\mathbf{f} \in \mathbb{R}^{m \times n}$ be the Jacobian matrix, and $\|\mathbf{f}\|_{lip}$ be the Lipschitz constant of \mathbf{f} induced by the Euclidean norm. For any function $\psi : \mathbb{R} \rightarrow \mathbb{R}$, let $\psi(\mathbf{r}) = (\psi(r_i))_{i=1}^n \in \mathbb{R}^n$ for all $\mathbf{r} \in \mathbb{R}^n$, in other words $\psi : \mathbb{R} \rightarrow \mathbb{R}$ acts componentwise on vectors, and let $\|\psi\|_\infty = \sup_{x \in \mathbb{R}} |\psi(x)|$ be the sup norm. If we write $C = C_1(a, b, c)$, C is a constant depending on (a, b, c) only. If two random vectors \mathbf{x}, \mathbf{y} are independent, we write $\mathbf{x} \perp\!\!\!\perp \mathbf{y}$. For a sequence of random variables $(U_n)_{n \geq 1}$, we write $U_n \xrightarrow{p} U$ to denote convergence in probability to U and $U_n = o_P(1)$ if $U_n \xrightarrow{p} 0$. For a sequence of reals $r_n > 0$, we write $U_n = O_P(r_n)$ if for any $\epsilon > 0$ there exists K_ϵ such that $\sup_{n \geq 1} \mathbb{P}(|U_n| > K_\epsilon r_n) < \epsilon$

2. Estimation of the out-of-sample error

Throughout, we assume that $(y_i, \mathbf{x}_i, \epsilon_i)_{i=1}^n$ are independently distributed according to

$$\forall i \in [n], \quad y_i = \mathbf{x}_i^\top \boldsymbol{\beta}^* + \epsilon_i, \quad \mathbf{x}_i \sim \mathcal{N}(\mathbf{0}_p, \boldsymbol{\Sigma}), \quad \epsilon_i \sim F_\epsilon, \quad \mathbf{x}_i \perp\!\!\!\perp \epsilon_i$$

where $\boldsymbol{\beta}^* \in \mathbb{R}^p$ is an unknown regression vector, $\boldsymbol{\Sigma} \in \mathbb{R}^{p \times p}$ is some symmetric positive semi-definite matrix, and F_ϵ is a probability distribution. Since our interest is the out-of-sample error $\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_*)\|^2$, and the out-of-sample error for the unregularized M-estimator $\hat{\boldsymbol{\beta}} \in \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \sum_{i=1}^n \rho(y_i - \mathbf{x}_i^\top \boldsymbol{\beta})$ is invariant with respect to $(\boldsymbol{\beta}^*, \boldsymbol{\Sigma})$ (cf. El Karoui (2013) or Section 5), we do not require any assumptions on $(\boldsymbol{\beta}^*, \boldsymbol{\Sigma})$.

When taking a limit as $n \rightarrow \infty$, we implicitly assume that the number of features p and the sample size n are increasing such that $p/n \rightarrow \gamma \in (0, 1)$, while other quantities such ρ and F_ϵ are fixed. In this setting, it has been shown by El Karoui et al. (2013); El Karoui (2018, 2013); Donoho and Montanari (2016); Thrampoulidis et al. (2018) that if the nonlinear system of equations

$$\begin{aligned} \alpha^2 \gamma &= \mathbb{E}[\left((\alpha Z + W) - \text{prox}[\kappa \rho](\alpha Z + W)\right)^2] \\ \alpha \gamma &= \mathbb{E}[\left((\alpha Z + W) - \text{prox}[\kappa \rho](\alpha Z + W)\right) \cdot Z] \end{aligned} \quad \text{where } \begin{cases} Z \sim \mathcal{N}(0, 1), \\ W \sim F_\epsilon, Z \perp\!\!\!\perp W \end{cases} \quad (15)$$

admits a unique solution (α, κ) , the out-of-sample error $\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_*)\|^2$ converges to α^2 in probability. The existence of solutions to (15) was established in Donoho and Montanari (2016) for strongly convex ρ , and recently in our companion paper Bellec and Koriyama (2023) (see Appendix A for details).

In this section, we will argue that the random quantity \hat{R} defined by

$$\hat{R} = \frac{p \|\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})\|_2^2}{\text{tr}[\mathbf{V}]^2}, \quad \text{where } \mathbf{V} := \frac{\partial \psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})}{\partial \mathbf{y}} \in \mathbb{R}^{n \times n}, \quad \psi = \rho' : \mathbb{R} \rightarrow \mathbb{R}, \quad (16)$$

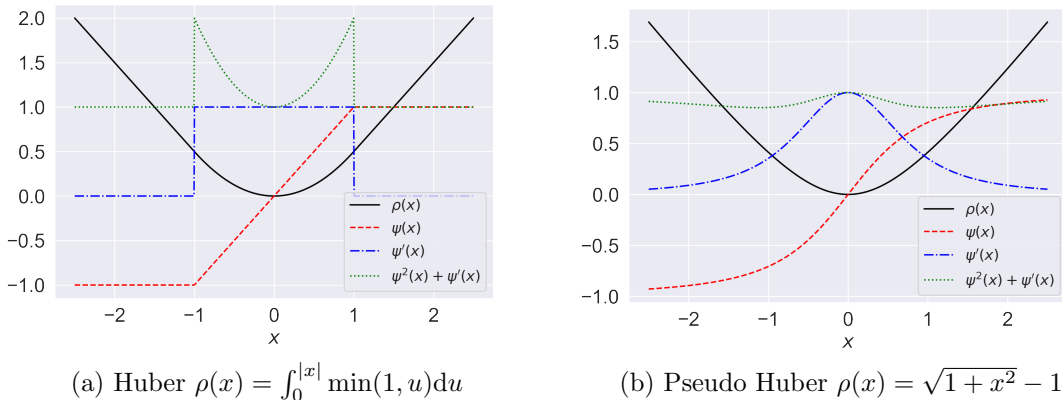


Figure 2: Example of loss satisfying Assumption 1

approximates well the out-of-sample error $\|\Sigma^{1/2}(\hat{\beta} - \beta_*)\|^2$. Note that $\mathbf{V} \in \mathbb{R}^{n \times n}$ is the Jacobian matrix of the map $\mathbf{y} \in \mathbb{R}^n \mapsto \psi(\mathbf{y} - \mathbf{X}\hat{\beta}) \in \mathbb{R}^n$. Before moving to the formal statement, we introduce our assumption on the loss ρ and noise distribution F_ϵ .

Assumption 1 ρ is convex, differentiable, and $\{0\} = \arg \min_x \rho(x)$, as well as $\psi = \rho' : \mathbb{R} \rightarrow \mathbb{R}$ satisfies

1. $\|\psi\|_\infty < +\infty$.
2. $\|\psi\|_{\text{lip}} = 1$.
3. $\exists \eta > 0$ such that $\psi(x)^2 / \|\psi\|_\infty^2 + \psi'(x) \geq \eta$ for almost every $x \in \mathbb{R}$.

Assumption 2 F_ϵ is a convolution of (F, \tilde{F}) , where \tilde{F} is arbitrary while F has some density $z \mapsto \exp(-\phi(z))$ for some twice continuously differentiable function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ with a bounded second derivative $\sup_{x \in \mathbb{R}} |\phi''(x)| < +\infty$.

The Huber loss $\rho(x) = \int_0^{|x|} \min(1, u) du$ and its smooth approximations, such as the pseudo Huber loss $\rho(x) = \sqrt{1 + x^2} - 1$, satisfy Assumption 1 (See Figure 2). Assumption 2 requires that each component ϵ_i of the noise in the linear model is equal in distribution to $z_i + \delta_i$ for two independent random variables $z_i \sim F$ and $\delta_i \sim \tilde{F}$, where z_i has a smooth density $z \mapsto \exp(-\phi(z))$. Typical example of the distribution F in Assumption 2 is the normal distribution $\mathcal{N}(0, \sigma^2)$ for some $\sigma > 0$. We emphasize that there is no assumption on \tilde{F} , so that the noise distribution $F_\epsilon = F * \tilde{F}$ can be heavy-tailed. For instance, we allow $F_\epsilon = \mathcal{N}(0, 1) * \text{Cauchy}(0, 1)$, in which F_ϵ has no finite moments.

Assumption 1(1) is the condition for Theorem 7 to be applicable; this is a sufficient condition to guarantee that the system (7) has a unique solution. Assumption 1(2) is mainly for the Jacobian \mathbf{V} of the map $\mathbf{y} \in \mathbb{R}^n \mapsto \psi(\mathbf{y} - \mathbf{X}\hat{\beta}) \in \mathbb{R}^n$ in (16) to be well-defined. Indeed, under this assumption, $\mathbf{y} \mapsto \psi(\mathbf{y} - \mathbf{X}\hat{\beta})$ is Lipschitz for almost every \mathbf{X} (cf. Proposition 3 or Proposition 4.1 in Bellec (2023)), so that Rademacher's theorem guarantees the existence of the Jacobian \mathbf{V} for almost every $(\mathbf{y}, \mathbf{X}) \in \mathbb{R}^n \times \mathbb{R}^{n \times p}$. The Lipschitz assumption in

Assumption 1(2) can be relaxed to $\|\psi\|_{\text{lip}} < +\infty$ since the unregularized M-estimator (1) remains invariant under a rescaling of the loss function, i.e., $\rho \mapsto \rho/\|\psi\|_{\text{lip}}$. For simplicity and to streamline the proof arguments presented in Section 5 and the Appendix, we assume $\|\psi\|_{\text{lip}} = 1$ throughout the paper. Assumption 1(3) is used to show that $\text{tr}[\mathbf{V}]$ in the denominator of \hat{R} in (16) is bounded from below by a positive constant times n with high probability.

Assumption 2 on F_ϵ is a technical condition to control $\text{tr}[\mathbf{V}]$; see the discussion surrounding (12). An equivalent formulation of Assumption 2 is that each entry ϵ_i of the noise is a sum $z_i + \delta_i$ of two independent random variables $z_i \sim F$ and $\delta_i \sim \tilde{F}$, where z_i has a smooth density $z \mapsto \exp(-\phi(z))$. Numerical simulations in Appendix D.1 suggest that our results still hold for some F_ϵ without the presence of the smooth noise component z_i , which suggests that it is an artifact of the proof. Some preliminary theoretical evidence that this assumption is not necessary is given in Lemma 6 where this assumption is replaced with $\epsilon_i = \sigma z_i + \delta_i$ with the same distributions as above but a vanishing amplitude σ for the smooth part z_i . However, our results in Section 3 require Assumption 2 (and cannot accommodate a vanishing σ) in Lemma 12.

The quantity $\text{tr}[\mathbf{V}]$ in (16) is observable and can be computed approximately by Monte Carlo schemes (see Section 2.11 in Bellec (2023) and the references therein) or off-the-shelf numerical methods to compute derivatives. For the Huber loss and the pseudo Huber loss, closed-form expressions for $\text{tr}[\mathbf{V}]$ are available: $\text{tr}[\mathbf{V}] = |\{i \in [n] : y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}} \in [-1, 1]\}| - p$ for the Huber loss and

$$\text{tr}[\mathbf{V}] = \sum_{i=1}^n \left[\psi'(\mathbf{x}_i^\top \hat{\boldsymbol{\beta}}) - \psi'(\mathbf{x}_i^\top \hat{\boldsymbol{\beta}})^2 \mathbf{x}_i^\top \left(\sum_{l=1}^n \mathbf{x}_l \psi'(\mathbf{x}_l^\top \hat{\boldsymbol{\beta}}) \mathbf{x}_l^\top \right)^{-1} \mathbf{x}_i \right]$$

for the pseudo Huber loss $\rho(x) = \sqrt{1+x^2} - 1$ or other twice-continuously differentiable loss functions with $\psi' = \rho''$ positive everywhere.

Remark 1 *Assumption 1-2 imply that the system (15) admits a unique solution. Indeed, according to Theorem 7, the sufficient conditions for the system to admit a unique solution are (1) ρ is convex, Lipschitz, and $\{0\} = \arg \min_{x \in \mathbb{R}} \rho(x)$, and (2) $\mathbb{P}(W \neq 0) > 0$ and $\inf_{\lambda > 0} \mathbb{E}[\text{dist}(G, \lambda \partial \rho(W))^2] > 1 - \gamma$ for independent $W \sim F_\epsilon$ and $G \sim \mathcal{N}(0, 1)$. Here, condition (1) is readily satisfied by Assumption 1. For condition (2), thanks to Assumption 2, F_ϵ does not have any point mass, so in particular $\mathbb{P}_{W \sim F_\epsilon}(W \neq 0) = 1$. Since ρ is differentiable by Assumption 1, $\partial \rho(W)$ is always the singleton $\{\rho'(W)\}$, which gives $\inf_{\lambda > 0} \mathbb{E}[\text{dist}(G, \lambda \partial \rho(W))^2] = \inf_{\lambda > 0} \mathbb{E}[(G - \lambda \rho'(W))^2] = 1 > 1 - \gamma$.*

Now we claim that the random quantity \hat{R} in (16) is a consistent estimate of the out-of-sample error $\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_*)\|^2$.

Theorem 1 *Assume that (ρ, F_ϵ) satisfy Assumption 1 and 2. Let $\psi = \rho' : \mathbb{R} \rightarrow \mathbb{R}$ be the derivative of the loss, and \mathbf{V} be the Jacobian matrix $(\partial/\partial \mathbf{y})\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \in \mathbb{R}^{n \times n}$. Then, as $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma \in (0, 1)$, we have*

$$\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|^2 = \hat{R} + o_P(1), \text{ where } \hat{R} := \frac{p \|\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})\|_2^2}{\text{tr}[\mathbf{V}]^2}. \quad (17)$$

An outline of the proof is given in Section 5 and the formal proof is given in Appendix B.4. We are able to relax Assumption 2 on the noise as the next proposition shows: the result (17) still holds if the smooth component of the noise has vanishing amplitude σ_n .

Proposition 1 *Let Assumption 1 be fulfilled, and assume that for each $i \in [n]$, the noise ϵ_i is equal in distribution to $\delta_i + \sigma_n z_i$ where $\delta_i \sim \tilde{F}$ and $z_i \sim F$ are independent and (F, \tilde{F}) are as in Assumption 2, and where $\sigma_n \rightarrow 0$ with $\sigma_n \geq n^{-1/8}$. Then (17) still holds as $n, p \rightarrow +\infty$ with $p/n \rightarrow \gamma \in (0, 1)$.*

The proof is given in Appendix B.5. Theorem 1 implies that the random quantity \hat{R} is a consistent estimate of the out-of-sample error. Importantly, $\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})$ and $\text{tr}[\mathbf{V}]$ are both observable, i.e., they can be computed by observed data (\mathbf{y}, \mathbf{X}) only. Thus, \hat{R} serves as a criterion to select different losses.

Corollary 1 *Assume either that F_ϵ satisfies Assumption 2, or that each iid component ϵ_i of the noise satisfies the relaxed condition in Proposition 1. For fixed integer K , consider K different loss functions ρ_1, \dots, ρ_K that satisfy Assumption 1. For each $k \in [K]$, let $\hat{\boldsymbol{\beta}}_k \in \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \sum_{i=1}^n \rho_k(y_i - \mathbf{x}_i^\top \boldsymbol{\beta})$ be the M-estimator computed by the k -th loss and \hat{R}_k be the corresponding criterion (16). Then, we have*

$$\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}}_{\hat{k}} - \boldsymbol{\beta}^*)\|^2 = \min_{k=1, \dots, K} \|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}^*)\|^2 + o_P(1), \quad \text{where } \hat{k} \in \arg \min_{k=1, \dots, K} \hat{R}_k.$$

See Appendix B.6 for the proof. Note that the left-hand side is the out-of-sample error of the M-estimator that minimizes the criterion $(\hat{R}_k)_{k=1}^K$ among K different losses, while $\min_{k=1, \dots, K} \|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}^*)\|^2$ on the right-hand side is the optimal out-of-sample error among the candidates. Thus, Corollary 1 implies that the M-estimator minimizing the criterion achieves the optimal out-of-sample error up to an error term that converges to 0 in probability.

3. Adaptive tuning of scale parameters

Let ρ be a fixed robust loss satisfying Assumption 1. For $\lambda > 0$, consider λ -scaled loss ρ_λ defined as

$$\forall x \in \mathbb{R}, \quad \rho_\lambda(x) := \lambda^2 \rho(x/\lambda).$$

This λ controls the sensitivity to outliers or heavy tails, and different λ lead to significantly different performance as seen in Figures 3a and 5a. When the base loss ρ is the Huber loss or the pseudo Huber loss (see Figure 2), it holds that

$$\forall x \in \mathbb{R}, \quad \rho_\lambda(x) = \begin{cases} x^2 \cdot (x/\lambda)^{-2} \rho(x/\lambda) & \rightarrow x^2/2 \quad \text{as } \lambda \rightarrow +\infty \\ \lambda x \cdot (x/\lambda)^{-1} \rho(x/\lambda) & \sim \lambda|x| \quad \text{as } \lambda \rightarrow 0+ \end{cases},$$

thanks to $\lim_{u \rightarrow 0} \rho(u)/u^2 = 1/2$ and $\lim_{u \rightarrow \pm\infty} \rho(u)/u = \pm 1$. Informally speaking, the scaled loss ρ_λ behaves like the square loss for large λ and like the absolute loss $|x|$ for small λ .

Previous sections have so far discussed the consistency of the estimate \hat{R} given by (16). In this section, we apply this consistency result to perform adaptive tuning of the parameter

λ . The goal of this section is to select some λ in a data-driven way so that the M-estimator $\hat{\beta}_\lambda$ defined with the scaled loss ρ_λ ,

$$\hat{\beta}_\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \rho_\lambda(y_i - \mathbf{x}_i^\top \beta) \quad \text{with} \quad \rho_\lambda(\cdot) := \lambda^2 \rho(\cdot/\lambda) \quad (18)$$

achieves an asymptotically optimal risk. Let us introduce some useful notation. For all $\lambda > 0$, let $\psi_\lambda : \mathbb{R} \rightarrow \mathbb{R}$ be the derivative of the λ -scaled loss ρ_λ :

$$\psi_\lambda : x \mapsto \rho'_\lambda(x) = \lambda \psi(x/\lambda) \quad \text{with} \quad \psi = \rho'.$$

Note that by construction, $\|\psi_\lambda\|_{\text{lip}}$ is invariant with respect to λ . We define the random functions R , \hat{R} , and the deterministic function α as

$$R : (0, \infty) \rightarrow \mathbb{R}, \quad \lambda \mapsto \|\Sigma^{1/2}(\hat{\beta}_\lambda - \beta^*)\|^2 \quad \text{with} \quad \hat{\beta}_\lambda \text{ being the M-estimator (18)} \quad (19)$$

$$\hat{R} : (0, \infty) \rightarrow \mathbb{R}, \quad \lambda \mapsto p \frac{\|\psi_\lambda(\mathbf{y} - \mathbf{X}\hat{\beta}_\lambda)\|^2}{\text{tr}[\mathbf{V}_\lambda]^2} \quad \text{with} \quad \mathbf{V}_\lambda = \frac{\partial \psi_\lambda(\mathbf{y} - \mathbf{X}\hat{\beta}_\lambda)}{\partial \mathbf{y}} \in \mathbb{R}^{n \times n} \quad (20)$$

$$\alpha : (0, \infty) \rightarrow \mathbb{R}, \quad \lambda \mapsto \alpha(\lambda) \quad (\text{the solution to (15) with } \rho(\cdot) = \rho_\lambda(\cdot)). \quad (21)$$

With the above notation, Donoho and Montanari (2016); Thrampoulidis et al. (2018) grants $R(\lambda) \rightarrow^p \alpha^2(\lambda)$ while Theorem 1 and Theorem 7 yield $\hat{R}(\lambda) \rightarrow^p \alpha^2$ and an explicit upper bound on $\alpha^2(\lambda)$ with respect to λ . We summarize these results in the following proposition.

Proposition 2 *Assume that (ρ, F_ϵ) satisfy Assumption 1 and 2. Then, the nonlinear system of equations (15) with $\rho = \rho_\lambda$ admits a unique solution for all $\lambda > 0$, so that the map $\lambda \mapsto \alpha^2(\lambda)$ in (21) is well-defined. Furthermore, as $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma \in (0, 1)$, we have*

$$\hat{R}(\lambda) \rightarrow^p \alpha^2(\lambda), \quad R(\lambda) \rightarrow^p \alpha^2(\lambda), \quad \alpha^2(\lambda) \leq C_2(\gamma, F_\epsilon, \rho)(\lambda^2 + 1).$$

for all $\lambda > 0$.

