

# HIGH-DIMENSIONAL ROBUST REGRESSION UNDER HEAVY-TAILED DATA: ASYMPTOTICS AND UNIVERSALITY

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We investigate the high-dimensional properties of robust regression estimators in the presence of heavy-tailed contamination of both the covariates and response functions. In particular, we provide a sharp asymptotic characterisation of M-estimators trained on a family of elliptical covariate and noise data distributions including cases where second and higher moments do not exist. We show that, despite being consistent, the Huber loss with optimally tuned location parameter  $\delta$  is suboptimal in the high-dimensional regime in the presence of heavy-tailed noise, highlighting the necessity of further regularisation to achieve optimal performance. This result also uncovers the existence of a curious transition in  $\delta$  as a function of the sample complexity and contamination. Moreover, we derive the decay rates for the excess risk of ridge regression. We show that, while it is both optimal and universal for noise distributions with finite second moment, its decay rate can be considerably faster when the covariates' second moment does not exist. Finally, we show that our formulas readily generalise to a richer family of models and data distributions, such as generalised linear estimation with arbitrary convex regularisation trained on mixture models.

## 1 INTRODUCTION

Consider the classical statistical problem of estimating a vector  $\beta_\star \in \mathbb{R}^d$  from  $n$  i.i.d. pairs of observations  $\mathcal{D} := \{(\mathbf{x}_i, y_i) \in \mathbb{R}^{d+1} : i \in [n]\}$  from a linear model:

$$y_i = \beta_\star^\top \mathbf{x}_i + \eta_i, \quad (1a)$$

where  $\mathbf{x}_i \in \mathbb{R}^d$  are the covariates and  $\eta_i \in \mathbb{R}$  is the label noise, which we assume to be a random quantity with zero mean and independent of the covariates. This manuscript is concerned with the characterisation of the following class of (regularised) M-estimators:

$$\hat{\beta}_\lambda := \arg \min_{\beta \in \mathbb{R}^d} \sum_{i=1}^n \rho(y_i - \beta^\top \mathbf{x}_i) + \frac{\lambda}{2} \|\beta\|_2^2, \quad \lambda \in \mathbb{R}_+ := [0, +\infty), \quad (1b)$$

with  $\rho: \mathbb{R} \rightarrow \mathbb{R}_+$  a convex objective function. A popular particular example is least-squares regression, where  $\rho(t) = t^2/2$  and  $\lambda = 0$ : in this case,  $\hat{\beta}_\lambda$  is the maximum likelihood estimator for  $\beta_\star$  if  $\eta_i \sim \mathcal{N}(0, 1)$ . It is well-known, however, that the least-squares estimator suffers in the presence of outliers in the data (Huber, 1973). Indeed, the fact that the gradient of the loss  $\rho'(t) = t$  is unbounded implies that an outlier can have an arbitrary influence over the solution of Eq. 1b. Tailoring the objective  $\rho$  to be insensitive (i.e., *robust*) to outliers in the training data  $\mathcal{D}$  is a classical statistical problem (Huber, 2004; Hampel et al., 2011; Rousseeuw and Leroy, 2005; Maronna et al., 2019). In his seminal work, Huber (1964) has shown that judiciously trimming the squared loss

$$\rho_\delta(t) = \begin{cases} t^2/2 & \text{if } |t| < \delta \\ \delta|t| - \delta^2/2 & \text{otherwise} \end{cases} \quad (2)$$

|                                     | $\varrho(\sigma)$                                       | $p(x) = \mathbb{E}_\sigma[\mathcal{N}(x; 0, \sigma^2)]$   | $k$ th moment exists if |
|-------------------------------------|---|---|-------------------------|
| Inverse-Gamma                       | $\frac{2b^a \exp(-b/\sigma^2)}{\Gamma(a)\sigma^{2a+1}}$ | $\frac{(2b)^a \Gamma(a+1/2)}{\sqrt{\pi}\Gamma(a)(2b+x^2)^{a+1/2}}$  | $k < 2a$                |
| for $a = b = 1/2$                   |   | $\frac{1}{\pi} \frac{1}{1+x^2}$ (Cauchy)  | none                    |
| for $2a = 2b = n$                   |   | $\frac{\Gamma(\frac{n+1}{2})}{\sqrt{\pi n}\Gamma(n/2)} \left(1 + \frac{x^2}{n}\right)^{-\frac{n+1}{2}}$ (Student-t) | $k < n$                 |
| for $a = 1 + b \rightarrow +\infty$ |   | $\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$ (Gaussian)   | all                     |
| Pareto                              | $\frac{2a\theta(\sigma-1)}{\sigma^{2a+1}}$              | $\frac{a2^a \gamma(a+1/2, 1/2x^2)}{\sqrt{\pi} x ^{2a+1}}$   | $k < 2a$                |
| for $a \rightarrow +\infty$         |   | $\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$ (Gaussian)   | all                     |

Table 1: Concrete examples of distributions  $p(x) = \mathbb{E}_\sigma[\mathcal{N}(x; 0, \sigma^2)]$  for different densities  $\varrho$  of  $\sigma$ . Here  $\gamma(a, x)$  is the lower incomplete Gamma function. Note that relevant distributions, such as the Cauchy and the Student’s  $t$ -distribution, appear as special cases. The last column shows the values of the parameters of  $\varrho$  for which the  $k$ th moment of the distribution is finite.

for  $\delta \geq 0$  provides a convenient solution to this problem while preserving the convexity of the task in Eq. 1b. Indeed, besides enjoying standard statistical guarantees in the classical limit of  $n \rightarrow \infty$ , such as consistency and asymptotic normality (Huber, 1973; van der Vaart, 2000), the so-called *Huber loss* has been shown to be optimal in different regards. For instance, Huber (1964) has shown it achieves minimax asymptotic variance under symmetric contamination of the normal distribution. It also has the smallest asymptotic variance among losses with bounded sensitivity to outliers, as it can be formalised by Hampel’s influence function (Hampel, 1974). However, while these guarantees are fit for a classical statistical regime where data is abundant ( $n \gg d$ ), they fall short in modern tasks where the number of features can be comparable to the quantity of data ( $n \approx d$ ). Investigating the properties of estimators in the proportional *high-dimensional regime* where  $n, d \rightarrow \infty$  at fixed sample complexity  $\alpha = n/d$  has been a major endeavor in the statistical literature in the past decade, where it has been shown that standard guarantees for the maximum likelihood estimator, such as unbiasedness (Javanmard and Montanari, 2018; Sur and Candès, 2019; 2020; Bellec et al., 2022; Zhao et al., 2022) and calibration (Bai et al., 2021; Clarté et al., 2023a;b) break down in this regime. However, the majority of these works have focused in the case of sub-Gaussian features and bounded noise variance. Our goal in this manuscript is to go beyond this assumption by providing a high-dimensional characterisation of M-estimators for a family of heavy-tailed distributions for both the covariates and noise distributions.

**Heavy-tailed data** — In the following, we consider a family of covariate distributions parametrised as:

$$\mathbf{x}_i = \sigma_i \mathbf{z}_i, \quad \mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, 1/d\mathbf{I}_d) \quad \sigma_i \sim \varrho, \quad (3)$$

where for each  $i \in [n]$ ,  $\sigma_i \in \mathbb{R}_+^* := (0, +\infty)$  is an independent random variable with probability density  $\varrho$  supported on the positive real line. This class of covariates appeared under different contexts in physics (Beck, 2003), statistics (El Karoui, 2018; Adomaityte et al., 2023) and signal processing (Wainwright and Simoncelli, 1999) where it has been shown that by judiciously choosing  $\varrho$  yields a large family of *power-law tailed distributions* for  $\mathbf{x}$ , see Table 1 for concrete examples. In particular, we are interested in investigating the impact of heavy-tail contamination of both the covariates and responses, and the sensitivity of M-estimators of the type in Eq. 1b to them. More concretely, in the following we consider the *Huber  $\epsilon$ -contamination model* for the covariates by assuming

$$p(\mathbf{x}) = \mathbb{E}[\mathcal{N}(\mathbf{x}; \mathbf{0}, \sigma^2/d\mathbf{I}_d)] \quad \text{with} \quad \sigma \sim \varrho = (1 - \epsilon_c)\delta_{\sigma,1} + \epsilon_c\varrho_0, \quad \epsilon_c \in [0, 1], \quad (4)$$

where  $\varrho_0$  is a density over  $\mathbb{R}_+^*$ . The distribution  $p(\mathbf{x})$  belongs therefore to the family defined by Eq. 3: the quantity  $\epsilon_c$  measures, therefore, the *contamination* of a pure Gaussian covariate distribution (recovered for  $\epsilon_c = 0$ ) with a possibly heavy-tailed law. Moreover, our results can be extended to the case of *mixtures*, i.e., the case of the covariates grouped in  $K$  different clouds. This generalisation is introduced and briefly discussed in Section 4.

**Main contributions** — The **key contributions** in this manuscript are:

- We provide an asymptotic characterisation of the statistics of the M-estimator  $\hat{\beta}_\lambda$  defined by Eqs. 1 with heavy-tailed covariates and general label noise distributions in the proportional high-dimensional regime.
- Further, we provide a similar high-dimensional characterisation for the performance of the optimal Bayesian estimator in this problem. These two results follow from an extension of the replica method for generalised linear estimation to covariate distributions in the family of Eq. 3, and hold for a broader mixture data model which we discuss in Section 4.
- Leveraging the characterisation above, we investigate the impact of heavy-tailed covariate and response contamination on the performance of M-estimators, including, for the first time for this model, analytic control of infinite-variance covariates. In particular, our result highlights the necessity of regularising the Huber loss in the high-dimensional regime to achieve optimal performance, as we show that it can be suboptimal in the presence of heavy-tailed response contamination even when the location parameter is optimally tuned, at odds with the optimality results of Huber (1964) and Hampel (1974) in the classical regime where  $n \gg d$ .
- We show that, despite the strong impact of heavy-tailed contamination in the high-dimensional regime where  $n/d = \Theta(1)$ , the error decay rates  $\|\hat{\beta} - \beta_\star\|_2 = \Theta(n^{-1/2})$  of optimally regularised Huber and least-squares regression are optimal provided the second moment of both the covariates and the noise are bounded. In contrast, we show that when the covariates’ second moment doesn’t exist (but under the assumption of a finite second moment for the noise), the rates explicitly depend on the tail behavior of the covariates’ distribution.

