



HAL
open science

Robust linear least squares regression

Jean-Yves Audibert, Olivier Catoni

► **To cite this version:**

| Jean-Yves Audibert, Olivier Catoni. Robust linear least squares regression. 2010. hal-00522534v1

HAL Id: hal-00522534

<https://hal.science/hal-00522534v1>

Preprint submitted on 30 Sep 2010 (v1), last revised 17 Sep 2011 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ROBUST LINEAR LEAST SQUARES REGRESSION

BY JEAN-YVES AUDIBERT^{*,†}, AND OLIVIER CATONI^{‡,§}

We consider the problem of robustly predicting as well as the best linear combination of d given functions in least squares regression, and variants of this problem including constraints on the parameters of the linear combination. For the ridge estimator and the ordinary least squares estimator, and their variants, we provide new risk bounds of order d/n without logarithmic factor unlike some standard results, where n is the size of the training data. We also provide a new estimator with better deviations in presence of heavy-tailed noise. It is based on truncating differences of losses in a min-max framework and satisfies a d/n risk bound both in expectation and in deviations. The key common surprising factor of these results is the absence of exponential moment condition on the output distribution while achieving exponential deviations. All risk bounds are obtained through a PAC-Bayesian analysis on truncated differences of losses. Experimental results strongly back up our truncated min-max estimator.

CONTENTS

1	Introduction	2
	Our statistical task	2
	Why should we be interested in this task	5
	Outline and contributions	6
2	Ridge regression and empirical risk minimization	7
3	A min-max estimator for robust estimation	9
	3.1 The min-max estimator and its theoretical guarantee	10
	3.2 The value of the uncentered kurtosis coefficient χ	12
	3.3 Computation of the estimator	14
	3.4 Synthetic experiments	17

^{*}Université Paris-Est, LIGM, Imagine, 6 avenue Blaise Pascal, 77455 Marne-la-Vallée, France, E-mail: audibert@imagine.enpc.fr

[†]CNRS/École Normale Supérieure/INRIA, LIENS, Willow – UMR 8548, 23 avenue d’Italie, 75214 Paris cedex 13, France.

[‡]École Normale Supérieure, CNRS – UMR 8553, Département de Mathématiques et Applications, 45 rue d’Ulm, 75230 Paris cedex 05, France, E-mail: olivier.catoni@ens.fr

[§] INRIA Paris-Rocquencourt - CLASSIC team.

AMS 2000 subject classifications: 62J05, 62J07

Keywords and phrases: Linear regression, Generalization error, Shrinkage, PAC-Bayesian theorems, Risk bounds, Robust statistics, Resistant estimators, Gibbs posterior distributions, Randomized estimators, Statistical learning theory

3.4.1	Noise distributions	17
3.4.2	Independent normalized covariates (INC(n, d))	18
3.4.3	Highly correlated covariates (HCC(n, d))	18
3.4.4	Trigonometric series (TS(n, d))	18
3.4.5	Experiments	19
4	Proofs	20
4.1	Main ideas of the proofs	20
4.1.1	Sub-exponential tails under a non-exponential moment assumption via truncation	20
4.1.2	Localized PAC-Bayesian inequalities to eliminate a logarithm factor	22
4.2	Proofs of Theorems 2.1 and 2.2	23
4.2.1	Proof of Theorem 2.1	30
4.2.2	Proof of Theorem 2.2	31
4.3	Proof of Theorem 3.1	33
	References	41
A	Experimental results for the min-max truncated estimator (Section 3.3)	42

1. Introduction.

Our statistical task. Let $Z_1 = (X_1, Y_1), \dots, Z_n = (X_n, Y_n)$ be $n \geq 2$ pairs of input-output and assume that each pair has been independently drawn from the same unknown distribution P . Let \mathcal{X} denote the input space and let the output space be the set of real numbers \mathbb{R} , so that P is a probability distribution on the product space $\mathcal{Z} \triangleq \mathcal{X} \times \mathbb{R}$. The target of learning algorithms is to predict the output Y associated with an input X for pairs $Z = (X, Y)$ drawn from the distribution P . The quality of a (prediction) function $f : \mathcal{X} \rightarrow \mathbb{R}$ is measured by the least squares *risk*:

$$R(f) \triangleq \mathbb{E}_{Z \sim P} \{[Y - f(X)]^2\}.$$

Through the paper, we assume that the output and all the prediction functions we consider are square integrable. Let Θ be a closed convex set of \mathbb{R}^d , and $\varphi_1, \dots, \varphi_d$ be d prediction functions. Consider the regression model

$$\mathcal{F} = \left\{ f_\theta = \sum_{j=1}^d \theta_j \varphi_j; (\theta_1, \dots, \theta_d) \in \Theta \right\}.$$

The best function f^* in \mathcal{F} is defined by

$$f^* = \sum_{j=1}^d \theta_j^* \varphi_j \in \operatorname{argmin}_{f \in \mathcal{F}} R(f).$$

Such a function always exists but is not necessarily unique. Besides it is unknown since the probability generating the data is unknown.

We will study the problem of predicting (at least) as well as function f^* . In other words, we want to deduce from the observations Z_1, \dots, Z_n a function \hat{f} having with high probability a risk bounded by the minimal risk $R(f^*)$ on \mathcal{F} plus a small remainder term, which is typically of order d/n up to a possible logarithmic factor. Except in particular settings (e.g., Θ is a simplex and $d \geq \sqrt{n}$), it is known that the convergence rate d/n cannot be improved in a minimax sense (see [11], and [12] for related results).

More formally, the target of the paper is to develop estimators \hat{f} for which the excess risk is controlled *in deviations*, i.e., such that for an appropriate constant $\kappa > 0$, for any $\varepsilon > 0$, with probability at least $1 - \varepsilon$,

$$(1.1) \quad R(\hat{f}) - R(f^*) \leq \kappa \frac{d + \log(\varepsilon^{-1})}{n}.$$

Note that by integrating the deviations (using the identity $\mathbb{E}W = \int_0^{+\infty} \mathbb{P}(W > t) dt$ which holds true for any nonnegative random variable W), Inequality (1.1) implies

$$(1.2) \quad \mathbb{E}R(\hat{f}) - R(f^*) \leq \kappa \frac{d + 1}{n}.$$

In this work, we do not assume that the function

$$f^{(\text{reg})} : x \mapsto \mathbb{E}[Y|X = x],$$

which minimizes the risk R among all possible measurable functions, belongs to the model \mathcal{F} . So we might have $f^* \neq f^{(\text{reg})}$ and in this case, bounds of the form

$$(1.3) \quad \mathbb{E}R(\hat{f}) - R(f^{(\text{reg})}) \leq C[R(f^*) - R(f^{(\text{reg})})] + \kappa \frac{d}{n},$$

with a constant C larger than 1 do not even ensure that $\mathbb{E}R(\hat{f})$ tends to $R(f^*)$ when n goes to infinity. This kind of bounds with $C > 1$ have been developed to analyze nonparametric estimators using linear approximation spaces, in which case the dimension d is a function of n chosen so that the bias term $R(f^*) - R(f^{(\text{reg})})$ has the order d/n of the estimation term (see [6] and references within). Here we intend to assess the generalization ability of the estimator even when the model is misspecified (namely when $R(f^*) > R(f^{(\text{reg})})$). Moreover we do not assume either that $Y - f^{(\text{reg})}(X)$ and X are independent or that Y has a low-tailed distribution: for the moment,

we just assume that $Y - f^*(X)$ admits a finite second order moment in order that the risk of f^* is finite.

Several risk bounds with $C = 1$ can be found in the litterature. A survey on these bounds is given in [1, Section 1]. Let us mention here the closest bound to what we are looking for. From the work of Birgé and Massart [3], we may derive the following risk bound for the empirical risk minimizer on a L^∞ ball (see Appendix B of [1]).

THEOREM 1.1. *Assume that \mathcal{F} has a diameter H for L^∞ -norm, i.e., for any f_1, f_2 in \mathcal{F} , $\sup_{x \in \mathcal{X}} |f_1(x) - f_2(x)| \leq H$ and there exists a function $f_0 \in \mathcal{F}$ satisfying the exponential moment condition:*

$$(1.4) \quad \text{for any } x \in \mathcal{X}, \quad \mathbb{E} \left\{ \exp \left[A^{-1} |Y - f_0(X)| \right] \mid X = x \right\} \leq M,$$

for some positive constants A and M . Let

$$\tilde{B} = \inf_{\phi_1, \dots, \phi_d} \sup_{\theta \in \mathbb{R}^d - \{0\}} \frac{\| \sum_{j=1}^d \theta_j \phi_j \|_\infty^2}{\| \theta \|_\infty^2}$$

where the infimum is taken with respect to all possible orthonormal basis of \mathcal{F} for the dot product $(f_1, f_2) \mapsto \mathbb{E} f_1(X) f_2(X)$ (when the set \mathcal{F} admits no basis with exactly d functions, we set $\tilde{B} = +\infty$). Then the empirical risk minimizer satisfies for any $\varepsilon > 0$, with probability at least $1 - \varepsilon$:

$$R(\hat{f}^{(\text{erm})}) - R(f^*) \leq \kappa(A^2 + H^2) \frac{d \log[2 + (\tilde{B}/n) \wedge (n/d)] + \log(\varepsilon^{-1})}{n},$$

where κ is a positive constant depending only on M .

The theorem gives exponential deviation inequalities of order at worse $d \log(n/d)/n$, and asymptotically, when n goes to infinity, of order d/n . This work will provide similar results under weaker assumption on the output distribution.

Notation. When $\Theta = \mathbb{R}^d$, the function f^* and the space \mathcal{F} will be written f_{lin}^* and \mathcal{F}_{lin} to emphasize that \mathcal{F} is the whole linear space spanned by $\varphi_1, \dots, \varphi_d$:

$$\mathcal{F}_{\text{lin}} = \text{span}\{\varphi_1, \dots, \varphi_d\} \quad \text{and} \quad f_{\text{lin}}^* \in \underset{f \in \mathcal{F}_{\text{lin}}}{\text{argmin}} R(f).$$

The Euclidean norm will simply be written as $\| \cdot \|$, and $\langle \cdot, \cdot \rangle$ will be its associated inner product. We will consider the vector valued function $\varphi : \mathcal{X} \rightarrow \mathbb{R}^d$ defined by $\varphi(X) = [\varphi_k(X)]_{k=1}^d$, so that for any $\theta \in \Theta$, we have

$$f_\theta(X) = \langle \theta, \varphi(X) \rangle.$$

The Gram matrix is the $d \times d$ -matrix $Q = \mathbb{E}[\varphi(X)\varphi(X)^T]$, and its smallest and largest eigenvalues will respectively be written as q_{\min} and q_{\max} . The empirical risk of a function f is

$$r(f) = \frac{1}{n} \sum_{i=1}^n [f(X_i) - Y_i]^2$$

and for $\lambda \geq 0$, the ridge regression estimator on \mathcal{F} is defined by $\hat{f}^{(\text{ridge})} = f_{\hat{\theta}^{(\text{ridge})}}$ with

$$\hat{\theta}^{(\text{ridge})} \in \arg \min_{\theta \in \Theta} \{r(f_\theta) + \lambda \|\theta\|^2\},$$

where λ is some nonnegative real parameter. In the case when $\lambda = 0$, the ridge regression $\hat{f}^{(\text{ridge})}$ is nothing but the empirical risk minimizer $\hat{f}^{(\text{erm})}$. Besides, the empirical risk minimizer when $\Theta = \mathbb{R}^d$ is also called the ordinary least squares estimator, and will be denoted $\hat{f}^{(\text{ols})}$.

In the same way, we introduce the optimal ridge function optimizing the expected ridge risk: $\tilde{f} = f_{\tilde{\theta}}$ with

$$(1.5) \quad \tilde{\theta} \in \arg \min_{\theta \in \Theta} \{R(f_\theta) + \lambda \|\theta\|^2\}.$$

Finally, let $Q_\lambda = Q + \lambda I$ be the ridge regularization of Q , where I is the identity matrix.

Why should we be interested in this task. There are four main reasons. First we intend to provide a non-asymptotic analysis of the parametric linear least squares method. Secondly, the task is central in nonparametric estimation for linear approximation spaces (piecewise polynomials based on a regular partition, wavelet expansions, trigonometric polynomials. . .)

Thirdly, it naturally arises in two-stage model selection. Precisely, when facing the data, the statistician has often to choose several models which are likely to be relevant for the task. These models can be of similar structures (like embedded balls of functional spaces) or on the contrary of very different nature (e.g., based on kernels, splines, wavelets or on parametric approaches). For each of these models, we assume that we have a learning scheme which produces a 'good' prediction function in the sense that it predicts as well as the best function of the model up to some small additive term. Then the question is to decide on how we use or combine/aggregate these schemes. One possible answer is to split the data into two groups, use the first group to train the prediction function associated with each model, and finally use the second group to build a prediction function which is as good as (i) the best of the previously learnt prediction functions, (ii) the best

convex combination of these functions or (iii) the best linear combination of these functions. This point of view has been introduced by Nemirovski in [8] and optimal rates of aggregation are given in [11] and references within. This paper focuses more on the linear aggregation task (even if (ii) enters in our setting), assuming implicitly here that the models are given in advance and are beyond our control and that the goal is to combine them appropriately.

Finally, in practice, the noise distribution often departs from the normal distribution. In particular, it can exhibit much heavier tails, and consequently induce highly non Gaussian residuals. It is then natural to ask whether classical estimators such as the ridge regression and the ordinary least squares estimators are sensitive to this type of noise, and whether we can design estimators robust to this type of noise.

Outline and contributions. Section 2 provides a new analysis of the ridge estimator and the ordinary least squares estimator, and their variants. Theorem 2.1 provides an asymptotic result for the ridge estimator while Theorem 2.2 gives a non asymptotic risk bound of the empirical risk minimizer, which is complementary to the theorems put in the survey section. In particular, the result has the benefit to hold for the ordinary least squares estimator and for heavy-tailed outputs. We show quantitatively that the ridge penalty leads to an implicit reduction of the input space dimension. Section 3 shows a non asymptotic d/n exponential deviation risk bound under weak moment conditions on the output Y and on the d -dimensional input representation $\varphi(X)$.

The main contribution of this paper is to show through a PAC-Bayesian analysis on truncated differences of losses that the output distribution does not need to have bounded conditional exponential moments in order for the excess risk of appropriate estimators to concentrate exponentially. Our results tend to say that truncation leads to more robust algorithms. Local robustness to contamination is usually invoked to advocate the removal of outliers, claiming that estimators should be made insensitive to small amounts of spurious data. Our work leads to a different theoretical explanation. The observed points having unusually large outputs when compared with the (empirical) variance should be down-weighted in the estimation of the mean, since they contain less information than noise. In short, huge outputs should be truncated because of their low signal-to-noise ratio.

2. Ridge regression and empirical risk minimization. We recall the definition

$$\mathcal{F} = \left\{ f_\theta = \sum_{j=1}^d \theta_j \varphi_j; (\theta_1, \dots, \theta_d) \in \Theta \right\},$$

where Θ is a closed convex set, not necessarily bounded (so that $\Theta = \mathbb{R}^d$ is allowed). In this section, we provide exponential deviation inequalities for the empirical risk minimizer and the ridge regression estimator on \mathcal{F} under weak conditions on the tail of the output distribution.

The most general theorem which can be obtained from the route followed in this section is Theorem 4.5 (page 29) stated along with the proof. It is expressed in terms of a series of empirical bounds. The first deduction we can make from this technical result is of asymptotic nature. It is stated under weak hypotheses, taking advantage of the weak law of large numbers.

THEOREM 2.1. *For $\lambda \geq 0$, let \tilde{f} be its associated optimal ridge function (see (1.5)). Let us assume that*

$$(2.1) \quad \mathbb{E}[\|\varphi(X)\|^4] < +\infty,$$

$$(2.2) \quad \text{and} \quad \mathbb{E}\left\{\|\varphi(X)\|^2 [\tilde{f}(X) - Y]^2\right\} < +\infty.$$

Let ν_1, \dots, ν_d be the eigenvalues of the Gram matrix $Q = \mathbb{E}[\varphi(X)\varphi(X)^T]$, and let $Q_\lambda = Q + \lambda I$ be the ridge regularization of Q . Let us define the effective ridge dimension

$$D = \sum_{i=1}^d \frac{\nu_i}{\nu_i + \lambda} \mathbb{1}(\nu_i > 0) = \text{Tr}[(Q + \lambda I)^{-1}Q] = \mathbb{E}[\|Q_\lambda^{-1/2}\varphi(X)\|^2].$$

When $\lambda = 0$, D is equal to the rank of Q and is otherwise smaller. For any $\varepsilon > 0$, there is n_ε , such that for any $n \geq n_\varepsilon$, with probability at least $1 - \varepsilon$,

$$\begin{aligned} R(\hat{f}^{(\text{ridge})}) + \lambda \|\hat{\theta}^{(\text{ridge})}\|^2 &\leq \min_{\theta \in \Theta} \{R(f_\theta) + \lambda \|\theta\|^2\} \\ &\quad + \frac{30 \mathbb{E}\{\|Q_\lambda^{-1/2}\varphi(X)\|^2 [\tilde{f}(X) - Y]^2\}}{\mathbb{E}\{\|Q_\lambda^{-1/2}\varphi(X)\|^2\}} \frac{D}{n} \\ &\quad + 1000 \sup_{v \in \mathbb{R}^d} \frac{\mathbb{E}\left[\langle v, \varphi(X) \rangle^2 [\tilde{f}(X) - Y]^2\right]}{\mathbb{E}(\langle v, \varphi(X) \rangle^2) + \lambda \|v\|^2} \frac{\log(3\varepsilon^{-1})}{n} \\ &\leq \min_{\theta \in \Theta} \{R(f_\theta) + \lambda \|\theta\|^2\} \end{aligned}$$

$$+ \text{ess sup } \mathbb{E}\{[Y - \tilde{f}(X)]^2 | X\} \frac{30D + 1000 \log(3\varepsilon^{-1})}{n}.$$

PROOF. See Section 4.2 (page 23). \square

This theorem shows that the ordinary least squares estimator (obtained when $\Theta = \mathbb{R}^d$ and $\lambda = 0$), as well as the empirical risk minimizer on any closed convex set, asymptotically reaches a d/n speed of convergence under very weak hypotheses. It shows also the regularization effect of the ridge regression. There emerges an *effective dimension* D , where the ridge penalty has a threshold effect on the eigenvalues of the Gram matrix.