See Appendix C.1 for the proof. Proposition 2 suggests that the M-estimator $\hat{\beta}_{\hat{\lambda}}$ in (18) with $\hat{\lambda}$ minimizing the criterion $\hat{R}(\lambda)$ over a discrete grid achieves the optimal risk limit, i.e.,

$$\alpha^2(\hat{\lambda}) \approx \min_{\lambda} \alpha^2(\lambda) \quad \text{with} \quad \hat{\lambda} \in \arg \min_{\lambda} \hat{R}(\lambda),$$

as long as the map $\lambda \mapsto \alpha^2(\lambda)$ is smooth enough and the grid fine enough, so that at least one element λ of the grid is sufficiently close to the optimal parameter (that may for instance fall between two consecutive elements of the grid). To ensure such smoothness condition on the function $\alpha^2(\cdot)$, we introduce an additional assumption on the loss ρ via $\psi = \rho'$.

Assumption 3 $\psi = \rho' : \mathbb{R} \rightarrow \mathbb{R}$ satisfies the following:

$$\sup_{\lambda, \tilde{\lambda} > 0, \lambda \neq \tilde{\lambda}} \sup_{x \in \mathbb{R}} \frac{|\lambda \psi(x/\lambda) - \tilde{\lambda} \psi(x/\tilde{\lambda})|}{|\lambda - \tilde{\lambda}|} < +\infty.$$

Note that the Huber loss $\rho(x) = \int_0^{|x|} \min(1, u) du$ satisfies Assumption 3. If ρ is twice continuously differentiable, the sufficient condition for Assumption 3 is

$$\sup_{u \in \mathbb{R}} |\psi(u) - u\psi'(u)| < +\infty \quad (22)$$

Indeed, for all $\lambda, \tilde{\lambda} \in (0, \infty)$ with $\lambda \neq \tilde{\lambda}$ and for all $x \in \mathbb{R}$, there exists $\lambda_x > 0$ by the mean value theorem, such that

$$\frac{|\lambda\psi(x/\lambda) - \tilde{\lambda}\psi(x/\tilde{\lambda})|}{|\lambda - \tilde{\lambda}|} = \left| \frac{d}{d\lambda} \left(\lambda\psi\left(\frac{x}{\lambda}\right) \right) \right|_{\lambda=\lambda_x} = \left| \psi\left(\frac{x}{\lambda_x}\right) - \frac{x}{\lambda_x}\psi'\left(\frac{x}{\lambda_x}\right) \right| \leq \sup_{u \in \mathbb{R}} |\psi(u) - u\psi'(u)|,$$

which implies that (22) is a sufficient condition for Assumption 3. We can easily check that (22) is satisfied by the pseudo Huber loss $\rho(x) = \sqrt{1+x^2} - 1$.

Now we will claim that the map $\lambda \mapsto \alpha^2(\lambda)$ in (21) is locally 1/2-Hölder continuous. See Appendix C.2 for the proof.

Theorem 2 *Assume that (ρ, F_ϵ) satisfy Assumption 1, 2, and 3. For each $\lambda > 0$, let $\alpha^2(\lambda)$ be the solution to the nonlinear system of equations (15) with $\rho(\cdot) = \lambda^2\rho(\cdot/\lambda)$. Then, the map $\lambda \mapsto \alpha^2(\lambda)$ is locally 1/2-Hölder continuous in the sense that*

$$\sup_{\lambda_{\min} \leq \lambda \leq \lambda_{\max}} \frac{|\alpha^2(\lambda) - \alpha^2(\tilde{\lambda})|}{|\lambda - \tilde{\lambda}|^{1/2}} \leq C_3(\lambda_{\min}, \lambda_{\max}, \rho, F_\epsilon, \gamma)$$

for all $\lambda_{\min}, \lambda_{\max} \in (0, \infty)$ with $\lambda_{\min} < \lambda_{\max}$.

Remark 2 *Similar result to Theorem 2 is known for the regularization parameter λ of the Lasso $\lambda\|\beta\|_1$ (see Celentano et al. (2023); Miolane and Montanari (2021)). This line of work shows a suitable smoothness of the associated nonlinear system with respect to λ , using the implicit function theorem. We proceed differently here for the proof of Theorem 2: we show the map*

$$\lambda \in (0, \infty) \mapsto \hat{R}(\lambda) = \frac{\|\psi_\lambda(\mathbf{y} - \mathbf{X}\hat{\beta}_\lambda)\|^2}{\text{tr}[\mathbf{V}_\lambda]^2}$$

is Hölder continuous. This Hölder continuity of $\hat{R}(\cdot)$ implies that of $\alpha^2(\cdot)$ thanks to the approximation $\alpha^2(\lambda) = \hat{R}(\lambda) + o_P(1)$ by Proposition 2. Hölder continuity of $\hat{R}(\cdot)$ follows if $\lambda \mapsto \|\psi_\lambda(\mathbf{y} - \mathbf{X}\hat{\beta}_\lambda)\|^2$ and $\lambda \mapsto \text{tr}[\mathbf{V}_\lambda]^2$ are both Hölder continuous. The proof in Appendix C.2 shows the Hölder continuity of the map $\lambda \mapsto \|\psi_\lambda(\mathbf{y} - \mathbf{X}\hat{\beta}_\lambda)\|^2$ using Assumption 3 and the KKT conditions. For the map $\lambda \mapsto \text{tr}[\mathbf{V}_\lambda]^2$, we show that $\text{tr}[\mathbf{V}_\lambda]$ inherits the Hölder continuity from $\psi_\lambda(\mathbf{y} - \mathbf{x}_i^\top \hat{\beta}_\lambda)$, by leveraging the assumption that the noise distribution is sufficiently smooth; see Theorem 6 and the discussion around (12) on how this smoothness assumption is also used for results in Section 2.

Finally, we apply Theorem 2 to the adaptive tuning of the scale parameter λ . Let us take λ_{\min} and λ_{\max} such that $0 < \lambda_{\min} < \lambda_{\max}$, and take the closed interval $I = [\lambda_{\min}, \lambda_{\max}]$. For this fixed interval I , we consider a finite grid I_N of $(N+1)$ points equispaced in log-scale, i.e.,

$$I_N := \left\{ \lambda_{\min} \left(\frac{\lambda_{\max}}{\lambda_{\min}} \right)^{i/N} : i = 0, \dots, N \right\} \subset I = [\lambda_{\min}, \lambda_{\max}]. \quad (23)$$

for all $N \in \mathbb{N}$. We define the oracle optimal parameter $\lambda_{\text{opt}} \in I$ and the data-driven selected parameter $\hat{\lambda}_N \in I_N$ as

$$\lambda_{\text{opt}} \in \arg \min_{\lambda \in I} \alpha^2(\lambda), \quad \hat{\lambda}_N \in \arg \min_{\lambda \in I_N} \hat{R}(\lambda).$$

Note that λ_{opt} exists thanks to the continuity of $\lambda \mapsto \alpha^2(\lambda)$ from Theorem 2 and the compactness of I . Here, λ_{opt} is the optimal scale parameter in the closed interval I , while $\hat{\lambda}_N$ is the minimizer of the criterion $\hat{R}(\lambda)$ among the finite grid I_N of $(N + 1)$ points. With the above notation, we claim that $\hat{\lambda}_N$ achieves the theoretically optimal risk limit $\alpha^2(\lambda_{\text{opt}}) = \min_{\lambda \in I} \alpha^2(\lambda)$.

Theorem 3 *Assume that F_ϵ satisfies Assumption 2, and ρ satisfies Assumption 1 and 3. Let R , \hat{R} , and α^2 be the maps defined by (19), (20), and (21), respectively. Then, for any closed interval $I = [\lambda_{\min}, \lambda_{\max}] \subset (0, \infty)$, there exists a sequence of integers $(N_n)_{n=1}^\infty = (N_n(\gamma, \rho, F_\epsilon, \lambda_{\min}, \lambda_{\max}))_{n=1}^\infty$ such that as $n, p \rightarrow +\infty$ with $p/n \rightarrow \gamma \in (0, 1)$, we have*

$$R(\hat{\lambda}_{N_n}) = \alpha^2(\hat{\lambda}_{N_n}) + o_P(1) = \alpha^2(\lambda_{\text{opt}}) + o_P(1),$$

where $\lambda_{\text{opt}} \in \arg \min_{\lambda \in I} \alpha^2(\lambda)$ and $\hat{\lambda}_N \in \arg \min_{\lambda \in I_N} \hat{R}(\lambda)$ with I_N given by (23).

See Appendix C for the proof. Theorem 3 implies that if the grid I_N is fine enough, the M-estimator $\hat{\beta}(\hat{\lambda})$ with $\hat{\lambda}$ minimizing the criteria $\hat{R}(\lambda)$ over the grid achieves the theoretically optimal risk $\alpha^2(\lambda_{\text{opt}}) = \min_{\lambda \in I} \alpha^2(\lambda)$ in the fixed interval I of λ . We will verify Theorem 3 by numerical simulations in Section 4.

4. Numerical simulations

We focus on the Huber loss $\rho(x) = \int_0^{|x|} \min(1, t) dt$ and consider the adaptive tuning of the scale parameter λ for the scaled loss $\rho_\lambda(\cdot) = \lambda^2 \rho(\cdot/\lambda)$. Note that the scaled loss ρ_λ and its derivative $\psi_\lambda := \rho'_\lambda$ can be written explicitly as

$$\forall x \in \mathbb{R}, \quad \rho_\lambda(x) := \begin{cases} x^2/2 & |x| \leq \lambda \\ \lambda|x| - \lambda^2/2 & |x| \geq \lambda \end{cases}, \quad \psi_\lambda(x) = \begin{cases} x & |x| \leq \lambda \\ \lambda \text{sign}(x) & |x| \geq \lambda \end{cases}. \quad (24)$$

Below, we use the notation in (19), (20), and (21); let $\hat{\beta}_\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^n \rho_\lambda(y_i - \mathbf{x}_i^\top \beta)$ be the unregularized M-estimator computed by the scaled loss ρ_λ , and let $R(\lambda)$ and $\hat{R}(\lambda)$ be the out-of-sample error and our proposed estimate

$$R(\lambda) := \|\Sigma^{1/2}(\hat{\beta}_\lambda - \beta_*)\|, \quad \hat{R}(\lambda) := \frac{p \sum_{i=1}^n \psi_\lambda(y_i - \mathbf{x}_i^\top \hat{\beta}_\lambda)^2}{|\{i \in [n] : |y_i - \mathbf{x}_i^\top \hat{\beta}_\lambda| \leq \lambda\}|^2},$$

where we have used the simplification $\text{tr}[\mathbf{V}_\lambda] = \sum_{i=1}^n \mathbf{1}\{|y_i - \mathbf{x}_i^\top \hat{\beta}_\lambda| \leq \lambda\}$ in (20). Let $\alpha^2(\lambda)$ be the solution to the nonlinear system (15) with $\rho = \rho_\lambda$.

With the above notation, we first verify

$$\hat{R}(\lambda) \approx R(\lambda) \approx \alpha^2(\lambda). \quad (25)$$

We set $(n, p) = (4000, 1200)$, $F_\epsilon = \text{t-dist}(\text{df} = 2)$, $\Sigma = \mathbf{I}_p$ and $\beta^* = \mathbf{0}_p$. Once we generate (\mathbf{y}, \mathbf{X}) , we compute $(R(\lambda), \hat{R}(\lambda))$ for each λ in a finite grid. We repeat the above procedure 100 times and plot $(R(\lambda), \hat{R}(\lambda))$ in Figure 3a, and the relative error $|\hat{R}(\lambda)/R(\lambda) - 1|$ in Figure 3b. We also plot $\alpha^2(\lambda)$ in Figure 3a by solving the nonlinear system of equations (15) (see Remark 3 for details). Figure 1a in the introduction features the same experiment as Figure 3a, without the theoretical curve $\lambda \rightarrow \alpha^2(\lambda)$. These figures are consistent with (25). Next, we conduct the adaptive tuning of the scale parameter $\lambda > 0$. Let us take $I = [1, 10]$ and the finite grid I_N as $\{\lambda_i = 10^{i/100} : i \in 0, 1, \dots, 100\} \subset I$. Then, Corollary 1 and Theorem 3 implies

$$\min_{\lambda \in I_N} R(\lambda) \approx R(\hat{\lambda}) \approx \min_{\lambda \in I} \alpha^2(\lambda) \quad \text{where} \quad \hat{\lambda} \in \arg \min_{\lambda \in I_N} \hat{R}(\lambda). \quad (26)$$

Below, we verify (26) as we change the scale of noise distribution F_ϵ in the following way:

$$F_\epsilon := \sigma \cdot \text{t-dist}(\text{df} = 2) \quad \text{where} \quad \sigma \in [1, 3].$$

For each σ in a finite grid over $[1, 3]$, we generate dataset (\mathbf{X}, \mathbf{y}) and calculate $R(\hat{\lambda}_N)$ and $\min_{\lambda \in I_N} R(\lambda)$. We repeat the above procedure 100 times and plot $R(\hat{\lambda})$ and $\min_{\lambda \in I} R(\lambda)$ in Figure 4a. We also plot $\min_{\lambda \in I} \alpha(\lambda)^2$ in the same figure. Figure 1b in the introduction has the same experiment as Figure 4a, without the theoretical limit.

In Figure 4b, we plot the two relative errors: $|\frac{R(\hat{\lambda}_N)}{\min_{\lambda \in I_N} R(\lambda)} - 1|$ and $|\frac{R(\hat{\lambda}_N)}{\min_{\lambda \in I} \alpha^2(\lambda)} - 1|$. These Figures are consistent with (26).

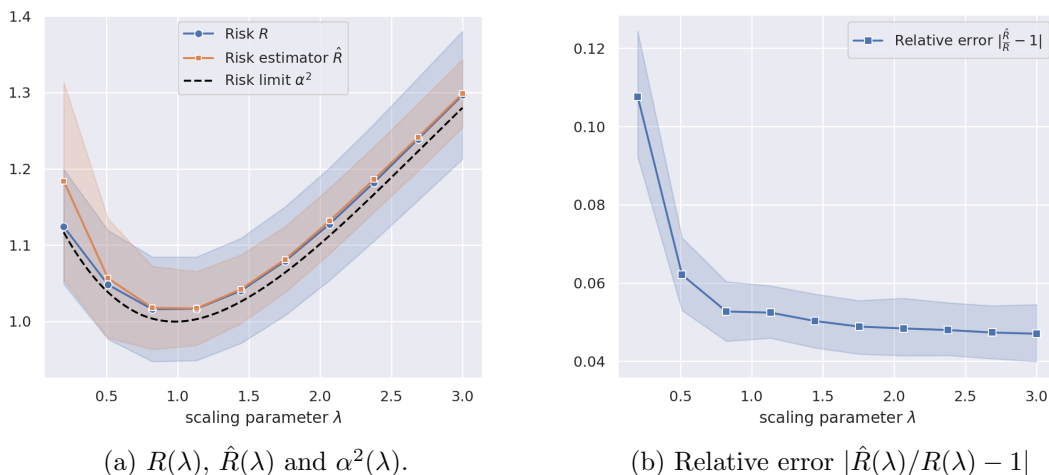


Figure 3: Plot of the out-of-sample error $R(\lambda)$ and estimator $\hat{R}(\lambda)$ over 100 repetitions, with $n = 4000$, $p = 1200$, for the Huber loss for different values of the scale parameters λ . The noise distribution is $\text{t-dist}(\text{df} = 2)$. $\alpha^2(\lambda)$ is the solution to the nonlinear system (15).

5. Outline of the proof

In this section, we give a sketch of proof of Theorem 1. See Appendix B for rigorous proofs. Let $\mathbf{G} = \mathbf{X}\Sigma^{-1/2} \in \mathbb{R}^{n \times p}$ so that \mathbf{G} has i.i.d $\mathcal{N}(0, 1)$ entries, and we denote by \mathbf{g}_i the i -th

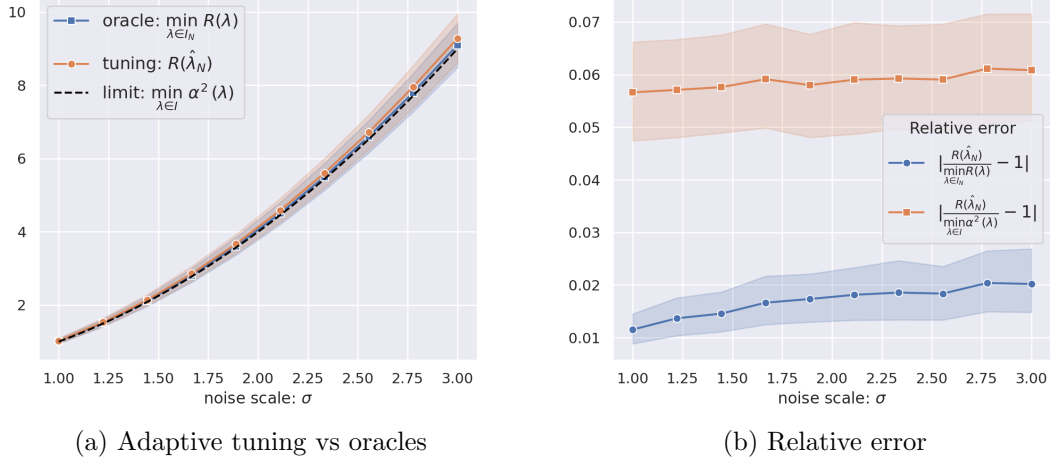


Figure 4: Adaptive tuning with the scale of noise σ changing as $F_\epsilon = \sigma \cdot \text{t-dist}(\text{df} = 2)$. Here, $I = [1, 10]$, $n = 4000$, $p = 1200$, and I_N is the uniform grid in log-scale of length 101. $R(\hat{\lambda}_N)$ is the out-of-sample error with λ selected by \hat{R} , $\min_{\lambda \in I_N} R(\lambda)$ is the optimal out-of-sample error among I_N , and $\min_{\lambda \in I} \alpha^2(\lambda)$ is the theoretically optimal risk limit. We repeat 100 times.

row of \mathbf{G} . By the change of variable $\boldsymbol{\beta} \mapsto \mathbf{h} = \boldsymbol{\Sigma}^{1/2}(\boldsymbol{\beta} - \boldsymbol{\beta}^*)$, we write the M-estimator $\hat{\boldsymbol{\beta}}$ in (1) by $\hat{\boldsymbol{\beta}} = \boldsymbol{\Sigma}^{-1/2}\hat{\mathbf{h}} + \boldsymbol{\beta}^*$, where $\hat{\mathbf{h}}$ is defined as

$$\hat{\mathbf{h}}(\boldsymbol{\epsilon}, \mathbf{G}) \in \arg \min_{\mathbf{h} \in \mathbb{R}^p} \frac{1}{n} \sum_i \rho(y_i - \mathbf{x}_i^\top (\boldsymbol{\Sigma}^{-1/2} \mathbf{h} + \boldsymbol{\beta}^*)) = \arg \min_{\mathbf{h} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \rho(\epsilon_i - \mathbf{g}_i^\top \mathbf{h}). \quad (27)$$

Throughout, we regard $\hat{\mathbf{h}}$ as a function of $(\boldsymbol{\epsilon}, \mathbf{G})$. Then, quantities of interest such as $\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|$, $\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})$, and $\mathbf{V} = (\partial/\partial \mathbf{y})\psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})$ can be expressed as functions of $(\boldsymbol{\epsilon}, \mathbf{G})$:

$$\|\boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\| = \|\hat{\mathbf{h}}\|, \quad \psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}) = \psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}) \in \mathbb{R}^n, \quad \mathbf{V} = (\partial/\partial \boldsymbol{\epsilon})\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}) \in \mathbb{R}^{n \times n}.$$

With the above notation, the statement of Theorem 1 is simplified to

$$\|\hat{\mathbf{h}}\|^2 = \frac{p}{\text{tr}[\mathbf{V}]^2} \|\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}})\|^2 + o_P(1). \quad (28)$$

To prove (28), we consider the ridge-regularized M-estimator $\hat{\mathbf{h}}_\mu$

$$\hat{\mathbf{h}}_\mu(\boldsymbol{\epsilon}, \mathbf{G}) := \arg \min_{\mathbf{h} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \rho(\epsilon_i - \mathbf{g}_i^\top \mathbf{h}) + \frac{\mu}{2} \|\mathbf{h}\|^2. \quad (29)$$

for a diminishing regularization parameter $\mu = \mu_n \rightarrow 0$. Note that $\hat{\mathbf{h}}_\mu$ coincides with $\hat{\mathbf{h}}$ defined by (27) when $\mu = 0$. The benefit of considering $\hat{\mathbf{h}}_\mu$ is its rich differentiability structures that will be explained in Section 5.1 ahead. Such derivative structure is helpful when we use probabilistic tools that involve derivatives, e.g., Stein's formula (Stein, 1981; Bellec and Zhang, 2021), the Gaussian Poincaré inequality (Boucheron et al., 2013, Theorem 3.20), and normal approximations (Chatterjee, 2009; Bellec and Zhang, 2023), etc.