**Related works** — Robust regression is a classical topic in statistics, with several books dedicated to the subject (Huber, 2004; Hampel et al., 2011; Rousseeuw and Leroy, 2005; Maronna et al., 2019). In contrast to the classical regime, literature on robust regression in the high-dimensional regime remains relatively scarce. Early works in this direction provided a characterisation of M-estimators in the proportional limit of  $n, d \rightarrow \infty$  with fixed  $\alpha = n/d$  for Gaussian (El Karoui et al., 2013; El Karoui, 2013; Donoho and Montanari, 2016) and sub-Gaussian (El Karoui, 2013) designs with bounded noise variance. In particular, it was shown that in this regime the maximum likelihood estimator is not necessarily the optimal choice of objective  $\rho$  (El Karoui, 2013). De-biasing and confidence intervals for high-dimensional M-estimation on Gaussian covariates were discussed by Bellec et al. (2022); Bellec (2023). A novel de-biasing framework for covariates with possibly heavy-tailed and asymmetric distributions was recently introduced by Li and Sur (2023). The family of distributions defined in Eq. 3 has previously appeared under different names and literatures, such as *superstatistics* in the context of statistical physics (Beck, 2003), *elliptical distributions* in the context of statistics (Couillet et al., 2015; El Karoui, 2018; Adomaityte et al., 2023) and *Gaussian scale mixture* in the context of signal processing (Wainwright and Simoncelli, 1999). It has been studied in the context of high-dimensional robust covariance estimation by Couillet et al. (2015; 2016). El Karoui (2018) considered (unregularised) robust regression for elliptically distributed covariates under boundedness conditions on the moments. Our work differs in many directions: (a) we relax the assumption of bounded moments and extend the analysis to the more general case of a *mixture* of elliptical distributions; (b) we derive the asymptotic performance of the *Bayes optimal estimator* for this family; (c) we consider a generic convex penalty; (d) we investigate the impact of Huber  $\epsilon$ -contamination. More recently, Vilucchio et al. (2023) studied the model Eq. 1a under a Gaussian design and a double Gaussian noise model for outliers in the proportional high-dimensional regime, showing that estimators of the type in Eq. 1b can fail to be consistent. Beyond (sub-)Gaussian designs, high-dimensional upper bounds on the regressor mean-squared error were obtained under different settings (Hsu and Sabato, 2016; Lugosi and Mendelson, 2019), including heavy-tailed noise (Sun et al., 2020) and heavy-tailed covariate contamination (Sasai, 2022; Pensia et al., 2021). Heavy-tailed covariates have also been studied in the context of kernel ridge regression in (Tomasini et al., 2022).

## 2 HIGH-DIMENSIONAL ASYMPTOTICS

In this section, we discuss our two main theoretical results: the high-dimensional asymptotic characterisation of the M-estimator defined in Eq. 1b and the corresponding Bayes-optimal error. In fact, the results in this section hold under the following slightly more general assumptions.

**Assumption 2.1** (Data). *The covariates  $\mathbf{x}_i \in \mathbb{R}^d$ ,  $i \in [n]$  are independently drawn from the family of “superstatistical” distributions  $\mathbf{x}_i \sim p(\mathbf{x}) := \mathbb{E}[\mathcal{N}(\mathbf{x}; \mathbf{0}, \sigma^2/d\mathbf{I}_d)]$ , where the expectation is over  $\sigma$  with generic distribution density  $\rho$  supported on the positive real line  $\mathbb{R}_+^* := (0, +\infty)$ . For each  $i \in [n]$ , the corresponding response  $y_i \in \mathcal{Y}$  is drawn from a conditional law  $P_0$  on  $\mathcal{Y}$ :*

$$y_i \sim P_0(\cdot | \beta_\star^\top \mathbf{x}_i), \quad (5)$$

with target weights  $\beta_\star \in \mathbb{R}^d$  having finite normalised norm  $\beta_\star^2 := \lim_{d \rightarrow \infty} 1/d \|\beta_\star\|_2^2$ .

Note that the model in Eq. 1a corresponds to the choice  $P_0(y|\tau) = p_\eta(\eta - \tau)$ .

**Assumption 2.2** (Predictor). *We consider the hypothesis class of generalised linear predictors  $\mathcal{H} = \{f_\beta(\mathbf{x}) := f(\beta^\top \mathbf{x}), \beta \in \mathbb{R}^d\}$ , where  $f: \mathbb{R} \rightarrow \mathcal{Y}$  is a generic activation function, and the weights  $\beta \in \mathbb{R}^d$  are obtained by minimising the following empirical risk:*

$$\hat{\beta}_\lambda := \arg \min_{\beta \in \mathbb{R}^d} \sum_{i=1}^n \rho(y_i - \beta^\top \mathbf{x}_i) + \frac{\lambda}{2} \|\beta\|_2^2, \quad \lambda \in \mathbb{R}_+, \quad (6)$$

for a convex objective function  $\rho: \mathbb{R} \rightarrow \mathbb{R}_+$ .

**Assumption 2.3** (Proportional regime). *We consider the proportional high-dimensional regime where both  $n, d \rightarrow \infty$  at a fixed ratio  $\alpha := n/d$ , known as the sample complexity.*

In particular, note that Assumption 2.1 covers the additive noise model in Eq. 1a, and that the standard robust regression setting is given by taking  $f(\beta^\top \mathbf{x}) = \beta^\top \mathbf{x}$ . In Appendix A and Appendix B we derive the following result.

**Result 2.4** (High-dimensional asymptotics). *Let  $\varphi: \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$  denote a test function, and define the following generalisation and training statistics:*

$$\mathcal{E}_g(\hat{\beta}) = \mathbb{E}_{(y, \mathbf{x})} [\varphi(y, \hat{\beta}^\top \mathbf{x})], \quad \mathcal{E}_t(\hat{\beta}) := \frac{1}{n} \sum_{i=1}^n \varphi(y_i, \hat{\beta}^\top \mathbf{x}_i). \quad (7)$$

Then, under Assumptions 2.1–2.3 we have:

$$\mathcal{E}_g(\hat{\beta}) \xrightarrow[n, d \rightarrow \infty]{\text{P}} \varepsilon_g(\alpha, \lambda, \beta_\star^2), \quad \mathcal{E}_t(\hat{\beta}) \xrightarrow[n, d \rightarrow \infty]{\text{P}} \varepsilon_t(\alpha, \lambda, \beta_\star^2), \quad (8)$$

explicitly given by:

$$\begin{aligned} \varepsilon_t(\alpha, \lambda, \beta_\star^2) &= \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma, \zeta} \left[ Z_0 \left( y, \frac{m\sigma}{\sqrt{q}} \zeta, \sigma^2 \beta_\star^2 - \frac{\sigma^2 m^2}{q} \right) \varphi \left( y, \sigma \sqrt{q} \zeta + v \sigma^2 f \right) \right], \\ \varepsilon_g(\alpha, \lambda, \beta_\star^2) &= \int_{\mathcal{Y}} dy \int d\eta \int d\tau P_0(y|\tau) \mathbb{E}_\sigma \left[ \mathcal{N} \left( \left( \frac{\tau}{\eta} \right); \mathbf{0}, \sigma^2 \begin{pmatrix} \beta_\star^2 & m \\ m & q \end{pmatrix} \right) \right] \varphi(y, \eta). \end{aligned} \quad (9)$$

where  $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ ,  $\zeta \sim \mathcal{N}(0, 1)$  are independent Gaussian variables. Similarly, the mean-squared error on the estimator is given by:

$$\varepsilon_{\text{est}}(\alpha, \lambda, \beta_\star^2) := \lim_{n, d \rightarrow +\infty} \frac{1}{d} \mathbb{E}_{\mathcal{D}} [\|\hat{\beta}_\lambda - \beta_\star\|_2^2] = \beta_\star^2 - 2m + q \xrightarrow{\alpha \rightarrow +\infty} \beta_\star^{-2} \lim_{\alpha \rightarrow +\infty} (m - \beta_\star^2)^2. \quad (10)$$

Moreover, the angle between the estimator  $\hat{\beta}_\lambda$  and  $\beta_\star$  is given by

$$\lim_{n, d \rightarrow +\infty} \mathbb{E}_{\mathcal{D}} [\text{angle}(\hat{\beta}_\lambda, \beta_\star)] = \frac{1}{\pi} \arccos \left( \frac{m}{\beta_\star \sqrt{q}} \right) \xrightarrow{\alpha \rightarrow +\infty} 0. \quad (11)$$

In the expressions above we have introduced the order parameters  $v$ ,  $q$ , and  $m$ . These quantities are found by solving the following set of self-consistent fixed-point equations:

$$\begin{aligned} q &= \frac{\hat{m}^2 \beta_\star^2 + \hat{q}}{(\lambda + \hat{v})^2} & \hat{q} &= \alpha \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0 \left( y, \frac{m\sigma}{\sqrt{q}} \zeta, \sigma^2 \beta_\star^2 - \frac{\sigma^2 m^2}{q} \right) f^2 \right], \\ m &= \frac{\beta_\star^2 \hat{m}}{\lambda + \hat{v}} & \hat{v} &= -\alpha \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0 \left( y, \frac{m\sigma}{\sqrt{q}} \zeta, \sigma^2 \beta_\star^2 - \frac{\sigma^2 m^2}{q} \right) \partial_\omega f \right], \\ v &= \frac{1}{\lambda + \hat{v}} & \hat{m} &= \alpha \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 \partial_\mu Z_0 \left( y, \frac{m\sigma}{\sqrt{q}} \zeta, \sigma^2 \beta_\star^2 - \frac{\sigma^2 m^2}{q} \right) f \right], \end{aligned} \quad (12a)$$

where the auxiliary function  $Z_0$  and proximal operator  $f$  are defined by:

$$\begin{aligned} Z_0(y, \mu, v) &:= \mathbb{E}_z [P_0(y|\mu + \sqrt{v}z)], \\ f &:= \arg \min_u \left[ \frac{\sigma^2 v u^2}{2} + \rho(y - \omega - \sigma^2 v u) \right] \in \mathbb{R}, \quad \omega = \sigma \sqrt{q} \zeta. \end{aligned} \quad (12b)$$

Result 2.4 reduces the high-dimensional optimisation problem defined by Eq. 2.2 to a self-consistent equation on the parameters  $\hat{q}, q, \hat{m}, m, \hat{v}, v \in \mathbb{R}$ . Despite being cumbersome, this low-dimensional problem can be efficiently solved numerically. Note that our result for  $\varepsilon_{\text{est}}$  implies that the estimator  $\hat{\beta}_\lambda$  is consistent if and only if  $\lim_{\alpha \rightarrow +\infty} m = \beta_\star^2$ . An interesting particular case, in which the convergence rates to zero of  $\varepsilon_{\text{est}}$  can be derived explicitly, is the case of square loss,  $\rho(t) = \frac{1}{2}t^2$ . In Appendix B.1 it is shown the following.

**Result 2.5** (Square loss rates). *Consider the linear model in Eq. 1a with covariates as in Eq. 3. Assume that  $\mathbb{E}[\eta_i] = 0$ ,  $\hat{\sigma}_0^2 := \mathbb{E}[\eta_i^2] \in \mathbb{R}_+^*$ . Let us also assume that in Eq. 3 we have  $\varrho(\sigma) \sim \frac{1}{\sigma^{2a+1}}$  for  $\sigma \gg 1$  with  $a > 0$ . Then, the ridge estimator  $\hat{\beta}_\lambda$  minimising Eq. 1b with  $\rho(t) = \frac{1}{2}t^2$  is consistent and*

$$\varepsilon_{\text{est}} \underset{\alpha \gg 1}{=} \begin{cases} \frac{\hat{\sigma}_0^2}{\sigma_0^2 \alpha} + o\left(\frac{1}{\alpha}\right) & \text{if } a > 1, \quad \text{with } \sigma_0^2 := \mathbb{E}[\sigma^2] = \lim_{x \rightarrow +\infty} x(1 - x S_{\sigma^2}(x)), \\ \frac{\hat{\sigma}_0^2}{\tilde{\sigma}_0^2 \alpha \ln \alpha} + o\left(\frac{1}{\alpha \ln \alpha}\right) & \text{if } a = 1, \quad \text{with } \tilde{\sigma}_0^2 = \lim_{x \rightarrow +\infty} \frac{x}{\ln x} (1 - x S_{\sigma^2}(x)), \\ \frac{\hat{\sigma}_0^2}{(\tilde{\sigma}_0^2 \alpha)^{1/a}} + o\left(\frac{1}{\alpha^{1/a}}\right) & \text{if } a \in (0, 1), \quad \text{with } \tilde{\sigma}_0^2 = \lim_{x \rightarrow +\infty} x^a (1 - x S_{\sigma^2}(x)), \end{cases} \quad (13)$$

where  $S_{\sigma^2}(x) := \mathbb{E}\left[\frac{1}{x + \sigma^2}\right]$  is the Stieltjes transform of the measure of  $\sigma^2$ . Finally, the asymptotic test and training error will depend on the noise distribution through its second moment only.