Let us remark that the second inequality stated in the theorem provides a simplified bound which makes sense only when

$$\text{ess sup } \mathbb{E}\{[Y - \tilde{f}(X)]^2 | X\} < +\infty,$$

implying that $\|\tilde{f} - f^{(\text{reg})}\|_\infty < +\infty$. We chose not to hide the first inequality in the proof, since it does not require such a tight relationship between \tilde{f} and $f^{(\text{reg})}$.

On the other hand, the weakness of this result is its asymptotic nature : n_ε may be arbitrarily large under such weak hypotheses, and this shows even in the simplest case of the estimation of the mean of a real valued random variable by its empirical mean (which is the case when $d = 1$ and $\varphi(X) \equiv 1$).

Let us now give some non asymptotic rate under stronger hypotheses and for the empirical risk minimizer (i.e., $\lambda = 0$).

THEOREM 2.2. *Assume that $\mathbb{E}\{[Y - f^*(X)]^4\} < +\infty$ and*

$$B = \sup_{f \in \text{span}\{\varphi_1, \dots, \varphi_d\} - \{0\}} \|f\|_\infty^2 / \mathbb{E}[f(X)^2] < +\infty.$$

Consider the (unique) empirical risk minimizer $\hat{f}^{(\text{erm})} = f_{\hat{\theta}^{(\text{erm})}} : x \mapsto \langle \hat{\theta}^{(\text{erm})}, \varphi(x) \rangle$ on \mathcal{F} for which $\hat{\theta}^{(\text{erm})} \in \text{span}\{\varphi(X_1), \dots, \varphi(X_n)\}$ ¹. For any values of ε and n such that $2/n \leq \varepsilon \leq 1$ and

$$n > 1280B^2 \left[3Bd + \log(2/\varepsilon) + \frac{16B^2d^2}{n} \right],$$

with probability at least $1 - \varepsilon$,

¹When $\mathcal{F} = \mathcal{F}_{\text{lin}}$, we have $\hat{\theta}^{(\text{erm})} = \mathbf{X}^+ \mathbf{Y}$, with $\mathbf{X} = (\varphi_j(X_i))_{1 \leq i \leq n, 1 \leq j \leq d}$, $\mathbf{Y} = [Y_j]_{j=1}^n$ and \mathbf{X}^+ is the Moore-Penrose pseudoinverse of \mathbf{X} .

$$(2.3) \quad R(\hat{f}^{(\text{erm})}) - R(f^*) \leq 1920 B \sqrt{\mathbb{E}[Y - f^*(X)]^4} \left[\frac{3Bd + \log(2\varepsilon^{-1})}{n} + \left(\frac{4Bd}{n} \right)^2 \right].$$

PROOF. See Section 4.2 (page 23). \square

It is quite surprising that the traditional assumption of uniform boundedness of the conditional exponential moments of the output can be replaced by a simple moment condition for reasonable confidence levels (i.e., $\varepsilon \geq 2/n$). For highest confidence levels, things are more tricky since we need to control with high probability a term of order $[r(f^*) - R(f^*)]d/n$ (see Theorem 4.6). The cost to pay to get the exponential deviations under only a fourth-order moment condition on the output is the appearance of the geometrical quantity B as a multiplicative factor.

To better understand the quantity B , let us consider two cases. First, consider that the input is uniformly distributed on $\mathcal{X} = [0, 1]$, and that the functions $\varphi_1, \dots, \varphi_d$ belongs to the Fourier basis. Then the quantity B behaves like a numerical constant. On the contrary, if we take $\varphi_1, \dots, \varphi_d$ as the first d elements of the wavelet expansions, the more localized wavelets induce high values of B , and B scales like \sqrt{d} , meaning that Theorem 2.2 fails to give a d/n -excess risk bound in this case. This limitation does not appear in Theorem 2.1.

To conclude, Theorem 2.2 is limited in at least four ways: it involves the quantity B , it applies only to uniformly bounded $\varphi(X)$, the output needs to have a fourth moment, and the confidence level should be as great as $\varepsilon \geq 2/n$. These limitations will be addressed in the next section by considering a more involved algorithm.

3. A min-max estimator for robust estimation. This section provides an alternative to the empirical risk minimizer with non asymptotic exponential risk deviations of order d/n for any confidence level. Moreover, we will assume only a second order moment condition on the output and cover the case of unbounded inputs, the requirement on $\varphi(X)$ being only a finite fourth order moment. On the other hand, we assume here that the set Θ of the vectors of coefficients is bounded. The computability of the proposed estimator and numerical experiments are discussed at the end of the section.

3.1. *The min-max estimator and its theoretical guarantee.* Let $\alpha > 0$, $\lambda \geq 0$, and consider the truncation function:

$$\psi(x) = \begin{cases} -\log(1 - x + x^2/2) & 0 \leq x \leq 1, \\ \log(2) & x \geq 1, \\ -\psi(-x) & x \leq 0, \end{cases}$$

For any $\theta, \theta' \in \Theta$, introduce

$$\mathcal{D}(\theta, \theta') = n\alpha\lambda(\|\theta\|^2 - \|\theta'\|^2) + \sum_{i=1}^n \psi\left(\alpha[Y_i - f_\theta(X_i)]^2 - \alpha[Y_i - f_{\theta'}(X_i)]^2\right).$$

We recall $\tilde{f} = f_{\tilde{\theta}}$ with $\tilde{\theta} \in \arg \min_{\theta \in \Theta} \{R(f_\theta) + \lambda\|\theta\|^2\}$, and the effective ridge dimension

$$D = \sum_{i=1}^d \frac{\nu_i}{\nu_i + \lambda} \mathbf{1}(\nu_i > 0) = \text{Tr}[(Q + \lambda I)^{-1}Q] = \mathbb{E}[\|Q_\lambda^{-1/2}\varphi(X)\|^2].$$

Let us assume in this section that

$$(3.1) \quad \mathbb{E}\{[Y - \tilde{f}(X)]^4\} < +\infty,$$

and for any $j \in \{1, \dots, d\}$,

$$(3.2) \quad \mathbb{E}[\varphi_j^4(X)] < +\infty.$$

Define

$$(3.3) \quad \mathcal{S} = \{f \in \mathcal{F}_{\text{lin}} : \mathbb{E}[f(X)^2] = 1\},$$

$$(3.4) \quad \sigma = \sqrt{\mathbb{E}\{[Y - \tilde{f}(X)]^2\}} = \sqrt{R(\tilde{f})},$$

$$(3.5) \quad \chi = \max_{f \in \mathcal{S}} \sqrt{\mathbb{E}[f(X)^4]},$$

$$(3.6) \quad \kappa = \frac{\sqrt{\mathbb{E}\{[\varphi(X)^T Q_\lambda^{-1} \varphi(X)]^2\}}}{\mathbb{E}[\varphi(X)^T Q_\lambda^{-1} \varphi(X)]},$$

$$(3.7) \quad \kappa' = \frac{\sqrt{\mathbb{E}\{[Y - \tilde{f}(X)]^4\}}}{\mathbb{E}\{[Y - \tilde{f}(X)]^2\}} = \frac{\sqrt{\mathbb{E}\{[Y - \tilde{f}(X)]^4\}}}{\sigma^2},$$

$$(3.8) \quad T = \max_{\theta \in \Theta, \theta' \in \Theta} \sqrt{\lambda\|\theta - \theta'\|^2 + \mathbb{E}[f_\theta(X) - f_{\theta'}(X)]^2}.$$

THEOREM 3.1. *Let us assume that (3.1) and (3.2) hold. For some numerical constants c and c' , for*

$$n > c\kappa\chi D,$$

by taking

$$(3.9) \quad \alpha = \frac{1}{2\chi[2\sqrt{\kappa'}\sigma + \sqrt{\chi}T]^2} \left(1 - \frac{c\kappa\chi D}{n}\right),$$

for any estimator $f_{\hat{\theta}}$ satisfying $\hat{\theta} \in \Theta$ a.s., for any $\varepsilon > 0$ and any $\lambda \geq 0$, with probability at least $1 - \varepsilon$, we have

$$\begin{aligned} R(f_{\hat{\theta}}) + \lambda\|\hat{\theta}\|^2 &\leq \min_{\theta \in \Theta} \{R(f_{\theta}) + \lambda\|\theta\|^2\} \\ &\quad + \frac{1}{n\alpha} \left(\max_{\theta_1 \in \Theta} \mathcal{D}(\hat{\theta}, \theta_1) - \inf_{\theta \in \Theta} \max_{\theta_1 \in \Theta} \mathcal{D}(\theta, \theta_1) \right) \\ &\quad + \frac{c\kappa\kappa' D\sigma^2}{n} + \frac{8\chi \left(\frac{\log(\varepsilon^{-1})}{n} + \frac{c'\kappa^2 D^2}{n^2} \right) [2\sqrt{\kappa'}\sigma + \sqrt{\chi}T]^2}{1 - \frac{c\kappa\chi D}{n}}. \end{aligned}$$

PROOF. See Section 4.3 (page 33). \square

By choosing an estimator such that

$$\max_{\theta_1 \in \Theta} \mathcal{D}(\hat{\theta}, \theta_1) < \inf_{\theta \in \Theta} \max_{\theta_1 \in \Theta} \mathcal{D}(\theta, \theta_1) + \sigma^2 \frac{D}{n},$$

Theorem 3.1 provides a non asymptotic bound for the excess (ridge) risk with a D/n convergence rate and an exponential tail even when neither the output Y nor the input vector $\varphi(X)$ has exponential moments. This stronger non asymptotic bound compared to the bounds of the previous section comes at the price of replacing the empirical risk minimizer by a more involved estimator. Section 3.3 provides a way of computing it approximately.

Theorem 3.1 requires a fourth order moment condition on the output. In fact, one can replace (3.1) by the following second order moment condition on the output: for any $j \in \{1, \dots, d\}$,

$$\mathbb{E}\{\varphi_j(X)^2[Y - \tilde{f}(X)]^2\} < +\infty,$$

and still obtain a D/n excess risk bound. This comes at the price of a more lengthy formula, where terms with κ' become terms involving the quantities $\max_{f \in \mathcal{S}} \mathbb{E}\{f(X)^2[Y - \tilde{f}(X)]^2\}$ and $\mathbb{E}\{\varphi(X)^T Q^{-1} \varphi(X)[Y - \tilde{f}(X)]^2\}$. (This can be seen by not using Cauchy-Schwarz's inequality in (4.28) and (4.29).)

3.2. *The value of the uncentered kurtosis coefficient χ .* Let us discuss here the value of the constant χ , which plays a critical role in the speed of convergence of our bound. With the convention $\frac{0}{0} = 0$, we have

$$\chi = \sup_{u \in \mathbb{R}^d} \frac{\mathbb{E}(\langle u, \varphi(X) \rangle^4)^{1/2}}{\mathbb{E}(\langle u, \varphi(X) \rangle^2)}.$$

Let us first examine the case when $\varphi_1(X) \equiv 1$ and $\varphi_2(X), \dots, \varphi_d(X)$ are independent. To compute χ , we can assume without loss of generality that $\varphi_2(X), \dots, \varphi_d(X)$ are centered and of unit variance. In this situation, introducing

$$\chi_* = \max_{j=1, \dots, d} \frac{\mathbb{E}[\varphi_j(X)^4]^{1/2}}{\mathbb{E}[\varphi_j(X)^2]},$$

we see that for any $u \in \mathbb{R}^d$ with $\|u\| = 1$, we have

$$\begin{aligned} \mathbb{E}(\langle u, \varphi(X) \rangle^4) &= \sum_{i=1}^d u_i^4 \mathbb{E}(\varphi_i(X)^4) + 6 \sum_{1 \leq i < j \leq d} u_i^2 u_j^2 \mathbb{E}[\varphi_i(X)^2] \mathbb{E}[\varphi_j(X)^2] \\ &\quad + 4 \sum_{i=2}^d u_1 u_i^3 \mathbb{E}[\varphi_i(X)^3] \\ &\leq \chi_*^2 \sum_{i=1}^d u_i^4 + 6 \sum_{i < j} u_i^2 u_j^2 + 4 \chi_*^{3/2} \sum_{i=2}^d |u_1 u_i|^3 \\ &\leq \sup_{u \in \mathbb{R}_+^d, \|u\|=1} (\chi_*^2 - 3) \sum_{i=1}^d u_i^4 + 3 \left(\sum_{i=1}^d u_i^2 \right)^2 + 4 \chi_*^{3/2} u_1 \sum_{i=2}^d u_i^3 \\ &\leq \frac{3^{3/2}}{4} \chi_*^{3/2} + \begin{cases} \chi_*^2, & \chi_*^2 \geq 3, \\ 3 + \frac{\chi_*^2 - 3}{d}, & 1 \leq \chi_*^2 < 3. \end{cases} \end{aligned}$$

Thus in this case

$$\chi \leq \begin{cases} \chi_* \left(1 + \frac{3^{3/2}}{4\sqrt{\chi_*}} \right)^{1/2}, & \chi_* \geq \sqrt{3}, \\ \left(3 + \frac{3^{3/2}}{4} \chi_*^{3/2} + \frac{\chi_*^2 - 3}{d} \right)^{1/2}, & 1 \leq \chi_* < \sqrt{3}. \end{cases}$$

If moreover the random variables $\varphi_2(X), \dots, \varphi_d(X)$ are not skewed, in the sense that $\mathbb{E}[\varphi_j(X)^3] = 0$, $j = 2, \dots, d$, then

$$\begin{cases} \chi = \chi_*, & \chi_* \geq \sqrt{3}, \\ \chi \leq \left(3 + \frac{\chi_*^2 - 3}{d} \right)^{1/2}, & 1 \leq \chi_* < \sqrt{3}. \end{cases}$$

In particular in the case when $\varphi_2(X), \dots, \varphi_d(X)$ are Gaussian random variables, we have $\chi = \chi_* = \sqrt{3}$ (as could be seen in a more straightforward way, since in this case $\langle u, \varphi(X) \rangle$ is also Gaussian!).

In particular, this situation arises in compress sensing using random projections on Gaussian vectors. Specifically, assume that we want to recover a signal $f \in \mathbb{R}^M$ that we know to be well approximated by a linear combination of d basis vectors f_1, \dots, f_d . We measure $n \ll M$ projections of the signal f on i.i.d. M -dimensional standard normal random vectors X_1, \dots, X_n : $Y_i = \langle f, X_i \rangle$, $i = 1, \dots, n$. Then, recovering the coefficient $\theta_1, \dots, \theta_d$ such that $f = \sum_{j=1}^d \theta_j f_j$ is associated to the least squares regression problem $Y \approx \sum_{j=1}^d \theta_j \varphi_j(X)$, with $\varphi_j(x) = \langle f_j, x \rangle$, and X having a M -dimensional standard normal distribution.

Let us discuss now a bound which is suited to the case when we are using a partial basis of regression functions. The functions φ_j are usually bounded (think of the Fourier basis, wavelet bases, histograms, splines ...).

Let us assume that for some positive constant A and any $u \in \mathbb{R}^d$,

$$\|u\| \leq A \mathbb{E}[\langle u, \varphi(X) \rangle^2]^{1/2}.$$

This appears as some stability property of the partial basis φ_j with respect to the \mathbb{L}_2 -norm, since it can also be written as

$$\sum_{j=1}^d u_j^2 \leq A^2 \mathbb{E} \left[\left(\sum_{j=1}^d u_j \varphi_j(X) \right)^2 \right], \quad u \in \mathbb{R}^d.$$

This will be the case if φ_j is nearly orthogonal in the sense that

$$\mathbb{E}[\varphi_j(X)^2] \geq 1, \quad \text{and} \quad \left| \mathbb{E}[\varphi_j(X)\varphi_k(X)] \right| \leq \frac{1 - A^{-2}}{d - 1}.$$

In this situation, by using

$$\mathbb{E}[\langle u, \varphi(X) \rangle^4] \leq \|u\|^2 \text{ess sup} \|\varphi(X)\|^2 \mathbb{E}[\langle u, \varphi(X) \rangle^2],$$

one can check that

$$\chi \leq A \left\| \left(\sum_{j=1}^d \varphi_j^2 \right)^{1/2} \right\|_{\infty}.$$

Therefore, if X is the uniform random variable on the unit interval and φ_j , $j = 1, \dots, d$ are any functions from the Fourier basis (meaning that they are of the form $\sqrt{2} \cos(2k\pi X)$ or $\sqrt{2} \sin(2k\pi X)$), then $\chi \leq \sqrt{2d}$ (because they form an orthogonal system, so that $A = 1$).