5.1 Differentiability of M-estimators

In this section, we introduce a useful differentiable structure of the regularized M-estimator $\hat{\mathbf{h}}_\mu$ defined in (29). To begin with, we discuss the differentiability of $\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu)$ with respect to $\boldsymbol{\epsilon}$.

Proposition 3 (Proposition 4.1 in Bellec (2023)) *Suppose that ρ is convex and differentiable and that $\psi = \rho'$ is 1-Lipschitz. Then, for any M-estimator of the form $\hat{\boldsymbol{\beta}} \in \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \sum_{i=1}^n \rho(\mathbf{y}_i - \mathbf{x}_i^\top \boldsymbol{\beta}) + g(\boldsymbol{\beta})$ where $g : \mathbb{R}^p \rightarrow \mathbb{R}$ is a convex penalty, the map $\mathbf{y} \in \mathbb{R}^n \mapsto \psi(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}) \in \mathbb{R}^n$ is 1-Lipschitz for almost every $\mathbf{X} \in \mathbb{R}^{n \times p}$.*

Since $\psi = \rho'$ is 1-Lipschitz, Proposition 3 implies that the map $\boldsymbol{\epsilon} \mapsto \psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu(\boldsymbol{\epsilon}, \mathbf{G}))$ is 1-Lipschitz for all $\mu \geq 0$. As a consequence, Rademacher's theorem gives the following:

$$\forall \mu \geq 0, \quad \mathbf{V}_\mu := (\partial/\partial \boldsymbol{\epsilon})\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu) \in \mathbb{R}^{n \times n} \text{ exists almost everywhere and } \|\mathbf{V}_\mu\|_{op} \leq 1. \quad (30)$$

We emphasize that (30) also holds for $\mu = 0$, so the original Jacobian matrix $\mathbf{V} = (\partial/\partial \boldsymbol{\epsilon})\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}})$ is also well-defined and $\|\mathbf{V}\|_{op} \leq 1$.

The next theorem introduces a detailed differentiable structure of $\hat{\mathbf{h}}_\mu$ with respect to \mathbf{G} when μ is strictly positive.

Theorem 4 (Theorem 1 in Bellec and Shen (2022)) *Suppose $\mu > 0$. Then, for almost every $(\boldsymbol{\epsilon}, \mathbf{G})$, the map $(\boldsymbol{\epsilon}, \mathbf{G}) \mapsto \hat{\mathbf{h}}_\mu(\boldsymbol{\epsilon}, \mathbf{G})$ defined by (29) is differentiable at almost every $(\boldsymbol{\epsilon}, \mathbf{G})$, and the derivatives are given by*

$$\begin{aligned} \forall i \in [n], \quad \frac{\partial \hat{\mathbf{h}}_\mu}{\partial \epsilon_i}(\boldsymbol{\epsilon}, \mathbf{G}) &= \hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{e}_i \psi'(r_i), \\ \forall i \in [n], \forall j \in [p], \quad \frac{\partial \hat{\mathbf{h}}_\mu}{\partial g_{ij}}(\boldsymbol{\epsilon}, \mathbf{G}) &= \hat{\mathbf{A}}_\mu \mathbf{e}_j \psi(r_i) - \hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{e}_i \psi'(r_i) \mathbf{e}_j^\top \hat{\mathbf{h}}_\mu, \end{aligned}$$

where $\hat{\mathbf{A}}_\mu = (\mathbf{G}^\top \text{diag}\{\psi'(\mathbf{r})\} \mathbf{G} + n\mu I_p)^{-1} \in \mathbb{R}^{p \times p}$ and $\mathbf{r} = \boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu \in \mathbb{R}^n$.

See Bellec and Shen (2022) for the differentiability of regularized M-estimators for general strongly convex penalties. Thanks to the strict positiveness of μ , the matrix $\hat{\mathbf{A}}_\mu$ is always well-defined and its operator norm is bounded from above by $(n\mu)^{-1}$. By Theorem 4 and the chain rule, we find that $\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu)$ is also differentiable at $(\boldsymbol{\epsilon}, \mathbf{G})$ when $\mu > 0$, with the derivatives given by

$$\begin{aligned} (\partial/\partial g_{ij})\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu) &= -\text{diag}\{\psi'(\mathbf{r})\} \mathbf{G} \hat{\mathbf{A}}_\mu \mathbf{e}_j \psi(r_i) - \mathbf{V}_\mu \mathbf{e}_i \mathbf{e}_j^\top \hat{\mathbf{h}}_\mu \in \mathbb{R}^n, \\ \mathbf{V}_\mu = (\partial/\partial \boldsymbol{\epsilon})\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu) &= \text{diag}\{\psi'(\mathbf{r})\} - \text{diag}\{\psi'(\mathbf{r})\} \mathbf{G} \hat{\mathbf{A}}_\mu \mathbf{G}^\top \text{diag}\{\psi'(\mathbf{r})\} \in \mathbb{R}^{n \times n}. \end{aligned}$$

Note that the derivative formula in Theorem 4 is only guaranteed for the smoothed M-estimator $\hat{\mathbf{h}}_\mu$ with a positive ridge parameter μ , and not for the original M-estimator $\hat{\mathbf{h}}$ ($\mu = 0$). This is the benefit of considering the smoothed M-estimator $\hat{\mathbf{h}}_\mu$.

Next, we will use those derivative formulae to prove the interplay between $\|\hat{\mathbf{h}}_\mu\|^2$, $\text{tr}[\mathbf{V}_\mu]$, and $\|\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu)\|^2$.

5.2 Consistency of the risk estimate for the smoothed M-estimator

We have discussed the differentiable structure of the regularized M-estimator $\hat{\mathbf{h}}_\mu$. In this section, we derive the key relationship (31) below between $\|\hat{\mathbf{h}}_\mu\|^2$, $\text{tr}[\mathbf{V}_\mu]$, and $\|\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu)\|^2$, which can be seen as an extension of (8) to M-estimators with a Ridge penalty function that vanishes polynomially in n as $n \rightarrow +\infty$.

Lemma 1 *Let Assumption 1 and 2 be fulfilled. Let $\hat{\mathbf{h}}_\mu$ be the regularized M-estimator with a ridge penalty $\mu\|\mathbf{h}\|^2/2$ defined in (29). If $\mu = n^{-c}$ for some $c \in (0, 1/4]$, we have*

$$\frac{p}{n}\|\boldsymbol{\psi}_\mu\|^2 - \frac{\text{tr}[\mathbf{V}_\mu]^2}{n}\|\hat{\mathbf{h}}_\mu\|^2 = O_P(n^{1-c}), \quad (31)$$

where $\boldsymbol{\psi}_\mu$ is the vector $\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu) \in \mathbb{R}^n$ and \mathbf{V}_μ is the Jacobian matrix $(\partial/\partial\boldsymbol{\epsilon})\psi(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu) \in \mathbb{R}^{n \times n}$.

The proof is given in Appendix B.1. It is based on the following moment inequality.

Theorem 5 (Theorem 7.2 in Bellec (2023)) *Assume that $\mathbf{G} = (g_{ij})_{ij} \in \mathbb{R}^{n \times p}$ has i.i.d $\mathcal{N}(0, 1)$ entries, and that $\mathbf{f} : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^n$ is a differentiable function with $\|\mathbf{f}\|^2 \leq 1$. Then, there exists an absolute constant $C > 0$ such that*

$$\mathbb{E}\left[\left|p\|\mathbf{f}\|^2 - \sum_{j=1}^p(\mathbf{f}^\top \mathbf{G} \mathbf{e}_j - \sum_{i=1}^n \frac{\partial f_i}{\partial g_{ij}})^2\right|\right] \leq C\left[\sqrt{p}\mathbb{E}[(1 + \sum_{ij} \|\frac{\partial \mathbf{f}}{\partial g_{ij}}\|^2)^{1/2}] + \mathbb{E}[\sum_{ij} \|\frac{\partial \mathbf{f}}{\partial g_{ij}}\|^2]\right] \quad (32)$$

Above, we have used $\sum_{ij} = \sum_{i=1}^n \sum_{j=1}^p$ for brevity. The proof is based on the Gaussian Poincaré inequality (cf. (Boucheron et al., 2013, Theorem 3.20)) and the second order Stein's formula (see Bellec and Zhang (2021) or Theorem 6).

Below, we describe a sketch of proof of Lemma 1. To begin with, we take $\mathbf{f} = \boldsymbol{\psi}_\mu/(n^{1/2}\|\boldsymbol{\psi}\|_\infty)$ in Theorem 5 so that $\|\mathbf{f}\|^2 \leq 1$. By the derivative formula in Theorem 4 with $\|\hat{\mathbf{A}}_\mu\|_{op} = \|(\mathbf{G}^\top \text{diag}\{\psi'(\mathbf{r})\}\mathbf{G} + n\mu I_p)^{-1}\|_{op} \leq (n\mu)^{-1} = n^{-1+c}$, we can show that the right-hand side of (32) is negligible compared to the left-hand side. As a consequence, we have

$$\frac{p}{n}\|\boldsymbol{\psi}_\mu\|^2 \approx \frac{1}{n} \sum_{j=1}^p \left(\boldsymbol{\psi}_\mu \mathbf{G}^\top \mathbf{e}_j - \sum_{i=1}^n \mathbf{e}_i^\top \frac{\partial \boldsymbol{\psi}_\mu}{\partial g_{ij}} \right)^2.$$

Using Theorem 4 and the KKT condition $\mathbf{G}^\top \boldsymbol{\psi}_\mu = n\mu \hat{\mathbf{h}}_\mu$, the terms inside the square can be written explicitly by

$$\boldsymbol{\psi}_\mu \mathbf{G}^\top \mathbf{e}_j - \sum_{i=1}^n \mathbf{e}_i^\top \frac{\partial \boldsymbol{\psi}_\mu}{\partial g_{ij}} = \left(n\mu \hat{\mathbf{h}}_\mu^\top + \boldsymbol{\psi}_\mu^\top \text{diag}\{\psi'(\mathbf{r})\}\mathbf{G} \hat{\mathbf{A}}_\mu + \text{tr}(\mathbf{V}_\mu) \hat{\mathbf{h}}_\mu^\top \right) \mathbf{e}_j,$$

for all $j \in [p]$, where $\hat{\mathbf{A}}_\mu = (\mathbf{G}^\top \text{diag}\{\psi'(\mathbf{r})\}\mathbf{G} + n\mu I_p)^{-1}$. Combining the above two displays, we obtain

$$\frac{p}{n}\|\boldsymbol{\psi}_\mu\|^2 \approx \frac{1}{n} \left\| n\mu \hat{\mathbf{h}}_\mu + \hat{\mathbf{A}}_\mu^\top \mathbf{G}^\top \text{diag}\{\psi'(\mathbf{r})\}\boldsymbol{\psi}_\mu + \text{tr}(\mathbf{V}_\mu) \hat{\mathbf{h}}_\mu \right\|^2.$$

Thanks to $\mu = n^{-c}$ and $\|\hat{\mathbf{A}}_\mu\|_{op} \leq (n\mu)^{-1}$, $\|n\mu \hat{\mathbf{h}}_\mu\|$ and $\|\hat{\mathbf{A}}_\mu^\top \mathbf{G}^\top \text{diag}\{\psi'(\mathbf{r})\}\boldsymbol{\psi}_\mu\|$ are negligible compared to $\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu$, so that we obtain $p/n \cdot \|\boldsymbol{\psi}_\mu\|^2 \approx n^{-1} \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2$ as desired.

5.3 Back to the original M-estimator

Finally, we derive the target (28) from Lemma 1. Toward that, we will show that the quantities of smoothed M-estimator, i.e., $(\hat{\mathbf{h}}_\mu, \boldsymbol{\psi}_\mu, \text{tr}[\mathbf{V}_\mu])$, are close to $(\hat{\mathbf{h}}, \boldsymbol{\psi}, \text{tr}[\mathbf{V}])$ under $\mu = \mu_n = n^{-c}$ for some $c \in (0, 1/4]$. More precisely, we prove the following:

$$\|\hat{\mathbf{h}}_\mu - \hat{\mathbf{h}}\|^2 = o_P(1), \quad \|\hat{\mathbf{h}}\|^2 = O_P(1), \quad (33)$$

$$\|\boldsymbol{\psi}_\mu\|^2 - \|\boldsymbol{\psi}\|^2 = o_P(n^{1-\frac{c}{2}}), \quad \|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\|_2 = O_P(n^{\frac{1-c}{2}}), \quad (34)$$

$$\text{tr}[\mathbf{V}_\mu]^2 - \text{tr}[\mathbf{V}]^2 = o_P(n^{2-\frac{c}{2}}). \quad (35)$$

Now we assume that the above displays are justified. Then, combined with Lemma 1, we have

$$\begin{aligned} \frac{p}{n^2} \|\boldsymbol{\psi}\|^2 - \frac{\text{tr}[\mathbf{V}]^2}{n^2} \|\hat{\mathbf{h}}\|^2 &= \left(\frac{p}{n^2} \|\boldsymbol{\psi}_\mu\|^2 - \frac{\text{tr}[\mathbf{V}_\mu]^2}{n^2} \|\hat{\mathbf{h}}\|^2 \right) + \frac{p}{n^2} (\|\boldsymbol{\psi}\|_2^2 - \|\boldsymbol{\psi}_\mu\|^2) \\ &\quad + \frac{\text{tr}[\mathbf{V}_\mu]^2 - \text{tr}[\mathbf{V}]^2}{n^2} \|\hat{\mathbf{h}}_\mu\|^2 - \frac{\text{tr}[\mathbf{V}]^2}{n^2} (\|\hat{\mathbf{h}}\|^2 - \|\hat{\mathbf{h}}_\mu\|^2) \\ &= O_P(n^{-c}) + o_P(n^{-\frac{c}{2}}) + o_P(n^{-\frac{c}{2}}) - \frac{\text{tr}[\mathbf{V}]^2}{n^2} (\|\hat{\mathbf{h}}\|^2 - \|\hat{\mathbf{h}}_\mu\|^2) \\ &= o_P(n^{-\frac{c}{2}}) - \frac{\text{tr}[\mathbf{V}]^2}{n^2} o_P(1). \end{aligned}$$

Furthermore, we will argue that $n^2/\text{tr}[\mathbf{V}]^2 = O_P(1)$, i.e., $\text{tr}(\mathbf{V})/n$ is not too small (see Lemma 9). Thus, dividing by $\text{tr}[\mathbf{V}]^2/n^2$ on the above display, we obtain the target (28).

Below, we describe a sketch of proof of (33), (34), and (35). For (33), we will show that $\|\hat{\mathbf{h}}\|^2$ and $\|\hat{\mathbf{h}}_\mu\|^2$ converge to the same finite limit by the main theorem in Thrampoulidis et al. (2018) and Theorem 7. (34) easily follows from the two KKT conditions and the Lipschitz property of $\boldsymbol{\psi}$. The most technical part is (35). Thanks to Assumption 2, the noise vector $\boldsymbol{\epsilon}$ can be represented as

$$\boldsymbol{\epsilon} = \mathbf{z} + \boldsymbol{\delta}, \quad \mathbf{z} \perp \boldsymbol{\delta}, \quad (z_i)_{i=1}^n \stackrel{\text{iid}}{\sim} F, \quad (\delta_i)_{i=1}^n \stackrel{\text{iid}}{\sim} \tilde{F},$$

where F has the density $z \mapsto \exp(-\phi(z))$. By the chain rule, $\mathbf{V} - \mathbf{V}_\mu$ can be written as

$$\mathbf{V} - \mathbf{V}_\mu = \frac{\partial}{\partial \boldsymbol{\epsilon}} (\boldsymbol{\psi} - \boldsymbol{\psi}_\mu) = \frac{\partial \mathbf{z}}{\partial \boldsymbol{\epsilon}} \cdot \frac{\partial}{\partial \mathbf{z}} (\boldsymbol{\psi} - \boldsymbol{\psi}_\mu) = \mathbf{I}_n \cdot \frac{\partial}{\partial \mathbf{z}} (\boldsymbol{\psi} - \boldsymbol{\psi}_\mu) = \frac{\partial}{\partial \mathbf{z}} (\boldsymbol{\psi} - \boldsymbol{\psi}_\mu). \quad (36)$$

Here, by the second order Stein's formula extended to the density of the form $\exp(-\phi(z))$, we will show that $\text{tr}[\mathbf{V}] - \text{tr}[\mathbf{V}_\mu]$ concentrates around $\phi'(\mathbf{z})^\top (\boldsymbol{\psi} - \boldsymbol{\psi}_\mu)$.

Theorem 6 (Equation (2.15) in Bellec and Zhang (2021)) *Assume that the random vector $\mathbf{z} \in \mathbb{R}^n$ has density $\mathbf{z} \mapsto \exp(-\phi(\mathbf{z}))$ where $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is twice continuously differentiable function with Hessian $\mathbf{H}(\mathbf{z}) = (\partial/\partial \mathbf{z})^2 \phi(\mathbf{z}) \in \mathbb{R}^{n \times n}$. Then, for any differentiable function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, we have*

$$\mathbb{E} \left[\left(\frac{\partial \phi(\mathbf{z})}{\partial \mathbf{z}} \mathbf{f}(\mathbf{z}) - \text{tr} \left[\frac{\partial \mathbf{f}(\mathbf{z})}{\partial \mathbf{z}} \right] \right)^2 \right] = \mathbb{E} \left[\mathbf{f}(\mathbf{z})^\top \mathbf{H}(\mathbf{z}) \mathbf{f}(\mathbf{z}) + \text{tr} \left[\left(\frac{\partial \mathbf{f}(\mathbf{z})}{\partial \mathbf{z}} \right)^2 \right] \right],$$

provided that the expectation on the right-hand side exists.

Applying Theorem 6 with $\phi(\mathbf{z}) = \sum_{i=1}^n \phi(z_i)$ and $\mathbf{f} = \boldsymbol{\psi} - \boldsymbol{\psi}_\mu$, using the identity (36) and $\mathbf{H}(\mathbf{z}) = \text{diag}\{\phi''(z_i) : i \in [n]\}$, we have

$$\mathbb{E}[(\phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi}) - \text{tr}[\mathbf{V} - \mathbf{V}_\mu])^2] = \mathbb{E}\left[\sum_{i=1}^n \phi''(z_i)(\psi_i - (\psi_\mu)_i)^2 + \text{tr}[(\mathbf{V} - \mathbf{V}_\mu)^2]\right],$$

and we can easily show that the right-hand side is $O(n)$. Hence, by the Markov inequality, we have

$$\text{tr}[\mathbf{V}_\mu] - \text{tr}[\mathbf{V}] = \phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi}) + O_P(n^{1/2}).$$

By the Cauchy–Schwarz inequality, we bound $|\phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi})|$ from above as

$$|\phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi})| \leq \|\phi'(\mathbf{z})\| \|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\| \stackrel{(*)}{=} O_P(n^{1/2}) \cdot O_P(n^{\frac{1-c}{2}}) = O_P(n^{1-\frac{c}{2}}),$$

where $(*)$ follows from $\|\phi'(\mathbf{z})\| = O_P(n^{1/2})$ (see Lemma 8) and $\|\boldsymbol{\psi} - \boldsymbol{\psi}_\mu\| = O_P(n^{1/2-c})$ by (34). As a consequence, we obtain

$$\text{tr}[\mathbf{V}_\mu] - \text{tr}[\mathbf{V}] = O_P(n^{1-\frac{c}{2}}) + O_P(n^{1/2}) = O_P(n^{1-\frac{c}{2}})$$

thanks to $c \leq 1$. Finally, $\|\mathbf{V}\|_{op}$ and $\|\mathbf{V}_\mu\|_{op} \leq 1$ lead to

$$|\text{tr}[\mathbf{V}_\mu]^2 - \text{tr}[\mathbf{V}]^2| \leq |\text{tr}[\mathbf{V}_\mu] + \text{tr}[\mathbf{V}]| |\text{tr}[\mathbf{V}_\mu] - \text{tr}[\mathbf{V}]| \leq 2nO_P(n^{1-\frac{c}{2}}) = O_P(n^{2-\frac{c}{2}}),$$

which is exactly (35).

Acknowledgments

P. C. Bellec’s research was partially supported by the National Science Foundation award DMS-1945428.

References

- Mohsen Bayati and Andrea Montanari. The LASSO risk for Gaussian matrices. *IEEE Transactions on Information Theory*, 58(4):1997–2017, 2011.
- Derek Bean, Peter J Bickel, Nouredine El Karoui, and Bin Yu. Optimal M-estimation in high-dimensional regression. *Proceedings of the National Academy of Sciences*, 110(36):14563–14568, 2013.
- Pierre C Bellec. Out-of-sample error estimation for M-estimators with convex penalty. *Information and Inference: A Journal of the IMA*, 12(4):2782–2817, 2023.
- Pierre C Bellec. Observable adjustments in single-index models for regularized M-estimators. *The Annals of Statistics*, forthcoming, 2025.
- Pierre C. Bellec and Takuya Koriyama. Existence of solutions to the nonlinear equations characterizing the precise error of M-estimators. *arXiv preprint arXiv:2312.13254*, 2023.