By definition, the mean-squared error  $\varepsilon_{\text{est}}$  of the M-estimator in Eq. 10 is lower-bounded by the minimum mean-squared error achieved by the Bayes-optimal estimator  $\hat{\beta}$ , which is given by the expectation over the Bayesian posterior  $\hat{\beta} = \mathbb{E}[\beta|\mathcal{D}]$ . Our second main contribution in this manuscript, derived in Appendix A.2, is to provide a sharp asymptotic formula for the minimum mean-squared error under the assumption that  $\beta_\star \sim \mathcal{N}(\mathbf{x}; \mathbf{0}, \beta_\star^2 \mathbf{I}_d)$ .

**Result 2.6** (Bayes optimal error). *Consider the minimum mean-squared error for the data model defined by 2.1 with  $\beta_\star \sim \mathcal{N}(\mathbf{x}; \mathbf{0}, \beta_\star^2 \mathbf{I}_d)$ :*

$$\mathcal{E}_{\text{BO}} := \min_{\beta \in \mathbb{R}^d} \frac{1}{d} \mathbb{E}[\|\beta_\star - \beta\|_2^2] = \frac{1}{d} \mathbb{E}[\|\beta_\star - \hat{\beta}\|_2^2]. \quad (14)$$

Then, in the high-dimensional limit defined by Assumption 2.3:

$$\mathcal{E}_{\text{BO}} \xrightarrow[n, d \rightarrow \infty]{\text{P}} \varepsilon_{\text{BO}} = \beta_\star^2 - \mathbf{q}, \quad (15)$$

where  $\mathbf{q}$  is given by the solution of the following self-consistent equations:

$$\hat{\mathbf{q}} = \alpha \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0(y, \mu, V) (\partial_\mu \ln Z_0(y, \mu, V))^2 \Big|_{\substack{\mu = \sigma \sqrt{q} \zeta \\ V = \sigma^2 (\beta_\star^2 - \mathbf{q})}} \right], \quad \mathbf{q} = \frac{\beta_\star^4 \hat{\mathbf{q}}}{1 + \beta_\star^2 \hat{\mathbf{q}}} \quad (16)$$

with  $\zeta \sim \mathcal{N}(0, 1)$  and  $Z_0$  given in Eq. 12b.

Similarly to Result 2.4, Result 2.6 reduces the problem of estimating the Bayes optimal error from the computationally intractable evaluation of the high-dimensional posterior marginals to a simple two-dimensional set of self-consistent equations on the variables  $\hat{\mathbf{q}}, \mathbf{q} \in \mathbb{R}$ .

A detailed derivation of both Results 2.4 & 2.6 is provided in Appendix A using the replica method. Despite being non-rigorous, Results 2.4 & 2.6 provide a natural extension of rigorous results in the established literature of high-dimensional asymptotics for generalised linear estimation on Gaussian covariates to the elliptical family defined in Assumption 2.1 (El Karoui, 2013; Donoho and Montanari, 2016; Thrampoulidis et al., 2018; Barbier et al., 2019; Loureiro et al., 2022). Indeed, by taking  $\varrho(\sigma) = \delta_{\sigma, \bar{\sigma}}$  for some  $\bar{\sigma} \in \mathbb{R}_+^*$ , Result 2.4 reduces to the rigorous formulas proven by Thrampoulidis et al. (2018), and Result 2.6 reduces to the rigorous formulas proven by Barbier et al. (2019). Nevertheless, as we will see in the next section, they produce an excellent agreement with moderately finite-size simulations, across different choices of penalty  $\rho$  and covariate distribution, including cases in which the variance is infinite, see Fig. 2 for instance.

### 3 DISCUSSION

In this section, we investigate the consequences of Results 2.4 – 2.6 in the context of robust regression with heavy-tail contamination as in Eq. 4 within the model in Eq. 1a. Our theoretical results are given for *any* noise distribution  $p_\eta$ , but, to exemplify our findings, we will focus on the case of heavy-tailed contamination of standard Gaussian noise, in a form similar to Eq. 4:

$$p_\eta(\eta) = \mathbb{E}[\mathcal{N}(\eta; 0, \sigma^2)], \quad \text{with } \sigma \sim \hat{\varrho} = (1 - \epsilon_n)\delta_{\sigma,1} + \epsilon_n\hat{\varrho}_0, \quad \epsilon_n \in [0, 1]. \quad (17)$$

As in the case of covariates,  $\epsilon_n$  measures the level of contamination of the purely Gaussian noise. For concreteness, in most of the plots, we take  $\beta_\star \sim \mathcal{N}(0, \beta_\star^2 \mathbf{I}_d)$  and use for  $\varrho_0$  and  $\hat{\varrho}_0$  one of the distributions in Table 1, which provide convenient parametric families of distributions where the existence of moments is easily tunable.

#### 3.1 LABEL CONTAMINATION

We start by discussing the impact of *response contamination* in high-dimensional robust regression. Since the Gaussian assumption for the noise label is widespread in statistics, we will focus our discussion on an  $\epsilon$ -contamination of the standard normal distribution using, for  $p_\eta$ , an expression as in Eq. 17, by adopting a proper  $\hat{\varrho}_0$  that generates fat tails. We recall the reader that when the loss function is the negative log-likelihood of the noise  $-\log p_\eta$ , strong guarantees hold for the M-estimator in the classical limit  $n \gg d$ . In particular, the estimator is unbiased and asymptotically normal, with estimation error (defining the variance)  $\|\beta_\star - \hat{\beta}_\lambda\|_2^2 = \Theta(n^{-1})$  (van der Vaart, 2000). Since the square loss corresponds to the negative log-likelihood of standard Gaussian noise, in the absence of contamination  $\epsilon_n = 0$  the square loss attains an optimal rate as  $n \rightarrow \infty$ . We discuss the insights provided by our exact high-dimensional characterisation in high dimensions, i.e., when  $n = \Theta(d)$ , assuming contaminated labels as in Eq. 17. On the other hand, we assume for now purely Gaussian covariates, i.e.,  $\epsilon_c = 0$  in Eq. 4.




Figure 1: Value of  $\epsilon_{\text{est}}$  obtained using a regularised Huber at given  $\lambda = 10^{-3}$  as a function of  $\delta$  for different values of  $\alpha$ . Here the contamination level is  $\epsilon_n = 0.5$ .

**The importance of regularisation** — Fig. 2 illustrates two distinct scenarios, i.e., contamination of the labels with an infinite-variance noise (top), so that  $\mathbb{E}[\eta^2] = +\infty$ , and contamination with a finite-variance noise (bottom), so that  $\mathbb{E}[\eta^2] = 1$ . While in the former case, the performance degrades as a function of the contamination, counter-intuitively, it improves when the contamination variance is bounded. As expected, in the absence of contamination  $\epsilon_n = 0$ , optimally regularised ridge regression achieves the minimum mean-squared error for this model. Indeed, in this case, the square loss is not only the maximum likelihood estimator but also coincides with the Bayes-optimal estimator, since from a Bayesian perspective the  $\ell_2$  penalty corresponds to the optimal Bayesian prior when optimally tuned. As contamination is introduced ( $\epsilon_n > 0$ ), this is no longer the case, and indeed, in the  $\mathbb{E}[\eta^2] < +\infty$  scenario, optimally regularised ridge regression is observed to be suboptimal. More surprising, perhaps, is that, at small values of the regularisation  $\lambda$ , the Huber loss in Eq. 2 with optimally chosen location parameter  $\delta^*$  is also suboptimal in the high-dimensional regime  $n = \Theta(d)$ . This is to be contrasted with the well-known optimality results in the classical regime  $n \gg d$ , which are recovered at  $\alpha \rightarrow \infty$ . Interestingly, at small but finite  $\lambda$ , the improvement of performance for the Huber loss with  $\alpha$  coincides with a sharp transition of the optimal location parameter  $\delta^*$  from  $O(\lambda)$  to  $O(1)$  at a given value of  $\alpha$ : at fixed  $(\lambda, \alpha)$ , the value of  $\epsilon_{\text{est}}$  develops indeed two minima as a function of  $\delta$ , one in  $\delta = O(\lambda)$  and another in  $\delta = O(1)$ , whose relative depth changes with  $\alpha$ , exhibiting an inversion in favour of the  $O(1)$  minimum at large  $\alpha$ , see Fig. 1. For  $\lambda \rightarrow 0$ , the minimum in  $\delta = O(\lambda)$  disappears. This phenomenology is observed for various choices of  $\hat{\varrho}_0$ , but






Figure 2: Response contamination of Gaussian covariates as a function of the sample complexity  $\alpha = n/d$  for different noise contamination levels  $\epsilon_n \in [0, 1]$  as in Eq. 17. Here we adopt  $\hat{\varrho}_0(\sigma) \propto \sigma^{-2a-1} \exp(-b\sigma^{-2})$  for the noise contamination, see Table 1. The theoretical predictions are compared with the results of numerical experiments (dots) obtained averaging over 20 instances with  $d = 10^3$ . **(Left)** Purely Gaussian noise ( $\epsilon_n = 0$ ). **(Top)**. Case  $b = 1$  and  $a = 4/5 < 1$ , implying  $\mathbb{E}[\eta^2] = +\infty$ . The performance degrades as the contamination level  $\epsilon_n$  is increased. Optimally regularised Huber (red) achieves near-optimal Bayesian performance (solid), while by fixing the regularisation, the Huber has suboptimal performances (purple). **(Bottom)**. Case  $a = 1 + b = 1 + 1/10$ , corresponding to  $\mathbb{E}[\eta^2] = 1$ . The performance uniformly improves as the contamination  $\epsilon_n$  grows. Optimally regularised Huber (red) achieves the optimal Bayesian performance (solid), while both Huber with untuned regularisation (purple) and optimally regularised ridge (blue) are suboptimal.

it also persists in other settings, as we mention later in the text. In general, adding a  $\ell_2$  regularisation and cross-validating significantly improve the performance of the Huber loss, bringing it very close to Bayes optimality at all contamination levels. Our results highlight the necessity of properly regularising robust estimators in order to achieve optimality in the high-dimensional regime. In particular, optimally regularising Huber is crucial to achieve near-Bayes-optimal performances in the case of infinite-variance label noise.




Figure 3: Estimation error  $\epsilon_{\text{est}}$  as in Eq. 10 at large- $\alpha$  using regularised optimal Huber loss (solid lines). The covariates are Gaussian, whereas the label noise is obtained as in Eq. 17 with  $\varrho(\sigma) \sim \sigma^{-2a-1}$ ,  $a > 0$ . The results are compared with the Bayes-optimal performance (squares). The dotted line shows a scaling of  $\alpha^{-1}$ .