On the other hand, a localized basis like the evenly spaced histogram basis of the unit interval

$$\varphi_j(x) = \sqrt{d}\mathbf{1}\left(x \in [(j-1)/d, j/d[\right),$$

will also be such that $\chi \leq \sqrt{d}$. Similar computations could be made for other local bases, like wavelet bases. Note that when χ is of order \sqrt{d} , Theorem 3.1 means that the excess risk of the min-max truncated estimator \hat{f} is upper bounded by $C\frac{d}{n}$ provided that $n \geq Cd^2$ for a large enough constant C .

Let us discuss the case when X is some observed random variable whose distribution is only approximately known. Namely let us assume that $(\varphi_j)_{j=1}^d$ is some basis of functions in $\mathbb{L}_2[\tilde{\mathbb{P}}]$ with some known coefficient $\tilde{\chi}$, where $\tilde{\mathbb{P}}$ is an approximation of the true distribution of X in the sense that the density of the true distribution \mathbb{P} of X with respect to the distribution $\tilde{\mathbb{P}}$ is in the range $(\eta^{-1/2}, \eta)$. In this situation, the coefficient χ satisfies the inequality $\chi \leq \eta\tilde{\chi}$. Indeed

$$\begin{aligned} \mathbb{E}_{X \sim \mathbb{P}}[\langle u, \varphi(X) \rangle^4] &\leq \eta \mathbb{E}_{X \sim \tilde{\mathbb{P}}}[\langle u, \varphi(X) \rangle^4] \\ &\leq \eta \tilde{\chi}^2 \mathbb{E}_{X \sim \tilde{\mathbb{P}}}[\langle u, \varphi(X) \rangle^2]^2 \leq \eta^2 \tilde{\chi}^2 \mathbb{E}_{X \sim \mathbb{P}}[\langle u, \varphi(X) \rangle^2]^2. \end{aligned}$$

Let us conclude this section with some scenario for the case when X is a real-valued random variable. Let us consider the distribution function of $\tilde{\mathbb{P}}$

$$\tilde{F}(x) = \tilde{\mathbb{P}}(X \leq x).$$

Then, if $\tilde{\mathbb{P}}$ has no atoms, the distribution of $\tilde{F}(X)$ is uniform in $(0, 1)$. Starting from some suitable partial basis $(\varphi_j)_{j=1}^d$ of $\mathbb{L}_2[(0, 1), \mathbb{U}]$ where \mathbb{U} is the uniform distribution, like the ones discussed above, we can build a basis for our problem as

$$\tilde{\varphi}_j(X) = \varphi_j[\tilde{F}(X)].$$

Moreover, if \mathbb{P} is absolutely continuous with respect to $\tilde{\mathbb{P}}$ with density g , then $\mathbb{P} \circ \tilde{F}^{-1}$ is absolutely continuous with respect to $\tilde{\mathbb{P}} \circ \tilde{F}^{-1}$, with density $g \circ \tilde{F}^{-1}$, and of course, the fact that g takes values in $(\eta^{-1/2}, \eta)$ implies the same property for $g \circ \tilde{F}^{-1}$. Thus, if $\tilde{\chi}$ is the coefficient corresponding to $\varphi_j(U)$ when U is the uniform random variable on the unit interval, then the true coefficient χ (corresponding to $\tilde{\varphi}_j(X)$) will be such that $\chi \leq \eta\tilde{\chi}$.

3.3. Computation of the estimator. For ease of description of the algorithm, we will write X for $\varphi(X)$, which is equivalent to considering without loss of generality that the input space is \mathbb{R}^d and that the functions $\varphi_1, \dots, \varphi_d$

are the coordinate functions. Therefore, the function f_θ maps an input x to $\langle \theta, x \rangle$. Let us introduce

$$\bar{L}_i(\theta) = \alpha (\langle \theta, X_i \rangle - Y_i)^2.$$

For any subset of indices $I \subset \{1, \dots, n\}$, let us define

$$r_I(\theta) = \lambda \|\theta\|^2 + \frac{1}{\alpha |I|} \sum_{i \in I} \bar{L}_i(\theta).$$

We suggest the following heuristics to compute an approximation of

$$\arg \min_{\theta \in \Theta} \sup_{\theta' \in \Theta} \mathcal{D}(\theta, \theta').$$

- Start from $I_1 = \{1, \dots, n\}$ with the ordinary least squares estimate

$$\hat{\theta}_1 = \arg \min_{\mathbb{R}^d} r_{I_1}.$$

- At step number k , compute

$$\hat{Q}_k = \frac{1}{|I_k|} \sum_{i \in I_k} X_i X_i^T.$$

- Consider the sets

$$J_{k,1}(\eta) = \left\{ i \in I_k : \bar{L}_i(\hat{\theta}_k) X_i^T \hat{Q}_k^{-1} X_i \left(1 + \sqrt{1 + [\bar{L}_i(\hat{\theta}_k)]^{-1}} \right)^2 < \eta \right\},$$

where \hat{Q}_k^{-1} is the (pseudo-)inverse of the matrix \hat{Q}_k .

- Let us define

$$\theta_{k,1}(\eta) = \arg \min_{\mathbb{R}^d} r_{J_{k,1}(\eta)},$$

$$J_{k,2}(\eta) = \left\{ i \in I_k : |\bar{L}_i(\theta_{k,1}(\eta)) - \bar{L}_i(\hat{\theta}_k)| \leq 1 \right\},$$

$$\theta_{k,2}(\eta) = \arg \min_{\mathbb{R}^d} r_{J_{k,2}(\eta)},$$

$$(\eta_k, \ell_k) = \arg \min_{\eta \in \mathbb{R}_+, \ell \in \{1, 2\}} \max_{j=1, \dots, k} \mathcal{D}(\theta_{k,\ell}(\eta), \hat{\theta}_j),$$

$$I_{k+1} = J_{k,\ell_k}(\eta_k),$$

$$\hat{\theta}_{k+1} = \theta_{k,\ell_k}(\eta_k).$$

- Stop when

$$\max_{j=1,\dots,k} \mathcal{D}(\widehat{\theta}_{k+1}, \widehat{\theta}_j) \geq 0,$$

and set $\widehat{\theta} = \widehat{\theta}_k$ as the final estimator of $\tilde{\theta}$.

Note that there will be at most n steps, since $I_{k+1} \subsetneq I_k$ and in practice much less in this iterative scheme. Let us give some justification for this proposal. Let us notice first that

$$\begin{aligned} \mathcal{D}(\theta + h, \theta) &= n\alpha\lambda(\|\theta + h\|^2 - \|\theta\|^2) \\ &\quad + \sum_{i=1}^n \psi\left(\alpha[2\langle h, X_i \rangle(\langle \theta, X_i \rangle - Y_i) + \langle h, X_i \rangle^2]\right). \end{aligned}$$

Hopefully, $\tilde{\theta} = \arg \min_{\theta \in \mathbb{R}^d} (R(f\theta) + \lambda\|\theta\|^2)$ is in some small neighbourhood of $\widehat{\theta}_k$ already, according to the distance defined by $Q \simeq \widehat{Q}_k$. So we may try to look for improvements of $\widehat{\theta}_k$ by exploring neighbourhoods of $\widehat{\theta}_k$ of increasing sizes with respect to some approximation of the relevant norm $\|\theta\|_Q^2 = \mathbb{E}[\langle \theta, X \rangle^2]$.

Since the truncation function ψ is constant on $(-\infty, -1]$ and $[1, +\infty)$, the map $\theta \mapsto \mathcal{D}(\theta, \widehat{\theta}_k)$ induces a decomposition of the parameter space into cells corresponding to different sets I of examples. Indeed, such a set I is associated to the set \mathcal{C}_I of θ such that $\overline{L}_i(\theta) - \overline{L}_i(\widehat{\theta}_k) < 1$ if and only if $i \in I$. Although this may not be the case, we will do as if the map $\theta \mapsto \mathcal{D}(\theta, \widehat{\theta}_k)$ restricted to the cell \mathcal{C}_I reached its minimum at some interior point of \mathcal{C}_I , and approximates this minimizer by the minimizer of r_I .

The idea is to remove first the examples which will become inactive in the closest cells to the current estimate $\widehat{\theta}_k$. The cells for which the contribution of example number i is constant are delimited by at most four parallel hyperplanes.

It is easy to see that the square of the inverse of the distance of $\widehat{\theta}_k$ to the closest of these hyperplanes is equal to

$$\frac{1}{\alpha} X_i^T \widehat{Q}_k^{-1} X_i \overline{L}_i(\widehat{\theta}_k) \left(1 + \sqrt{1 + \frac{1}{\overline{L}_i(\widehat{\theta}_k)}} \right)^2.$$

Indeed, this distance is the infimum of $\|\widehat{Q}_k^{1/2} h\|$, where h is a solution of

$$\langle h, X_i \rangle^2 + 2\langle h, X_i \rangle(\langle \widehat{\theta}_k, X_i \rangle - Y_i) = \frac{1}{\alpha}.$$

It is computed by considering h of the form $h = \xi \|\widehat{Q}_k^{-1/2} X_i\|^{-1} \widehat{Q}_k^{-1} X_i$ and solving an equation of order two in ξ .

This explains the proposed choice of $J_{k,1}(\eta)$. Then a first estimate $\theta_{k,1}(\eta)$ is computed on the basis of this reduced sample, and the sample is readjusted to $J_{k,2}(\eta)$ by checking which constraints are really activated in the computation of $\mathcal{D}(\theta_{k,1}(\eta), \widehat{\theta}_k)$. The estimated parameter is then readjusted taking into account the readjusted sample (this could as a variant be iterated more than once). Now that we have some new candidates $\theta_{k,\ell}(\eta)$, we check the minimax property against them to elect I_{k+1} and $\widehat{\theta}_{k+1}$. Since we did not check the minimax property against the whole parameter set $\Theta = \mathbb{R}^d$, we have no theoretical warranty for this simplified algorithm. Nonetheless, similar computations to what we did could prove that we are close to solving $\min_{j=1,\dots,k} R(f_{\widehat{\theta}_j})$, since we checked the minimax property on the reduced parameter set $\{\widehat{\theta}_j, j = 1, \dots, k\}$. Thus the proposed heuristics is capable of improving on the performance of the ordinary least squares estimator, while being guaranteed not to degrade its performance significantly.

3.4. Synthetic experiments. In Section 3.4.1, we detail the different kinds of noises we work with. Then, Sections 3.4.2, 3.4.3 and 3.4.4 describe the three types of functional relationships between the input, the output and the noise involved in our experiments. A motivation for choosing these input-output distributions was the ability to compute exactly the excess risk, and thus to compare easily estimators. Section 3.4.5 provides details about the implementation, its computational efficiency and the main conclusions of the numerical experiments. Figures and tables are postponed to Appendix A.

3.4.1. Noise distributions. In our experiments, we consider different types of noise that are centered and with unit variance:

- the standard Gaussian noise: $W \sim \mathcal{N}(0, 1)$,
- a heavy-tailed noise defined by: $W = \text{sign}(V)/|V|^{1/q}$, with $V \sim \mathcal{N}(0, 1)$ a standard Gaussian random variable and $q = 2.01$ (the real number q is taken strictly larger than 2 as for $q = 2$, the random variable W would not admit a finite second moment).
- an asymmetric heavy-tailed noise defined by:

$$W = \begin{cases} |V|^{-1/q} & \text{if } V > 0, \\ -\frac{q}{q-1} & \text{otherwise,} \end{cases}$$

with $q = 2.01$ with $V \sim \mathcal{N}(0, 1)$ a standard Gaussian random variable.

- a mixture of a Dirac random variable with a low-variance Gaussian random variable defined by: with probability p , $W = \sqrt{(1-\rho)/p}$, and

with probability $1 - p$, W is drawn from

$$\mathcal{N}\left(-\frac{\sqrt{p(1-\rho)}}{1-p}, \frac{\rho}{1-p} - \frac{p(1-\rho)}{(1-p)^2}\right).$$

The parameter $\rho \in [p, 1]$ characterizes the part of the variance of W explained by the Gaussian part of the mixture. Note that this noise admits exponential moments, but for n of order $1/p$, the Dirac part of the mixture generates low signal-to-noise points.

3.4.2. Independent normalized covariates (INC(n, d)). In INC(n, d), we consider $\varphi(X) = X$, and the input-output pair is such that

$$Y = \langle \theta^*, X \rangle + \sigma W,$$

where the components of X are independent standard normal distributions, $\theta^* = (10, \dots, 10)^T \in \mathbb{R}^d$, and $\sigma = 10$.

3.4.3. Highly correlated covariates (HCC(n, d)). In HCC(n, d), we consider $\varphi(X) = X$, and the input-output pair is such that

$$Y = \langle \theta^*, X \rangle + \sigma W,$$

where X is a multivariate centered normal Gaussian with covariance matrix Q obtained by drawing a (d, d) -matrix A of uniform random variables in $[0, 1]$ and by computing $Q = AA^T$, $\theta^* = (10, \dots, 10)^T \in \mathbb{R}^d$, and $\sigma = 10$. So the only difference with the setting of Section 3.4.2 is the correlation between the covariates.

3.4.4. Trigonometric series (TS(n, d)). Let X be a uniform random variable on $[0, 1]$. Let d be an even number. In TS(n, d), we consider

$$\varphi(X) = (\cos(2\pi X), \dots, \cos(d\pi X), \sin(2\pi X), \dots, \sin(d\pi X))^T,$$

and the input-output pair is such that

$$Y = 20X^2 - 10X - \frac{5}{3} + \sigma W,$$

with $\sigma = 10$. One can check that this implies

$$\theta^* = \left(\frac{20}{\pi^2}, \dots, \frac{20}{\pi^2(\frac{d}{2})^2}, -\frac{10}{\pi}, \dots, -\frac{10}{\pi(\frac{d}{2})}\right)^T \in \mathbb{R}^d.$$

3.4.5. Experiments.

Choice of the parameters and implementation details. The min-max truncated algorithm has two parameters α and λ . In the subsequent experiments, we set the ridge parameter λ to the natural default choice for it: $\lambda = 0$. For the truncation parameter α , according to our analysis (see (3.9)), it roughly should be of order $1/\sigma^2$ up to kurtosis coefficients. By using the ordinary least squares estimator, we roughly estimate this value, and test values of α in a geometric grid (of 8 points) around it (with ratio 3). Cross-validation can be used to select the final α . Nevertheless, it is computationally expensive and is significantly outperformed in our experiments by the following simple procedure: start with the smallest α in the geometric grid and increase it as long as $\hat{\theta} = \theta_1$, that is as long as we stop at the end of the first iteration and output the empirical risk minimizer.

To compute $\theta_{k,1}(\eta)$ or $\theta_{k,2}(\eta)$, one needs to determine a least squares estimate (for a modified sample). To reduce the computational burden, we do not want to test all possible values of η (note that there are at most n values leading to different estimates). Our experiments show that testing only three levels of η is sufficient. Precisely, we sort the quantity

$$\bar{L}_i(\hat{\theta}_k) X_i^T \hat{Q}_k^{-1} X_i \left(1 + \sqrt{1 + [\bar{L}_i(\hat{\theta}_k)]^{-1}} \right)^2$$

by decreasing order and consider η being the first, 5-th and 25-th value of the ordered list. Overall, in our experiments, the computational complexity is approximately fifty times larger than the one of computing the ordinary least squares estimator.

Results. The tables and figures have been gathered in Appendix A. Tables 1 and 2 give the results for the mixture noise. Tables 3, 4 and 5 provide the results for the heavy-tailed noise and the standard Gaussian noise. Each line of the tables has been obtained after 1000 generations of the training set. These results show that the min-max truncated estimator is often equal to the ordinary least squares estimator $\hat{f}^{(\text{ols})}$, while it ensures impressive consistent improvements when it differs from $\hat{f}^{(\text{ols})}$. In this latter case, the number of points that are not considered in \hat{f} , i.e. the number of points with low signal-to-noise ratio, varies a lot from 1 to 150 and is often of order 30. Note that not only the points that we expect to be considered as outliers (i.e. very large output points) are erased, and that these points seem to be taken out by local groups: see Figures 1 and 2 in which the erased points are marked by surrounding circles.

Besides, the heavier the noise tail is (and also the larger the variance of the noise is), the more often the truncation modifies the initial ordinary

least squares estimator, and the more improvements we get from the min-max truncated estimator, which also becomes much more robust than the ordinary least squares estimator (see the confidence intervals in the tables).

Finally, we have also tested more traditional methods in robust regression: namely, the M-estimators with Huber’s loss, L_1 -loss, and Tukey’s bisquare influence function, and also the least trimmed squares estimator, the S-estimator and the MM-estimator (see [9, 13] and references within). These methods rely on diminishing the influence of points having “unreasonably” large residuals. They were developed to handle training sets containing true outliers, i.e. points (X, Y) not generated by the distribution P . This is not the case in our estimation framework. By overweighting points having reasonably small residuals, these methods are often biased even in settings where the noise is symmetric and the regression function $f^{(\text{reg})} : x \mapsto \mathbb{E}[Y|X = x]$ belongs to \mathcal{F}_{lin} (i.e., $f^{(\text{reg})} = f_{\text{lin}}^*$), and also even when there is no noise (but $f^{(\text{reg})} \notin f_{\text{lin}}^*$).