- Pierre C Bellec and Yiwei Shen. Derivatives and residual distribution of regularized M-estimators with application to adaptive tuning. In *Conference on Learning Theory*, pages 1912–1947. PMLR, 2022.
- Pierre C Bellec and Cun-Hui Zhang. Second-order Stein: SURE for SURE and other applications in high-dimensional inference. *The Annals of Statistics*, 49(4):1864–1903, 2021.
- Pierre C Bellec and Cun-Hui Zhang. Debiasing convex regularized estimators and interval estimation in linear models. *The Annals of Statistics*, 51(2):391–436, 2023.
- Pierre C Bellec, Yiwei Shen, and Cun-Hui Zhang. Asymptotic normality of robust M-estimators with convex penalty. *Electronic Journal of Statistics*, 16(2):5591–5622, 2022.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Michael Celentano and Andrea Montanari. Fundamental barriers to high-dimensional regression with convex penalties. *The Annals of Statistics*, 50(1):170–196, 2022.
- Michael Celentano, Andrea Montanari, and Yuting Wei. The lasso with general gaussian designs with applications to hypothesis testing. *The Annals of Statistics*, 51(5):2194–2220, 2023.
- Sourav Chatterjee. Fluctuations of eigenvalues and second order Poincaré inequalities. *Probability Theory and Related Fields*, 143(1-2):1–40, 2009.
- David Donoho and Andrea Montanari. High Dimensional Robust M-estimation: Asymptotic Variance via Approximate Message Passing. *Probability Theory and Related Fields*, 166: 935–969, 2016.
- Noureddine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170:95–175, 2018.
- Noureddine El Karoui, Derek Bean, Peter J Bickel, Chingway Lim, and Bin Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 110(36):14557–14562, 2013.
- Noureddine El El Karoui. Asymptotic behavior of unregularized and ridge-regularized high-dimensional robust regression estimators: rigorous results. *arXiv preprint arXiv:1311.2445*, 2013.
- Peter J Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- Sophie Lambert-Lacroix and Laurent Zwald. Robust regression through the Huber’s criterion and adaptive lasso penalty. *Electronic Journal of Statistics*, 5(none):1015 – 1053, 2011.

- Po-Ling Loh. Scale calibration for high-dimensional robust regression. *Electronic Journal of Statistics*, 15(2):5933–5994, 2021.
- Po-Ling Loh. A theoretical review of modern robust statistics. *Annual Review of Statistics and Its Application*, 12, 2024.
- Bruno Loureiro, Cedric Gerbelot, Hugo Cui, Sebastian Goldt, Florent Krzakala, Marc Mezard, and Lenka Zdeborová. Learning curves of generic features maps for realistic datasets with a teacher-student model. *Advances in Neural Information Processing Systems*, 34:18137–18151, 2021.
- Bruno Loureiro, Cédric Gerbelot, Maria Refinetti, Gabriele Sicuro, and Florent Krzakala. Fluctuations, bias, variance & ensemble of learners: Exact asymptotics for convex losses in high-dimension. In *International Conference on Machine Learning*, pages 14283–14314. PMLR, 2022.
- Ricardo A Maronna, R Douglas Martin, Victor J Yohai, and Matías Salibián-Barrera. *Robust statistics: theory and methods (with R)*. John Wiley & Sons, 2019.
- Léo Miolane and Andrea Montanari. The distribution of the lasso: Uniform control over sparse balls and adaptive parameter tuning. *The Annals of Statistics*, 49(4):2313–2335, 2021.
- Charles M Stein. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, pages 1135–1151, 1981.
- Qiang Sun, Wen-Xin Zhou, and Jianqing Fan. Adaptive huber regression. *Journal of the American Statistical Association*, 115(529):254–265, 2020.
- Pragya Sur and Emmanuel J Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525, 2019.
- Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Precise error analysis of regularized M -estimators in high dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, 2018.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- Lili Wang, Chao Zheng, Wen Zhou, and Wen-Xin Zhou. A new principle for tuning-free Huber regression. *Statistica Sinica*, 31(4):2153–2177, 2021.

SUPPLEMENT

Appendix A. Existence and uniqueness of solutions to (15)

The following theorem from our companion paper Bellec and Koriyama (2023) guarantees the uniqueness and the existence of the solution (α, κ) to (15), with an explicit upper bound of α . We include it here for convenience.

Theorem 7 (Section 2 Bellec and Koriyama (2023)) *Suppose that $\gamma \in (0, 1)$ and (ρ, F_ϵ) satisfy the conditions*

1. ρ is convex, Lipschitz, and $\{0\} = \arg \min_{x \in \mathbb{R}} \rho(x)$.
2. $\mathbb{P}_{W \sim F_\epsilon}(W \neq 0) > 0$ and $\inf_{\lambda > 0} \mathbb{E}_{W \sim F_\epsilon}[\text{dist}(G, \lambda \partial \rho(W))^2] > 1 - \gamma$.

Then, the nonlinear system of equations (15) has a unique solution $(\alpha, \kappa) \in \mathbb{R}_{>0}^2$, and this α is bounded from above as

$$\alpha \leq \frac{Q_{F_\epsilon}(r^2 c_\gamma)}{r c_\gamma} + \frac{b}{c_\gamma} \tag{37}$$

where c_γ is a positive constant depending on γ only, $Q_{F_\epsilon}(x) := \inf\{q > 0 : \mathbb{P}_{W \sim F_\epsilon}(|W| \geq q) \leq x\}$, and $r \in (0, 1]$ and $b \geq 0$ are any constant such that the coercivity condition

$$\|\rho_\lambda\|_{\text{lip}}^{-1}(\rho(x) - \rho(0)) \geq r(|x| - b) \quad \text{for all } x \in \mathbb{R}$$

is satisfied.

The upper bound (37) is helpful for Section 3, in which we will consider the scaled loss $\rho_\lambda : \mathbb{R} \rightarrow \mathbb{R}$, $x \mapsto \lambda^2 \rho(x/\lambda)$ for some scale parameter $\lambda > 0$. Let $\alpha(\lambda)$ be the solution to the nonlinear system (15) with $\rho = \rho_\lambda$. Then, (37) with $\rho = \rho_\lambda$ gives the following explicit upper bound of $\alpha(\lambda)$ with respect to $\lambda > 0$:

$$\forall \lambda > 0, \quad \alpha^2(\lambda) \leq C_4(\rho, F_\epsilon, \gamma)(\lambda^2 + 1).$$

See Proposition 2 and its proof for its derivation. We observe that the upper bound explodes as $\lambda \rightarrow +\infty$, which reflects the fact that the out-of-sample error of the ordinary least square (OLS) estimator is unbounded when the noise distribution has no second moment.

Remark 3 *Suppose that the loss ρ is the Huber loss ρ_λ , parametrized by a scale parameter $\lambda > 0$. This loss function was introduced in (24) and used for numerical simulations in Section 4 and Appendix D. For convenience, we recall its definition:*

$$\rho_\lambda(x) := \begin{cases} x^2/2 & |x| \leq \lambda \\ \lambda|x| - \lambda^2/2 & |x| > \lambda \end{cases}$$

In this case, the proximal operator of $\rho = \rho_\lambda$ takes a closed form and satisfies

$$x - \text{prox}[\kappa \rho_\lambda](x) = \kappa \lambda \zeta\left(\frac{x}{\lambda(1 + \kappa)}\right) \quad \text{where} \quad \zeta(u) := \begin{cases} u & |u| \leq 1 \\ \text{sign}(u) & |u| > 1 \end{cases}$$

for all $\kappa > 0$ and $x \in \mathbb{R}$. Thus, the nonlinear system (15) with $\rho = \rho_\lambda$ can be written as

$$\alpha^2 \gamma - \kappa^2 \lambda^2 \mathbb{E} \left[\zeta \left(\frac{aG + W}{\lambda(1 + \kappa)} \right)^2 \right] = 0, \quad \alpha \gamma - \kappa \lambda \mathbb{E} \left[\zeta \left(\frac{aG + W}{\lambda(1 + \kappa)} \right) \cdot Z \right] = 0$$

with positive unknown (α, κ) . In the numerical simulations presented in Section 4 and Appendix D, this system was solved using the solver `scipy.optimize.fsolve` from `scipy` (Virtanen et al., 2020).

Appendix B. Proof of Theorem 1

B.1 Proof of Lemma 1

The claim of Lemma 1 is recalled here for convenience:

$$\frac{p}{n} \|\boldsymbol{\psi}_\mu\|^2 - \frac{\text{tr}[\mathbf{V}_\mu]^2}{n} \|\hat{\mathbf{h}}_\mu\|^2 = O_P(n^{1-c}) \text{ under } \mu = n^{-c} \text{ for all } c \in (0, 1/4].$$

To prove this, we use a variant of Theorem 5:

Lemma 2 (Theorem 7.4 in Bellec and Shen (2022)) *Let $h : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^p$ and $\psi : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^n$ be locally Lipschitz functions. Suppose $\mathbf{G} \in \mathbb{R}^{n \times p}$ has i.i.d $\mathcal{N}(0, 1)$ entries, and we denote $h(\mathbf{G})$ by \mathbf{h} and $\psi(\mathbf{G})$ by $\boldsymbol{\psi}$. Then, we have*

$$\mathbb{E} \left[\frac{\left| \frac{p}{n} \|\boldsymbol{\psi}\|^2 - \frac{1}{n} \sum_{j=1}^p (\boldsymbol{\psi}^\top \mathbf{G} \mathbf{e}_j - \sum_{i=1}^n \frac{\partial \psi_i}{\partial g_{ij}})^2 \right|}{\|\mathbf{h}\|^2 + \|\boldsymbol{\psi}\|^2/n} \right] \leq C(\sqrt{n+p}(1 + \Xi^{1/2}) + \Xi), \quad (38)$$

where $\Xi = \mathbb{E} \left[\frac{1}{\|\mathbf{h}\|^2 + \|\boldsymbol{\psi}\|^2/n} \sum_{i=1}^n \sum_{j=1}^p (\|\frac{\partial \mathbf{h}}{\partial g_{ij}}\|^2 + \frac{1}{n} \|\frac{\partial \boldsymbol{\psi}}{\partial g_{ij}}\|^2) \right]$

Lemma 2 is obtained from Theorem 5 with $\mathbf{f} = n^{-1/2} \boldsymbol{\psi} / (\|\boldsymbol{\psi}\|/n + \|\mathbf{h}\|)$. Below, we prove Lemma 1 using Lemma 2 with $(\boldsymbol{\psi}, \mathbf{h}) = (\boldsymbol{\psi}_\mu, \mathbf{h}_\mu) = (\boldsymbol{\psi}(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu), \hat{\mathbf{h}}_\mu)$. First, we bound Ξ in (38). By the derivative formula in Theorem 4, the derivatives inside the expectation can be written as

$$\begin{aligned} \sum_{i=1}^n \sum_{j=1}^p \left\| \frac{\partial \hat{\mathbf{h}}_\mu}{\partial g_{ij}} \right\|^2 &= \sum_{i=1}^n \sum_{j=1}^p \left\| \hat{\mathbf{A}}_\mu \mathbf{e}_j \mathbf{e}_i^\top \boldsymbol{\psi}_\mu - \hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu \mathbf{e}_i \mathbf{e}_j^\top \hat{\mathbf{h}}_\mu \right\|^2, \\ \sum_{i=1}^n \sum_{j=1}^p \left\| \frac{\partial \boldsymbol{\psi}_\mu}{\partial g_{ij}} \right\|^2 &= \sum_{i=1}^n \sum_{j=1}^p \left\| -\mathbf{D}_\mu \mathbf{G} \hat{\mathbf{A}}_\mu \mathbf{e}_j \mathbf{e}_i^\top \boldsymbol{\psi}_\mu - \mathbf{V}_\mu \mathbf{e}_i \mathbf{e}_j^\top \hat{\mathbf{h}}_\mu \right\|^2, \end{aligned}$$

where $\mathbf{D}_\mu = \text{diag}\{\psi'(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu)\}$, $\hat{\mathbf{A}}_\mu = (\mathbf{G}^\top \text{diag}\{\psi'(\mathbf{r})\} \mathbf{G} + n\mu \mathbf{I}_p)^{-1}$, and $\mathbf{V}_\mu = (\partial/\partial \boldsymbol{\epsilon}) \boldsymbol{\psi}(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu)$. Note that the operator norm of these matrices is bounded from above as

$$\begin{aligned} \|\hat{\mathbf{A}}_\mu\|_{op} &\leq (n\mu)^{-1} && \text{by } \psi'(r_i) \geq 0 \text{ for all } i \in [n] \\ \|\mathbf{D}_\mu\|_{op} &\leq 1 && \text{by } \|\boldsymbol{\psi}\|_{\text{lip}} = 1 \\ \|\mathbf{V}_\mu\|_{op} &\leq 1 && \text{by (30)} \end{aligned}$$

Then, $\frac{1}{2} \sum_{ij} \|\frac{\partial \hat{\mathbf{h}}_\mu}{\partial g_{ij}}\|^2$ can be bounded from above as

$$\begin{aligned}
 \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^p \|\frac{\partial \hat{\mathbf{h}}_\mu}{\partial g_{ij}}\|^2 &\leq \sum_{ij} \|\hat{\mathbf{A}}_\mu \mathbf{e}_j \mathbf{e}_i^\top \boldsymbol{\psi}_\mu\|^2 + \|\hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu \mathbf{e}_i \mathbf{e}_j^\top \hat{\mathbf{h}}_\mu\|^2 && \text{by } (a+b)^2 \leq 2(a^2 + b^2) \\
 &= \|\hat{\mathbf{A}}_\mu\|_F^2 \|\boldsymbol{\psi}_\mu\|^2 + \|\hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu\|_F^2 \|\hat{\mathbf{h}}_\mu\|^2 && \text{by } \sum_i \|\mathbf{M} \mathbf{e}_i\|^2 = \|\mathbf{M}\|_F^2 \\
 &\leq \|\hat{\mathbf{A}}_\mu\|_F^2 (\|\boldsymbol{\psi}_\mu\|^2 + \|\mathbf{G}^\top \mathbf{D}_\mu\|_{op}^2 \|\hat{\mathbf{h}}_\mu\|^2) && \text{by } \|\hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu\|_F^2 \leq \|\hat{\mathbf{A}}_\mu\|_F^2 \|\mathbf{G}^\top \mathbf{D}_\mu\|_{op}^2 \\
 &\leq \|\hat{\mathbf{A}}_\mu\|_F^2 (\|\boldsymbol{\psi}_\mu\|^2 + \|\mathbf{G}\|_{op}^2 \|\hat{\mathbf{h}}_\mu\|^2) && \text{by } \|\mathbf{D}_\mu\|_{op} \leq 1 \\
 &\leq \|\hat{\mathbf{A}}_\mu\|_F^2 (n + \|\mathbf{G}\|_{op}^2) (\|\boldsymbol{\psi}_\mu\|^2/n + \|\hat{\mathbf{h}}_\mu\|^2) \\
 &\leq p(n\mu)^{-2} (n + \|\mathbf{G}\|_{op}^2) (\|\boldsymbol{\psi}_\mu\|^2/n + \|\hat{\mathbf{h}}_\mu\|^2) && \text{by } \|\hat{\mathbf{A}}_\mu\|_F^2 \leq p \|\hat{\mathbf{A}}_\mu\|_{op}^2 \leq p(n\mu)^{-2}.
 \end{aligned}$$

By the same algebra, we bound $\frac{1}{2n} \sum_{ij} \|\frac{\partial \hat{\boldsymbol{\psi}}_\mu}{\partial g_{ij}}\|^2$ from above as

$$\begin{aligned}
 \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^p \|\frac{\partial \hat{\boldsymbol{\psi}}_\mu}{\partial g_{ij}}\|^2 &\leq \frac{1}{n} \sum_{ij} \left(\|\mathbf{D}_\mu \mathbf{G} \hat{\mathbf{A}}_\mu \mathbf{e}_j \mathbf{e}_i^\top \boldsymbol{\psi}_\mu\|^2 + \|\mathbf{V}_\mu \mathbf{e}_i \mathbf{e}_j^\top \hat{\mathbf{h}}_\mu\|^2 \right) \\
 &= n^{-1} \|\mathbf{D}_\mu \mathbf{G} \hat{\mathbf{A}}_\mu\|_F^2 \|\boldsymbol{\psi}_\mu\|^2 + n^{-1} \|\mathbf{V}_\mu\|_F^2 \|\hat{\mathbf{h}}_\mu\|^2 \\
 &\leq \|\hat{\mathbf{A}}_\mu\|_F^2 \|\mathbf{G}\|_{op}^2 \|\boldsymbol{\psi}_\mu\|^2/n + \|\hat{\mathbf{h}}_\mu\|^2 \\
 &\leq (\|\hat{\mathbf{A}}_\mu\|_F^2 \|\mathbf{G}\|_{op}^2 + 1) (\|\boldsymbol{\psi}_\mu\|^2/n + \|\hat{\mathbf{h}}_\mu\|^2) \\
 &\leq (p(n\mu)^{-2} \|\mathbf{G}\|_{op}^2 + 1) (\|\boldsymbol{\psi}_\mu\|^2/n + \|\hat{\mathbf{h}}_\mu\|^2)
 \end{aligned}$$

By the above displays, we obtain the upper bound of $\Xi = \mathbb{E}[\frac{1}{\|\mathbf{h}\|^2 + \|\boldsymbol{\psi}\|^2/n} \sum_{ij} (\|\frac{\partial \mathbf{h}}{\partial g_{ij}}\|^2 + \frac{1}{n} \|\frac{\partial \boldsymbol{\psi}}{\partial g_{ij}}\|^2)]$:

$$\Xi \leq \mathbb{E} [2p(n\mu)^{-2} (n + \|\mathbf{G}\|_{op}^2) + 2(p(n\mu)^{-2} \|\mathbf{G}\|_{op}^2 + 1)] \stackrel{(*)}{=} O(n^{2c}),$$

where $(*)$ follows from $\mu = n^{-c}$ with $c > 0$ and $\mathbb{E}[\|\mathbf{G}\|_{op}^2] = O(n)$ for the matrix $\mathbf{G} \in \mathbb{R}^{n \times p}$ with iid standard normal entries as $p/n \rightarrow \gamma \in (0, \infty)$. Substituting this bound to the right-hand side of Lemma 2, we have

$$\mathbb{E} \left[\frac{|\frac{p}{n} \|\boldsymbol{\psi}_\mu\|^2 - \frac{1}{n} \sum_{j=1}^p (\boldsymbol{\psi}_\mu^\top \mathbf{G} \mathbf{e}_j - \sum_{i=1}^n \frac{\partial \mathbf{e}_i^\top \boldsymbol{\psi}_\mu}{\partial g_{ij}})^2|}{\|\boldsymbol{\psi}_\mu\|^2/n + \|\hat{\mathbf{h}}_\mu\|^2} \right] \leq C(\sqrt{n+p}(1 + \Xi^{1/2}) + \Xi) \stackrel{(**)}{=} O(n^{c+1/2}), \tag{39}$$

where we have used $0 < c < 1/2$ and $\Xi = O(n^{2c})$ for $(**)$. It remains to bound the error $\sum_{j=1}^p (\boldsymbol{\psi}_\mu^\top \mathbf{G} \mathbf{e}_j - \sum_{i=1}^n (\partial/\partial g_{ij}) \mathbf{e}_i^\top \boldsymbol{\psi}_\mu)^2 - \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2$ inside the expectation on the left-hand side. Now the derivatives formula and the KKT condition $\mathbf{G}^\top \boldsymbol{\psi}_\mu = n\mu \hat{\mathbf{h}}_\mu$ yield

$$\sum_{j=1}^p (\boldsymbol{\psi}_\mu^\top \mathbf{G} \mathbf{e}_j - \sum_{i=1}^n \frac{\partial \mathbf{e}_i^\top \boldsymbol{\psi}_\mu}{\partial g_{ij}})^2 = \|n\mu \hat{\mathbf{h}}_\mu + \hat{\mathbf{A}}_\mu^\top \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu + \text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2,$$

so that we have

$$\begin{aligned}
 & \frac{1}{2} \left| \sum_{j=1}^p (\boldsymbol{\psi}_\mu^\top \mathbf{G} e_j - \sum_{i=1}^n \frac{\partial e_i^\top \boldsymbol{\psi}_\mu}{\partial g_{ij}})^2 - \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2 \right| \\
 &= 2^{-1} \|n\mu \hat{\mathbf{h}}_\mu + \hat{\mathbf{A}}_\mu^\top \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu + \text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2 - \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2 \\
 &\leq 2^{-1} \|n\mu \hat{\mathbf{h}}_\mu + \hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\|^2 + \|n\mu \hat{\mathbf{h}}_\mu + \mathbf{A}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\| \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\| && \text{by } |(a+b)^2 - b^2| \leq a^2 + 2|ab| \\
 &\leq n^2 \mu^2 \|\hat{\mathbf{h}}_\mu\|^2 + \|\mathbf{A}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\|^2 + |\text{tr}[\mathbf{V}_\mu]| \|n\mu \hat{\mathbf{h}}_\mu + \mathbf{A}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\| \|\hat{\mathbf{h}}_\mu\| && \text{by } 2^{-1}(a+b)^2 \leq a^2 + b^2 \\
 &\leq n^2 \mu^2 \|\hat{\mathbf{h}}_\mu\|^2 + \|\mathbf{A}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\|^2 + n \|n\mu \hat{\mathbf{h}}_\mu + \mathbf{A}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\| \|\hat{\mathbf{h}}_\mu\| && \text{by } |\text{tr}[\mathbf{V}_\mu]| \leq n \|\mathbf{V}_\mu\|_{op} \leq n \\
 &\leq n^2 \mu^2 \|\hat{\mathbf{h}}_\mu\|^2 + \|\hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\|^2 + n^2 \mu \|\hat{\mathbf{h}}_\mu\|^2 + \|\hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\| \|\hat{\mathbf{h}}_\mu\| && \text{by triangle inequality} \\
 &\leq (n^2 \mu^2 + n^2 \mu + 2^{-1}) \|\hat{\mathbf{h}}_\mu\|^2 + (1 + 2^{-1}) \|\hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu \boldsymbol{\psi}_\mu\|^2 && \text{by } ab \leq 2^{-1}(a^2 + b^2) \\
 &\leq (\|\hat{\mathbf{h}}_\mu\|^2 + n^{-1} \|\boldsymbol{\psi}_\mu\|^2) (n^2 \mu^2 + n^2 \mu + 2^{-1} + 2^{-1} 3n \|\hat{\mathbf{A}}_\mu \mathbf{G}^\top \mathbf{D}_\mu\|_{op}^2) \\
 &\leq (\|\hat{\mathbf{h}}_\mu\|^2 + n^{-1} \|\boldsymbol{\psi}_\mu\|^2) (n^2 \mu^2 + n^2 \mu + 2^{-1} + 2^{-1} 3n (n\mu)^{-2} \|\mathbf{G}\|_{op}^2) && \text{by } \|\hat{\mathbf{A}}_\mu\|_{op} \leq (n\mu)^{-1} \\
 & && \text{and } \|\mathbf{D}_\mu\|_{op} \leq 1.
 \end{aligned}$$