**Convergence rates** — In Fig. 3 we consider the case of full contamination  $\epsilon_n = 1$  of the labels with both finite- and infinite-variance noise in the large- $\alpha$  regime: this corresponds to the classical limit  $n \gg d$ . In particular, we focus on the regularised Huber loss, whose location parameter and regularisation strength have been optimised to reach the minimum test error. A fatter tail in the noise hurts performance in terms of the estimation error  $\epsilon_{\text{est}}$ , but does not affect its  $\alpha^{-1}$  scaling, which remains universal even when  $\mathbb{E}[\eta^2] = +\infty$  and is also observed in the Bayes optimal lower bound. We will see below that this will not be the case if a contamination of the covariates is considered.

### 3.2 COVARIATE CONTAMINATION IN THE PRESENCE OF HEAVY-TAILED LABEL NOISE

We now move to discuss the impact of *covariates' contamination*, focusing on the case in which covariates are generated with a density as in Eq. 4 with  $\epsilon_c > 0$ . On top of this, we assume a non-Gaussian label noise with finite variance,  $\mathbb{E}[\eta^2] = 1$ : note that, as discussed above, in this setting regularised ridge regression is not optimal.

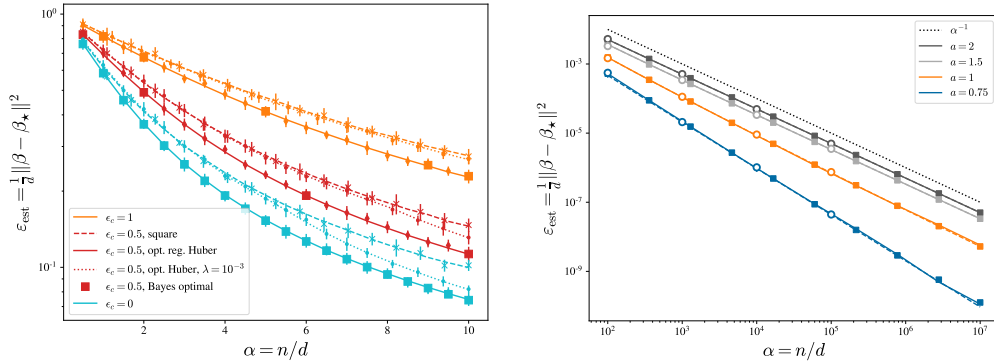


Figure 4: **(Left)** Estimation error as a function of sample complexity  $\alpha$ . Dots correspond to the average error of 50 numerical experiments in dimension  $d = 10^3$ . Covariates are contaminated as in Eq 4 using for  $\varrho$  an inverse-Gamma with  $b = a - 1 = 1/10$ . **(Right)** Estimation error  $\epsilon_{\text{est}}$  as in Eq. 10 at large- $\alpha$  obtained from our theory for the square loss (dashed line), regularised optimal Huber loss (solid line) and Bayes-optimal performance given by Result 2.6 (squares). The covariates’ noise is Pareto-distributed to have  $\varrho(\sigma) \sim \sigma^{-2a-1}$  with  $a > 0$ , and the label noise is Gaussian. White dots correspond to numerical experiments in dimension  $d = 50$ . The black dotted line shows a scaling of  $\alpha^{-1}$ .

Fig. 4 (left) shows that as the contamination level of the covariates  $\epsilon_c$  grows, the performance is negatively affected for all metrics. The phenomenology is similar to the one observed for response contamination. While optimally regularised square loss performs the worst amongst the considered losses, an optimally  $\delta$ -tuned Huber loss, ridge-regularised with small, fixed strength  $\lambda$ , performs halfway between the square and the Bayes-optimal bound, approaching the latter as  $\alpha$  grows. As observed in the previous section, the improvement in the  $\delta$ -tuned Huber loss occurs, for small but non-zero  $\lambda$ , with a sharp jump in the optimal  $\delta^*$  parameter from a value  $O(\lambda)$  to a value  $O(1)$  as  $\alpha$  increases, reflected in a kink in the curve of  $\epsilon_{\text{est}}$  as function of  $\alpha$ . Finally, an optimally  $(\delta, \lambda)$ -tuned regularised Huber loss achieves Bayes-optimal performance. The phenomenology is therefore analogous to the one described in section 3.1.

**Convergence rates** — As we show in Result 2.5, if the statistician adopts the square loss (hence choosing  $\rho(t) = t^2/2$ ) in the study of power-law tailed covariates, the estimation error rate might depend on the tail, under the assumption that the second moment of the noise is finite. In particular, Result 2.4 implies that, in this case, the estimation error rate depends only on the covariates’ second moment. As a consequence the rate  $\epsilon_{\text{est}} \sim n^{-1}$  is universal as long as the second moment of the covariates exists. Curiously, if the second moment does not exist, the estimation error decays *faster*, depending explicitly on the tail exponent of the distribution. Remarkably, this suggests that the presence of very fat tails in the covariates can actually improve the classical convergence rates of M-estimation when  $n \rightarrow \infty$ . Indeed, despite not having a closed form expression for the rates under different other penalties, we have numerically verified that the convergence rates reported in Result 2.4 are Bayes-optimal and are also attained optimally tuned and regularised Huber loss in different scenarios. This is illustrated in Fig. 4 (right), which clearly shows the dependence on the tail exponent of the covariate distribution when the latter has no finite second moment.

## 4 GENERALISED LINEAR ESTIMATION ON AN ELLIPTICAL MIXTURE

We conclude our work by stating the most general form of our exact high-dimensional asymptotic result, valid for an arbitrary mixture of  $K$  elliptical distributions, that we call *elliptical mixture model* (EMM). To be more precise, let us consider the following general Assumptions:



**Assumption 4.1** (EMM Data). *In the dataset  $\mathcal{D} := \{(y_i, \mathbf{x}_i)\}_{i \in [n]}$ , the covariates  $\mathbf{x}_i \in \mathbb{R}^d$ ,  $i \in [n]$  are independently drawn from the mixture  $\mathbf{x}_i \sim p(\mathbf{x}) := \sum_{c=1}^K p_c \mathbb{E}_{\sigma_c}[\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_c, \sigma_c^2/d\mathbf{I}_d)]$ , where, for each  $c \in [K]$ ,  $d\|\boldsymbol{\mu}_c\|_2^2 = \Theta(1/d)$ ,  $p_c \in [0, 1]$  and the expectation is over  $\sigma_c \sim \varrho_c$ , generic distribution density supported on  $\mathbb{R}_+^*$ . Moreover,  $\sum_{c=1}^K p_c = 1$ . For each  $\mathbf{x}_i$ ,  $i \in [n]$ , the corresponding response  $y_i \in \mathcal{Y}$  is drawn from a conditional law  $P_0(y|\boldsymbol{\beta}_*^\top \mathbf{x}_i)$  on  $\mathcal{Y}$ , with target weights  $\boldsymbol{\beta}_* \in \mathbb{R}^d$  having finite normalised norm  $\beta_*^2 := \lim_{d \rightarrow \infty} 1/d \|\boldsymbol{\beta}_*\|_2^2$ .*

**Assumption 4.2** (Predictor). *We consider the hypothesis class of generalised linear predictors  $\mathcal{H} = \{f_\beta(\mathbf{x}) := f(\boldsymbol{\beta}^\top \mathbf{x}), \boldsymbol{\beta} \in \mathbb{R}^d\}$ , where  $f: \mathbb{R} \rightarrow \mathcal{Y}$  is a generic activation function, and the weights  $\boldsymbol{\beta} \in \mathbb{R}^d$  are obtained by minimising the following empirical risk:*

$$\hat{\boldsymbol{\beta}}_\lambda := \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^d} \sum_{i=1}^n \ell(y_i, \boldsymbol{\beta}^\top \mathbf{x}_i) + \lambda r(\boldsymbol{\beta}), \quad \lambda \in \mathbb{R}_+, \quad (18)$$

for a convex loss function  $\ell: \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}_+$  and a convex regularisation function  $r: \mathbb{R} \rightarrow \mathbb{R}_+$ .

Under the further Assumption 2.3, in Appendix A it is shown that the following holds.

**Result 4.3** (High-dimensional asymptotics for the EMM). *Let us consider the independent random variables  $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ , and  $z, \zeta \sim \mathcal{N}(0, 1)$ . In the proportional asymptotic regime, the training and test error defined in Eq. 7 and associated with the  $M$ -estimator of the problem in Eq. 18 have a limit in probability as in Eq. 8, with*

$$\begin{aligned} \varepsilon_t(\alpha, \lambda, \beta_*^2) &= \sum_c p_c \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma_c, \zeta} \left[ Z_0 \left( y, \boldsymbol{\mu}_c^\top \boldsymbol{\beta}_* + \frac{\sigma_c m}{\sqrt{q}} \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) \varphi(y, \sigma_c \sqrt{q} \zeta + v \sigma_c^2 f_c) \right], \\ \varepsilon_g(\alpha, \lambda, \beta_*^2) &= \sum_c p_c \int_{\mathcal{Y}} dy \int d\eta \int d\tau P_0(y|\tau) \mathbb{E}_{\sigma_c} \left[ \mathcal{N} \left( (\tau); \left( \boldsymbol{\mu}_c^\top \boldsymbol{\beta}_* \right), \sigma_c^2 \left( \frac{\beta_*^2 m}{q} \right) \right) \right] \varphi(y, \eta). \end{aligned} \quad (19)$$

In particular, the estimation error is given by

$$\varepsilon_{\text{est}} := \lim_{d \rightarrow +\infty} \frac{1}{d} \mathbb{E}_{\mathcal{D}} [\|\hat{\boldsymbol{\beta}}_\lambda - \boldsymbol{\beta}_*\|_2^2] = \beta_*^2 - 2m + q \quad (20)$$