The worst results were obtained by the L_1 -loss, since estimating the (conditional) median is here really different from estimating the (conditional) mean. The MM-estimator and the M-estimators with Huber’s loss and Tukey’s bisquare influence function give good results as long as the signal-to-noise ratio is low. When the signal-to-noise ratio is high, a lack of consistency drastically appears in part of our simulations, showing that these methods are thus not suited for our estimation framework.

The S-estimator is almost consistently improving on the ordinary least squares estimator (in our simulations). However, when the signal-to-noise ratio is low (that is, in the setting of the forementioned simulations with $\sigma = 10$), the improvements are much less significant than the ones of the min-max truncated estimator.

4. Proofs.

4.1. *Main ideas of the proofs.* The goal of this section is to explain the key ingredients appearing in the proofs which both allows to obtain sub-exponential tails for the excess risk under a non-exponential moment assumption and get rid of the logarithmic factor in the excess risk bound.

4.1.1. *Sub-exponential tails under a non-exponential moment assumption via truncation.* Let us start with the idea allowing us to prove exponential inequalities under just a moment assumption (instead of the traditional exponential moment assumption). To understand it, we can consider the (apparently) simplistic 1-dimensional situation in which we have $\Theta = \mathbb{R}$ and the marginal distribution of $\varphi_1(X)$ is the Dirac distribution at 1. In

this case, the risk of the prediction function f_θ is $R(f_\theta) = \mathbb{E}(Y - \theta)^2 = \mathbb{E}(Y - \theta^*)^2 + (\mathbb{E}Y - \theta)^2$, so that the least squares regression problem boils down to the estimation of the mean of the output variable. If we only assume that Y admits a finite second moment, say $\mathbb{E}Y^2 \leq 1$, it is not clear whether for any $\varepsilon > 0$, it is possible to find $\hat{\theta}$ such that with probability at least $1 - 2\varepsilon$,

$$(4.1) \quad R(f_{\hat{\theta}}) - R(f^*) = (\mathbb{E}(Y) - \hat{\theta})^2 \leq c \frac{\log(\varepsilon^{-1})}{n},$$

for some numerical constant c . Indeed, from Chebyshev's inequality, the trivial choice $\hat{\theta} = \frac{\sum_{i=1}^n Y_i}{n}$ just satisfies: with probability at least $1 - 2\varepsilon$,

$$R(f_{\hat{\theta}}) - R(f^*) \leq \frac{1}{n\varepsilon},$$

which is far from the objective (4.1) for small confidence levels (consider $\varepsilon = \exp(-\sqrt{n})$ for instance). The key idea is thus to average (soft) *truncated* values of the outputs. This is performed by taking

$$\hat{\theta} = \frac{1}{n\lambda} \sum_{i=1}^n \log \left(1 + \lambda Y_i + \frac{\lambda^2 Y_i^2}{2} \right),$$

with $\lambda = \sqrt{\frac{2\log(\varepsilon^{-1})}{n}}$. Since we have

$$\log \mathbb{E} \exp(n\lambda\hat{\theta}) = n \log \left(1 + \lambda \mathbb{E}(Y) + \frac{\lambda^2}{2} \mathbb{E}(Y^2) \right) \leq n\lambda \mathbb{E}(Y) + n \frac{\lambda^2}{2},$$

the exponential Chebyshev's inequality guarantees that with probability at least $1 - \varepsilon$, we have $n\lambda(\hat{\theta} - \mathbb{E}(Y)) \leq n \frac{\lambda^2}{2} + \log(\varepsilon^{-1})$, hence

$$\hat{\theta} - \mathbb{E}(Y) \leq \sqrt{\frac{2\log(\varepsilon^{-1})}{n}}.$$

Replacing Y by $-Y$ in the previous argument, we obtain that with probability at least $1 - \varepsilon$, we have

$$n\lambda \left\{ \mathbb{E}(Y) + \frac{1}{n\lambda} \sum_{i=1}^n \log \left(1 - \lambda Y_i + \frac{\lambda^2 Y_i^2}{2} \right) \right\} \leq n \frac{\lambda^2}{2} + \log(\varepsilon^{-1}).$$

Since $-\log(1 + x + x^2/2) \leq \log(1 - x + x^2/2)$, this implies $\mathbb{E}(Y) - \hat{\theta} \leq \sqrt{\frac{2\log(\varepsilon^{-1})}{n}}$. The two previous inequalities imply Inequality (4.1) (for $c = 2$), showing that sub-exponential tails are achievable even when we only assume that the random variable admits a finite second moment (see [5] for more details on the robust estimation of the mean of a random variable).

4.1.2. *Localized PAC-Bayesian inequalities to eliminate a logarithm factor.* Let us first recall that the Kullback-Leibler divergence between distributions ρ and μ defined on \mathcal{F} is

$$(4.2) \quad K(\rho, \mu) \triangleq \begin{cases} \mathbb{E}_{f \sim \rho} \log \left[\frac{d\rho}{d\mu}(f) \right] & \text{if } \rho \ll \mu, \\ +\infty & \text{otherwise,} \end{cases}$$

where $\frac{d\rho}{d\mu}$ denotes as usual the density of ρ w.r.t. μ . For any real-valued (measurable) function h defined on \mathcal{F} such that $\int \exp[h(f)]\pi(df) < +\infty$, we define the distribution π_h on \mathcal{F} by its density:

$$(4.3) \quad \frac{d\pi_h}{d\pi}(f) = \frac{\exp[h(f)]}{\int \exp[h(f')]\pi(df')}.$$

The analysis of statistical inference generally relies on upper bounding the supremum of an empirical process χ indexed by the functions in a model \mathcal{F} . One central tool to obtain these bounds is the concentration inequalities. An alternative approach, called the PAC-Bayesian one, consists in using the entropic equality

$$(4.4) \quad \mathbb{E} \exp \left(\sup_{\rho \in \mathcal{M}} \left\{ \int \rho(df) \chi(f) - K(\rho, \pi') \right\} \right) = \int \pi'(df) \mathbb{E} \exp(\chi(f)).$$

where \mathcal{M} is the set of probability distributions on \mathcal{F} .

Let $\tilde{r} : \mathcal{F} \rightarrow \mathbb{R}$ be an observable process such that for any $f \in \mathcal{F}$, we have

$$\mathbb{E} \exp(\chi(f)) \leq 1$$

for $\chi(f) = \lambda[R(f) - \tilde{r}(f)]$ and some $\lambda > 0$. Then (4.4) leads to: for any $\varepsilon > 0$, with probability at least $1 - \varepsilon$, for any distribution ρ on \mathcal{F} , we have

$$(4.5) \quad \int \rho(df) R(f) \leq \int \rho(df) \tilde{r}(f) + \frac{K(\rho, \pi') + \log(\varepsilon^{-1})}{\lambda}.$$

The lefthand-side quantity represents the expected risk with respect to the distribution ρ . To get the smallest upper bound on this quantity, a natural choice of the (posterior) distribution ρ is obtained by minimizing the righthand-side, that is by taking $\rho = \pi'_{-\lambda\tilde{r}}$ (with the notation introduced in (4.3)). This distribution concentrates on functions $f \in \mathcal{F}$ for which $\tilde{r}(f)$ is small. Without prior knowledge, one may want to choose a prior distribution $\pi' = \pi$ which is rather “flat” (e.g., the one induced by the Lebesgue measure in the case of a model \mathcal{F} defined by a bounded parameter set in some

Euclidean space). Consequently the Kullback-Leibler divergence $K(\rho, \pi')$, which should be seen as the complexity term, might be excessively large.

To overcome the lack of prior information and the resulting high complexity term, one can alternatively use a more “localized” prior distribution. Here we will use Gaussian distributions centered at the function of interest (for instance, the function f^*), and with covariance matrix proportional to the inverse of the Gram matrix Q . The idea of using PAC-Bayesian inequalities with Gaussian prior and posterior distributions goes back to Langford and Shawe-Taylor [7] in the context of linear classification.

4.2. Proofs of Theorems 2.1 and 2.2. To shorten the formulae, we will write X for $\varphi(X)$, which is equivalent to considering without loss of generality that the input space is \mathbb{R}^d and that the functions $\varphi_1, \dots, \varphi_d$ are the coordinate functions. Therefore, the function f_θ maps an input x to $\langle \theta, x \rangle$. With a slight abuse of notation, $R(\theta)$ will denote the risk of this prediction function.

Let us first assume that the matrix $Q_\lambda = Q + \lambda I$ is positive definite. This indeed does not restrict the generality of our study, even in the case when $\lambda = 0$, as we will discuss later (Remark 4.1).

Consider the change of coordinates

$$\bar{X} = Q_\lambda^{-1/2} X.$$

Let us introduce

$$\bar{R}(\theta) = \mathbb{E}[(\langle \theta, \bar{X} \rangle - Y)^2],$$

so that

$$\bar{R}(Q_\lambda^{1/2} \theta) = R(\theta) = \mathbb{E}[(\langle \theta, X \rangle - Y)^2].$$

Let

$$\bar{\Theta} = \{Q_\lambda^{1/2} \theta; \theta \in \Theta\}.$$

Consider

$$(4.6) \quad r(\theta) = \frac{1}{n} \sum_{i=1}^n (\langle \theta, X_i \rangle - Y_i)^2,$$

$$(4.7) \quad \bar{r}(\theta) = \frac{1}{n} \sum_{i=1}^n (\langle \theta, \bar{X}_i \rangle - Y_i)^2,$$

$$(4.8) \quad \theta_0 = \arg \min_{\theta \in \bar{\Theta}} \bar{R}(\theta) + \lambda \|Q_\lambda^{-1/2} \theta\|^2,$$

$$(4.9) \quad \hat{\theta} \in \arg \min_{\theta \in \Theta} r(\theta) + \lambda \|\theta\|^2,$$

$$(4.10) \quad \theta_1 = Q_\lambda^{1/2} \hat{\theta} \in \arg \min_{\theta \in \Theta} \bar{r}(\theta) + \lambda \|Q_\lambda^{-1/2} \theta\|^2.$$

For $\alpha > 0$, let us introduce the notation

$$\begin{aligned} W_i(\theta) &= \alpha \left\{ (\langle \theta, \bar{X}_i \rangle - Y_i)^2 - (\langle \theta_0, \bar{X}_i \rangle - Y_i)^2 \right\}, \\ W(\theta) &= \alpha \left\{ (\langle \theta, \bar{X} \rangle - Y)^2 - (\langle \theta_0, \bar{X} \rangle - Y)^2 \right\}. \end{aligned}$$

For any $\theta_2 \in \mathbb{R}^d$ and $\beta > 0$, let us consider the Gaussian distribution centered at θ_2

$$\rho_{\theta_2}(d\theta) = \left(\frac{\beta}{2\pi} \right)^{d/2} \exp \left(-\frac{\beta}{2} \|\theta - \theta_2\|^2 \right) d\theta.$$

LEMMA 4.1. *For any $\eta > 0$ and $\alpha > 0$, with probability at least $1 - \exp(-\eta)$, for any $\theta_2 \in \mathbb{R}^d$,*

$$\begin{aligned} & -n \int \rho_{\theta_2}(d\theta) \log \left\{ 1 - \mathbb{E}[W(\theta)] + \mathbb{E}[W(\theta)^2]/2 \right\} \\ & \leq - \sum_{i=1}^n \left(\int \rho_{\theta_2}(d\theta) \log \left\{ 1 - W_i(\theta) + W_i(\theta)^2/2 \right\} \right) + \mathcal{K}(\rho_{\theta_2}, \rho_{\theta_0}) + \eta, \end{aligned}$$

where $\mathcal{K}(\rho_{\theta_2}, \rho_{\theta_0})$ is the Kullback-Leibler divergence function :

$$\mathcal{K}(\rho_{\theta_2}, \rho_{\theta_0}) = \int \rho_{\theta_2}(d\theta) \log \left[\frac{d\rho_{\theta_2}(\theta)}{d\rho_{\theta_0}(\theta)} \right].$$

PROOF.

$$\mathbb{E} \left(\int \rho_{\theta_0}(d\theta) \prod_{i=1}^n \frac{1 - W_i(\theta) + W_i(\theta)^2/2}{1 - \mathbb{E}[W(\theta)] + \mathbb{E}[W(\theta)^2]/2} \right) \leq 1,$$

thus with probability at least $1 - \exp(-\eta)$

$$\log \left(\int \rho_{\theta_0}(d\theta) \prod_{i=1}^n \frac{1 - W_i(\theta) + W_i(\theta)^2/2}{1 - \mathbb{E}[W(\theta)] + \mathbb{E}[W(\theta)^2]/2} \right) \leq \eta.$$

We conclude from the convex inequality (see [4, page 159])

$$\log \left(\int \rho_{\theta_0}(d\theta) \exp[h(\theta)] \right) \geq \int \rho_{\theta_0}(d\theta) h(\theta) - \mathcal{K}(\rho_{\theta_2}, \rho_{\theta_0}).$$

□

Let us compute some useful quantities

$$(4.11) \quad \begin{aligned} \mathcal{K}(\rho_{\theta_2}, \rho_{\theta_0}) &= \frac{\beta}{2} \|\theta_2 - \theta_0\|^2, \\ \int \rho_{\theta_2}(d\theta) [W(\theta)] &= \alpha \int \rho_{\theta_2}(d\theta) \langle \theta - \theta_2, \bar{X} \rangle^2 + W(\theta_2) \end{aligned}$$

$$(4.12) \quad = W(\theta_2) + \alpha \frac{\|\bar{X}\|^2}{\beta},$$

$$(4.13) \quad \int \rho_{\theta_2}(d\theta) \langle \theta - \theta_2, \bar{X} \rangle^4 = \frac{3\|\bar{X}\|^4}{\beta^2},$$

$$(4.14) \quad \begin{aligned} \int \rho_{\theta_2}(d\theta) [W(\theta)^2] &= \alpha^2 \int \rho_{\theta_2}(d\theta) \langle \theta - \theta_0, \bar{X} \rangle^2 (\langle \theta + \theta_0, \bar{X} \rangle - 2Y)^2 \\ &= \alpha^2 \int \rho_{\theta_2}(d\theta) \left[\langle \theta - \theta_2 + \theta_2 - \theta_0, \bar{X} \rangle (\langle \theta - \theta_2 + \theta_2 + \theta_0, \bar{X} \rangle - 2Y) \right]^2 \\ &= \int \rho_{\theta_2}(d\theta) \left[\alpha \langle \theta - \theta_2, \bar{X} \rangle^2 + 2\alpha \langle \theta - \theta_2, \bar{X} \rangle (\langle \theta_2, \bar{X} \rangle - Y) + W(\theta_2) \right]^2 \\ &= \int \rho_{\theta_2}(d\theta) \left[\alpha^2 \langle \theta - \theta_2, \bar{X} \rangle^4 + 4\alpha^2 \langle \theta - \theta_2, \bar{X} \rangle^2 (\langle \theta_2, \bar{X} \rangle - Y)^2 + W(\theta_2)^2 \right. \\ &\quad \left. + 2\alpha \langle \theta - \theta_2, \bar{X} \rangle^2 W(\theta_2) \right] \\ &= \frac{3\alpha^2 \|\bar{X}\|^4}{\beta^2} + \frac{2\alpha \|\bar{X}\|^2}{\beta} \left[2\alpha (\langle \theta_2, \bar{X} \rangle - Y)^2 + W(\theta_2) \right] + W(\theta_2)^2. \end{aligned}$$

Using the fact that

$$2\alpha (\langle \theta_2, \bar{X} \rangle - Y)^2 + W(\theta_2) = 2\alpha (\langle \theta_0, \bar{X} \rangle - Y)^2 + 3W(\theta_2),$$

and that for any real numbers a and b , $6ab \leq 9a^2 + b^2$, we get

LEMMA 4.2.

$$(4.15) \quad \begin{aligned} \int \rho_{\theta_2}(d\theta) [W(\theta)] &= W(\theta_2) + \alpha \frac{\|\bar{X}\|^2}{\beta}, \\ \int \rho_{\theta_2}(d\theta) [W(\theta)^2] &= W(\theta_2)^2 + \frac{2\alpha \|\bar{X}\|^2}{\beta} \left[2\alpha (\langle \theta_0, \bar{X} \rangle - Y)^2 + 3W(\theta_2) \right] \\ &\quad + \frac{3\alpha^2 \|\bar{X}\|^4}{\beta^2} \end{aligned}$$

$$(4.17) \quad \leq 10W(\theta_2)^2 + \frac{4\alpha^2 \|\bar{X}\|^2}{\beta} (\langle \theta_0, \bar{X} \rangle - Y)^2 + \frac{4\alpha^2 \|\bar{X}\|^4}{\beta^2},$$

and the same holds true when W is replaced with W_i and (\bar{X}, Y) with (\bar{X}_i, Y_i) .