Thus, we have

$$\begin{aligned}
 & \mathbb{E} \left[\frac{\frac{1}{n} \left| \sum_{j=1}^p (\boldsymbol{\psi}_\mu^\top \mathbf{G} e_j - \sum_{i=1}^n \frac{\partial e_i^\top \boldsymbol{\psi}_\mu}{\partial g_{ij}})^2 - \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2 \right|}{\|\boldsymbol{\psi}_\mu\|^2/n + \|\hat{\mathbf{h}}_\mu\|^2} \right] \\
 &\leq C_5 \mathbb{E} [n\mu^2 + n\mu + n^{-1} + (n\mu)^{-2} \|\mathbf{G}\|_{op}^2] \\
 &= O(n^{1-2c} + n^{1-c} + n^{-1} + n^{-1+2c}) && \text{by } \mu = n^{-c} \text{ and } \mathbb{E}[\|\mathbf{G}\|_{op}^2] = O(n), \\
 &= O(n^{1-c}) && \text{by } 0 < c \leq 1/4 < 2/3
 \end{aligned}$$

Combined with (39), by the triangle inequality,

$$\mathbb{E} \left[\frac{\left| \frac{p}{n} \|\boldsymbol{\psi}_\mu\|^2 - \frac{1}{n} \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2 \right|}{\|\hat{\mathbf{h}}_\mu\|^2 + \|\boldsymbol{\psi}_\mu\|^2/n} \right] = O(n^{c+1/2}) + O(n^{1-c}) = O(n^{1-c}) \quad (40)$$

by $c \in (0, 1/4]$. Thus, we obtain

$$\left| \frac{p}{n} \|\boldsymbol{\psi}_\mu\|^2 - \frac{1}{n} \|\text{tr}[\mathbf{V}_\mu] \hat{\mathbf{h}}_\mu\|^2 \right| = \left(\|\hat{\mathbf{h}}_\mu\|^2 + \frac{\|\boldsymbol{\psi}_\mu\|^2}{n} \right) O_P(n^{1-c}) \leq (\|\hat{\mathbf{h}}_\mu\|^2 + \|\boldsymbol{\psi}_\mu\|_\infty^2) O_P(n^{1-c}) = O_P(n^{1-c}),$$

where we have used $\|\hat{\mathbf{h}}_\mu\|_2^2 = O_P(1)$ (see Lemma 3). This finishes the proof.

B.2 Convergence of smoothed quantities

The lemmas below are the key to relating Lemma 1 to Theorem 1.

Lemma 3 *Suppose that (ρ, F_ϵ) satisfies Assumption 1 and Assumption 2. Then, under $\mu = n^{-c}$ for any $c > 0$, we have*

$$\|\hat{\mathbf{h}}\|^2 \xrightarrow{p} \alpha^2, \quad \|\hat{\mathbf{h}}_\mu\|^2 \xrightarrow{p} \alpha^2, \quad \alpha^2 < +\infty, \quad \|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\| = o_p(1)$$

where α is the unique solution to the nonlinear system (15).

Proof The definition of $\hat{\mathbf{h}}$ and $\hat{\mathbf{h}}_\mu$ are recalled here for convenience:

$$\hat{\mathbf{h}} \in \arg \min_{\mathbf{h} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \rho(\epsilon_i - \mathbf{g}_i^\top \mathbf{h}), \quad \hat{\mathbf{h}}_\mu \in \arg \min_{\mathbf{h} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \rho(\epsilon_i - \mathbf{g}_i^\top \mathbf{h}) + \frac{\mu \|\mathbf{h}\|^2}{2},$$

where $(\epsilon_i)_{i=1}^n \stackrel{\text{iid}}{\sim} F_\epsilon$, $(\mathbf{g}_i)_{i=1}^n \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_p, \mathbf{I}_p)$, and $(\epsilon_i)_{i=1}^n \perp\!\!\!\perp (\mathbf{g}_i)_{i=1}^n$. Since Assumption 1-2 imply the assumption in Theorem 7 (see Remark 1), the nonlinear system of equations (15) has a unique solution α and $\|\hat{\mathbf{h}}\|^2 \rightarrow^p \alpha^2$. It remains to show $\|\hat{\mathbf{h}}_\mu\|^2 \rightarrow^p \alpha^2$ under the diminishing regularization parameter $\mu = n^{-c}$. Define

$$f : \quad \mathbb{R}^p \rightarrow \mathbb{R}, \quad \mathbf{x} \mapsto n^{-c} \|\mathbf{x}\|^2 / 2$$

so that $\hat{\mathbf{h}}_\mu$ is the regularized M-estimator with the Lipschitz convex loss ρ and penalty f . Suppose the following condition is satisfied:

$$\forall a \in \mathbb{R}, \forall \tau > 0, \quad n^{-1} (e_f(a\mathbf{g}; \tau) - f(\mathbf{0}_p)) \rightarrow^p 0, \quad \text{where } \mathbf{g} \sim \mathcal{N}(\mathbf{0}_p, \mathbf{I}_p). \quad (41)$$

where $e_f(x; \tau) = \arg \min_u \frac{1}{2\tau} (x - u)^2 + f(u)$ is the Moreau envelope of f . Let $L(c, \tau) := \mathbb{E}[e_\rho(cZ + W; \tau) - \rho(W)]$ and assume $Z \sim N(0, 1)$, $W \sim F_\epsilon$, $Z \perp\!\!\!\perp W$. Then, by (Thrapoulidis et al., 2018, Theorem 3.1) (with $F = 0$ and taking for instance the signal distribution to be a single point mass at 0), if the convex-concave optimization

$$\inf_{\alpha \geq 0, \tau_g > 0} \sup_{\beta \geq 0, \tau_h > 0} \frac{\beta \tau_g}{2} + \frac{1}{\gamma} \cdot L(\alpha, \frac{\tau_g}{\beta}) - \frac{\alpha \tau_h}{2} - \frac{\alpha^2 \beta^2}{2\tau_h}, \quad (42)$$

admits a unique solution α_* , we have $\|\hat{\mathbf{h}}_\mu\|^2 \rightarrow^p \alpha_*^2$. Note that the stationary condition of the above convex-concave optimization is the nonlinear system of equations (15) (cf. (Thrapoulidis et al., 2018, Section 5.1)), while Theorem 7 guarantees the uniqueness and existence of the solution to (15). Therefore, we obtain $\|\hat{\mathbf{h}}_\mu\|_2^2 \rightarrow^p \alpha_*^2$ where α_* is the unique solution to (15) (we add a star to denote the unique solution here, to avoid confusion with the variable α of the minimization in (42)). It remains to show (41). Note that the Moreau envelope $e_f(\mathbf{x}; \tau)$ of the convex function $\mathbf{x} \in \mathbb{R}^p \mapsto n^{-c} \|\mathbf{x}\|^2 / 2$ is given by

$$\forall \mathbf{x} \in \mathbb{R}^p, \forall \tau > 0, \quad e_f(\mathbf{x}; \tau) = \min_{\mathbf{u} \in \mathbb{R}^p} \frac{\|\mathbf{x} - \mathbf{u}\|^2}{2\tau} + \frac{n^{-c}}{2} \|\mathbf{u}\|^2 = \frac{n^{-c}}{2(1 + n^{-c}\tau)} \|\mathbf{x}\|^2,$$

so that for all $a \in \mathbb{R}$ and $\tau > 0$, we have

$$\frac{e_{f_\mu}(a\mathbf{g}; \tau) - f_\mu(\mathbf{0}_p)}{n} = \frac{n^{-c}}{1 + n^{-c}\tau} \frac{a^2 \|\mathbf{g}\|^2}{n} = \frac{n^{-c}}{1 + n^{-c}\tau} \frac{a^2}{n} \sum_{i=1}^n g_i^2 \rightarrow^p 0 \cdot a^2 = 0,$$

where we have used the weak law of large number to the iid sum $n^{-1} \sum_{i=1}^n g_i^2$ with $\mathbb{E}[g_i^2] = 1$. This finishes the proof of (41).

It remains to show $\|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\| = o_p(1)$. Now we verify $\|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\|^2 \leq \|\hat{\mathbf{h}}\|^2 - \|\hat{\mathbf{h}}_\mu\|^2$. Letting $\mathcal{L} : \mathbb{R}^p \rightarrow \mathbb{R}$ be the convex function $\mathcal{L}(\mathbf{h}) := \sum_{i=1}^n \rho(\epsilon_i - \mathbf{g}_i^\top \mathbf{h})$ then $\hat{\mathbf{h}}$ and $\hat{\mathbf{h}}_\mu$ solve

$$\hat{\mathbf{h}} \in \arg \min_{\mathbf{h} \in \mathbb{R}^p} \mathcal{L}(\mathbf{h}), \quad \hat{\mathbf{h}}_\mu \in \arg \min_{\mathbf{h} \in \mathbb{R}^p} \mathcal{L}(\mathbf{h}) + \frac{\mu}{2} \|\mathbf{h}\|^2$$

Note in passing that $\hat{\mathbf{h}}_\mu$ is also a minimizer of the convex function $\mathcal{F} : \mathbb{R}^p \rightarrow \mathbb{R}$

$$\hat{\mathbf{h}}_\mu \in \arg \min_{\mathbf{h} \in \mathbb{R}^p} \mathcal{F}(\mathbf{h}) \quad \text{where} \quad \mathcal{F}(\mathbf{h}) := \mathcal{L}(\mathbf{h}) + \frac{\mu}{2} \left(\|\mathbf{h}\|^2 - \|\mathbf{h} - \hat{\mathbf{h}}_\mu\|^2 \right),$$

since the gradient of $\mathbf{h} \mapsto \|\mathbf{h} - \hat{\mathbf{h}}_\mu\|^2$ is $\mathbf{0}_p$ at $\hat{\mathbf{h}}_\mu$ and $\hat{\mathbf{h}}_\mu$ satisfies the new KKT condition $\nabla \mathcal{L}(\hat{\mathbf{h}}_\mu) = \mathbf{0}_p$. Then, $\mathcal{F}(\hat{\mathbf{h}}_\mu) \leq \mathcal{F}(\hat{\mathbf{h}})$ and $\mathcal{L}(\hat{\mathbf{h}}) \leq \mathcal{L}(\hat{\mathbf{h}}_\mu)$ yield

$$\begin{aligned} 0 &\geq \mathcal{F}(\hat{\mathbf{h}}_\mu) - \mathcal{F}(\hat{\mathbf{h}}) = \frac{\mu}{2} (\|\hat{\mathbf{h}}_\mu\|^2 - \|\hat{\mathbf{h}}\|^2 + \|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\|^2) + \mathcal{L}(\hat{\mathbf{h}}_\mu) - \mathcal{L}(\hat{\mathbf{h}}) \\ &\geq \frac{\mu}{2} (\|\hat{\mathbf{h}}_\mu\|^2 - \|\hat{\mathbf{h}}\|^2 + \|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\|^2) + 0, \end{aligned}$$

so that $\|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\|^2 \leq \|\hat{\mathbf{h}}\|^2 - \|\hat{\mathbf{h}}_\mu\|^2$. Combined with $\|\hat{\mathbf{h}}\|^2, \|\hat{\mathbf{h}}_\mu\|^2 \rightarrow^p \alpha^2 < +\infty$, we have $\|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\|^2 = o_p(1)$ and conclude the proof. \blacksquare

Lemma 4 *Suppose that (ρ, F_ϵ) satisfies Assumption 1-2. Then, under $\mu = n^{-c}$ for some $c > 0$, we have*

$$\|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\|_2 = o_P(n^{\frac{1-c}{2}}) \quad \text{and} \quad \|\boldsymbol{\psi}\|^2 - \|\boldsymbol{\psi}_\mu\|^2 = o_P(n^{1-\frac{c}{2}}),$$

where $\boldsymbol{\psi} = \boldsymbol{\psi}(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}})$ and $\boldsymbol{\psi}_\mu = \boldsymbol{\psi}(\boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu)$.

Proof Note that $\mathbf{G}^\top \boldsymbol{\psi}_\mu = n\mu \hat{\mathbf{h}}_\mu$ and $\mathbf{G}^\top \boldsymbol{\psi} = \mathbf{0}_p$ by the KKT conditions. Then, using Lemma 7 that is introduced later, $\|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\|^2$ can be bounded from above as

$$\begin{aligned} \|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\|^2 &\leq (\boldsymbol{\psi}_\mu - \boldsymbol{\psi})^\top (\mathbf{G}\hat{\mathbf{h}} - \mathbf{G}\hat{\mathbf{h}}_\mu) \quad \text{by Lemma 7 with } \mathbf{u} = \boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\mu \text{ and } \mathbf{v} = \boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}} \\ &\leq (n\mu \hat{\mathbf{h}}_\mu - \mathbf{0}_p)^\top (\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu) \quad \text{by } \mathbf{G}^\top \boldsymbol{\psi}_\mu = n\mu \hat{\mathbf{h}}_\mu \text{ and } \mathbf{G}^\top \boldsymbol{\psi} = \mathbf{0}_p \\ &\leq n\mu \|\hat{\mathbf{h}}_\mu\| \|\hat{\mathbf{h}} - \hat{\mathbf{h}}_\mu\| \quad \text{by the Cauchy-Schwarz inequality} \\ &= o_P(n^{1-c}) \quad \text{by Lemma 3 and } \mu = n^{-c}, \end{aligned} \tag{43}$$

which finishes the proof for $\|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\|^2$. For the bound of $\|\boldsymbol{\psi}_\mu\|^2 - \|\boldsymbol{\psi}\|^2$, the Cauchy-Schwarz inequality implies

$$|\|\boldsymbol{\psi}\|^2 - \|\boldsymbol{\psi}_\mu\|^2| = |(\boldsymbol{\psi} - \boldsymbol{\psi}_\mu)^\top (\boldsymbol{\psi} + \boldsymbol{\psi}_\mu)| \leq \|\boldsymbol{\psi} + \boldsymbol{\psi}_\mu\| \|\boldsymbol{\psi} - \boldsymbol{\psi}_\mu\| \leq 2\sqrt{n} \|\boldsymbol{\psi}\|_\infty \|\boldsymbol{\psi} - \boldsymbol{\psi}_\mu\|.$$

Since $\|\boldsymbol{\psi}\|_\infty < +\infty$ by Assumption 1 and we have shown that $\|\boldsymbol{\psi} - \boldsymbol{\psi}_\mu\|^2 = o_P(n^{1-c})$, we obtain $|\|\boldsymbol{\psi}\|^2 - \|\boldsymbol{\psi}_\mu\|^2| = o_P(n^{1-\frac{c}{2}})$. This finishes the proof. \blacksquare

Lemma 5 *Suppose that (ρ, F_ϵ) satisfies Assumption 1-2. Then, under $\mu = n^{-c}$ for some $c \in (0, 1)$, we have*

$$\text{tr}[\mathbf{V}_\mu]^2 - \text{tr}[\mathbf{V}]^2 = o_P(n^{2-\frac{c}{2}}),$$

where $\mathbf{V} = (\partial/\partial \boldsymbol{\epsilon})\boldsymbol{\psi}$ and $\mathbf{V}_\mu = (\partial/\partial \boldsymbol{\epsilon})\boldsymbol{\psi}_\mu$.

Proof By Assumption 2, we can write $\boldsymbol{\epsilon} = \mathbf{z} + \boldsymbol{\delta}$ where $\mathbf{z} \perp\!\!\!\perp \boldsymbol{\delta}$ and $(z_i)_{i=1}^n$ has i.i.d. density $\exp(-\phi(z))$. By the chain rule, $\mathbf{V} - \mathbf{V}_\mu$ can be written as

$$\mathbf{V} - \mathbf{V}_\mu = \frac{\partial}{\partial \boldsymbol{\epsilon}}(\boldsymbol{\psi} - \boldsymbol{\psi}_\mu) = \frac{\partial \mathbf{z}}{\partial \boldsymbol{\epsilon}} \cdot \frac{\partial}{\partial \mathbf{z}}(\boldsymbol{\psi} - \boldsymbol{\psi}_\mu) = \mathbf{I}_n \cdot \frac{\partial}{\partial \mathbf{z}}(\boldsymbol{\psi} - \boldsymbol{\psi}_\mu) = \frac{\partial}{\partial \mathbf{z}}(\boldsymbol{\psi} - \boldsymbol{\psi}_\mu).$$

Then, Theorem 6 applied with $\mathbf{f} = \boldsymbol{\psi} - \boldsymbol{\psi}_\mu$ yields

$$\begin{aligned} & \mathbb{E}[(\phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi}) - \text{tr}[\mathbf{V} - \mathbf{V}_\mu])^2] & (44) \\ &= \mathbb{E}\left[\sum_{i=1}^n \phi''(z_i)(\psi_i - (\psi_\mu)_i)^2 + \text{tr}[(\mathbf{V} - \mathbf{V}_\mu)^2]\right] & \text{by Theorem 6} \\ &\leq \|\phi''\|_\infty^2 \mathbb{E}[\|\boldsymbol{\psi} - \boldsymbol{\psi}_\mu\|^2] + n \mathbb{E}[\|\mathbf{V} - \mathbf{V}_\mu\|_{op}^2] & \text{using } \text{tr}[\mathbf{M}] \leq n\|\mathbf{M}\|_{op} \text{ for } \mathbf{M} \in \mathbb{R}^{n \times n} \\ &\leq 4\|\phi''\|_\infty^2 \|\boldsymbol{\psi}\|_\infty^2 n + 4n & \text{using } \|\mathbf{V}\|_{op}, \|\mathbf{V}_\mu\|_{op} \leq 1 \text{ from (30)} \\ &= C_6(\rho, F_\epsilon)n, & \text{since } \|\phi''\|_\infty, \|\boldsymbol{\psi}\|_\infty < +\infty \text{ by Assumption 1-2,} \end{aligned}$$

so that $|\text{tr}[\mathbf{V} - \mathbf{V}_\mu] - \phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi})| = O_P(n^{1/2})$. From this bound and the triangle inequality, it follows that

$$\begin{aligned} |\text{tr}[\mathbf{V} - \mathbf{V}_\mu]| &\leq |\phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi})| + |\text{tr}[\mathbf{V} - \mathbf{V}_\mu] - \phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\mu - \boldsymbol{\psi})| & (45) \\ &\leq \|\phi'(\mathbf{z})\| \|\boldsymbol{\psi} - \boldsymbol{\psi}_\mu\| + O_P(n^{1/2}) & \text{by the Cauchy-Schwarz} \\ &= O_P(n^{1/2}) o_P(n^{\frac{1-c}{2}}) + O_P(n^{1/2}) & \text{by Lemma 4 and 8} \\ &= o_P(n^{1-\frac{c}{2}}) & \text{by } c < 1. \end{aligned}$$

Finally, using $\|\mathbf{V}\|_{op}, \|\mathbf{V}_\mu\|_{op} \leq 1$ from (30), we have

$$|\text{tr}[\mathbf{V}]^2 - \text{tr}[\mathbf{V}_\mu]^2| = |\text{tr}[\mathbf{V}] - \text{tr}[\mathbf{V}_\mu]| \cdot |\text{tr}[\mathbf{V}] + \text{tr}[\mathbf{V}_\mu]| \leq o_P(n^{1-\frac{c}{2}}) \cdot 2n = o_P(n^{2-\frac{c}{2}}),$$

which concludes the proof. \blacksquare

In order to relax the assumption on the noise and prove Proposition 1, we now provide here a modification of this argument to allow for a vanishing smooth noise component z_i : the conclusion $|\text{tr}[\mathbf{V} - \mathbf{V}_\mu]|/n \xrightarrow{P} 0$ still holds if the linear model noise is $\epsilon_i = \sigma_n z_i + \delta_i$ with vanishing σ_n (depending on n) provided that $\sigma_n \geq \sqrt{\mu}$ and $\sigma_n \sqrt{n} \rightarrow +\infty$.