The required set of order parameters and proximals can be found by self-consistently solving the equations below

$$\begin{aligned} v &= \frac{1}{d\sqrt{q}} \mathbb{E}_{\boldsymbol{\xi}}[\mathbf{g}^\top \boldsymbol{\xi}], & \hat{q}_c &= \alpha p_c \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma_c, \zeta} \left[ \sigma_c^2 Z_0 \left( y, \boldsymbol{\mu}_c^\top \boldsymbol{\beta}_* + \frac{\sigma_c m}{\sqrt{q}} \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) f_c^2 \right], \\ q &= \frac{1}{d} \mathbb{E}_{\boldsymbol{\xi}}[\|\mathbf{g}\|_2^2], & \hat{v} &= -\alpha \sum_c p_c \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma_c, \zeta} \left[ \sigma_c^2 Z_0 \left( y, \boldsymbol{\mu}_c^\top \boldsymbol{\beta}_* + \frac{\sigma_c m}{\sqrt{q}} \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) \partial_\omega f_c \right], \\ m &= \frac{1}{d} \mathbb{E}_{\boldsymbol{\xi}}[\mathbf{g}^\top \boldsymbol{\beta}_*], & \hat{m}_c &= \alpha p_c \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma_c, \zeta} \left[ \sigma_c^2 \partial_\mu Z_0 \left( y, \boldsymbol{\mu}_c^\top \boldsymbol{\beta}_* + \frac{\sigma_c m}{\sqrt{q}} \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) f_c \right], \\ t_c &= \mathbb{E}_{\boldsymbol{\xi}}[\mathbf{g}^\top \boldsymbol{\mu}_c], & \hat{t}_c &= \alpha p_c \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma_c, \zeta} \left[ Z_0 \left( y, \boldsymbol{\mu}_c^\top \boldsymbol{\beta}_* + \frac{\sigma_c m}{\sqrt{q}} \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) f_c \right], \end{aligned} \quad (21a)$$

where, as before,  $Z_0(y, \mu, V) := \mathbb{E}_z[P_0(y|\mu + \sqrt{V}z)]$  and

$$\begin{aligned} \mathbf{g} &:= \arg \min_{\boldsymbol{\beta}} \left( \frac{\hat{v}\|\boldsymbol{\beta}\|_2^2}{2} - \sum_c \boldsymbol{\beta}^\top (\hat{m}_c \boldsymbol{\beta}_* + \hat{t}_c \boldsymbol{\mu}_c) - \sum_c \sqrt{\hat{q}_c} \boldsymbol{\xi}^\top \boldsymbol{\beta} + \lambda r(\boldsymbol{\beta}) \right) \in \mathbb{R}^d, \\ f_c &:= \arg \min_u \left[ \frac{vu^2 \sigma_c^2}{2} + \ell(y, t_c + \sigma_c \sqrt{q} \zeta + \sigma_c^2 vu) \right] \in \mathbb{R}. \end{aligned} \quad (21b)$$

The previous set of equations covers a wide range of distributions and possibly any power-law tail decay. The generality of this setting relies on the fact that any distribution can be approximated by a possibly uncountable superposition of Gaussians (Nestoridis et al., 2011; Alspach and Sorenson, 1972; Ghosh and Ramamoorthi, 2006). Moreover, in Appendices B & B.1 the following universality result, generalising results of Gerace et al. (2022) and Pesce et al. (2023) in the Gaussian setting, is given.

**Result 4.4** (Universality for uncorrelated target). *Assume that, for all  $c \in [K]$ ,  $\lim_{d \rightarrow +\infty} \boldsymbol{\mu}_c^\top \boldsymbol{\beta}_* = 0$  and that the labels are generated according to the model in Eq. 1a, the noise  $\eta$  having even distribution and zero mean. Then, if the loss  $\ell$  is even, the asymptotic training and generalisation errors are the same as if the covariates had distribution  $p(\mathbf{x}) = \mathbb{E}[\mathcal{N}(\mathbf{x}; \mathbf{0}, \sigma/d\mathbf{I}_d)]$  with  $\sigma \sim \sum_{c=1}^K p_c \varrho_c$ . Furthermore, if  $\ell(y, t) = \frac{1}{2}(y - t)^2$  and  $r(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ , under the assumption that  $\mathbb{E}[\eta^2] < +\infty$ , then the training loss is*

$$\lim_{\lambda \rightarrow 0^+} \frac{1}{2n} \sum_{i=1}^n (y_i - \boldsymbol{\beta}_\lambda^\top \mathbf{x}_i)^2 \xrightarrow[\alpha = n/d = \Theta(1)]{d \rightarrow +\infty} \frac{\mathbb{E}[\eta^2]}{2} \left(1 - \frac{1}{\alpha}\right)_+, \quad (x)_+ := x\theta(x), \quad (22)$$

for any distribution of the random variables  $\sigma_c$ .

## REFERENCES

- U. Adomaityte, G. Sicuro, and P. Vivo. Classification of superstatistical features in high dimensions. *arXiv:2304.02912*, 2023.
- D. Alspach and H. Sorenson. Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Trans. Autom. Control*, 17(4):439–448, 1972.
- Y. Bai, S. Mei, H. Wang, and C. Xiong. Don’t just blame over-parametrization for over-confidence: Theoretical analysis of calibration in binary classification. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *PMRL*, pages 566–576. PMLR, 2021.
- J. Barbier, F. Krzakala, N. Macris, L. Miolane, and L. Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proc. Natl. Acad. Sci. U.S.A.*, 116(12):5451–5460, 2019.
- C. Beck. Superstatistics: Theory and applications. *Contin. Mech. Thermodyn.*, 16, 2003.
- P.C. Bellec. Out-of-sample error estimate for robust m-estimators with convex penalty. *arXiv:2008.11840*, 2023.
- P.C. Bellec, Y. Shen, and C.-H. Zhang. Asymptotic normality of robust M-estimators with convex penalty. *Electron. J. Stat.*, 16(2):5591 – 5622, 2022. doi: 10.1214/22-EJS2065.
- L. Clarté, B. Loureiro, F. Krzakala, and L. Zdeborová. Theoretical characterization of uncertainty in high-dimensional linear classification. *Machine Learning: Science and Technology*, 4(2):025029, 2023a.
- L. Clarté, B. Loureiro, F. Krzakala, and L. Zdeborová. On double-descent in uncertainty quantification in overparametrized models. In F. Ruiz, J. Dy, and J.-W. van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *PMRL*, pages 7089–7125. PMLR, 2023b.
- R. Couillet, F. Pascal, and J.W. Silverstein. The random matrix regime of maronna’s m-estimator with elliptically distributed samples. *J. Multivar. Anal.*, 139:56–78, 2015.
- R. Couillet, A. Kammoun, and F. Pascal. Second order statistics of robust estimators of scatter. application to glrt detection for elliptical signals. *J. Multivar. Anal.*, 143:249–274, 2016. ISSN 0047-259X. doi: <https://doi.org/10.1016/j.jmva.2015.08.021>.
- D. Donoho and A. Montanari. High dimensional robust  $m$ -estimation: asymptotic variance via approximate message passing. *Probab. Theory Relat. Fields.*, 166(3):935–969, 2016.
- N. El Karoui. Asymptotic behavior of unregularized and ridge-regularized high-dimensional robust regression estimators : rigorous results, 2013.
- N. El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probab. Theory Relat. Fields*, 170, 2018.
- N. El Karoui, D. Bean, P.J. Bickel, C. Lim, and B. Yu. On robust regression with high-dimensional predictors. *Proc. Natl. Acad. Sci. U.S.A.*, 110(36):14557–14562, 2013.
- F. Gerace, F. Krzakala, B. Loureiro, L. Stephan, and L. Zdeborová. Gaussian universality of perceptrons with random labels. *arXiv:2205.13303*, 2022.
- J.K. Ghosh and R.V. Ramamoorthi. *Bayesian Nonparametrics*. Springer Series in Statistics. Springer New York, 2006.
- F.R. Hampel. The influence curve and its role in robust estimation. *J. Am. Stat. Assoc.*, 69(346):383–393, 1974.

- F.R. Hampel, E.M. Ronchetti, P.J. Rousseeuw, and W.A. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. Wiley Series in Probability and Statistics. Wiley, 2011.
- D. Hsu and S. Sabato. Loss minimization and parameter estimation with heavy tails. *J. Mach. Learn. Res.*, 17(18):1–40, 2016.
- P.J. Huber. Robust Estimation of a Location Parameter. *Ann. Math. Stat.*, 35(1):73 – 101, 1964.
- P.J. Huber. Robust Regression: Asymptotics, Conjectures and Monte Carlo. *Ann. Stat.*, 1(5):799 – 821, 1973.
- P.J. Huber. *Robust Statistics*. Wiley Series in Probability and Statistics - Applied Probability and Statistics Section Series. Wiley, 2004.
- A. Javanmard and A. Montanari. Debiasing the lasso: Optimal sample size for Gaussian designs. *Ann. Stat.*, 46(6A):2593 – 2622, 2018.
- Yufan Li and Pragya Sur. Spectrum-aware adjustment: A new debiasing framework with applications to principal component regression, 2023.
- B. Loureiro, G. Sicuro, C. Gerbelot, A. Pacco, F. Krzakala, and L. Zdeborová. Learning Gaussian Mixtures with Generalized Linear Models: Precise Asymptotics in High-dimensions. *Advances in Neural Information Processing Systems*, 34:10144–10157, 2021.
- B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mézard, and L. Zdeborová. Learning curves of generic features maps for realistic datasets with a teacher-student model. *J. Stat. Mech.: Theory Exp.*, 2022(11):114001, 2022.
- G. Lugosi and S. Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Found. Comput. Math.*, 19(5):1145–1190, 2019.
- R.A. Maronna, R.D. Martin, V.J. Yohai, and M. Salibián-Barrera. *Robust Statistics: Theory and Methods (with R)*. Wiley Series in Probability and Statistics. Wiley, 2019.
- V. Nestoridis, S. Schmutzhard, and V. Stefanopoulos. Universal series induced by approximate identities and some relevant applications. *J. Approx. Theory*, 163(12):1783–1797, 2011.
- A. Pensia, V. Jog, and P.-L. Loh. Robust regression with covariate filtering: Heavy tails and adversarial contamination. *arXiv:2009.12976*, 2021.
- L. Pesce, F. Krzakala, B. Loureiro, and L. Stephan. Are Gaussian data all you need? The extents and limits of universality in high-dimensional generalized linear estimation. In A. Krause et al., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 27680–27708. PMLR, 23–29 Jul 2023.
- P.J. Rousseeuw and A.M. Leroy. *Robust Regression and Outlier Detection*. Wiley Series in Probability and Statistics. Wiley, 2005.
- T. Sasai. Robust and sparse estimation of linear regression coefficients with heavy-tailed noises and covariates, 2022.
- Q. Sun, W.-X. Zhou, and J. Fan. Adaptive huber regression. *J. Am. Stat. Assoc.*, 115(529): 254–265, 2020.
- P. Sur and E.J. Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. 116(29):14516–14525, 2019.
- P. Sur and E.J. Candès. The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression. *Ann. Stat.*, 48(1):27 – 42, 2020. doi: 10.1214/18-AOS1789.

- C. Thrampoulidis, E. Abbasi, and B. Hassibi. Precise error analysis of regularized  $m$ -estimators in high dimensions. *IEEE Trans. Inf. Theory*, 64(8):5592–5628, 2018.
- U.M Tomasini, A. Scocchi, and M. Wyart. Failure and success of the spectral bias prediction for Laplace kernel ridge regression: the case of low-dimensional data. In K. Chaudhuri et al., editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 21548–21583. PMLR, 17–23 Jul 2022.
- A.W. van der Vaart. *Asymptotic Statistics*. Asymptotic Statistics. Cambridge University Press, 2000.
- M. Vilucchio, E. Troiani, V. Erba, and F. Krzakala. Asymptotic characterisation of robust empirical risk minimisation performance in the presence of outliers. *arXiv:2305.18974*, 2023.
- M.J. Wainwright and E. Simoncelli. Scale mixtures of gaussians and the statistics of natural images. In S. Solla, T. Leen, and K. Müller, editors, *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.
- Q. Zhao, P. Sur, and E.J. Candès. The asymptotic distribution of the MLE in high-dimensional logistic models: Arbitrary covariance. *Bernoulli*, 28(3):1835 – 1861, 2022.