Another important thing to realize is that

$$\begin{aligned}
\mathbb{E}[\|\bar{X}\|^2] &= \mathbb{E}[\text{Tr}(\bar{X}\bar{X}^T)] &&= \mathbb{E}[\text{Tr}(Q_\lambda^{-1/2}XX^TQ_\lambda^{-1/2})] \\
&= \mathbb{E}[\text{Tr}(Q_\lambda^{-1}XX^T)] &&= \text{Tr}[Q_\lambda^{-1}\mathbb{E}(XX^T)] \\
(4.18) \quad &= \text{Tr}(Q_\lambda^{-1}(Q_\lambda - \lambda I)) &&= d - \lambda \text{Tr}(Q_\lambda^{-1}) = D.
\end{aligned}$$

We can weaken Lemma 4.1 (page 24) noticing that for any real number x , $x \leq -\log(1-x)$ and

$$\begin{aligned}
-\log\left(1-x+\frac{x^2}{2}\right) &= \log\left(\frac{1+x+x^2/2}{1+x^4/4}\right) \\
&\leq \log\left(1+x+\frac{x^2}{2}\right) \leq x+\frac{x^2}{2}.
\end{aligned}$$

We obtain with probability at least $1 - \exp(-\eta)$

$$\begin{aligned}
n\mathbb{E}[W(\theta_2)] &+ \frac{n\alpha}{\beta}\mathbb{E}[\|\bar{X}\|^2] - 5n\mathbb{E}[W(\theta_2)^2] \\
&- \mathbb{E}\left\{\frac{2n\alpha^2\|\bar{X}\|^2}{\beta}(\langle\theta_0, \bar{X}\rangle - Y)^2 + \frac{2n\alpha^2\|\bar{X}\|^4}{\beta^2}\right\} \\
&\leq \sum_{i=1}^n \left\{W_i(\theta_2) + 5W_i(\theta_2)^2\right. \\
&+ \left.\frac{\alpha\|\bar{X}_i\|^2}{\beta} + \frac{2\alpha^2\|\bar{X}_i\|^2}{\beta}(\langle\theta_0, \bar{X}_i\rangle - Y)^2 + \frac{2\alpha^2\|\bar{X}_i\|^4}{\beta^2}\right\} \\
&\quad + \frac{\beta}{2}\|\theta_2 - \theta_0\|^2 + \eta.
\end{aligned}$$

Noticing that for any real numbers a and b , $4ab \leq a^2 + 4b^2$, we can then bound

$$\begin{aligned}
\alpha^{-2}W(\theta_2)^2 &= \langle\theta_2 - \theta_0, \bar{X}\rangle^2(\langle\theta_2 + \theta_0, \bar{X}\rangle - 2Y)^2 \\
&= \langle\theta_2 - \theta_0, \bar{X}\rangle^2\left[\langle\theta_2 - \theta_0, \bar{X}\rangle + 2(\langle\theta_0, \bar{X}\rangle - Y)\right]^2 \\
&= \langle\theta_2 - \theta_0, \bar{X}\rangle^4 + 4\langle\theta_2 - \theta_0, \bar{X}\rangle^3(\langle\theta_0, \bar{X}\rangle - Y) \\
&\quad + 4\langle\theta_2 - \theta_0, \bar{X}\rangle^2(\langle\theta_0, \bar{X}\rangle - Y)^2 \\
&\leq 2\langle\theta_2 - \theta_0, \bar{X}\rangle^4 + 8\langle\theta_2 - \theta_0, \bar{X}\rangle^2(\langle\theta_0, \bar{X}\rangle - Y)^2.
\end{aligned}$$

THEOREM 4.3. *Let us put*

$$\begin{aligned}\widehat{D} &= \frac{1}{n} \sum_{i=1}^n \|\overline{X}_i\|^2 && \text{(let us remind that } D = \mathbb{E}[\|\overline{X}\|^2] \text{ from (4.18))}, \\ B_1 &= 2\mathbb{E}\left[\|\overline{X}\|^2 (\langle \theta_0, \overline{X} \rangle - Y)^2\right], \\ \widehat{B}_1 &= \frac{2}{n} \sum_{i=1}^n \left[\|\overline{X}_i\|^2 (\langle \theta_0, \overline{X}_i \rangle - Y_i)^2\right], \\ B_2 &= 2\mathbb{E}\left[\|\overline{X}\|^4\right], \\ \widehat{B}_2 &= \frac{2}{n} \sum_{i=1}^n \|\overline{X}_i\|^4, \\ B_3 &= 40 \sup\left\{\mathbb{E}\left[\langle u, \overline{X} \rangle^2 (\langle \theta_0, \overline{X} \rangle - Y)^2\right] : u \in \mathbb{R}^d, \|u\| = 1\right\}, \\ \widehat{B}_3 &= \sup\left\{\frac{40}{n} \sum_{i=1}^n \langle u, \overline{X}_i \rangle^2 (\langle \theta_0, \overline{X}_i \rangle - Y_i)^2 : u \in \mathbb{R}^d, \|u\| = 1\right\}, \\ B_4 &= 10 \sup\left\{\mathbb{E}\left[\langle u, \overline{X} \rangle^4\right] : u \in \mathbb{R}^d, \|u\| = 1\right\}, \\ \widehat{B}_4 &= \sup\left\{\frac{10}{n} \sum_{i=1}^n \langle u, \overline{X}_i \rangle^4 : u \in \mathbb{R}^d, \|u\| = 1\right\}.\end{aligned}$$

With probability at least $1 - \exp(-\eta)$, for any $\theta_2 \in \mathbb{R}^d$,

$$\begin{aligned}n\mathbb{E}[W(\theta_2)] &- \left[n\alpha^2(B_3 + \widehat{B}_3) + \frac{\beta}{2}\right] \|\theta_2 - \theta_0\|^2 \\ &- n\alpha^2(B_4 + \widehat{B}_4) \|\theta_2 - \theta_0\|^4 \\ &\leq \sum_{i=1}^n W_i(\theta_2) + \frac{n\alpha}{\beta}(\widehat{D} - D) + \frac{n\alpha^2}{\beta}(B_1 + \widehat{B}_1) + \frac{n\alpha^2}{\beta^2}(B_2 + \widehat{B}_2) + \eta.\end{aligned}$$

Let us now assume that $\theta_2 \in \overline{\Theta}$ and let us use the fact that $\overline{\Theta}$ is a convex set and that $\theta_0 = \arg \min_{\theta \in \overline{\Theta}} \overline{R}(\theta) + \lambda \|Q_\lambda^{-1/2} \theta\|^2$. Introduce $\theta_* = \arg \min_{\theta \in \mathbb{R}^d} \overline{R}(\theta) + \lambda \|Q_\lambda^{-1/2} \theta\|^2$. As we have

$$\overline{R}(\theta) + \lambda \|Q_\lambda^{-1/2} \theta\|^2 = \|\theta - \theta_*\|^2 + \overline{R}(\theta_*) + \lambda \|Q_\lambda^{-1/2} \theta_*\|^2,$$

the vector θ_0 is uniquely defined as the projection of θ_* on $\overline{\Theta}$ for the Euclidean distance, and for any $\theta_2 \in \overline{\Theta}$

$$\begin{aligned}
(4.19) \quad & \alpha^{-1}\mathbb{E}[W(\theta_2)] + \lambda\|Q_\lambda^{-1/2}\theta_2\|^2 - \lambda\|Q_\lambda^{-1/2}\theta_0\|^2 \\
& = \overline{R}(\theta_2) - \overline{R}(\theta_0) + \lambda\|Q_\lambda^{-1/2}\theta_2\|^2 - \lambda\|Q_\lambda^{-1/2}\theta_0\|^2 \\
& = \|\theta_2 - \theta_*\|^2 - \|\theta_0 - \theta_*\|^2 \\
& = \|\theta_2 - \theta_0\|^2 + 2\langle\theta_2 - \theta_0, \theta_0 - \theta_*\rangle \geq \|\theta_2 - \theta_0\|^2.
\end{aligned}$$

This and the inequality

$$\alpha^{-1}\sum_{i=1}^n W_i(\theta_1) + n\lambda\|Q_\lambda^{-1/2}\theta_1\|^2 - n\lambda\|Q_\lambda^{-1/2}\theta_0\|^2 \leq 0$$

leads to the following result.

THEOREM 4.4. *With probability at least $1 - \exp(-\eta)$,*

$$\begin{aligned}
R(\hat{\theta}) + \lambda\|\hat{\theta}\|^2 - \inf_{\theta \in \Theta} [R(\theta) + \lambda\|\theta\|^2] \\
= \alpha^{-1}\mathbb{E}[W(\theta_1)] + \lambda\|Q_\lambda^{-1/2}\theta_1\|^2 - \lambda\|Q_\lambda^{-1/2}\theta_0\|^2
\end{aligned}$$

is not greater than the smallest positive non degenerate root of the following polynomial equation as soon as it has one

$$\begin{aligned}
& \left\{1 - \left[\alpha(B_3 + \widehat{B}_3) + \frac{\beta}{2n\alpha}\right]\right\}x - \alpha(B_4 + \widehat{B}_4)x^2 \\
& = \frac{1}{\beta} \max(\widehat{D} - D, 0) + \frac{\alpha}{\beta}(B_1 + \widehat{B}_1) + \frac{\alpha}{\beta^2}(B_2 + \widehat{B}_2) + \frac{\eta}{n\alpha}.
\end{aligned}$$

PROOF. Let us remark first that when the polynomial appearing in the theorem has two distinct roots, they are of the same sign, due to the sign of its constant coefficient. Let $\widehat{\Omega}$ be the event of probability at least $1 - \exp(-\eta)$ described in Theorem 4.3 (page 27). For any realization of this event for which the polynomial described in Theorem 4.4 does not have two distinct positive roots, the statement of Theorem 4.4 is void, and therefore fulfilled. Let us consider now the case when the polynomial in question has two distinct positive roots $x_1 < x_2$. Consider in this case the random (trivially nonempty) closed convex set

$$\widehat{\Theta} = \left\{\theta \in \Theta : R(\theta) + \lambda\|\theta\|^2 \leq \inf_{\theta' \in \Theta} [R(\theta') + \lambda\|\theta'\|^2] + \frac{x_1 + x_2}{2}\right\}.$$

Let $\theta_3 \in \arg \min_{\theta \in \widehat{\Theta}} r(\theta) + \lambda\|\theta\|^2$ and $\theta_4 \in \arg \min_{\theta \in \Theta} r(\theta) + \lambda\|\theta\|^2$. We see from Theorem 4.3 that

$$(4.20) \quad R(\theta_3) + \lambda\|\theta_3\|^2 < R(\theta_0) + \lambda\|\theta_0\|^2 + \frac{x_1 + x_2}{2},$$

because it cannot be larger from the construction of $\widehat{\Theta}$. On the other hand, since $\widehat{\Theta} \subset \Theta$, the line segment $[\theta_3, \theta_4]$ is such that $[\theta_3, \theta_4] \cap \widehat{\Theta} \subset \arg \min_{\theta \in \widehat{\Theta}} r(\theta) + \lambda \|\theta\|^2$. We can therefore apply equation (4.20) to any point of $[\theta_3, \theta_4] \cap \widehat{\Theta}$, which proves that $[\theta_3, \theta_4] \cap \widehat{\Theta}$ is an open subset of $[\theta_3, \theta_4]$. But it is also a closed subset by construction, and therefore, as it is non empty and $[\theta_3, \theta_4]$ is connected, it proves that $[\theta_3, \theta_4] \cap \widehat{\Theta} = [\theta_3, \theta_4]$, and thus that $\theta_4 \in \widehat{\Theta}$. This can be applied to any choice of $\theta_3 \in \arg \min_{\theta \in \widehat{\Theta}} r(\theta) + \lambda \|\theta\|^2$ and $\theta_4 \in \arg \min_{\theta \in \Theta} r(\theta) + \lambda \|\theta\|^2$, proving that $\arg \min_{\theta \in \Theta} r(\theta) + \lambda \|\theta\|^2 \subset \arg \min_{\theta \in \widehat{\Theta}} r(\theta) + \lambda \|\theta\|^2$ and therefore that any $\theta_4 \in \arg \min_{\theta \in \Theta} r(\theta) + \lambda \|\theta\|^2$ is such that

$$R(\theta_4) + \lambda \|\theta_4\|^2 \leq \inf_{\theta \in \Theta} [R(\theta) + \lambda \|\theta\|^2] + x_1.$$

because the values between x_1 and x_2 are excluded by Theorem 4.3. \square

The actual convergence speed of the least squares estimator $\widehat{\theta}$ on Θ will depend on the speed of convergence of the “empirical bounds” \widehat{B}_k towards their expectations. We can rephrase the previous theorem in the following more practical way:

THEOREM 4.5. *Let $\eta_0, \eta_1, \dots, \eta_5$ be positive real numbers. With probability at least*

$$1 - \mathbb{P}(\widehat{D} > D + \eta_0) - \sum_{k=1}^4 \mathbb{P}(\widehat{B}_k - B_k > \eta_k) - \exp(-\eta_5),$$

$R(\widehat{\theta}) + \lambda \|\widehat{\theta}\|^2 - \inf_{\theta \in \Theta} [R(\theta) + \lambda \|\theta\|^2]$ is smaller than the smallest non degenerate positive root of

$$(4.21) \quad \left\{ 1 - \left[\alpha(2B_3 + \eta_3) + \frac{\beta}{2n\alpha} \right] \right\} x - \alpha(2B_4 + \eta_4)x^2 \\ = \frac{\eta_0}{\beta} + \frac{\alpha}{\beta}(2B_1 + \eta_1) + \frac{\alpha}{\beta^2}(2B_2 + \eta_2) + \frac{\eta_5}{n\alpha},$$

where we can optimize the values of $\alpha > 0$ and $\beta > 0$, since this equation has non random coefficients. For example, taking for simplicity

$$\alpha = \frac{1}{8B_3 + 4\eta_3}, \\ \beta = \frac{n\alpha}{2},$$

we obtain

$$x - \frac{2B_4 + \eta_4}{4B_3 + 2\eta_3}x^2 = \frac{16\eta_0(2B_3 + \eta_3)}{n} + \frac{8B_1 + 4\eta_1}{n} + \frac{32(2B_3 + \eta_3)(2B_2 + \eta_2)}{n^2} + \frac{8\eta_5(2B_3 + \eta_3)}{n}.$$

4.2.1. *Proof of Theorem 2.1.* Let us now deduce Theorem 2.1 (page 7) from Theorem 4.5. Let us first remark that with probability at least $1 - \varepsilon/2$

$$\widehat{D} \leq D + \sqrt{\frac{B_2}{\varepsilon n}},$$

because the variance of \widehat{D} is less than $\frac{B_2}{2n}$. For a given $\varepsilon > 0$, let us take $\eta_0 = \sqrt{\frac{B_2}{\varepsilon n}}$, $\eta_1 = B_1$, $\eta_2 = B_2$, $\eta_3 = B_3$ and $\eta_4 = B_4$. We get that $R_\lambda(\hat{\theta}) - \inf_{\theta \in \Theta} R_\lambda(\theta)$ is smaller than the smallest positive non degenerate root of

$$x - \frac{B_4}{2B_3}x^2 = \frac{48B_3}{n} \sqrt{\frac{B_2}{n\varepsilon}} + \frac{12B_1}{n} + \frac{288B_2B_3}{n^2} + \frac{24 \log(3/\varepsilon)B_3}{n},$$

with probability at least

$$1 - \frac{5\varepsilon}{6} - \sum_{k=1}^4 \mathbb{P}(\widehat{B}_k > B_k + \eta_k).$$

According to the weak law of large numbers, there is n_ε such that for any $n \geq n_\varepsilon$,

$$\sum_{k=1}^4 \mathbb{P}(\widehat{B}_k > B_k + \eta_k) \leq \varepsilon/6.$$

Thus, increasing n_ε and the constants to absorb the second order terms, we see that for some n_ε and any $n \geq n_\varepsilon$, with probability at least $1 - \varepsilon$, the excess risk is less than the smallest positive root of

$$x - \frac{B_4}{2B_3}x^2 = \frac{13B_1}{n} + \frac{24 \log(3/\varepsilon)B_3}{n}.$$

Now, as soon as $ac < 1/4$, the smallest positive root of $x - ax^2 = c$ is $\frac{2c}{1 + \sqrt{1 - 4ac}}$. This means that for n large enough, with probability at least $1 - \varepsilon$,

$$R_\lambda(\hat{\theta}) - \inf_{\theta} R_\lambda(\theta) \leq \frac{15B_1}{n} + \frac{25 \log(3/\varepsilon)B_3}{n},$$

which is precisely the statement of Theorem 2.1 (page 7), up to some change of notation.