Lemma 6 (Vanishing noise component) *Assume that the noise $\boldsymbol{\epsilon}$ in the linear model has iid coordinates of the form $\epsilon_i = \sigma_n z_i + \delta_i$ where δ_i and z_i are independent and z_i has density $z \mapsto \exp(-\phi(z))$ with twice-continuously differentiable ϕ and $\sup_{x \in \mathbb{R}} |\phi''(z)| < +\infty$. If $\sigma_n \geq \sqrt{\mu}$, we have*

$$\mathbb{E}\left[\frac{|\text{tr}[\mathbf{V} - \mathbf{V}_\mu]|}{n}\right] \leq \mathbb{E}\left[\min\left(2, \frac{\|\phi'(\mathbf{z})\|}{\sqrt{n}} \|\mathbf{h}_\mu\| \|\mathbf{h} - \mathbf{h}_\mu\|\right)\right] + \frac{2\|\phi''\|_\infty \|\boldsymbol{\psi}\|_\infty}{\sigma_n \sqrt{n}} + \frac{2}{\sqrt{n}}. \quad (46)$$

If we take $\mu = n^{-c}$ with $c \in (0, 1)$ as in Lemma 3, then $\|\phi'(\mathbf{z})\| n^{-\frac{1}{2}} \|\mathbf{h}_\mu\| \|\mathbf{h} - \mathbf{h}_\mu\| \xrightarrow{P} 0$ by Lemma 3 and Lemma 8 for the first term, while $\sigma_n \sqrt{n} \geq \sqrt{n^{1-c}} \rightarrow +\infty$ for the second term. This implies that $|\text{tr}[\mathbf{V} - \mathbf{V}_\mu]|/n$ converges to 0 in L1 and in probability for all $\sigma_n \geq n^{-c/2}$.

Proof This is a modification of the argument of Lemma 5 to obtain a condition on how small the amplitude σ_n of the smooth noise z_i is allowed. Similarly to (44), by Theorem 6 we have

$$\begin{aligned} \mathbb{E}[(\phi'(\mathbf{z})^\top (\frac{\boldsymbol{\psi}_\mu - \boldsymbol{\psi}}{\sigma_n}) - \text{tr}[\mathbf{V} - \mathbf{V}_\mu])^2] &= \mathbb{E}\left[\sum_{i=1}^n \phi''(z_i) \frac{(\psi_i - (\psi_\mu)_i)^2}{\sigma_n^2} + \text{tr}[(\mathbf{V} - \mathbf{V}_\mu)^2]\right] \\ &\leq 4n(\|\phi''\|_\infty^2 \|\boldsymbol{\psi}\|_\infty^2 / \sigma_n^2 + 1) \end{aligned}$$

using $\|\mathbf{V}\|_{op}, \|\mathbf{V}_\mu\|_{op} \leq 1$ from (30). By the triangle inequality combined with $0 \leq \text{tr}[\mathbf{V}] \leq n$ and $0 \leq \text{tr}[\mathbf{V}_\mu] \leq n$,

$$|\text{tr}[\mathbf{V} - \mathbf{V}_\mu]| \leq \min\left(2n, \frac{\|\phi'(\mathbf{z})\| \|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\|}{\sigma_n}\right) + \left|\phi'(\mathbf{z})^\top (\frac{\boldsymbol{\psi}_\mu - \boldsymbol{\psi}}{\sigma_n}) - \text{tr}[\mathbf{V} - \mathbf{V}_\mu]\right|$$

Recalling (43) we have $\|\boldsymbol{\psi}_\mu - \boldsymbol{\psi}\|^2 \leq n\mu \|\mathbf{h}_\mu\| \|\mathbf{h} - \mathbf{h}_\mu\|$ if $\sigma_n \geq \sqrt{\mu}$ we obtain

$$\mathbb{E}[|\text{tr}[\mathbf{V} - \mathbf{V}_\mu]|/n] \leq \mathbb{E}[\min(2, (\|\phi'(\mathbf{z})\| n^{-1/2}) \|\mathbf{h}_\mu\| \|\mathbf{h} - \mathbf{h}_\mu\|)] + 2(\|\phi''\|_\infty^2 \|\boldsymbol{\psi}\|_\infty^2 / \sigma_n^2 + 1)^{1/2} / \sqrt{n}.$$

Using $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, we obtain (46). ■

Lemma 7 *If $\psi : \mathbb{R} \rightarrow \mathbb{R}$ is 1-Lipschitz and nondecreasing, it holds that*

$$\|\psi(\mathbf{u}) - \psi(\mathbf{v})\|^2 \leq (\psi(\mathbf{u}) - \psi(\mathbf{v}))^\top (\mathbf{u} - \mathbf{v})$$

for all $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$.

Proof We have $0 \leq \psi(x) - \psi(y) \leq (x - y)$ for all $x, y \in \mathbb{R}$ such that $x \leq y$. Switching the role of x and y in this inequality, we have $(\psi(x) - \psi(y))^2 \leq (\psi(x) - \psi(y))(x - y)$ for all $x, y \in \mathbb{R}$. Applying this inequality with $(x, y) = (u_i, v_i)$ for each $i \in [n]$, we conclude the proof. ■

Lemma 8 *Suppose $(z_i)_{i=1}^n$ has iid density $\exp(-\phi(z))$ where $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is twice continuously differentiable and $\|\phi''\|_\infty < +\infty$. Then $\mathbb{E}[\phi'(z_i)^2] = \mathbb{E}[\phi''(z_i)]$ and $\lim_{n \rightarrow \infty} \mathbb{P}(\|\phi'(\mathbf{z})\|^2 \leq n(\|\phi''\|_\infty + 1)) = 1$.*

Proof Note $\mathbb{E}[\phi'(z_i)^2] = \mathbb{E}[\phi''(z_i)] \leq \|\phi''\|_\infty < +\infty$ by integration part. By the weak law of large number, we have $n^{-1} \|\phi'(\mathbf{z})\|^2 = n^{-1} \sum_{i=1}^n \phi'(z_i)^2 \xrightarrow{p} \mathbb{E}[\phi'(z)^2] = \mathbb{E}[\phi''(z_i)]$. Thus, we obtain

$$\mathbb{P}(n^{-1} \|\phi'(\mathbf{z})\|^2 > \|\phi''\|_\infty + 1) \leq \mathbb{P}(n^{-1} \|\phi'(\mathbf{z})\|^2 > \mathbb{E}[\phi''(z_i)] + 1) \rightarrow 0,$$

which completes the proof. ■

B.3 Lower bound on the trace of \mathbf{V}

In this section, we derive a lower bound of $\text{tr}[\mathbf{V}]/n$.

Lemma 9 *Suppose that (ρ, F_ϵ) satisfy Assumption 1-2. Let $\alpha(\rho, F_\epsilon, \gamma)$ be the unique solution to the nonlinear system of equations (15) for $(\rho, F_\epsilon, \gamma)$ and $\eta = \eta(\rho)$ be some positive constant such that $\psi(x)^2/\|\psi\|_\infty^2 + \psi'(x) \geq \eta^2$ for almost every $x \in \mathbb{R}$ (we can always take such $\eta > 0$ thanks to Assumption 1(3)). Then, we have*

$$\mathbb{P}\left(\frac{\text{tr}[\mathbf{V}]^2}{n^2} \geq B(\gamma, \rho, F_\epsilon)\right) \rightarrow 1, \text{ where } B(\gamma, \rho, F_\epsilon) = C(\gamma) \cdot \min(\eta(\rho)^4, \frac{\|\psi\|_\infty^2 \cdot \eta(\rho)^2}{\alpha^2(\rho, F_\epsilon, \gamma)}) > 0.$$

Proof Recall that we have shown in Lemma 5 that

$$\text{tr}(\mathbf{V})^2/n^2 - \text{tr}(\mathbf{V}_\mu)^2/n^2 = o_P(n^{-\frac{c}{2}}) = o_P(1) \text{ under } \mu = n^{-c} \text{ for all } c \in (0, 1).$$

Thus, it suffices to show the lower bound for $\text{tr}[\mathbf{V}_\mu]^2/n^2$ for some $c \in (0, 1)$, say $c = 1/4$. Define the matrices \mathbf{D} and $\tilde{\mathbf{D}}_\mu$ as

$$\mathbf{D}_\mu := \text{diag}\{\psi'(\epsilon - \mathbf{G}\hat{\mathbf{h}}_\mu)\}, \quad \tilde{\mathbf{D}}_\mu := \|\psi\|_\infty^{-2} \text{diag}\{\psi^2(\epsilon - \mathbf{G}\hat{\mathbf{h}}_\mu)\}.$$

Thanks to $\|\psi\|_{\text{lip}} \leq 1$ and $\psi(x)^2/\|\psi\|_\infty^2 + \psi'(x) \geq \eta^2$, we have

$$\text{tr}(\mathbf{D}_\mu) \leq n, \quad \tilde{\mathbf{D}}_\mu + \mathbf{D}_\mu \succeq \eta^2 \mathbf{I}_n$$

where \succeq is the positive semi-definite order. By the derivative formula (Theorem 4), we have

$$\begin{aligned} \mathbf{V}_\mu &= \mathbf{D}_\mu - \mathbf{D}_\mu \mathbf{G} (\mathbf{G}^\top \mathbf{D}_\mu \mathbf{G} + n\mu \mathbf{I}_p)^{-1} \mathbf{G}^\top \mathbf{D}_\mu \\ &= \mathbf{D}_\mu^{1/2} (\mathbf{I}_n - \mathbf{D}_\mu^{1/2} \mathbf{G} (\mathbf{G}^\top \mathbf{D}_\mu \mathbf{G} + n\mu \mathbf{I}_p)^{-1} \mathbf{G}^\top \mathbf{D}_\mu^{1/2}) \mathbf{D}_\mu \\ &= \mathbf{D}_\mu^{1/2} \mathbf{H}_\mu \mathbf{D}_\mu^{1/2}, \end{aligned}$$

where $\mathbf{H}_\mu := \mathbf{I}_n - \mathbf{D}_\mu^{1/2} \mathbf{G} (\mathbf{G}^\top \mathbf{D}_\mu \mathbf{G} + n\mu \mathbf{I}_p)^{-1} \mathbf{G}^\top \mathbf{D}_\mu^{1/2}$ is positive semi-definite, and satisfies $\|\mathbf{H}_\mu\|_{op} \leq 1$ and $\text{tr}(\mathbf{H}_\mu) \geq n - p$. Then, simple algebra yields

$$\begin{aligned} \text{tr}[\mathbf{V}_\mu] &= \text{tr}(\mathbf{D}_\mu^{1/2} \mathbf{H}_\mu \mathbf{D}_\mu^{1/2}) \\ &= \text{tr}(\mathbf{H}_\mu^{1/2} \mathbf{D}_\mu \mathbf{H}_\mu^{1/2}) && \text{by } \text{tr}(\mathbf{D}_\mu^{1/2} \mathbf{H}_\mu \mathbf{D}_\mu^{1/2}) = \text{tr}(\mathbf{H}_\mu \mathbf{D}_\mu) = \text{tr}(\mathbf{H}_\mu^{1/2} \mathbf{D}_\mu \mathbf{H}_\mu^{1/2}) \\ &\geq \text{tr}(\mathbf{H}_\mu^{1/2} (\eta^2 \mathbf{I}_n - \tilde{\mathbf{D}}_\mu) \mathbf{H}_\mu^{1/2}) && \text{by } \tilde{\mathbf{D}}_\mu + \mathbf{D}_\mu \succeq \eta^2 \mathbf{I}_n \\ &= \eta^2 \text{tr}(\mathbf{H}_\mu) - \text{tr}(\mathbf{H}_\mu^{1/2} \tilde{\mathbf{D}}_\mu \mathbf{H}_\mu^{1/2}) \\ &\geq \eta^2 (n - p) - \text{tr}(\tilde{\mathbf{D}}_\mu) \|\mathbf{H}_\mu\|_{op} && \text{by } \text{tr}[\mathbf{H}_\mu] \geq n - p \text{ and } \text{tr}[\mathbf{AB}] \leq \text{tr}[\mathbf{A}] \cdot \|\mathbf{B}\|_{op} \\ &\geq \eta^2 (n - p) - \|\psi\|_\infty^{-2} \|\boldsymbol{\psi}_\mu\|^2 \cdot 1. \end{aligned}$$

Thus, under the event $\Omega := \{\|\boldsymbol{\psi}_\mu\|^2 \leq n\|\psi\|_\infty^2 \eta^2 (1 - \gamma)/2\}$, we have

$$\mathbf{1}_\Omega \cdot n^{-2} \text{tr}[\mathbf{V}_\mu]^2 \geq \mathbf{1}_\Omega \cdot 4^{-1} \eta^4 (1 - \gamma)^2 \quad (47)$$

It remains to bound $n^{-2} \text{tr}[\mathbf{V}_\mu]^2$ from below under the complement Ω^c . Lemma 1 implies

$$n^{-2} \text{tr}(\mathbf{V}_\mu)^2 \|\hat{\mathbf{h}}_\mu\|^2 - n^{-2} p \|\boldsymbol{\psi}_\mu\|^2 = O_P(n^{-c}) = o_P(1),$$

where $n^{-2}\text{tr}(\mathbf{V}_\mu)^2\|\hat{\mathbf{h}}_\mu\|^2 = n^{-2}\text{tr}(\mathbf{V}_\mu)^2\alpha^2 + o_P(1)$ from $\text{tr}[\mathbf{V}_\mu]^2/n^2 \leq \|\mathbf{V}_\mu\|_{op}^2 \leq 1$ and $\|\hat{\mathbf{h}}_\mu\|^2 \xrightarrow{P} \alpha^2 > 0$ by Lemma 3. Thus, we have $n^{-2}\text{tr}[\mathbf{V}_\mu]^2\alpha^2 - \gamma n^{-1}\|\boldsymbol{\psi}_\mu\|^2 = o_P(1)$, so that

$$\mathbb{P}\left(\frac{\text{tr}[\mathbf{V}_\mu]^2}{n^2} - \frac{\gamma}{n\alpha^2}\|\boldsymbol{\psi}_\mu\|^2 \geq -\frac{\|\psi\|_\infty^2\eta^2(1-\gamma)\gamma}{4\alpha^2}\right) \rightarrow 1.$$

Thereby, under $\Omega^c = \{\|\boldsymbol{\psi}_\mu\|^2 \geq n\|\psi\|_\infty^2\eta^2(1-\gamma)/2\}$, we have

$$\mathbb{P}\left(\mathbf{1}_{\Omega^c} \frac{\text{tr}(\mathbf{V}_\mu)^2}{n^2} \geq \mathbf{1}_{\Omega^c} \left(\frac{\|\psi\|_\infty^2\eta^2(1-\gamma)\gamma}{2\alpha^2} - \frac{\|\psi\|_\infty^2\eta^2(1-\gamma)\gamma}{4\alpha^2}\right) = \mathbf{1}_{\Omega^c} \frac{\|\psi\|_\infty^2\eta^2(1-\gamma)\gamma}{4\alpha^2}\right) \rightarrow 1 \quad (48)$$

Consequently, (47) and (47) yield

$$\mathbb{P}\left(\frac{\text{tr}(\mathbf{V}_\mu)^2}{n^2} \geq \min\left\{\frac{\eta^4(1-\gamma)^2}{4}, \frac{\|\psi\|_\infty^2\eta^2(1-\gamma)\gamma}{4\alpha^2}\right\}\right) \rightarrow 1,$$

which completes the proof. \blacksquare

B.4 Proof of Theorem 1

Recall that the goal is to show (28). Putting Lemma 1, Lemma 3, and Lemma 4, Lemma 5 together, we have

$$\begin{aligned} & \frac{p}{n^2}\|\boldsymbol{\psi}\|^2 - \frac{\text{tr}[\mathbf{V}]^2}{n^2}\|\hat{\mathbf{h}}\|^2 \quad (49) \\ &= \left(\frac{p}{n^2}\|\boldsymbol{\psi}_\mu\|^2 - \frac{\text{tr}[\mathbf{V}_\mu]^2}{n^2}\|\hat{\mathbf{h}}\|^2\right) + \frac{p}{n^2}(\|\boldsymbol{\psi}\|_2^2 - \|\boldsymbol{\psi}_\mu\|^2) + \frac{\text{tr}[\mathbf{V}_\mu]^2 - \text{tr}[\mathbf{V}]^2}{n^2}\|\hat{\mathbf{h}}_\mu\|^2 - \frac{\text{tr}[\mathbf{V}]^2}{n^2}(\|\hat{\mathbf{h}}\|^2 - \|\hat{\mathbf{h}}_\mu\|^2) \\ &= O_P(n^{-c}) + o_P(n^{-c/2}) + o_P(n^{-c/2}) - \frac{\text{tr}[\mathbf{V}]^2}{n^2}o_P(1) \\ &= o_P(n^{-c/2}) - \frac{\text{tr}[\mathbf{V}]^2}{n^2}o_P(1). \end{aligned}$$

Multiplying $(\text{tr}(\mathbf{V})^2/n^2)^{-1}$, which is $O_P(1)$ by Lemma 9, we conclude the proof.

B.5 Proof of Proposition 1

The strategy is exactly the same as for Theorem 1, with Lemma 5 replaced by Lemma 6 to allow for a vanishing smooth component in the noise. More precisely, set $c = 1/4$ so that $\mu = n^{-1/4}$. The conclusions of Lemma 1, Lemma 3 and Lemma 4 still hold for $\epsilon_i = \sigma_n z_i + \delta_i$ with $\sigma_n \rightarrow 0$, where α is now the solution to the system (7) with $W \sim \tilde{F}$, since the contribution of terms involving $\sigma_n z_i$ in (42) is 0 due to $\sigma_n \rightarrow 0$ and the fact that the loss is Lipschitz by Assumption 1. The condition $\sigma_n \geq \sqrt{\mu}$ of Lemma 6 is satisfied thanks to the assumption $\sigma_n \geq n^{-1/8}$ in Proposition 1. The argument and conclusion of Lemma 9 still hold as well, with $B(\gamma, \rho, F_\epsilon)$ replaced by $B(\gamma, \rho, \tilde{F})$, again because α is now the solution to the system (7) with $W \sim \tilde{F}$. The decomposition (49), and the same argument as for Theorem 1, thus yield the conclusion of Proposition 1.