# SUPPLEMENTARY INFORMATION

## A REPLICIA DERIVATION OF THE FIXED-POINT EQUATIONS

In this Appendix, we will derive the fixed point equations for the order parameters following the analysis by Loureiro et al. (2021); Pesce et al. (2023); Adomaityte et al. (2023) in the most general setting discussed in Result 4.3: the case in Result 2.4 is obtained by fixing  $K = 1$  and  $\boldsymbol{\mu}_1 \equiv \mathbf{0}$  below, and by assuming a ridge regularisation. The dataset  $\mathcal{D} := \{(\mathbf{x}_i, y_i)\}_{i \in [n]}$  consists of  $n$  independent datapoints  $\mathbf{x}_i \in \mathbb{R}^d$  each associated to a label  $y_i \in \mathcal{Y}$ . The elements of the dataset are independently generated by using a law  $P(\mathbf{x}, y)$  which we assume can be put in the form of a superstatistical mixture model (SMM) involving  $K$  clusters  $\mathcal{C} = \{1, \dots, K\}$ ,

$$P(\mathbf{x}, y) \equiv P_0(y|\boldsymbol{\beta}_*^\top \mathbf{x}) \sum_{c \in \mathcal{C}} p_c \mathbb{E}_{\sigma_c} [\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_c, \sigma_c^2/d \mathbf{I}_d)], \quad (23)$$

and  $P_0(\bullet|\tau)$  is the distribution of the scalar label  $y$  produced via the ‘‘teacher’’  $\boldsymbol{\beta}_*$ . In the following, we assume that  $\beta_*^2 = 1/d \|\boldsymbol{\beta}_*\|_2^2 = \Theta(1)$ . In the equation above,  $\forall c \in \mathcal{C}$ ,  $p_c \in [0, 1]$  and  $\boldsymbol{\mu}_c \in \mathbb{R}^d$  with  $\|\boldsymbol{\mu}_c\|_2^2 = \Theta(1/d)$ . It is assumed that  $\sum_c p_c = 1$ . The expectation is intended over  $\sigma_c$ , a positive random variable with density  $\varrho_c$ . We will perform our regression task searching for a set of *weights*  $\hat{\boldsymbol{\beta}}_\lambda$ , that will allow us to construct an estimator via a certain classifier  $f: \mathbb{R} \rightarrow \mathcal{Y}$ :

$$\mathbf{x} \mapsto f(\hat{\boldsymbol{\beta}}_\lambda^\top \mathbf{x}) = y, \quad (24)$$

which will provide us with our prediction for a datapoint  $\mathbf{x}$ . The weights will be chosen by minimising an empirical risk function in the form

$$\mathcal{R}(\boldsymbol{\beta}) \equiv \sum_{\nu=1}^n \ell(y_i, \boldsymbol{\beta}^\top \mathbf{x}_i) + \lambda r(\boldsymbol{\beta}), \quad (25)$$

i.e., they are given by

$$\hat{\boldsymbol{\beta}}_\lambda := \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^d} \mathcal{R}(\boldsymbol{\beta}). \quad (26)$$

We will assume that  $\ell$  is a convex loss function with respect to its second argument, and  $r$  is a strictly convex regularisation function: the parameter  $\lambda \geq 0$  will tune the strength of the regularisation. Note that this setting is slightly more general than the one given in the main text. The starting point is to reformulate the problem as an optimisation problem by introducing a Gibbs measure over the parameters  $\boldsymbol{\beta}$  depending on a positive parameter  $\beta$ ,

$$\mu_\beta(\boldsymbol{\beta}) \propto e^{-\beta \mathcal{R}(\boldsymbol{\beta})} = \underbrace{e^{-\beta r(\boldsymbol{\beta})}}_{P_w} \prod_{i=1}^n \underbrace{\exp[-\beta \ell(y_i, \boldsymbol{\beta}^\top \mathbf{x}_i)]}_{P_\ell}, \quad (27)$$

so that, in the  $\beta \rightarrow +\infty$  limit, the Gibbs measure concentrates on  $\hat{\boldsymbol{\beta}}_\lambda$ . The functions  $P_y$  and  $P_w$  can be interpreted as (unnormalised) likelihood and prior distribution respectively. Our analysis will go through the computation of the average free energy density associated with such Gibbs measure in a specific proportional limit, i.e.,

$$f_\beta := - \lim_{\substack{n, d \rightarrow +\infty \\ n/d = \alpha}} \mathbb{E}_{\mathcal{D}} \left[ \frac{\ln \mathcal{Z}_\beta}{d\beta} \right] = \lim_{\substack{n, d \rightarrow +\infty \\ n/d = \alpha}} \lim_{s \rightarrow 0} \frac{1 - \mathbb{E}_{\mathcal{D}}[\mathcal{Z}_\beta^s]}{sd\beta}, \quad (28)$$

where  $\mathbb{E}_{\mathcal{D}}[\bullet]$  is the average over the training dataset, and we have introduced the partition function

$$\mathcal{Z}_\beta := \int e^{-\beta \mathcal{R}(\boldsymbol{\beta})} d\boldsymbol{\beta}. \quad (29)$$

### A.1 REPLICIA APPROACH.

In our replica approach, we need to evaluate

$$\mathbb{E}_{\mathcal{D}}[\mathcal{Z}_\beta^s] = \prod_{a=1}^s \int d\boldsymbol{\beta}^a P_w(\boldsymbol{\beta}^a) \left( \mathbb{E}_{(\mathbf{x}, y)} \left[ \prod_{a=1}^s P_\ell(y|\mathbf{x}^\top \boldsymbol{\beta}^a) \right] \right)^n. \quad (30)$$

Let us take the inner average introducing a new set of *local fields*  $\eta^a$  and  $\tau$ ,

$$\begin{aligned} \mathbb{E}_{(\mathbf{x}, y)} \left[ \prod_{a=1}^s P_\ell(y | \mathbf{x}^\top \beta^a) \right] &= \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \int_{\mathbf{y}} dy \int_{\mathbb{R}^d} d\mathbf{x} P_0(y | \mathbf{x}^\top \beta_\star) \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_c, \sigma_c^2/d \mathbf{I}_d) \prod_{a=1}^s P_\ell(y | \mathbf{x}^\top \beta^a) \right] \\ &= \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \int d\boldsymbol{\eta} \int d\tau \int_{\mathbf{y}} dy P_0(y | \tau) \prod_{a=1}^s P_\ell(y | \eta^a) \mathcal{N}\left(\left(\frac{\tau}{\boldsymbol{\eta}}\right); \left(\frac{\boldsymbol{\mu}_c^\top \beta_\star}{\boldsymbol{\mu}_c^\top \beta^a}\right), \frac{\sigma_c^2}{d} \begin{pmatrix} d\beta_\star^2 & \beta_\star^\top \beta^b \\ \beta_\star^\top \beta^a & \beta^{a\top} \beta^b \end{pmatrix}\right) \right]. \end{aligned} \quad (31)$$

We can write then

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} [\mathcal{Z}_\beta^s] &= \prod_{a=1}^s \int d\beta^a P_w(\beta^a) \times \\ &\left( \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \int d\boldsymbol{\eta} \int d\tau \int_{\mathbf{y}} dy P_0(y | \tau) \prod_{a=1}^s P_\ell(y | \eta^a) \mathcal{N}\left(\left(\frac{\tau}{\boldsymbol{\eta}}\right); \left(\frac{\boldsymbol{\mu}_c^\top \beta_\star}{\boldsymbol{\mu}_c^\top \beta^a}\right), \frac{\sigma_c^2}{d} \begin{pmatrix} d\beta_\star^2 & \beta_\star^\top \beta^b \\ \beta_\star^\top \beta^a & \beta^{a\top} \beta^b \end{pmatrix}\right) \right] \right)^n \\ &= \prod_c \left( \prod_{a \leq b} \int \mathcal{D}Q^{ab} \mathcal{D}\hat{Q}^{ab} \right) \left( \prod_a \int \mathcal{D}M^a \mathcal{D}\hat{M}^a \right) \left( \prod_a \int dt^a d\hat{t}^a \right) e^{-d\beta\Phi^{(s)}}. \end{aligned} \quad (32)$$

In the equation above we introduced the *order parameters*

$$Q_c^{ab} = \frac{\sigma_c^2}{d} \beta^{a\top} \beta^b \in \mathbb{R}, \quad a, b = 1, \dots, s, \quad (33)$$

$$M_c^a = \frac{\sigma_c^2}{d} \beta_\star^\top \beta^a \in \mathbb{R}, \quad a = 1, \dots, s, \quad (34)$$

$$t_c^a = \boldsymbol{\mu}_c^\top \beta^a \in \mathbb{R}, \quad a = 1, \dots, s, \quad (35)$$

whilst the integration is over all possible order parameters,  $Q_c^{ab}$  and  $m_c^a$  to be intended as functions of  $\sigma_c$ . In the equation, we have also denoted the replicated free-energy

$$\begin{aligned} \beta\Phi^{(s)}(\mathbf{Q}, \mathbf{M}, \hat{\mathbf{Q}}, \hat{\mathbf{M}}) &= \sum_c \sum_a \mathbb{E}_{\sigma_c} [\hat{M}_c^a M_c^a] + \sum_c \sum_{a \leq b} \mathbb{E}_{\sigma_c} [\hat{Q}_c^{ab} Q_c^{ab}] + \frac{1}{d} \sum_{c,a} \hat{t}_c^a t_c^a \\ &- \frac{1}{d} \ln \prod_{a=1}^s \int P_w(\beta^a) d\beta^a \prod_c \exp \left( \sum_{a \leq b} \mathbb{E}_{\sigma_c} [\sigma_c^2 \hat{Q}_c^{ab}] \beta^{a\top} \beta^b + \sum_a \mathbb{E}_{\sigma_c} [\sigma_c^2 \hat{M}_c^a] \beta^{a\top} \beta_\star + \sum_a \hat{t}_c^a \beta^{a\top} \boldsymbol{\mu}_c \right) \\ &- \alpha \ln \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \int d\boldsymbol{\eta} \int d\tau \int_{\mathbf{y}} dy P_0(y | \tau) \prod_{a=1}^s P_\ell(y | \eta^a) \mathcal{N}\left(\left(\frac{\tau}{\boldsymbol{\eta}}\right); \left(\frac{t_c^0}{t_c^a}\right) \begin{pmatrix} \sigma_c^2 \beta_\star^2 & M_c^b \\ M_c^a & Q_c^{ab} \end{pmatrix}\right) \right], \end{aligned} \quad (36)$$

where, for the sake of brevity,  $t_c^0 := \boldsymbol{\mu}_c^\top \beta_\star$ . At this point, the free energy  $f_\beta$  should be computed functionally extremizing with respect to all the order parameters by virtue of the Laplace approximation,

$$f_\beta = \lim_{s \rightarrow 0} \text{Extr}_{\substack{\mathbf{M}, \hat{\mathbf{M}}, \mathbf{t} \\ \mathbf{Q}, \hat{\mathbf{Q}}, \hat{\mathbf{t}}}} \frac{\Phi^{(s)}(\mathbf{Q}, \mathbf{M}, \hat{\mathbf{Q}}, \hat{\mathbf{M}}, \mathbf{t}, \hat{\mathbf{t}})}{s}. \quad (37)$$