4.2.2. *Proof of Theorem 2.2.* Let us now weaken Theorem 4.4 in order to make a more explicit non asymptotic result and obtain Theorem 2.2. From now on, we will assume that $\lambda = 0$. We start by giving bounds on the quantity defined in Theorem 4.3 in terms of

$$B = \sup_{f \in \text{span}\{\varphi_1, \dots, \varphi_d\} - \{0\}} \|f\|_\infty^2 / \mathbb{E}[f(X)]^2.$$

Since we have

$$\|\bar{X}\|^2 = \|Q_\lambda^{-1/2} X\|^2 \leq dB,$$

we get

$$\hat{d} = \frac{1}{n} \sum_{i=1}^n \|\bar{X}_i\|^2 \leq dB,$$

$$B_1 = 2\mathbb{E} \left[\|\bar{X}\|^2 (\langle \theta_0, \bar{X} \rangle - Y)^2 \right] \leq 2dB R(f^*),$$

$$\hat{B}_1 = \frac{2}{n} \sum_{i=1}^n \left[\|\bar{X}_i\|^2 (\langle \theta_0, \bar{X}_i \rangle - Y_i)^2 \right] \leq 2dB r(f^*),$$

$$B_2 = 2\mathbb{E} \left[\|\bar{X}\|^4 \right] \leq 2d^2 B^2,$$

$$\hat{B}_2 = \frac{2}{n} \sum_{i=1}^n \|\bar{X}_i\|^4 \leq 2d^2 B^2,$$

$$B_3 = 40 \sup \left\{ \mathbb{E} \left[\langle u, \bar{X} \rangle^2 (\langle \theta_0, \bar{X} \rangle - Y)^2 \right] : u \in \mathbb{R}^d, \|u\| = 1 \right\} \leq 40B R(f^*),$$

$$\hat{B}_3 = \sup \left\{ \frac{40}{n} \sum_{i=1}^n \langle u, \bar{X}_i \rangle^2 (\langle \theta_0, \bar{X}_i \rangle - Y_i)^2 : u \in \mathbb{R}^d, \|u\| = 1 \right\} \leq 40B r(f^*),$$

$$B_4 = 10 \sup \left\{ \mathbb{E} \left[\langle u, \bar{X} \rangle^4 \right] : u \in \mathbb{R}^d, \|u\| = 1 \right\} \leq 10B^2,$$

$$\hat{B}_4 = \sup \left\{ \frac{10}{n} \sum_{i=1}^n \langle u, \bar{X}_i \rangle^4 : u \in \mathbb{R}^d, \|u\| = 1 \right\} \leq 10B^2.$$

Let us put

$$a_0 = \frac{2dB + 4dB\alpha[R(f^*) + r(f^*)] + \eta}{\alpha n} + \frac{16B^2 d^2}{\alpha n^2},$$

$$a_1 = 3/4 - 40\alpha B[R(f^*) + r(f^*)],$$

and

$$a_2 = 20\alpha B^2.$$

Theorem 4.4 applied with $\beta = n\alpha/2$ implies that with probability at least $1 - \eta$ the excess risk $R(\hat{f}^{(\text{erm})}) - R(f^*)$ is upper bounded by the smallest positive root of $a_1x - a_2x^2 = a_0$ as soon as $a_1^2 > 4a_0a_2$. In particular, setting $\varepsilon = \exp(-\eta)$ when (4.22) holds, we have

$$R(\hat{f}^{(\text{erm})}) - R(f^*) \leq \frac{2a_0}{a_1 + \sqrt{a_1^2 - 4a_0a_2}} \leq \frac{2a_0}{a_1}.$$

We conclude that

THEOREM 4.6. *For any $\alpha > 0$ and $\varepsilon > 0$, with probability at least $1 - \varepsilon$, if the inequality*

$$(4.22) \quad 80 \left(\frac{(2 + 4\alpha[R(f^*) + r(f^*)])Bd + \log(\varepsilon^{-1})}{n} + \left(\frac{4Bd}{n} \right)^2 \right) < \left(\frac{3}{4B} - 40\alpha[R(f^*) + r(f^*)] \right)^2$$

holds, then we have

$$(4.23) \quad R(\hat{f}^{(\text{erm})}) - R(f^*) \leq \mathcal{J} \left(\frac{(2 + 4\alpha[R(f^*) + r(f^*)])Bd + \log(\varepsilon^{-1})}{n} + \left(\frac{4Bd}{n} \right)^2 \right),$$

where $\mathcal{J} = 8/(3\alpha - 160\alpha^2 B[R(f^*) + r(f^*)])$

Now, the Bienaymé-Chebyshev inequality implies

$$\mathbb{P}(r(f^*) - R(f^*) \geq t) \leq \frac{\mathbb{E}(r(f^*) - R(f^*))^2}{t^2} \leq \mathbb{E}[Y - f^*(X)]^4 / nt^2.$$

Under the finite moment assumption of Theorem 2.2, we obtain that for any $\varepsilon \geq 1/n$, with probability at least $1 - \varepsilon$,

$$r(f^*) < R(f^*) + \sqrt{\mathbb{E}[Y - f^*(X)]^4}.$$

From Theorem 4.6 and a union bound, by taking

$$\alpha = \left(80B[2R(f^*) + \sqrt{\mathbb{E}[Y - f^*(X)]^4}] \right)^{-1},$$

we get that with probability $1 - 2\varepsilon$,

$$R(\hat{f}^{(\text{erm})}) - R(f^*) \leq \mathcal{J}_1 B \left(\frac{3Bd + \log(\varepsilon^{-1})}{n} + \left(\frac{4Bd}{n} \right)^2 \right),$$

with $\mathcal{J}_1 = 640 \left(2R(f^*) + \sqrt{\mathbb{E}\{[Y - f^*(X)]^4\}} \right)$. This concludes the proof of Theorem 2.2.

REMARK 4.1. *Let us indicate now how to handle the case when Q is degenerate. Let us consider the linear subspace S of \mathbb{R}^d spanned by the eigenvectors of Q corresponding to positive eigenvalues. Then almost surely $\text{Span}\{X_i, i = 1, \dots, n\} \subset S$. Indeed for any θ in the kernel of Q , $\mathbb{E}(\langle \theta, X \rangle^2) = 0$ implies that $\langle \theta, X \rangle = 0$ almost surely, and considering a basis of the kernel, we see that $X \in S$ almost surely, S being orthogonal to the kernel of Q . Thus we can restrict the problem to S , as soon as we choose*

$$\hat{\theta} \in \text{span}\{X_1, \dots, X_n\} \cap \arg \min_{\theta} \sum_{i=1}^n (\langle \theta, X_i \rangle - Y_i)^2,$$

or equivalently with the notation $\mathbf{X} = (\varphi_j(X_i))_{1 \leq i \leq n, 1 \leq j \leq d}$ and $\mathbf{Y} = [Y_j]_{j=1}^n$,

$$\hat{\theta} \in \text{im } \mathbf{X}^T \cap \arg \min_{\theta} \|\mathbf{X}\theta - \mathbf{Y}\|^2$$

This proves that the results of this section apply to this special choice of the empirical least squares estimator. Since we have $\mathbb{R}^d = \ker \mathbf{X} \oplus \text{im } \mathbf{X}^T$, this choice is unique. Finally, we also have that (2.3) still holds by replacing d by $\text{rank}(Q)$.

4.3. *Proof of Theorem 3.1.* We use a similar notation as in Section 4.2: we write X for $\varphi(X)$. Therefore, the function f_{θ} maps an input x to $\langle \theta, x \rangle$. We consider the change of coordinates

$$\bar{X} = Q_{\lambda}^{-1/2} X.$$

Thus, from (4.18), we have $\mathbb{E}[\|\bar{X}\|^2] = D$. We will use

$$\bar{R}(\theta) = \mathbb{E}[(\langle \theta, \bar{X} \rangle - Y)^2],$$

so that $\bar{R}(Q^{1/2}\theta) = \mathbb{E}[(\langle \theta, X \rangle - Y)^2] = R(f_{\theta})$. Let

$$\bar{\Theta} = \{Q_{\lambda}^{1/2}\theta; \theta \in \Theta\}.$$

Consider

$$\theta_0 = \arg \min_{\theta \in \bar{\Theta}} \left\{ \bar{R}(\theta) + \lambda \|Q_{\lambda}^{-1/2}\theta\|^2 \right\}.$$

We thus have $\tilde{\theta} = Q_{\lambda}^{-1/2}\theta_0$, and

$$\sigma = \sqrt{\mathbb{E}[(\langle \theta_0, \bar{X} \rangle - Y)^2]},$$

$$\begin{aligned}
\chi &= \sup_{u \in \mathbb{R}^d} \frac{\mathbb{E}(\langle u, \bar{X} \rangle^4)^{1/2}}{\mathbb{E}(\langle u, \bar{X} \rangle^2)}, \\
\kappa &= \frac{\mathbb{E}(\|\bar{X}\|^4)^{1/2}}{\mathbb{E}(\|\bar{X}\|^2)} = \frac{\mathbb{E}(\|\bar{X}\|^4)^{1/2}}{D}, \\
\kappa' &= \frac{\mathbb{E}[(\langle \theta_0, \bar{X} \rangle - Y)^4]^{1/2}}{\sigma^2}, \\
T &= \|\bar{\Theta}\| = \max_{\theta, \theta' \in \bar{\Theta}} \|\theta - \theta'\|.
\end{aligned}$$

For $\alpha > 0$, we introduce

$$\begin{aligned}
J_i(\theta) &= \langle \theta, \bar{X}_i \rangle - Y_i, & J(\theta) &= \langle \theta, \bar{X} \rangle - Y \\
\bar{L}_i(\theta) &= \alpha(\langle \theta, \bar{X}_i \rangle - Y_i)^2, & \bar{L}(\theta) &= \alpha(\langle \theta, \bar{X} \rangle - Y)^2 \\
W_i(\theta) &= \bar{L}_i(\theta) - \bar{L}_i(\theta_0), & W(\theta) &= \bar{L}(\theta) - \bar{L}(\theta_0),
\end{aligned}$$

and

$$r'(\theta, \theta') = \lambda(\|Q_\lambda^{-1/2}\theta\|^2 - \|Q_\lambda^{-1/2}\theta'\|^2) + \frac{1}{n\alpha} \sum_{i=1}^n \psi(\bar{L}(\theta) - \bar{L}(\theta')).$$

Let $\bar{\theta} = Q_\lambda^{1/2}\hat{\theta} \in \bar{\Theta}$. We have

$$(4.24) \quad -r'(\theta_0, \bar{\theta}) = r'(\bar{\theta}, \theta_0) \leq \max_{\theta_1 \in \bar{\Theta}} r'(\bar{\theta}, \theta_1) \leq \gamma + \max_{\theta_1 \in \bar{\Theta}} r'(\theta_0, \theta_1),$$

where $\gamma = \max_{\theta_1 \in \bar{\Theta}} r'(\bar{\theta}, \theta_1) - \inf_{\theta \in \bar{\Theta}} \max_{\theta_1 \in \bar{\Theta}} r'(\theta, \theta_1)$ is a quantity which can be made arbitrary small by choice of the estimator. By using an upper bound $r'(\theta_0, \theta_1)$ that holds uniformly in θ_1 , we will control both left and right hand sides of (4.24).

To achieve this, we will upper bound

$$(4.25) \quad r'(\theta_0, \theta_1) = \lambda(\|Q_\lambda^{-1/2}\theta_0\|^2 - \|Q_\lambda^{-1/2}\theta_1\|^2) + \frac{1}{n\alpha} \sum_{i=1}^n \psi[-W_i(\theta_1)]$$

by the expectation of a distribution depending on θ_1 of a *quantity that does not depend on θ_1* , and then use the PAC-Bayesian argument to control this expectation uniformly in θ_1 . The distribution depending on θ_1 should therefore be taken such that for any $\theta_1 \in \bar{\Theta}$, its Kullback-Leibler divergence with respect to some fixed distribution is small (at least when θ_1 is close to θ_0).

Let us start with the following result.

LEMMA 4.7. *Let $f, g : \mathbb{R} \rightarrow \mathbb{R}$ be two Lebesgue measurable functions such that $f(x) \leq g(x)$, $x \in \mathbb{R}$. Let us assume that there exists $h \in \mathbb{R}$ such that $x \mapsto g(x) + h\frac{x^2}{2}$ is convex. Then for any probability distribution μ on the real line,*

$$f\left(\int x\mu(dx)\right) \leq \int g(x)\mu(dx) + \min\left\{\sup f - \inf f, \frac{h}{2}\text{Var}(\mu)\right\}.$$

PROOF. Let us put $x_0 = \int x\mu(dx)$ The function

$$x \mapsto g(x) + \frac{h}{2}(x - x_0)^2$$

is convex. Thus, by Jensen's inequality

$$f(x_0) \leq g(x_0) \leq \int \mu(dx) \left[g(x) + \frac{h}{2}(x - x_0)^2 \right] = \int g(x)\mu(dx) + \frac{h}{2}\text{Var}(\mu).$$

On the other hand

$$\begin{aligned} f(x_0) &\leq \sup f \leq \sup f + \int [g(x) - \inf f]\mu(dx) \\ &= \int g(x)\mu(dx) + \sup f - \inf f. \end{aligned}$$

The lemma is a combination of these two inequalities. \square

The above lemma will be used with $f = g = \psi$, where ψ is the increasing influence function

$$\psi(x) = \begin{cases} -\log(2), & x \leq -1, \\ \log(1 + x + x^2/2), & -1 \leq x \leq 0, \\ -\log(1 - x + x^2/2), & 0 \leq x \leq 1, \\ \log(2), & x \geq 1. \end{cases}$$

Since we have for any $x \in \mathbb{R}$

$$-\log\left(1 - x + \frac{x^2}{2}\right) = \log\left(\frac{1 + x + \frac{x^2}{2}}{1 + \frac{x^4}{4}}\right) < \log\left(1 + x + \frac{x^2}{2}\right),$$

the function ψ satisfies for any $x \in \mathbb{R}$

$$-\log\left(1 - x + \frac{x^2}{2}\right) < \psi(x) < \log\left(1 + x + \frac{x^2}{2}\right).$$

Moreover

$$\psi'(x) = \frac{1-x}{1-x+\frac{x^2}{2}}, \quad \psi''(x) = \frac{x(x-2)}{2(1-x+\frac{x^2}{2})^2} \geq -2, \quad 0 \leq x \leq 1,$$

showing (by symmetry) that the function $x \mapsto \psi(x) + 2x^2$ is convex on the real line.

For any $\theta' \in \mathbb{R}^d$ and $\beta > 0$, we consider the Gaussian distribution with mean θ' and covariance $\beta^{-1}I$:

$$\rho_{\theta'}(d\theta) = \left(\frac{\beta}{2\pi}\right)^{d/2} \exp\left(-\frac{\beta}{2}\|\theta - \theta'\|^2\right) d\theta.$$

From Lemmas 4.2 and 4.7 (with μ the distribution of $-W_i(\theta) + \frac{\alpha\|\bar{X}_i\|^2}{\beta}$ when θ is drawn from ρ_{θ_1} and for a fixed pair (X_i, Y_i)), we can see that

$$\begin{aligned} \psi[-W_i(\theta_1)] &= \psi\left\{\int \rho_{\theta_1}(d\theta) \left[-W_i(\theta) + \frac{\alpha\|\bar{X}_i\|^2}{\beta}\right]\right\} \\ &\leq \int \rho_{\theta_1}(d\theta) \psi\left[-W_i(\theta) + \frac{\alpha\|\bar{X}_i\|^2}{\beta}\right] \\ &\quad + \min\left\{\log(4), \text{Var}_{\rho_{\theta_1}}[\bar{L}_i(\theta)]\right\}. \end{aligned}$$

Let us compute

$$\begin{aligned} \frac{1}{\alpha^2} \text{Var}_{\rho_{\theta_1}}[\bar{L}_i(\theta)] &= \text{Var}_{\rho_{\theta_1}}[J_i^2(\theta) - J_i^2(\theta_1)] \\ &= \int \rho_{\theta_1}(d\theta) [J_i^2(\theta) - J_i^2(\theta_1)]^2 - \frac{\|\bar{X}_i\|^4}{\beta^2} \\ &= \int \rho_{\theta_1}(d\theta) \left[\langle \theta - \theta_1, \bar{X}_i \rangle^2 + 2\langle \theta - \theta_1, \bar{X}_i \rangle J_i(\theta_1)\right]^2 - \frac{\|\bar{X}_i\|^4}{\beta^2} \\ (4.26) \quad &= \frac{2\|\bar{X}_i\|^4}{\beta^2} + \frac{4\bar{L}_i(\theta_1)\|\bar{X}_i\|^2}{\alpha\beta}. \end{aligned}$$

Let $\xi \in (0, 1)$. Now we can remark that

$$\bar{L}_i(\theta_1) \leq \frac{\bar{L}_i(\theta)}{\xi} + \frac{\alpha\langle \theta - \theta_1, \bar{X}_i \rangle^2}{1-\xi}.$$

We get

$$\min\left\{\log(4), \text{Var}_{\rho_{\theta_1}}[\bar{L}_i(\theta)]\right\}$$

$$\begin{aligned}
&= \min\left\{\log(4), \frac{4\alpha\|\bar{X}_i\|^2\bar{L}_i(\theta_1)}{\beta} + \frac{2\alpha^2\|\bar{X}_i\|^4}{\beta^2}\right\} \\
&\leq \int \rho_{\theta_1}(d\theta) \min\left\{\log(4), \right. \\
&\quad \left. \frac{4\alpha\|\bar{X}_i\|^2\bar{L}_i(\theta)}{\beta\xi} + \frac{2\alpha^2\|\bar{X}_i\|^4}{\beta^2} + \frac{4\alpha^2\|\bar{X}_i\|^2\langle\theta - \theta_1, \bar{X}_i\rangle^2}{\beta(1-\xi)}\right\} \\
&\leq \int \rho_{\theta_1}(d\theta) \min\left\{\log(4), \frac{4\alpha\|\bar{X}_i\|^2\bar{L}_i(\theta)}{\beta\xi} + \frac{2\alpha^2\|\bar{X}_i\|^4}{\beta^2}\right\} \\
&\quad + \min\left\{\log(4), \frac{4\alpha^2\|\bar{X}_i\|^4}{\beta^2(1-\xi)}\right\}.
\end{aligned}$$