B.6 Proof of Corollary 1

Fix $\epsilon > 0$. Letting $R_k = \|\Sigma^{1/2}(\hat{\beta}_k - \beta^*)\|^2$ for each $k \in [K] := \{1, \dots, K\}$, Theorem 1 or Proposition 1 imply $\mathbb{P}(|R_k - \hat{R}_k| \geq \epsilon/2) \rightarrow 0$ for all k . Note that $\hat{R}_{\hat{k}} < \hat{R}_k$ for all k by the definition of \hat{k} . Then, we have

$$\begin{aligned}
 \mathbb{P}(R_{\hat{k}} > \min_{k \in [K]} R_k + \epsilon) &\leq \mathbb{P}(\exists k \in [K], R_{\hat{k}} > R_k + \epsilon) \\
 &\leq \sum_{k=1}^K \mathbb{P}(R_{\hat{k}} > R_k + \epsilon) && \text{by the union bound} \\
 &= \sum_{k=1}^K \mathbb{P}(R_{\hat{k}} > R_k + \epsilon, \hat{R}_k \geq \hat{R}_{\hat{k}}) && \text{by definition of } \hat{k} \\
 &\leq \sum_{k=1}^K \sum_{l=1}^K \mathbb{P}(R_l > R_k + \epsilon, \hat{R}_k > \hat{R}_l) && \text{by the union bound.}
 \end{aligned}$$

Here, $\mathbb{P}(R_l > R_k + \epsilon, \hat{R}_k > \hat{R}_l) \rightarrow 0$ as $n \rightarrow \infty$ by the following argument:

$$\begin{aligned}
 \mathbb{P}(R_l > R_k + \epsilon, \hat{R}_k > \hat{R}_l) &\leq \mathbb{P}(R_l - \hat{R}_l - (R_k - \hat{R}_k) > \hat{R}_k - \hat{R}_l + \epsilon > \epsilon) \\
 &\leq \mathbb{P}(|R_l - \hat{R}_l| > \epsilon/2) + \mathbb{P}(|R_k - \hat{R}_k| > \epsilon/2) \rightarrow 0.
 \end{aligned}$$

Since K is finite, we conclude $\mathbb{P}(R_{\hat{k}} > \min_{k \in [K]} R_k + \epsilon) \rightarrow 0$. Since $\epsilon > 0$ is taken arbitrarily, this finishes the proof.

Appendix C. Proofs for Section 3

Throughout of this section, we fix (ρ, F_ϵ) that satisfy Assumption 1, Assumption 2 and Assumption 3. For all $\lambda > 0$, define ρ_λ and ψ_λ as

$$\forall \lambda > 0, \quad \rho_\lambda(\cdot) := \lambda^2 \rho(\cdot/\lambda), \quad \psi_\lambda(\cdot) := \rho'_\lambda(\cdot) = \lambda \psi(\cdot/\lambda).$$

Note in passing that $\|\psi_\lambda\|_{\text{lip}} = \|\psi\|_{\text{lip}} = 1$ and $\|\psi_\lambda\|_\infty = \lambda \|\psi\|_\infty < +\infty$. Define $R(\lambda)$, $\hat{R}(\lambda)$, and $\alpha(\lambda)$ as

$$\begin{aligned}
 R(\lambda) &:= \|\Sigma^{1/2}(\hat{\beta}_\lambda - \beta^*)\|^2 && \text{with } \hat{\beta}_\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^n \rho(y_i - \mathbf{x}_i^\top \beta) \\
 \hat{R}(\lambda) &:= p \frac{\|\psi_\lambda\|^2}{\text{tr}[\mathbf{V}_\lambda]^2} && \text{with } \mathbf{V}_\lambda = \frac{\partial \psi_\lambda}{\partial \mathbf{y}} \in \mathbb{R}^{n \times n}, \quad \psi_\lambda := \psi_\lambda(\mathbf{y} - \mathbf{X} \hat{\beta}_\lambda) \\
 \alpha(\lambda) &:= \alpha(\rho_\lambda, F_\epsilon, \gamma) && \text{(the solution to (15) with } \rho = \rho_\lambda).
 \end{aligned}$$

C.1 Proof of Proposition 2

$\hat{R}(\lambda) = R(\lambda) + o_P(1)$ by Theorem 1, while $R(\lambda) \rightarrow^p \alpha^2(\lambda)$ by Thrampoulidis et al. (2018) and Theorem 7. It remains to show the upper bound of $\alpha^2(\lambda)$. Inequality (37) in Theorem 7 with $\rho = \rho_\lambda$ implies that $\alpha(\lambda)$ is bounded from above as

$$\alpha(\lambda) \leq \frac{Q_{F_\epsilon}(r^2 c_\gamma)}{r c_\gamma} + \frac{b}{c_\gamma}$$

where $r \in (0, 1]$ and $b \geq 0$ are any constant such that the coercivity condition

$$\|\rho_\lambda\|_{\text{lip}}^{-1}(\rho_\lambda(x) - \rho_\lambda(0)) \geq r(|x| - b) \quad \text{for all } x \in \mathbb{R}$$

is satisfied. Now we verify that

$$r = \frac{\min\{|\psi(\pm 1)|\}}{\|\rho\|_{\text{lip}}} \quad \text{and} \quad b = \lambda \cdot \frac{\max\{|(\mp 1)\psi(\pm 1) - \rho(\pm 1) + \rho(0)|\}}{\min\{|\psi(\pm 1)|\}}$$

satisfy the coercivity condition. By the convexity, evaluating the derivative at λ , using $\rho_\lambda(\cdot) = \lambda^2 \rho(\cdot/\lambda)$ and $\psi_\lambda(\cdot) = \lambda \psi(\cdot/\lambda)$, we have

$$\rho_\lambda(x) \geq \rho_\lambda(\lambda) + (x - \lambda)\psi_\lambda(\lambda) = \lambda\psi(1)x - \lambda^2\psi(1) + \lambda^2\rho(1)$$

for all $x \in \mathbb{R}$. By the same argument, evaluating the derivative at $-\lambda$ gives

$$\rho_\lambda(x) \geq \rho_\lambda(-\lambda) + (x + \lambda)\psi_\lambda(-\lambda) = \lambda\psi(-1)x + \lambda^2\psi(-1) + \lambda^2\rho(-1).$$

Thanks to $\{0\} = \arg \min_x \rho(x)$ by Assumption 1, $\psi(-1) < 0 < \psi(1)$ holds, and hence,

$$\rho_\lambda(x) - \rho_\lambda(0) \geq \lambda \min(|\psi(1)|, |\psi(-1)|)|x| - \lambda^2 \max(|\psi(1) - \rho(1) + \rho(0)|, |-\psi(-1) - \rho(-1) + \rho(0)|).$$

Dividing the both sides by $\|\rho_\lambda\|_{\text{lip}} = \lambda\|\rho\|_{\text{lip}}$, we obtain

$$\frac{\rho_\lambda(x) - \rho_\lambda(0)}{\|\rho_\lambda\|_{\text{lip}}} \geq \frac{\min(|\psi(\pm 1)|)}{\|\rho\|_{\text{lip}}}|x| - \lambda \frac{\max(|(\mp 1)\psi(\pm 1) - \rho(\pm 1) + \rho(0)|)}{\|\rho\|_{\text{lip}}}.$$

This means that the pair of (r, b) specified above satisfies the coercivity condition. Substituting this to the upper bound of $\alpha(\lambda)$, since the dependence on λ only comes from b , we obtain

$$\alpha(\lambda) \leq C_7(\rho, F_\epsilon, \gamma)(1 + \lambda),$$

which finishes the proof.

C.2 Proof of Theorem 2

Thanks to Proposition 2, we have

$$\alpha^2(\lambda) - \alpha^2(\tilde{\lambda}) = \hat{R}(\lambda) - \hat{R}(\tilde{\lambda}) + o_P(1),$$

so it suffices to show a suitable smoothness of the map $\lambda \mapsto \hat{R}(\lambda)$.

Lemma 10 *We have*

$$\forall \lambda > 0, \quad |\text{tr}[\mathbf{V}_\lambda]| \leq n, \quad \mathbb{P}(n^2 \cdot [\text{tr} \mathbf{V}_\lambda]^{-2} \leq C_8(\rho, F_\epsilon, \gamma)(1 + \lambda^{-2})) \rightarrow 1.$$

Proof $\text{tr}[\mathbf{V}_\lambda]/n \leq 1$ immediately follows from Proposition 3 and $\|\psi_\lambda\|_{\text{lip}} = \|\psi\|_{\text{lip}} = 1$. Note in passing that $\psi_\lambda(\cdot) = \lambda\psi(\cdot/\lambda)$ satisfies $\|\psi_\lambda\|_\infty = \lambda\|\psi\|_\infty < +\infty$, and for almost every x ,

$$\begin{aligned} \psi_\lambda^2(x)/\|\psi_\lambda\|_\infty^2 + \psi'_\lambda(x) &= \psi(x/\lambda)^2/\|\psi\|_\infty + \psi'(x/\lambda) \\ &\geq \eta^2 > 0 \end{aligned}$$

thanks to Assumption 1(3) where η is a constant that only depends on ρ . Thus, Lemma 9 with $\rho = \rho_\lambda$ implies

$$\mathbb{P}\left(\frac{\text{tr}[\mathbf{V}_\lambda]^2}{n^2} \geq C(\gamma) \min\left(\eta^4, \frac{\lambda^2 \|\psi\|_\infty^2 \eta^2}{\alpha^2(\lambda)}\right)\right) \rightarrow 1.$$

Putting this lower bound and $\alpha^2(\lambda) \leq C'(\gamma, \rho, F_\epsilon)(\lambda^2 + 1)$ by Proposition 2 together, we have

$$\frac{\text{tr}[\mathbf{V}_\lambda]^2}{n^2} \geq C(\gamma) \min\left(\eta^4, \frac{\lambda^2 \|\psi\|_\infty^2 \eta^2}{C'(\rho, F_\epsilon, \gamma)(1 + \lambda^2)}\right) \geq C(\gamma) \min\left(\eta^4, \frac{\eta^2 \|\psi\|_\infty^2}{C'(\rho, F_\epsilon, \gamma)}\right) \frac{\lambda^2}{1 + \lambda^2}$$

with high probability. This finishes the proof. \blacksquare

Lemma 11 *We have the followings for all $\lambda, \tilde{\lambda} > 0$:*

$$\begin{aligned} \mathbb{P}\left(\|\boldsymbol{\psi}_\lambda - \boldsymbol{\psi}_{\tilde{\lambda}}\| \leq \sqrt{n} C_9(\gamma, \rho, F_\epsilon) |\lambda - \tilde{\lambda}|^{1/2} (1 + \lambda^{1/2} + \tilde{\lambda}^{1/2})\right) &\rightarrow 1, \\ \mathbb{P}\left(\left|\|\boldsymbol{\psi}_\lambda\|^2 - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|^2\right| \leq n C_{10}(\gamma, \rho, F_\epsilon) |\lambda - \tilde{\lambda}|^{1/2} (1 + \lambda^{3/2} + \tilde{\lambda}^{3/2})\right) &\rightarrow 1. \end{aligned}$$

Proof Note that the second display immediately follows from the first display since

$$\begin{aligned} \left|\|\boldsymbol{\psi}_\lambda\|^2 - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|^2\right| &= \left|\|\boldsymbol{\psi}_\lambda\| - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|\right| (\|\boldsymbol{\psi}_\lambda\| + \|\boldsymbol{\psi}_{\tilde{\lambda}}\|) \\ &\leq \left|\|\boldsymbol{\psi}_\lambda\| - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|\right| \sqrt{n} (\|\psi_\lambda\|_\infty + \|\psi_{\tilde{\lambda}}\|_\infty) \\ &= \left|\|\boldsymbol{\psi}_\lambda\| - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|\right| \sqrt{n} \|\psi\|_\infty (\lambda + \tilde{\lambda}) \quad \text{by } \|\boldsymbol{\psi}_\lambda\|_\infty = \lambda \|\psi\|_\infty \end{aligned}$$

It remains to bound $\left|\|\boldsymbol{\psi}_\lambda\| - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|\right|$. Below, we write $\mathbf{r}_\lambda = \boldsymbol{\epsilon} - \mathbf{G}\hat{\mathbf{h}}_\lambda$ with $\hat{\mathbf{h}}_\lambda = \boldsymbol{\Sigma}^{1/2}(\hat{\boldsymbol{\beta}}_\lambda - \boldsymbol{\beta}_*)$ so that $\boldsymbol{\psi}_\lambda = \psi_\lambda(\mathbf{r}_\lambda)$ and $\|\hat{\mathbf{h}}_\lambda\|^2 = R(\lambda)$. By Assumption 3, there exists some constant $c(\rho) > 0$ that only depends on ρ such that for all $\lambda, \tilde{\lambda} > 0$,

$$\sup_{\mathbf{u} \in \mathbb{R}^n} \|\psi_\lambda(\mathbf{u}) - \psi_{\tilde{\lambda}}(\mathbf{u})\|_2 \leq \sqrt{n} \cdot \sup_{x \in \mathbb{R}} |\psi_\lambda(x) - \psi_{\tilde{\lambda}}(x)| \leq \sqrt{n} c(\rho) |\lambda - \tilde{\lambda}|,$$

so that $\|\boldsymbol{\psi}_\lambda - \boldsymbol{\psi}_{\tilde{\lambda}}\|$ can be upper bounded from above as

$$\|\boldsymbol{\psi}_\lambda - \boldsymbol{\psi}_{\tilde{\lambda}}\| \leq \|\psi_\lambda(\mathbf{r}_\lambda) - \psi_\lambda(\mathbf{r}_{\tilde{\lambda}})\| + \|\psi_\lambda(\mathbf{r}_{\tilde{\lambda}}) - \psi_{\tilde{\lambda}}(\mathbf{r}_{\tilde{\lambda}})\| \leq \|\psi_\lambda(\mathbf{r}_\lambda) - \psi_\lambda(\mathbf{r}_{\tilde{\lambda}})\| + \sqrt{n} c(\rho) |\lambda - \tilde{\lambda}|,$$

Here, using the KKT conditions $\mathbf{G}^\top \boldsymbol{\psi}_\lambda(\mathbf{r}_\lambda) = \mathbf{G}^\top \boldsymbol{\psi}_{\tilde{\lambda}}(\mathbf{r}_{\tilde{\lambda}}) = \mathbf{0}_p$, we bound the first from above as

$$\begin{aligned} \|\psi_\lambda(\mathbf{r}_\lambda) - \psi_\lambda(\mathbf{r}_{\tilde{\lambda}})\|^2 &\leq (\psi_\lambda(\mathbf{r}_\lambda) - \psi_\lambda(\mathbf{r}_{\tilde{\lambda}}))^\top (\mathbf{r}_\lambda - \mathbf{r}_{\tilde{\lambda}}) && \text{by Lemma 7 and } \|\psi_\lambda\|_{\text{lip}} = 1 \\ &= (\psi_\lambda(\mathbf{r}_\lambda) - \psi_\lambda(\mathbf{r}_{\tilde{\lambda}}))^\top \mathbf{G}(\hat{\mathbf{h}}_{\tilde{\lambda}} - \hat{\mathbf{h}}_\lambda) \\ &= (\psi_{\tilde{\lambda}}(\mathbf{r}_{\tilde{\lambda}}) - \psi_\lambda(\mathbf{r}_{\tilde{\lambda}}))^\top \mathbf{G}(\hat{\mathbf{h}}_{\tilde{\lambda}} - \hat{\mathbf{h}}_\lambda) && \text{by } \mathbf{G}^\top \boldsymbol{\psi}_\lambda(\mathbf{r}_\lambda) = \mathbf{G}^\top \boldsymbol{\psi}_{\tilde{\lambda}}(\mathbf{r}_{\tilde{\lambda}}) = \mathbf{0}_p \\ &\leq \|\psi_{\tilde{\lambda}}(\mathbf{r}_{\tilde{\lambda}}) - \psi_\lambda(\mathbf{r}_{\tilde{\lambda}})\| \|\mathbf{G}\|_{op} (\|\hat{\mathbf{h}}_\lambda\| + \|\hat{\mathbf{h}}_{\tilde{\lambda}}\|) \\ &\leq \sqrt{n} c(\rho) |\lambda - \tilde{\lambda}| \|\mathbf{G}\|_{op} (\|\hat{\mathbf{h}}_\lambda\| + \|\hat{\mathbf{h}}_{\tilde{\lambda}}\|) \end{aligned}$$

Therefore, using $\|\hat{\mathbf{h}}_\lambda\| \rightarrow^p \alpha(\lambda)$ and $\alpha(\lambda) \leq C_{11}(\rho, \gamma F_\epsilon)(1 + \lambda)$ by Proposition 2, we have

$$\begin{aligned} \|\boldsymbol{\psi}_\lambda - \boldsymbol{\psi}_{\tilde{\lambda}}\|^2 &\leq 2\sqrt{nc}(\rho)|\lambda - \tilde{\lambda}|\|\mathbf{G}\|_{op}(\|\hat{\mathbf{h}}_\lambda\| + \|\hat{\mathbf{h}}_{\tilde{\lambda}}\|) + 2nc(\rho)^2|\lambda - \tilde{\lambda}|^2 \\ &\leq C_{12}(\rho) \cdot n \cdot (n^{-1/2}\|\mathbf{G}\|_{op} + 1)(|\lambda - \tilde{\lambda}|(\|\hat{\mathbf{h}}_\lambda\| + \|\hat{\mathbf{h}}_{\tilde{\lambda}}\|) + |\lambda - \tilde{\lambda}|^2) \\ &\leq C_{13}(\rho, \gamma, F_\epsilon) \cdot n(|\lambda - \tilde{\lambda}|(1 + \lambda + \tilde{\lambda}) + |\lambda - \tilde{\lambda}|^2) \\ &\leq C_{14}(\rho, \gamma, \epsilon) \cdot n|\lambda - \tilde{\lambda}| \cdot (1 + \lambda + \tilde{\lambda}) \end{aligned}$$

with high probability. This finishes the proof. \blacksquare

Lemma 12 *Suppose that (ρ, F_ϵ) satisfy Assumption 1, Assumption 2 and Assumption 3. Then, we have for all $\lambda, \tilde{\lambda} > 0$ that*

$$\mathbb{P}\left(n^{-2}|\text{tr}[\mathbf{V}_\lambda]^2 - \text{tr}[\mathbf{V}_{\tilde{\lambda}}]^2| \leq C_{15}(\rho, F_\epsilon, \gamma)|\lambda - \tilde{\lambda}|^{1/2}(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2})\right) \rightarrow 1,$$

Proof By the same argument in the proof of Lemma 5, the variant of second order Stein's formula leads to

$$\text{tr}[\mathbf{V}(\lambda)] - \text{tr}[\mathbf{V}(\tilde{\lambda})] = \partial_{\mathbf{z}}(\boldsymbol{\psi}_\lambda - \boldsymbol{\psi}_{\tilde{\lambda}}) = \phi'(\mathbf{z})^\top(\boldsymbol{\psi}_\lambda - \boldsymbol{\psi}_{\tilde{\lambda}}) + O_P(n^{1/2}),$$

where ϕ is the density function in Assumption 2 such that $\|\phi''\|_\infty < +\infty$. By Lemma 11, Lemma 8, and Cauchy–Schwarz, we have

$$|\text{tr}[\mathbf{V}_\lambda] - \text{tr}[\mathbf{V}_{\tilde{\lambda}}]| \leq \|\phi'(\mathbf{z})\| \|\boldsymbol{\psi}_\lambda - \boldsymbol{\psi}_{\tilde{\lambda}}\| + O_P(\sqrt{n}) \leq nC_{16}(\rho, F_\epsilon, \gamma)|\lambda - \tilde{\lambda}|^{1/2}(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2})$$

with high probability. Combined with $\text{tr}[\mathbf{V}_\lambda] \in [0, n]$ from Lemma 10, we conclude the proof. \blacksquare

Proof [Proof of Theorem 2] Lemma 10, Lemma 11, and Lemma 12 imply that the followings events hold simultaneously with high probability:

$$\begin{aligned} n^2 \text{tr}[\mathbf{V}_{\lambda'}]^{-2} &\leq C_{17}(\rho, F_\epsilon, \gamma)(1 + \lambda'^{-2}) && \text{for } \lambda' \in \{\lambda, \tilde{\lambda}\}, \\ n^{-1}|\|\boldsymbol{\psi}_\lambda\|^2 - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|^2| &\leq C_{18}(\rho, F_\epsilon, \gamma)|\lambda - \tilde{\lambda}|^{1/2}(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2}), \\ n^{-2}|\text{tr}[\mathbf{V}_\lambda]^2 - \text{tr}[\mathbf{V}_{\tilde{\lambda}}]^2| &\leq C_{19}(\rho, F_\epsilon, \gamma)|\lambda - \tilde{\lambda}|^{1/2}(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2}). \end{aligned}$$

Thus, we have

$$\begin{aligned} |\hat{R}(\lambda) - \hat{R}(\tilde{\lambda})| &= \left| \frac{p\|\boldsymbol{\psi}_\lambda\|^2}{\text{tr}[\mathbf{V}_\lambda]^2} - \frac{p\|\boldsymbol{\psi}_{\tilde{\lambda}}\|^2}{\text{tr}[\mathbf{V}_{\tilde{\lambda}}]^2} \right| \\ &= \frac{n^2}{\text{tr}[\mathbf{V}_\lambda]^2} \frac{n}{p} \left| \frac{\|\boldsymbol{\psi}_\lambda\|^2}{n} - \frac{\text{tr}[\mathbf{V}_\lambda]^2}{\text{tr}[\mathbf{V}_{\tilde{\lambda}}]^2} \frac{\|\boldsymbol{\psi}_{\tilde{\lambda}}\|^2}{n} \right| \\ &\leq \frac{n^2}{\text{tr}[\mathbf{V}_\lambda]^2} \frac{n}{p} \left(\frac{|\|\boldsymbol{\psi}_\lambda\|^2 - \|\boldsymbol{\psi}_{\tilde{\lambda}}\|^2|}{n} + \frac{n^2}{\text{tr}[\mathbf{V}_{\tilde{\lambda}}]^2} \frac{|\text{tr}[\mathbf{V}_\lambda]^2 - \text{tr}[\mathbf{V}_{\tilde{\lambda}}]^2|}{n^2} \frac{\|\boldsymbol{\psi}_{\tilde{\lambda}}\|^2}{n} \right) \\ &\leq C(\rho, \gamma, F_\epsilon)|\lambda - \tilde{\lambda}|^{1/2}(1 + \lambda^{-2}) \left(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2} + (1 + \tilde{\lambda}^{-2})(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2})\tilde{\lambda} \right) \end{aligned}$$

Combined with $\hat{R}(\lambda) \xrightarrow{p} \alpha^2(\lambda)$ for each $\lambda \in (0, \infty)$, we have the following for all $\lambda, \tilde{\lambda} > 0$:

$$|\alpha^2(\lambda) - \alpha^2(\tilde{\lambda})| \leq C(\rho, \gamma, F_\epsilon) |\lambda - \tilde{\lambda}|^{1/2} L(\lambda, \tilde{\lambda}),$$

where $L(\lambda, \tilde{\lambda}) := (1 + \lambda^{-2})(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2} + (1 + \tilde{\lambda}^{-2})(1 + \lambda^{1/2} + \tilde{\lambda}^{1/2})\tilde{\lambda})$. Since $\sup_{\lambda, \tilde{\lambda} \in [\lambda_{\min}, \lambda_{\max}]} L(\lambda, \tilde{\lambda})$ is bounded from above by some positive constant $C(\lambda_{\min}, \lambda_{\max})$, we conclude the proof. ■

C.3 Proof of Theorem 3

Lemma 13 *Fix $I = [\lambda_{\min}, \lambda_{\max}]$ for some $0 < \lambda_{\min} < \lambda_{\max} < +\infty$. For all $N \in \mathbb{N}$, let $I_N = \{\lambda_{\min}^{i/N} \cdot \lambda_{\max}^{1-i/N} : i = 0, 1, \dots, N\}$ be the finite grid over I and take $\hat{\lambda}_N \in \arg \min_{\lambda \in I_N} \hat{R}(\lambda)$. Then, for all $\delta > 0$, there exists $N_\delta = N(\rho, F_\epsilon, \gamma, \lambda_{\min}, \lambda_{\max}, \delta) \in \mathbb{N}$ such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}(|\alpha^2(\hat{\lambda}_{N_\delta}) - \alpha^2(\lambda_{\text{opt}})| < \delta) = \lim_{n \rightarrow \infty} \mathbb{P}(|R(\hat{\lambda}_{N_\delta}) - \alpha^2(\lambda_{\text{opt}})| < \delta) \rightarrow 1,$$

where $\hat{\lambda}_{N_\delta} \in \arg \min_{\lambda \in I_{N_\delta}} \hat{R}(\lambda)$ and $\lambda_{\text{opt}} \in \arg \min_{\lambda \in I} \alpha^2(\lambda)$.