**Replica symmetric ansatz.** Before taking the  $s \rightarrow 0$  limit we make the replica symmetric assumptions

$$\begin{aligned} Q_c^{aa} &= \begin{cases} R_c, & a = b \\ Q_c & a \neq b \end{cases} & \hat{Q}_c^{aa} &= \begin{cases} -\frac{1}{2} \hat{R}_c, & a = b \\ \hat{Q}_c & a \neq b \end{cases} \\ M_c^a &= M_c & \hat{M}_c^a &= \hat{M}_c \quad \forall a \\ t_c^a &= t_c & \hat{t}_c^a &= \hat{t}_c \quad \forall a \end{aligned} \quad (38)$$

If we denote  $V_c := R_c - Q_c$  we obtain, after some work we obtain



$$\begin{aligned} & \ln \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \int d\boldsymbol{\eta} \int d\tau \int_y dy P_0(y|\tau) \prod_{a=1}^s P_\ell(y|\eta^a) \mathcal{N}\left(\begin{pmatrix} \tau \\ \boldsymbol{\eta} \end{pmatrix}; \begin{pmatrix} t_c^0 \\ t_c \mathbf{1}_s \end{pmatrix}, \begin{pmatrix} \sigma_c^2 \beta_*^2 & M_c \mathbf{1}_s^\top \\ M_c \mathbf{1}_s & Q_c \mathbf{1}_s \times \mathbf{1}_s \end{pmatrix}\right) \right] \\ & = s \sum_c p_c \mathbb{E}_{\sigma_c, \zeta} \left[ \int_y dy Z_0\left(y, t_c^0 + \frac{M_c \zeta}{\sqrt{Q_c}}, \sigma_c^2 \beta_*^2 - \frac{M_c^2}{Q_c}\right) \ln Z_\ell\left(y, t_c + \sqrt{Q_c} \zeta, V_c\right) \right] + o(s), \end{aligned} \quad (39)$$

with  $\zeta \sim \mathcal{N}(0, 1)$  is normally distributed and we have introduced the function

$$Z_\bullet(y, \mu, V) := \int \frac{d\tau P_\bullet(y|\tau)}{\sqrt{2\pi V}} e^{-\frac{(\tau-\mu)^2}{2V}}, \quad \bullet \in \{0, \ell\}. \quad (40)$$

On the other hand, denoted by  $\hat{V}_c = \hat{R}_c + \hat{Q}_c$ , and introducing  $\hat{q}_c := \mathbb{E}_{\sigma_c}[\sigma_c^2 \hat{Q}_c]$ ,  $\hat{v}_c := \mathbb{E}_{\sigma_c}[\sigma_c^2 \hat{V}_c]$ , and  $\hat{m}_c := \mathbb{E}_{\sigma_c}[\sigma_c^2 \hat{M}_c]$

$$\begin{aligned} & \frac{1}{d} \ln \prod_{a=1}^s \left( \int P_w(\boldsymbol{\beta}^a) d\boldsymbol{\beta}^a \prod_c e^{-\frac{\hat{v}_c}{2} \|\boldsymbol{\beta}^a\|_2^2 + \boldsymbol{\beta}^{a\top} (\hat{m}_c \boldsymbol{\beta}_* + \hat{t}_c \boldsymbol{\mu}_c)} \prod_{b,c} e^{\frac{1}{2} \hat{q}_c \boldsymbol{\beta}^{a\top} \boldsymbol{\beta}^b} \right) = \\ & = \frac{s}{d} \mathbb{E}_\xi \ln \left[ \int P_w(\boldsymbol{\beta}) d\boldsymbol{\beta} \prod_c \exp\left(-\frac{\hat{v}_c \|\boldsymbol{\beta}\|_2^2}{2} + \boldsymbol{\beta}^\top (\hat{m}_c \boldsymbol{\beta}_* + \hat{t}_c \boldsymbol{\mu}_c) + \sqrt{\hat{q}_c} \boldsymbol{\xi}^\top \boldsymbol{\beta}\right) \right] + o(s). \end{aligned} \quad (41)$$

In the expression above we have introduced  $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ . Therefore, the (replicated) *replica symmetric* free-energy is given by

$$\lim_{s \rightarrow 0} \frac{\beta}{s} \Phi_{\text{RS}}^{(s)} = \frac{1}{d} \sum_c \hat{t}_c t_c + \sum_c \hat{M}_c M_c + \frac{\sum_c \mathbb{E}_{\sigma_c} [\hat{V}_c Q_c - \hat{Q}_c V_c - \hat{V}_c V_c]}{2} - \alpha \beta \Psi_\ell(M, Q, V) - \beta \Psi_w(\hat{m}, \hat{q}, \hat{v}) \quad (42)$$

where we have defined two contributions

$$\begin{aligned} \Psi_\ell(M, Q, V) & := \frac{1}{\beta} \sum_c p_c \mathbb{E}_{\sigma_c, \zeta} \left[ \int_y dy Z_0\left(y, t_c^0 + \frac{M_c \zeta}{\sqrt{Q_c}}, \sigma_c^2 \beta_*^2 - \frac{M_c^2}{Q_c}\right) \ln Z_\ell\left(y, t_c + \sqrt{Q_c} \zeta, V_c\right) \right], \\ \Psi_w(\hat{m}, \hat{q}, \hat{v}) & := \frac{1}{\beta d} \mathbb{E}_\xi \ln \left[ \int P_w(\boldsymbol{\beta}) d\boldsymbol{\beta} \prod_c \exp\left(-\frac{\hat{v}_c \|\boldsymbol{\beta}\|_2^2}{2} + \boldsymbol{\beta}^\top (\hat{m}_c \boldsymbol{\beta}_* + \hat{t}_c \boldsymbol{\mu}_c) + \sqrt{\hat{q}_c} \boldsymbol{\xi}^\top \boldsymbol{\beta}\right) \right]. \end{aligned} \quad (43)$$

Note that we have separated the contribution coming from the chosen loss (the so-called *channel* part  $\Psi_\ell$ ) from the contribution depending on the regularisation (the *prior* part  $\Psi_w$ ). To write down the saddle-point equations in the  $\beta \rightarrow +\infty$  limit, let us first rescale our order parameters as  $\hat{M}_c \mapsto \beta \hat{M}_c$ ,  $\hat{t}_c \mapsto d\beta \hat{t}_c$ ,  $\hat{Q}_c \mapsto \beta^2 \hat{Q}_c$ ,  $\hat{V}_c \mapsto \beta \hat{V}_c$  and  $V_c \mapsto \beta^{-1} V_c$ . Also, for future convenience, let us rescale  $Q_c \mapsto \sigma_c^2 q_c$ ,  $M_c \mapsto \sigma_c^2 m_c$ ,  $V_c \mapsto \sigma_c^2 v_c$ . For  $\beta \rightarrow +\infty$  the channel part is

$$\begin{aligned} \Psi_\ell(m, q, v, t) & = \\ & = - \sum_c p_c \mathbb{E}_{\sigma_c, \zeta} \left[ \int_y dy Z_0\left(y, t_c^0 + \frac{\sigma_c m \zeta}{\sqrt{q_c}}, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m_c^2}{q_c}\right) \left( \frac{(h_c - t_c - \sigma_c \sqrt{q_c} \zeta)^2}{2\sigma_c^2 v_c} + \ell(y, h_c) \right) \right]. \end{aligned} \quad (44)$$

where we have written  $\Psi_\ell$  of a Moreau envelope, i.e., in terms of a proximal

$$h_c := \arg \min_u \left[ \frac{(u - \omega_c)^2}{2\sigma_c^2 v_c} + \ell(y, u) \right] \quad \text{with } \omega_c = t_c + \sigma_c \sqrt{q_c} \zeta. \quad (45)$$

A similar expression can be obtained for  $\Psi_w$ . Introducing the proximal

$$\mathbf{g} = \arg \min_{\boldsymbol{\beta}} \left( \frac{\|\boldsymbol{\beta}\|_2^2 \sum_c \hat{v}_c}{2} - \boldsymbol{\beta}^\top \sum_c (\hat{m}_c \boldsymbol{\beta}_* + \hat{t}_c \boldsymbol{\mu}_c) - \boldsymbol{\xi}^\top \boldsymbol{\beta} \sum_c \sqrt{\hat{q}_c} + \lambda r(\boldsymbol{\beta}) \right) \in \mathbb{R}^d \quad (46)$$

We can rewrite the prior contribution  $\Psi_w$  as

$$\Psi_w(\hat{m}, \hat{q}, \hat{v}, \hat{t}) = -\frac{1}{d} \mathbb{E}_\xi \left[ \frac{\|\mathbf{g}\|_2^2}{2} \sum_c \hat{v}_c - \mathbf{g}^\top \sum_c (\hat{m}_c \boldsymbol{\beta}_* + \hat{t}_c \boldsymbol{\mu}_c) - \boldsymbol{\xi}^\top \mathbf{g} \sum_c \sqrt{\hat{q}_c} + \lambda r(\mathbf{g}) \right] \quad (47)$$

The parallelism between the two contributions is evident, aside from the different dimensionality of the involved objects. The replica symmetric free energy in the  $\beta \rightarrow +\infty$  limit is computed by extremising with respect to the introduced order parameters,

$$f_{\text{RS}} = \text{Extr} \left[ \sum_c \mathbb{E}_{\sigma_c} [\sigma_c^2 \hat{M}_c m_c] + \frac{1}{2} \sum_c \mathbb{E}_{\sigma_c} \left[ \sigma_c^2 (\hat{V}_c q_c - \hat{Q}_c v_c) \right] + \sum_c t_c \hat{t}_c \right. \\ \left. - \alpha \Psi_\ell(m, q, v, t) - \Psi_w(\hat{m}, \hat{q}, \hat{v}, \hat{t}) \right]. \quad (48)$$

To do so, we have to write down a set of saddle-point equations and solve them.

**Saddle-point equations.** The saddle-point equations are derived straightforwardly from the obtained free energy functionally extremising with respect to all parameters. It is easily seen that  $v_c$ ,  $q_c$  and  $m_c$  are independent from  $\sigma_c$ , and that it is possible to reduce the number of variables by introducing  $\hat{v} = \sum_c \hat{v}_c$ , so that we can write

$$v_c = \frac{\mathbb{E}_\xi [\mathbf{g}^\top \boldsymbol{\xi}]}{d\sqrt{\hat{q}_c}}, \quad (49a)$$

$$q = \frac{\mathbb{E}_\xi [\|\mathbf{g}\|_2^2]}{d}, \quad (49b)$$

$$m = \frac{\mathbb{E}_\xi [\mathbf{g}^\top \boldsymbol{\beta}_*]}{d}, \quad (49c)$$

$$t_c = \mathbb{E}_\xi [\mathbf{g}^\top \boldsymbol{\mu}_c]. \quad (49d)$$

and the remaining equations can be rewritten as

$$\hat{q}_c = \alpha p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} \left[ \sigma_c^2 Z_0 \left( y, t_c^0 + \frac{\sigma_c m}{\sqrt{q}}, \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) f_c^2 \right], \quad (50a)$$

$$\hat{v} = -\alpha \sum_c p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} \left[ \sigma_c^2 Z_0 \left( y, t_c^0 + \frac{\sigma_c m}{\sqrt{q}}, \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) \partial_w f_c \right], \quad (50b)$$

$$\hat{m}_c = \alpha p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} \left[ \sigma_c^2 \partial_\mu Z_0 \left( y, t_c^0 + \frac{\sigma_c m}{\sqrt{q}}, \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) f_c \right], \quad (50c)$$

$$\hat{t}_c = \alpha p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} \left[ Z_0 \left( y, t_c^0 + \frac{\sigma_c m}{\sqrt{q}}, \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) f_c \right] \quad (50d)$$

with

$$f_c := \arg \min_u \left[ \frac{\sigma_c^2 v_c u^2}{2} + \ell(y, \omega_c + \sigma_c^2 v_c u) \right], \quad \omega_c = t_c + \sigma_c \sqrt{q} \zeta, \quad (51)$$

$$\mathbf{g} = \arg \min_{\boldsymbol{\beta}} \left( \frac{\|\boldsymbol{\beta}\|_2^2 \hat{v}}{2} - \boldsymbol{\beta}^\top \sum_c (\hat{m}_c \boldsymbol{\beta}_* + d \hat{t}_c \boldsymbol{\mu}_c) - \boldsymbol{\xi}^\top \boldsymbol{\beta} \sum_c \sqrt{\hat{q}_c} + \lambda r(\boldsymbol{\beta}) \right).$$

To obtain the replica symmetric free energy, therefore, the given set of equations has to be solved, and the result is then plugged in Eq. 48. The obtained saddle-point equations correspond to the ones given in the Result 4.3.