Let us now put $a = \frac{3}{\log(4)} < 2.17$, $b = a + a^2 \log(4) < 8.7$ and let us remark that

$$\begin{aligned}
&\min\{\log(4), x\} + \min\{\log(4), y\} \\
&\leq \log[1 + a \min\{\log(4), x\}] + \log(1 + ay) \\
&\leq \log(1 + ax + by), \quad x, y \in \mathbb{R}_+.
\end{aligned}$$

Thus

$$\begin{aligned}
&\min\left\{\log(4), \text{Var}_{\rho_{\theta_1}}[\bar{L}_i(\theta)]\right\} \\
&\leq \int \rho_{\theta_1}(d\theta) \log\left[1 + \frac{4a\alpha\|\bar{X}_i\|^2\bar{L}_i(\theta)}{\beta\xi} + \frac{2\alpha^2\|\bar{X}_i\|^4}{\beta^2}\left(a + \frac{2b}{1-\xi}\right)\right].
\end{aligned}$$

We can then remark that

$$\begin{aligned}
&\psi(x) + \log(1 + y) = \log[\exp[\psi(x)] + y \exp[\psi(x)]] \\
&\leq \log[\exp[\psi(x)] + 2y] \leq \log\left(1 + x + \frac{x^2}{2} + 2y\right), \quad x \in \mathbb{R}, y \in \mathbb{R}_+.
\end{aligned}$$

Thus, putting $c_0 = a + \frac{2b}{1-\xi}$, we get

$$(4.27) \quad \psi[-W_i(\theta_1)] \leq \int \rho_{\theta_1}(d\theta) \log[A_i(\theta)],$$

with

$$A_i(\theta) = 1 - W_i(\theta) + \frac{\alpha\|\bar{X}_i\|^2}{\beta} + \frac{1}{2}\left(-W_i(\theta) + \frac{\alpha\|\bar{X}_i\|^2}{\beta}\right)^2$$

$$+ \frac{8a\alpha\|\bar{X}_i\|^2\bar{L}_i(\theta)}{\beta\xi} + \frac{4c_0\alpha^2\|\bar{X}_i\|^4}{\beta^2}.$$

Similarly, we define $A(\theta)$ by replacing (X_i, Y_i) by (X, Y) . Since we have

$$\mathbb{E} \exp \left(\sum_{i=1}^n \log[A_i(\theta)] - n \log[\mathbb{E}A(\theta)] \right) = 1,$$

from the usual PAC-Bayesian argument, we have with probability at least $1 - \varepsilon$, for any $\theta_1 \in \mathbb{R}^d$,

$$\begin{aligned} \int \rho_{\theta_1}(d\theta) \left(\sum_{i=1}^n \log[A_i(\theta)] \right) - n \int \rho_{\theta_1}(d\theta) \log[A(\theta)] &\leq K(\rho_{\theta_1}, \rho_{\theta_0}) + \log(\varepsilon^{-1}) \\ &\leq \frac{\beta\|\theta_1 - \theta_0\|^2}{2} + \log(\varepsilon^{-1}) \end{aligned}$$

From (4.25) and (4.27), with probability at least $1 - \varepsilon$, for any $\theta_1 \in \mathbb{R}^d$, we get

$$\begin{aligned} r'(\theta_0, \theta_1) &\leq \frac{1}{\alpha} \log \left\{ 1 + \mathbb{E} \left[\int \rho_{\theta_1}(d\theta) \left(-W(\theta) + \frac{\alpha\|\bar{X}\|^2}{\beta} \right. \right. \right. \\ &\quad \left. \left. + \frac{1}{2} \left(-W(\theta) + \frac{\alpha\|\bar{X}\|^2}{\beta} \right)^2 + \frac{8a\alpha\|\bar{X}\|^2\bar{L}(\theta)}{\beta\xi} + \frac{4c_0\alpha^2\|\bar{X}\|^4}{\beta^2} \right) \right] \right\} \\ &\quad + \frac{\beta\|\theta_1 - \theta_0\|^2}{2n\alpha} + \frac{\log(\varepsilon^{-1})}{n\alpha} + \lambda(\|Q_\lambda^{-1/2}\theta_0\|^2 - \|Q_\lambda^{-1/2}\theta_1\|^2). \end{aligned}$$

Now from (4.26) and $\frac{\alpha\|\bar{X}\|^2}{\beta} = -\bar{L}(\theta_1) + \int \rho_{\theta_1}(d\theta)\bar{L}(\theta)$, we have

$$\begin{aligned} \int \rho_{\theta_1}(d\theta) \left(-W(\theta) + \frac{\alpha\|\bar{X}\|^2}{\beta} \right)^2 &= \text{Var}_{\rho_{\theta_1}}[\bar{L}(\theta)] + W(\theta_1)^2 \\ &= W(\theta_1)^2 + \frac{4\alpha\bar{L}(\theta_1)\|\bar{X}\|^2}{\beta} + \frac{2\alpha^2\|\bar{X}\|^4}{\beta^2}. \end{aligned}$$

PROPOSITION 4.8. *With probability at least $1 - \varepsilon$, for any $\theta_1 \in \mathbb{R}^d$,*

$$\begin{aligned} r'(\theta_0, \theta_1) &\leq \frac{1}{\alpha} \log \left\{ 1 + \mathbb{E} \left[-W(\theta_1) + \frac{W(\theta_1)^2}{2} + \frac{(2 + 8a/\xi)\alpha\|\bar{X}\|^2\bar{L}(\theta_1)}{\beta} \right. \right. \\ &\quad \left. \left. + \frac{(1 + 8a/\xi + 4c_0)\alpha^2\|\bar{X}\|^4}{\beta^2} \right] \right\} + \frac{\beta\|\theta_1 - \theta_0\|^2}{2n\alpha} + \frac{\log(\varepsilon^{-1})}{n\alpha} \\ &\quad + \lambda(\|Q_\lambda^{-1/2}\theta_0\|^2 - \|Q_\lambda^{-1/2}\theta_1\|^2) \end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{E} \left[J(\theta_0)^2 - J(\theta_1)^2 + \frac{1}{2\alpha} W(\theta_1)^2 + \frac{(2 + 8a/\xi) \|\bar{X}\|^2 \bar{L}(\theta_1)}{\beta} \right. \\
&\quad \left. + \frac{(1 + 8a/\xi + 4c_0)\alpha \|\bar{X}\|^4}{\beta^2} \right] + \frac{\beta \|\theta_1 - \theta_0\|^2}{2n\alpha} + \frac{\log(\varepsilon^{-1})}{n\alpha} \\
&\quad + \lambda (\|Q_\lambda^{-1/2} \theta_0\|^2 - \|Q_\lambda^{-1/2} \theta_1\|^2).
\end{aligned}$$

By using the triangular inequality and Cauchy-Schwarz's inequality, we get

$$\begin{aligned}
(4.28) \quad \frac{1}{\alpha^2} \mathbb{E}[W(\theta_1)^2] &= \mathbb{E} \left\{ [\langle \theta_1 - \theta_0, \bar{X} \rangle^2 + 2\langle \theta_1 - \theta_0, \bar{X} \rangle J(\theta_0)]^2 \right\} \\
&\leq \left\{ \mathbb{E}[\langle \theta_1 - \theta_0, \bar{X} \rangle^4]^{1/2} + 2\mathbb{E}[\langle \theta_1 - \theta_0, \bar{X} \rangle^4]^{1/4} \mathbb{E}[J(\theta_0)^4]^{1/4} \right\}^2 \\
&\leq \left\{ \chi \|\theta_1 - \theta_0\|^2 \mathbb{E} \left[\left\langle \frac{\theta_1 - \theta_0}{\|\theta_1 - \theta_0\|}, \bar{X} \right\rangle^2 \right] \right. \\
&\quad \left. + 2\|\theta_1 - \theta_0\| \sigma \sqrt{\kappa' \chi} \sqrt{\mathbb{E} \left[\left\langle \frac{\theta_1 - \theta_0}{\|\theta_1 - \theta_0\|}, \bar{X} \right\rangle^2 \right]} \right\}^2 \\
&\leq \frac{\chi q_{\max}}{q_{\max} + \lambda} \|\theta_1 - \theta_0\|^2 \left\{ \|\theta_1 - \theta_0\| \sqrt{\frac{\chi q_{\max}}{q_{\max} + \lambda}} + 2\sigma \sqrt{\kappa'} \right\}^2,
\end{aligned}$$

and

$$\begin{aligned}
(4.29) \quad \frac{1}{\alpha} \mathbb{E}[\|\bar{X}\|^2 \bar{L}(\theta_1)] &= \mathbb{E} \left\{ [\|\bar{X}\| \langle \theta_1 - \theta_0, \bar{X} \rangle + \|\bar{X}\| J(\theta_0)]^2 \right\} \\
&\leq \mathbb{E}[\|\bar{X}\|^4]^{1/2} \left\{ \mathbb{E}[\langle \theta_1 - \theta_0, \bar{X} \rangle^4]^{1/4} + \mathbb{E}[J(\theta_0)^4]^{1/4} \right\}^2 \\
&\leq \kappa D \left\{ \|\theta_1 - \theta_0\| \sqrt{\frac{\chi q_{\max}}{q_{\max} + \lambda}} + 2\sigma \sqrt{\kappa'} \right\}^2,
\end{aligned}$$

Let us put

$$\begin{aligned}
\tilde{R}(\theta) &= \bar{R}(\theta) + \lambda \|Q_\lambda^{-1/2} \theta\|^2, \\
c_1 &= 4(2 + 8a/\xi), \\
c_2 &= 4(1 + 8a/\xi + 4c_0), \\
\delta &= \frac{c_1 \kappa \kappa' D \sigma^2}{n} + \frac{2\chi \left(\frac{\log(\varepsilon^{-1})}{n} + \frac{c_2 \kappa^2 D^2}{n^2} \right) [2\sqrt{\kappa'} \sigma + \|\bar{\Theta}\| \sqrt{\chi}]^2}{1 - \frac{4c_1 \kappa \chi D}{n}}.
\end{aligned}$$

We have proved the following result.

PROPOSITION 4.9. *With probability at least $1 - \varepsilon$, for any $\theta_1 \in \mathbb{R}^d$,*

$$\begin{aligned} r'(\theta_0, \theta_1) &\leq \tilde{R}(\theta_0) - \tilde{R}(\theta_1) + \frac{\alpha}{2} \chi \|\theta_1 - \theta_0\|^2 [2\sqrt{\kappa'}\sigma + \|\theta_1 - \theta_0\|\sqrt{\chi}]^2 \\ &\quad + \frac{c_1\alpha}{4\beta} \kappa D [\sqrt{\kappa'}\sigma + \|\theta_1 - \theta_0\|\sqrt{\chi}]^2 + \frac{c_2\alpha\kappa^2 D^2}{4\beta^2} \\ &\quad + \frac{\beta\|\theta_1 - \theta_0\|^2}{2n\alpha} + \frac{\log(\varepsilon^{-1})}{n\alpha}. \end{aligned}$$

Let us assume from now on that $\theta_1 \in \bar{\Theta}$, our convex bounded parameter set. In this case, as seen in (4.19), we have $\|\theta_0 - \theta_1\|^2 \leq \tilde{R}(\theta_1) - \tilde{R}(\theta_0)$. We can also use the fact that

$$[\sqrt{\kappa'}\sigma + \|\theta_1 - \theta_0\|\sqrt{\chi}]^2 \leq 2\kappa'\sigma^2 + 2\chi\|\theta_1 - \theta_0\|^2.$$

We deduce from these remarks that with probability at least $1 - \varepsilon$,

$$\begin{aligned} r'(\theta_0, \theta_1) &\leq \left\{ -1 + \frac{\alpha\chi}{2} [2\sqrt{\kappa'}\sigma + \|\bar{\Theta}\|\sqrt{\chi}]^2 + \frac{\beta}{2n\alpha} + \frac{c_1\alpha\kappa D\chi}{2\beta} \right\} [\tilde{R}(\theta_1) - \tilde{R}(\theta_0)] \\ &\quad + \frac{c_1\alpha\kappa D\kappa'\sigma^2}{2\beta} + \frac{c_2\alpha\kappa^2 D^2}{4\beta^2} + \frac{\log(\varepsilon^{-1})}{n\alpha}. \end{aligned}$$

Let us assume that $n > 4c_1\kappa\chi D$ and let us choose

$$\begin{aligned} \beta &= \frac{n\alpha}{2}, \\ \alpha &= \frac{1}{2\chi[2\sqrt{\kappa'}\sigma + \|\bar{\Theta}\|\sqrt{\chi}]^2} \left(1 - \frac{4c_1\kappa\chi D}{n} \right), \end{aligned}$$

to get

$$r'(\theta_0, \theta_1) \leq -\frac{\tilde{R}(\theta_1) - \tilde{R}(\theta_0)}{2} + \delta.$$

Plugging this into (4.24), we get

$$\frac{\tilde{R}(\bar{\theta}) - \tilde{R}(\theta_0)}{2} - \delta \leq r'(\bar{\theta}, \theta_0) \leq \max_{\theta_1 \in \bar{\Theta}} \left(\frac{\tilde{R}(\theta_0) - \tilde{R}(\theta_1)}{2} \right) + \gamma + \delta = \gamma + \delta,$$

hence

$$\tilde{R}(\bar{\theta}) - \tilde{R}(\theta_0) \leq 2\gamma + 4\delta.$$

Computing the numerical values of the constants when $\xi = 0.8$ gives $c_1 < 95$ and $c_2 < 1511$.

References.

- [1] J.-Y. Audibert and O. Catoni. Robust linear regression through PAC-Bayesian truncation, 2010.
- [2] Y. Baraud. Model selection for regression on a fixed design. *Probability Theory and Related Fields*, 117(4):467–493, 2000.
- [3] L. Birgé and P. Massart. Minimum contrast estimators on sieves: exponential bounds and rates of convergence. *Bernoulli*, 4(3):329–375, 1998.
- [4] O. Catoni. *Statistical Learning Theory and Stochastic Optimization, Lectures on Probability Theory and Statistics, École d’Été de Probabilités de Saint-Flour XXXI – 2001*, volume 1851 of *Lecture Notes in Mathematics*. Springer, 2004. Pages 1–269.
- [5] O. Catoni. Challenging the empirical mean and empirical variance: a deviation study, 2010. arXiv:1009.2048v1 [math.ST].
- [6] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer, 2004.
- [7] J. Langford and J. Shawe-Taylor. PAC-bayes & margins. In *Advances in Neural Information Processing Systems*, pages 423–430, 2002.
- [8] A. Nemirovski. *Lectures on probability theory and statistics. Topics in non-parametric statistics. Ecole d’Été de Probabilités de Saint-Flour XXVIII-1998*. Springer-Verlag, 2000.
- [9] P.J. Rousseeuw and V.J. Yohai. Robust regression by means of S-estimators. *Robust and nonlinear time series analysis*, 26:256–272, 1984.
- [10] M. Sauvé. Piecewise polynomial estimation of a regression function. *IEEE Trans. Inf. Theor.*, 56(1):597–613, 2010.
- [11] A.B. Tsybakov. Optimal rates of aggregation. In B.Scholkopf and M.Warmuth, editors, *Computational Learning Theory and Kernel Machines, Lecture Notes in Artificial Intelligence*, volume 2777, pages 303–313. Springer, Heidelberg, 2003.
- [12] Y. Yang. Aggregating regression procedures for a better performance. *Bernoulli*, 10:25–47, 2004.
- [13] V.J. Yohai. High breakdown-point and high efficiency robust estimates for regression. *The Annals of Statistics*, 15(2):642–656, 1987.