Proof [Proof of Lemma 13] First, we prove $\mathbb{P}(|\alpha^2(\hat{\lambda}_{N_\delta}) - \alpha^2(\lambda_{\text{opt}})| < \delta) \rightarrow 1$. Let us take some $N \in \mathbb{N}$ to be specified later. Consider the decomposition

$$0 < \alpha^2(\hat{\lambda}_N) - \alpha^2(\lambda_{\text{opt}}) = \alpha^2(\hat{\lambda}_N) - \min_{\lambda \in I_N} \alpha^2(\lambda) + \min_{\lambda \in I_N} \alpha^2(\lambda) - \alpha^2(\lambda_{\text{opt}}).$$

Now, the cardinality of I_N is finite and $\hat{R}(\lambda) = \alpha^2(\lambda) + o_P(1)$ for all $\lambda > 0$ by Proposition 2. Then, by the same argument of Appendix B.6, where the role of R is replaced by α^2 , we have

$$\alpha^2(\hat{\lambda}_N) - \min_{\lambda \in I_N} \alpha^2(\lambda) = o_P(1),$$

so in particular $\lim_{n \rightarrow \infty} \mathbb{P}(\alpha^2(\hat{\lambda}_N) - \min_{\lambda \in I_N} \alpha^2(\lambda) < \delta/2) = 1$. Below, we show $\min_{\lambda \in I_N} \alpha^2(\lambda) - \alpha^2(\lambda_{\text{opt}}) \leq \delta/2$ for some $N = N_\delta$. By the definition of I_N , we can take $\lambda_N^* \in I_N$ such that $\lambda_N^* \leq \lambda_{\text{opt}} < (\lambda_{\max}/\lambda_{\min})^{1/N} \lambda_N^*$, so that

$$|\lambda_{\text{opt}} - \lambda_N^*| \leq [(\lambda_{\max}/\lambda_{\min})^{1/N} - 1] \lambda_N^* \leq [(\lambda_{\max}/\lambda_{\min})^{1/N} - 1] \lambda_{\max}.$$

On the other hand, Theorem 2 implies that

$$\text{for all } \lambda, \tilde{\lambda} \in I = [\lambda_{\min}, \lambda_{\max}], \quad |\alpha^2(\lambda) - \alpha^2(\tilde{\lambda})| \leq L |\lambda - \tilde{\lambda}|^{1/2},$$

where $L = L(\rho, F_\epsilon, \gamma, \lambda_{\min}, \lambda_{\max}) > 0$ is some positive constant. From the above displays, it follows that

$$\begin{aligned} 0 < \min_{\lambda \in I_N} \alpha^2(\lambda) - \alpha^2(\lambda_{\text{opt}}) & \quad \text{by } \lambda_{\text{opt}} \in \arg \min_{\lambda \in I} \alpha^2(\lambda) \text{ and } I_N \subset I \\ & \leq \alpha^2(\lambda_N^*) - \alpha^2(\lambda_{\text{opt}}) & \quad \text{by } \lambda_N^* \in I_N \\ & \leq L |\lambda_N^* - \lambda_{\text{opt}}|^{1/2} & \quad \text{by } \lambda_N^*, \lambda_{\text{opt}} \in I \\ & \leq L \{ \lambda_{\max} [(\lambda_{\max}/\lambda_{\min})^{1/N} - 1] \}^{1/2} & \quad \text{by } |\lambda_{\text{opt}} - \lambda_N^*| \leq [(\lambda_{\max}/\lambda_{\min})^{1/N} - 1] \lambda_{\max} \\ & \leq \delta/2 & \quad \text{if } (\lambda_{\max}/\lambda_{\min})^{1/N} \leq 1 + \delta^2/(4L\lambda_{\max}). \end{aligned}$$

Therefore, if we take $N = N_\delta := \lceil \log(\lambda_{\max}/\lambda_{\min})/\log\{1 + \delta^2/(4L\lambda_{\max})\} \rceil$, we obtain $\mathbb{P}(|\alpha^2(\hat{\lambda}_{N_\delta}) - \alpha^2(\lambda_{\text{opt}})| < \delta) \rightarrow 1$.

Next, we prove $\mathbb{P}(|R(\hat{\lambda}_N) - \alpha^2(\lambda_{\text{opt}})| \leq \epsilon) \rightarrow 1$. By the triangle inequality, it follows that

$$|R(\hat{\lambda}_N) - \alpha^2(\lambda_{\text{opt}})| \leq |R(\hat{\lambda}_N) - \min_{\lambda \in I_N} R(\lambda)| + |\min_{\lambda \in I_N} R(\lambda) - \min_{\lambda \in I_N} \alpha^2(\lambda)| + |\min_{\lambda \in I_N} \alpha^2(\lambda) - \alpha^2(\lambda_{\text{opt}})|.$$

The first term is $o_P(1)$ by Corollary 1, while the second term is $o_P(1)$ as $R(\lambda) \rightarrow^P \alpha^2(\lambda)$ for all $\lambda > 0$ and the cardinality of I_N is finite. For the third term, we have shown that it is less than $\delta/2$ for $N = N_\delta$. This completes the proof of $\mathbb{P}(|R(\hat{\lambda}_{N_\delta}) - \alpha^2(\lambda_{\text{opt}})| \leq \delta) \rightarrow 1$. ■

Proof [Proof of Theorem 3] Our goal is to construct an array $(N_n)_{n=1}^\infty$ of integers such that

$$\forall \epsilon > 0, \quad \lim_{n \rightarrow \infty} \mathbb{P}(|\alpha^2(\hat{\lambda}_{N_n}) - \alpha^2(\lambda_{\text{opt}})| > \epsilon) + \mathbb{P}(|R(\hat{\lambda}_{N_n}) - \alpha^2(\lambda_{\text{opt}})| > \epsilon) = 0. \quad (50)$$

By Lemma 13 with $\delta = 1/k$, there exists a function $N : \mathbb{N} \rightarrow \mathbb{N}$ such that

$$\forall k \in \mathbb{N}, \quad \lim_{n \rightarrow \infty} \underbrace{\mathbb{P}(|\alpha^2(\hat{\lambda}_{N(k)}) - \alpha^2(\lambda_{\text{opt}})| \geq 1/k) + \mathbb{P}(|R(\hat{\lambda}_{N(k)}) - \alpha^2(\lambda_{\text{opt}})| \geq 1/k)}_{:= p_{n,k}} = 0. \quad (51)$$

Thus, if we can find some function $\varphi : \mathbb{N} \rightarrow \mathbb{N}$ such that $\lim_{n \rightarrow \infty} \varphi(n) \rightarrow \infty$ and $\lim_{n \rightarrow \infty} p_{n, \varphi(n)} = 0$, the array $\{N_n\}_{n=1}^\infty := \{N \circ \varphi(n)\}_{n=1}^\infty$ satisfies (50), thereby completing the proof. Below we construct such $n \mapsto \varphi(n)$.

By (51), there exists an array of integers $\{n_k\}_{k=1}^\infty$ such that

$$\forall k \in \mathbb{N}, \quad \forall n \geq n_k, \quad |p_{n,k}| \leq 1/k.$$

Here, we assume without loss of generality that $\{n_k\}_k$ is strictly increasing; otherwise redefine recursively as $\tilde{n}_k := \max(\tilde{n}_{k-1}, n_k) + 1$. For this $\{n_k\}_{k=1}^\infty$, define the step function $\varphi : \mathbb{N} \rightarrow \mathbb{N}$ as $\varphi(n) := \sum_{k=1}^\infty k \mathbf{1}\{n_k \leq n < n_{k+1}\}$. Then, we have $\lim_{n \rightarrow \infty} \varphi(n) = +\infty$, and $|p_{n,k}| \leq 1/k$ holds for $k = \varphi(n)$ since $n \geq n_k$ by construction. Therefore, we obtain

$$\forall n \geq n_1, \quad |p_{n, \varphi(n)}| \leq 1/\varphi(n).$$

Since $\varphi(n) \rightarrow +\infty$ as $n \rightarrow \infty$, we have $\lim_{n \rightarrow \infty} p_{n, \varphi(n)} = 0$. This finishes the proof. ■

Appendix D. Additional numerical simulation

D.1 Relaxing Assumption 2

In the same setting as Section 4, we change the noise distribution to be the following discrete distribution so that Assumption 2 is not satisfied:

$$\epsilon_i \stackrel{\text{iid}}{\sim} 3 \cdot \lceil \text{t-dist}(\text{df} = 2) \rceil$$

where $\lceil x \rceil = \max\{n \in \mathbb{Z} : n \leq x\}$. Figure 5 implies that our result still holds in this setting.

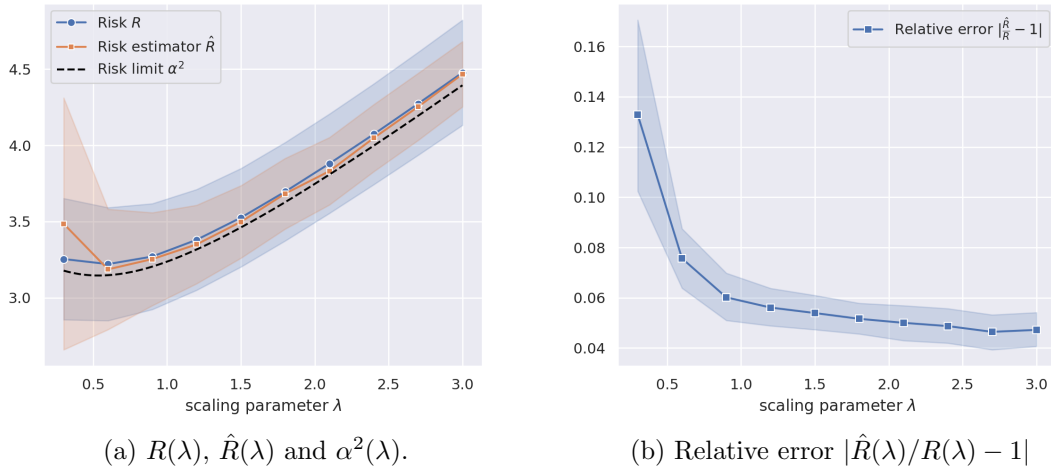


Figure 5: Plot of the out-of-sample error $R(\lambda)$ and estimator $\hat{R}(\lambda)$ over 100 repetitions, with $n = 4000$, $p = 1200$, the Huber loss for different values of the scale parameters λ . The noise distribution is $3[\text{t-dist}(\text{df} = 2)]$. $\alpha^2(\lambda)$ is the theoretical limit given by the nonlinear system (7).

D.2 Robustness of risk estimation across covariate distributions

We vary the distribution of the covariate \mathbf{X} from Gaussian to Rademacher, uniform and Student’s t-distribution in Figure 6. In each of these settings, we confirm the effectiveness of the risk estimator, suggesting that our results extend to a broader class of covariate distributions.

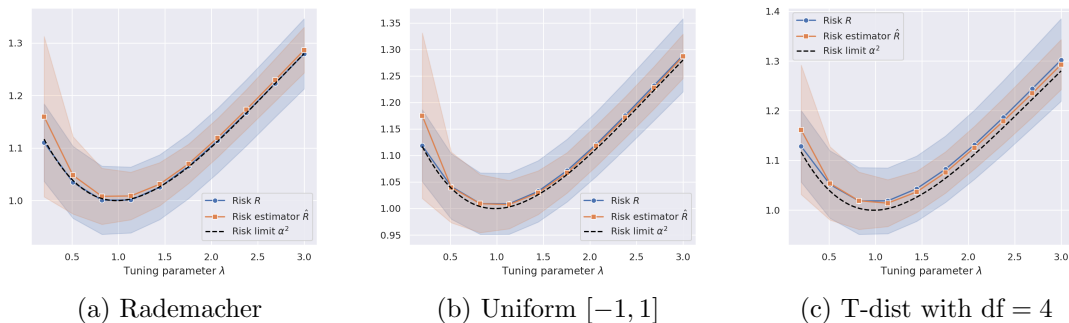


Figure 6: We change the distribution of the design matrix \mathbf{X} to Uniform, Student’s T, and Rademacher distributions with proper normalization so that the variance is 1. Simulation parameter: the noise distribution $F_\epsilon = \text{t-dist}(\text{df} = 2)$, sample size $n = 4000$, feature $p = 1200$, scaling parameter $\lambda = 1$, 100 repetition.

D.3 Risk behavior as dimensionality ratio $\gamma = p/n$ approaches 1

We plot the risk and its estimator as $\gamma = p/n$ varies in Figure 7. The results show that the risk estimation becomes increasingly unstable as $\gamma \rightarrow 1^-$.

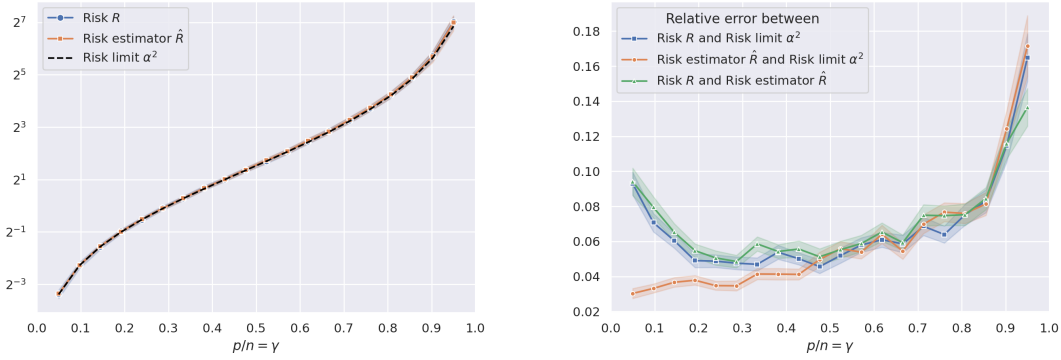


Figure 7: We plot the risk and its estimator for different ratio $\gamma = p/n$. Simulation parameters are as follows: sample size $n = 4000$, noise distribution $F_\epsilon = \text{t-dist}(\text{df} = 2)$, scaling parameter $\lambda = 1$, with 100 repetition. The relative error between A and B is $|A/B - 1|$.

D.4 Convergence of risk estimation variance with increasing sample size

We plot the risk and its estimator as n varies for different values of $\gamma = p/n$: $\gamma = 0.25$ in Figure 8 and $\gamma = 0.5$ in Figure 9. In both cases, the variances of the risk and risk estimator decrease as $n \rightarrow +\infty$. Interestingly, these figures suggest that the variances decays in \sqrt{n} -rate.

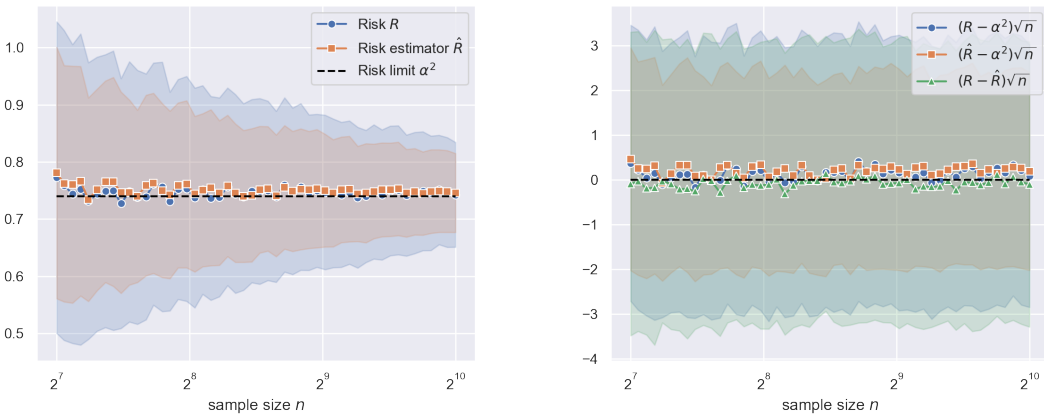


Figure 8: We plot the risk and its estimator for different n . Simulation parameters are as follows: $p/n = \gamma = 0.25$, $F_\epsilon = \text{t-dist}(\text{df} = 2)$, scaling parameter $\lambda = 1$, with 1000 repetition.

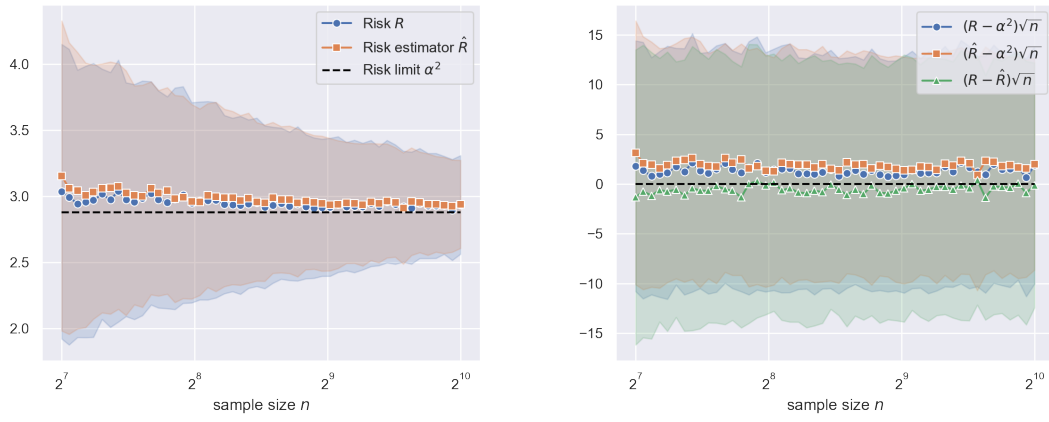


Figure 9: We plot the risk and its estimator for different n . Simulation parameters are as follows: $p/n = \gamma = 0.5$, $F_\epsilon = \text{t-dist}(\text{df} = 2)$, scaling parameter $\lambda = 1$, with 1000 repetition.