APPENDIX A: EXPERIMENTAL RESULTS FOR THE MIN-MAX
TRUNCATED ESTIMATOR (SECTION 3.3)

TABLE 1

Comparison of the min-max truncated estimator \hat{f} with the ordinary least squares estimator $\hat{f}^{(ols)}$ for the mixture noise (see Section 3.4.1) with $\rho = 0.1$ and $p = 0.005$. In parenthesis, the 95%-confidence intervals for the estimated quantities.

	nb of iterations	nb of iter. with $R(\hat{f}) \neq R(\hat{f}^{(ols)})$	nb of iter. with $R(\hat{f}) < R(\hat{f}^{(ols)})$	$\mathbb{E}R(\hat{f}^{(ols)}) - R(f^*)$	$\mathbb{E}R(\hat{f}) - R(f^*)$	$\mathbb{E}R[(\hat{f}^{(ols)} f \neq \hat{f}^{(ols)})] - R(f^*)$	$\mathbb{E}[R(\hat{f}) f \neq \hat{f}^{(ols)}] - R(f^*)$
INC(n=200,d=1)	1000	419	405	0.567(±0.083)	0.178(±0.025)	1.191(±0.178)	0.262(±0.052)
INC(n=200,d=2)	1000	506	498	1.055(±0.112)	0.271(±0.030)	1.884(±0.193)	0.334(±0.050)
HCC(n=200,d=2)	1000	502	494	1.045(±0.103)	0.267(±0.024)	1.866(±0.174)	0.316(±0.032)
TS(n=200,d=2)	1000	561	554	1.069(±0.089)	0.310(±0.027)	1.720(±0.132)	0.367(±0.036)
INC(n=1000,d=2)	1000	402	392	0.204(±0.015)	0.109(±0.008)	0.316(±0.029)	0.081(±0.011)
INC(n=1000,d=10)	1000	950	946	1.030(±0.041)	0.228(±0.016)	1.051(±0.042)	0.207(±0.014)
HCC(n=1000,d=10)	1000	942	942	0.980(±0.038)	0.222(±0.015)	1.008(±0.039)	0.203(±0.015)
TS(n=1000,d=10)	1000	976	973	1.009(±0.037)	0.228(±0.017)	1.018(±0.038)	0.217(±0.016)
INC(n=2000,d=2)	1000	209	207	0.104(±0.007)	0.078(±0.005)	0.206(±0.021)	0.082(±0.012)
HCC(n=2000,d=2)	1000	184	183	0.099(±0.007)	0.076(±0.005)	0.196(±0.023)	0.070(±0.010)
TS(n=2000,d=2)	1000	172	171	0.101(±0.007)	0.080(±0.005)	0.206(±0.020)	0.083(±0.012)
INC(n=2000,d=10)	1000	669	669	0.510(±0.018)	0.206(±0.012)	0.572(±0.023)	0.117(±0.009)
HCC(n=2000,d=10)	1000	669	669	0.499(±0.018)	0.207(±0.013)	0.561(±0.023)	0.125(±0.011)
TS(n=2000,d=10)	1000	754	753	0.516(±0.018)	0.195(±0.013)	0.558(±0.022)	0.131(±0.011)

TABLE 2

Comparison of the min-max truncated estimator \hat{f} with the ordinary least squares estimator $\hat{f}^{(\text{ols})}$ for the mixture noise (see Section 3.4.1) with $\rho = 0.4$ and $p = 0.005$. In parenthesis, the 95%-confidence intervals for the estimated quantities.

	nb of iterations	nb of iter. with $R(\hat{f}) \neq R(\hat{f}^{(\text{ols})})$	nb of iter. with $R(\hat{f}) < R(\hat{f}^{(\text{ols})})$	$\mathbb{E}R(\hat{f}^{(\text{ols})}) - R(f^*)$	$\mathbb{E}R(\hat{f}) - R(f^*)$	$\mathbb{E}R[(\hat{f}^{(\text{ols})}) \hat{f} \neq \hat{f}^{(\text{ols})}] - R(f^*)$	$\mathbb{E}[R(\hat{f}) \hat{f} \neq \hat{f}^{(\text{ols})}] - R(f^*)$
INC(n=200,d=1)	1000	234	211	0.551(±0.063)	0.409(±0.042)	1.211(±0.210)	0.606(±0.110)
INC(n=200,d=2)	1000	195	186	1.046(±0.088)	0.788(±0.061)	2.174(±0.293)	0.848(±0.118)
HCC(n=200,d=2)	1000	222	215	1.028(±0.079)	0.748(±0.051)	2.157(±0.243)	0.897(±0.112)
TS(n=200,d=2)	1000	291	268	1.053(±0.079)	0.805(±0.058)	1.701(±0.186)	0.851(±0.093)
INC(n=1000,d=2)	1000	127	117	0.201(±0.013)	0.181(±0.012)	0.366(±0.053)	0.207(±0.035)
INC(n=1000,d=10)	1000	262	249	1.023(±0.035)	0.902(±0.030)	1.238(±0.081)	0.777(±0.054)
HCC(n=1000,d=10)	1000	201	192	0.991(±0.033)	0.902(±0.031)	1.235(±0.088)	0.790(±0.067)
TS(n=1000,d=10)	1000	171	162	1.009(±0.033)	0.951(±0.031)	1.166(±0.098)	0.825(±0.071)
INC(n=2000,d=2)	1000	80	77	0.105(±0.007)	0.099(±0.006)	0.214(±0.042)	0.135(±0.029)
HCC(n=2000,d=2)	1000	44	42	0.102(±0.007)	0.099(±0.007)	0.187(±0.050)	0.120(±0.034)
TS(n=2000,d=2)	1000	47	47	0.101(±0.007)	0.099(±0.007)	0.147(±0.032)	0.103(±0.026)
INC(n=2000,d=10)	1000	116	113	0.511(±0.016)	0.491(±0.016)	0.611(±0.052)	0.437(±0.042)
HCC(n=2000,d=10)	1000	110	105	0.500(±0.016)	0.481(±0.015)	0.602(±0.056)	0.430(±0.044)
TS(n=2000,d=10)	1000	101	98	0.511(±0.016)	0.499(±0.016)	0.601(±0.054)	0.486(±0.051)

TABLE 3
 Comparison of the min-max truncated estimator \hat{f} with the ordinary least squares estimator $\hat{f}^{(\text{ols})}$ with the heavy-tailed noise (see Section 3.4.1).

	nb of iterations	nb of iter. with $R(\hat{f}) \neq R(\hat{f}^{(\text{ols})})$	nb of iter. with $R(\hat{f}) < R(\hat{f}^{(\text{ols})})$	$\mathbb{E}R(\hat{f}^{(\text{ols})}) - R(f^*)$	$\mathbb{E}R(\hat{f}) - R(f^*)$	$\mathbb{E}R[\hat{f}^{(\text{ols})} \hat{f} \neq \hat{f}^{(\text{ols})}] - R(f^*)$	$\mathbb{E}[R(\hat{f}) \hat{f} \neq \hat{f}^{(\text{ols})}] - R(f^*)$
INC(n=200,d=1)	1000	163	145	7.72(±3.46)	3.92(±0.409)	30.52(±20.8)	7.20(±1.61)
INC(n=200,d=2)	1000	104	98	22.69(±23.14)	19.18(±23.09)	45.36(±14.1)	11.63(±2.19)
HCC(n=200,d=2)	1000	120	117	18.16(±12.68)	8.07(±0.718)	99.39(±105)	15.34(±4.41)
TS(n=200,d=2)	1000	110	105	43.89(±63.79)	39.71(±63.76)	48.55(±18.4)	10.59(±2.01)
INC(n=1000,d=2)	1000	104	100	3.98(±2.25)	1.78(±0.128)	23.18(±21.3)	2.03(±0.56)
INC(n=1000,d=10)	1000	253	242	16.36(±5.10)	7.90(±0.278)	41.25(±19.8)	7.81(±0.69)
HCC(n=1000,d=10)	1000	220	211	13.57(±1.93)	7.88(±0.255)	33.13(±8.2)	7.28(±0.59)
TS(n=1000,d=10)	1000	214	211	18.67(±11.62)	13.79(±11.52)	30.34(±7.2)	7.53(±0.58)
INC(n=2000,d=2)	1000	113	103	1.56(±0.41)	0.89(±0.059)	6.74(±3.4)	0.86(±0.18)
HCC(n=2000,d=2)	1000	105	97	1.66(±0.43)	0.95(±0.062)	7.87(±3.8)	1.13(±0.23)
TS(n=2000,d=2)	1000	101	95	1.59(±0.64)	0.88(±0.058)	8.03(±6.2)	1.04(±0.22)
INC(n=2000,d=10)	1000	259	255	8.77(±4.02)	4.23(±0.154)	21.54(±15.4)	4.03(±0.39)
HCC(n=2000,d=10)	1000	250	242	6.98(±1.17)	4.13(±0.127)	15.35(±4.5)	3.94(±0.25)
TS(n=2000,d=10)	1000	238	233	8.49(±3.61)	5.95(±3.486)	14.82(±3.8)	4.17(±0.30)

TABLE 4
 Comparison of the min-max truncated estimator \hat{f} with the ordinary least squares estimator $\hat{f}^{(\text{ols})}$ with the asymmetric heavy-tailed noise (see Section 3.4.1).

	nb of iterations	nb of iter. with $R(\hat{f}) \neq R(\hat{f}^{(\text{ols})})$	nb of iter. with $R(\hat{f}) < R(\hat{f}^{(\text{ols})})$	$\mathbb{E}R(\hat{f}^{(\text{ols})}) - R(f^*)$	$\mathbb{E}R(\hat{f}) - R(f^*)$	$\mathbb{E}R(\hat{f}^{(\text{ols})}) f \neq \hat{f}^{(\text{ols})} - R(f^*)$	$\mathbb{E}R(\hat{f}) f \neq \hat{f}^{(\text{ols})} - R(f^*)$
INC(n=200,d=1)	1000	87	77	5.49(±3.07)	3.00(±0.330)	35.44(±34.7)	6.85(±2.48)
INC(n=200,d=2)	1000	70	66	19.25(±23.23)	17.4(±23.2)	37.95(±13.1)	11.05(±2.87)
HCC(n=200,d=2)	1000	67	66	7.19(±0.88)	5.81(±0.397)	31.52(±10.5)	10.87(±2.64)
TS(n=200,d=2)	1000	76	68	39.80(±64.09)	37.9(±64.1)	34.28(±14.8)	9.21(±2.05)
INC(n=1000,d=2)	1000	101	92	2.81(±2.21)	1.31(±0.106)	16.76(±21.8)	1.88(±0.69)
INC(n=1000,d=10)	1000	211	195	10.71(±4.53)	5.86(±0.222)	29.00(±21.3)	6.03(±0.71)
HCC(n=1000,d=10)	1000	197	185	8.67(±1.16)	5.81(±0.177)	20.31(±5.59)	5.79(±0.43)
TS(n=1000,d=10)	1000	258	233	13.62(±11.27)	11.3(±11.2)	14.68(±2.45)	5.60(±0.36)
INC(n=2000,d=2)	1000	106	92	1.04(±0.37)	0.64(±0.042)	4.54(±3.45)	0.79(±0.16)
HCC(n=2000,d=2)	1000	99	90	0.90(±0.11)	0.66(±0.042)	3.23(±0.93)	0.82(±0.16)
TS(n=2000,d=2)	1000	84	81	1.11(±0.66)	0.60(±0.042)	6.80(±7.79)	0.69(±0.17)
INC(n=2000,d=10)	1000	238	222	6.32(±4.18)	3.07(±0.147)	16.84(±17.5)	3.18(±0.51)
HCC(n=2000,d=10)	1000	221	203	4.49(±0.98)	2.98(±0.091)	9.76(±4.39)	2.93(±0.22)
TS(n=2000,d=10)	1000	412	350	5.93(±3.51)	4.59(±3.44)	6.07(±1.76)	2.84(±0.16)

TABLE 5
 Comparison of the min-max truncated estimator \hat{f} with the ordinary least squares estimator $\hat{f}^{(ols)}$ for standard Gaussian noise.

	nb of iter.	nb of iter. with $R(\hat{f}) \neq R(\hat{f}^{(ols)})$	nb of iter. with $R(\hat{f}) < R(\hat{f}^{(ols)})$	$\mathbb{E}R(\hat{f}^{(ols)}) - R(f^*)$	$\mathbb{E}R(\hat{f}) - R(f^*)$	$\mathbb{E}R(\hat{f}^{(ols)} \hat{f} \neq \hat{f}^{(ols)}) - R(f^*)$	$\mathbb{E}[R(\hat{f}) \hat{f} \neq \hat{f}^{(ols)}] - R(f^*)$
INC(n=200,d=1)	1000	20	8	0.541(±0.048)	0.541(±0.048)	0.401(±0.168)	0.397(±0.167)
INC(n=200,d=2)	1000	1	0	1.051(±0.067)	1.051(±0.067)	2.566	2.757
HCC(n=200,d=2)	1000	1	0	1.051(±0.067)	1.051(±0.067)	2.566	2.757
TS(n=200,d=2)	1000	0	0	1.068(±0.067)	1.068(±0.067)	-	-
INC(n=1000,d=2)	1000	0	0	0.203(±0.013)	0.203(±0.013)	-	-
INC(n=1000,d=10)	1000	0	0	1.023(±0.029)	1.023(±0.029)	-	-
HCC(n=1000,d=10)	1000	0	0	1.023(±0.029)	1.023(±0.029)	-	-
TS(n=1000,d=10)	1000	0	0	0.997(±0.028)	0.997(±0.028)	-	-
INC(n=2000,d=2)	1000	0	0	0.112(±0.007)	0.112(±0.007)	-	-
HCC(n=2000,d=2)	1000	0	0	0.112(±0.007)	0.112(±0.007)	-	-
TS(n=2000,d=2)	1000	0	0	0.098(±0.006)	0.098(±0.006)	-	-
INC(n=2000,d=10)	1000	0	0	0.517(±0.015)	0.517(±0.015)	-	-
HCC(n=2000,d=10)	1000	0	0	0.517(±0.015)	0.517(±0.015)	-	-
TS(n=2000,d=10)	1000	0	0	0.501(±0.015)	0.501(±0.015)	-	-

FIG 1. Surrounding points are the points of the training set generated several times from $TS(1000, 10)$ (with the mixture noise with $p = 0.005$ and $\rho = 0.4$) that are not taken into account in the min-max truncated estimator (to the extent that the estimator would not change by removing simultaneously all these points). The min-max truncated estimator $x \mapsto \hat{f}(x)$ appears in dash-dot line, while $x \mapsto \mathbb{E}(Y|X = x)$ is in solid line. In these six simulations, it outperforms the ordinary least squares estimator.

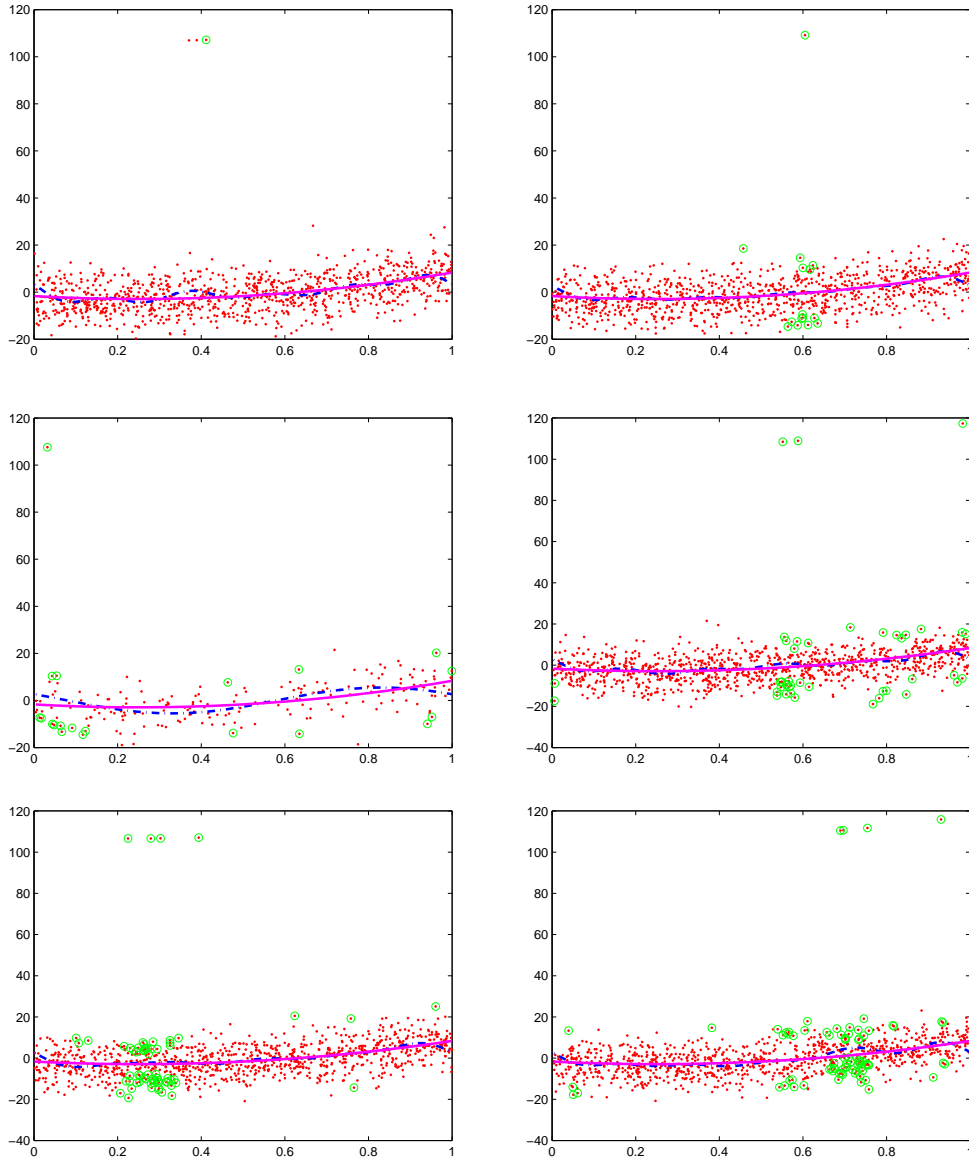


FIG 2. *Surrounding points are the points of the training set generated several times from $TS(200, 2)$ (with the heavy-tailed noise) that are not taken into account in the min-max truncated estimator (to the extent that the estimator would not change by removing these points). The min-max truncated estimator $x \mapsto \hat{f}(x)$ appears in dash-dot line, while $x \mapsto \mathbb{E}(Y|X = x)$ is in solid line. In these six simulations, it outperforms the ordinary least squares estimator. Note that in the last figure, it does not consider 64 points among the 200 training points.*