**Training and test errors.** Let us show how to use the previous result to estimate the training loss and the generalisation error. Let us start from estimating

$$\varepsilon_\ell := \lim_{n \rightarrow +\infty} \frac{1}{n} \sum_{i=1}^n \ell(y_i, \hat{\boldsymbol{\beta}}_\lambda^\top \mathbf{x}_i). \quad (52)$$

The best way to proceed is to observe that

$$\varepsilon_\ell = - \lim_{\beta \rightarrow +\infty} \partial_\beta (\beta \Psi_\ell) = \sum_c p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} \left[ Z_0 \left( y, t_c^0 + \frac{\sigma_c m}{\sqrt{q}}, \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right) \ell(y, h_c) \right]. \quad (53)$$

This concentration result holds for a generic function  $\varphi: \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$ , so that more generally, under Assumption 2.3,

$$\frac{1}{n} \sum_{i=1}^n \varphi(y_i, \hat{\beta}_\lambda^\top \mathbf{x}_i) \xrightarrow[n, d \rightarrow +\infty]{\text{P}} \sum_c p_c \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma_c, \zeta} \left[ Z_0 \left( y, t_c^0 + \frac{\sigma_c m}{\sqrt{q}} \zeta, \sigma_c^2 \beta_\star^2 - \frac{\sigma_c^2 m^2}{q} \right) \varphi(y, h_c) \right]. \quad (54)$$

The expressions above hold in general, but, as anticipated, important simplifications can occur in the set of saddle-point equations Eq. 50 and Eq. 49 depending on the choice of the loss  $\ell$  and of the regularization function  $r$ . The population's expectation (e.g., in the computation of the test error) of a performance function  $\varphi: \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$  is given instead by

$$\varepsilon_g := \lim_{n \rightarrow +\infty} \mathbb{E}_{(y, \mathbf{x})} \left[ \varphi(y, \hat{\beta}_\lambda^\top \mathbf{x}) \right], \quad (55)$$

where the expectation has to be taken on a newly sampled datapoint  $(y, \mathbf{x}) \notin \mathcal{D}$ . This expression can be rewritten as

$$\begin{aligned} & \mathbb{E}_{y|\mathbf{x}} \left[ \int d\eta \varphi(y, \eta) \mathbb{E}_{\mathbf{x}} \left[ \delta \left( \eta - \hat{\beta}_\lambda^\top \mathbf{x} \right) \right] \right] \\ & \xrightarrow[n, d \rightarrow +\infty]{\text{P}} \int d\eta \int d\tau \int_{\mathcal{Y}} dy P_0(y|\tau) \varphi(y, \eta) \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \mathcal{N} \left( \begin{pmatrix} \tau \\ \eta \end{pmatrix}; \begin{pmatrix} t_c^0 \\ t_c \end{pmatrix}, \sigma_c^2 \begin{pmatrix} \beta_\star^2 & m \\ m & q \end{pmatrix} \right) \right]. \end{aligned} \quad (56)$$

This can be easily computed numerically once the order parameters are given. Finally, another relevant quantity for our investigations is the estimation mean-squared error

$$\varepsilon_{\text{est}} := \lim_{d \rightarrow +\infty} \frac{1}{d} \mathbb{E}_{\mathcal{D}} \left[ \|\hat{\beta}_\lambda - \beta_\star\|_2^2 \right] = \beta_\star^2 - 2m + q. \quad (57)$$

## A.2 BAYES-OPTIMAL ESTIMATOR FOR $K = 1$

A derivation similar to the one discussed above can be repeated to obtain information on the statistical properties of the Bayes optimal estimator presented in Result 2.6. Given a dataset  $\mathcal{D}$  of observation, we have that

$$P(\beta|\mathcal{D}) = \frac{P(\beta)P(\mathcal{D}|\beta)}{\mathcal{Z}(\mathcal{D})} = \frac{P(\beta)}{\mathcal{Z}(\mathcal{D})} \prod_{i=1}^n P_0(y_i|\beta^\top \mathbf{x}_i) \quad (58)$$

where  $P(\beta)$  is the prior on the teacher that we assume to be Gaussian,  $P(\beta) = \mathcal{N}(\beta; \mathbf{0}, \beta_\star^2 \mathbf{I}_d)$ , and

$$\mathcal{Z}(\mathcal{D}) := \int d\beta P(\beta) \prod_{i=1}^n P_0(y_i|\beta^\top \mathbf{x}_i) = \frac{1}{(2\pi)^{n/2}} \int d\beta \exp \left( -\frac{\|\beta\|_2^2}{2\beta_\star^2} + \sum_{i=1}^n \ln P_0(y_i|\beta^\top \mathbf{x}_i) \right). \quad (59)$$

The calculation of  $\mathcal{Z}(\mathcal{D})$  gives access in particular to the free entropy  $\phi(\mathcal{D}) := \lim_n \frac{1}{n} \ln \mathcal{Z}(\mathcal{D})$ . The computation of  $\phi$ , which has an information-theoretical interpretation as mutual information, provides us the statistics of  $\beta$  according to the true posterior  $P(\beta|\mathcal{D})$ . By assuming a concentration in the large  $n$  limit, the calculation is performed on  $\mathbb{E}_{\mathcal{D}}[\ln \mathcal{Z}(\mathcal{D})]$ . Using the replica trick as before, the calculation can be repeated almost identically. For the sake of simplicity, we focus on the case in which only one cluster is present, centered in the origin. It is found that the statistics of a Bayes optimal estimator can be characterised therefore by an order parameter  $\mathbf{q}$  satisfying a self-consistent equation not different from the one presented above (we will use below a different font to stress that we are currently analysing the Bayes optimal setting)

$$\hat{\mathbf{q}} = \alpha \int_{\mathcal{Y}} dy \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0(y, \mu, V) \left( \partial_\mu \ln Z_0(y, \mu, V) \right)^2 \Big|_{\substack{\mu = \sigma \sqrt{q} \zeta \\ V = \sigma^2 (\beta_\star^2 - q)}} \right], \quad \mathbf{q} = \frac{\beta_\star^4 \hat{\mathbf{q}}}{1 + \beta_\star^2 \hat{\mathbf{q}}}. \quad (60)$$

with  $Z_0(y, \mu, V) := \mathbb{E}_z [P_0(y|\mu + \sqrt{V}z)]$  with  $z \sim \mathcal{N}(0, 1)$ . We can then compute the Bayes optimal estimation error for the Bayes optimal estimator  $\hat{\beta}_{\text{BO}} = \mathbb{E}_{\beta|\mathcal{D}}[\beta]$  as

$$\varepsilon_{\text{BO}} = \lim_{d \rightarrow +\infty} \frac{1}{d} \|\beta_\star - \hat{\beta}_{\text{BO}}\|_2^2 = \beta_\star^2 - \mathbf{q}. \quad (61)$$

## B ASYMPTOTIC RESULTS FOR RIDGE-REGULARISED LOSSES

Let us fix now  $r(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ . In this case, the computation of  $\Psi_w$  can be performed explicitly via a Gaussian integration, and the saddle-point equations can take a more compact form that is particularly suitable for a numerical solution. In particular, the prior proximal is found as

$$\mathbf{g} = \frac{\sum_c (\hat{m}_c \boldsymbol{\beta}_* + d \hat{t}_c \boldsymbol{\mu}_c) + \sum_c \sqrt{\hat{q}_c} \boldsymbol{\xi}}{\lambda + \hat{v}} \quad (62)$$

so that the prior saddle-point equations obtained from  $\Psi_w$  become

$$\begin{aligned} q &= \frac{1}{d} \left( \sum_c \frac{\hat{m}_c \boldsymbol{\beta}_* + d \hat{t}_c \boldsymbol{\mu}_c}{\lambda + \hat{v}} \right)^2 + \left( \sum_c \frac{\sqrt{\hat{q}_c}}{\lambda + \hat{v}} \right)^2 & \hat{q}_c &= \alpha p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} [\sigma_c^2 Z_0 f_c^2], \\ m &= \frac{\sum_c (\beta_*^2 \hat{m}_c + t_c^0 \hat{t}_c)}{\lambda + \hat{v}} & \hat{v} &= -\alpha \sum_c p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} [\sigma_c^2 Z_0 \partial_\omega f_c], \\ v_c &= \frac{1}{\lambda + \hat{v}} \sum_{c'} \sqrt{\frac{\hat{q}_{c'}}{\hat{q}_c}} & \hat{m}_c &= \alpha p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} [\sigma_c^2 \partial_\mu Z_0 f_c], \\ t_c &= \frac{\sum_{c'} (\hat{t}_{c'} \mu_{c'} + t_c^0 \hat{m}_{c'})}{\lambda + \hat{v}}, & \hat{t}_c &= \alpha p_c \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma_c, \zeta} [Z_0 f_c] \end{aligned} \quad (63)$$

We have used the shorthand notation  $Z_0 \equiv Z_0 \left( y, t_c^0 + \frac{\sigma_c m}{\sqrt{q}} \zeta, \sigma_c^2 \beta_*^2 - \frac{\sigma_c^2 m^2}{q} \right)$  and  $\boldsymbol{\mu}_{c'} := d \boldsymbol{\mu}_{c'}^\top \boldsymbol{\mu}_c$ .

**Regression on one cloud: consistency** In the special case in which  $|\mathcal{C}| = 1$  and the cloud is centered in the origin, we have  $t_1 = \hat{t}_1 = 0$  and, dropping the subscript referring to the cluster,

$$\begin{aligned} q &= \frac{\beta_*^2 \hat{m}^2 + \hat{q}}{(\lambda + \hat{v})^2} & \hat{q} &= \alpha \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0 \left( y, \frac{\sigma m}{\sqrt{q}} \zeta, \sigma^2 \beta_*^2 - \frac{\sigma^2 m^2}{q} \right) f^2 \right], \\ m &= \frac{\beta_*^2 \hat{m}}{\lambda + \hat{v}} & \hat{v} &= -\alpha \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0 \left( y, \frac{\sigma m}{\sqrt{q}} \zeta, \sigma^2 \beta_*^2 - \frac{\sigma^2 m^2}{q} \right) \partial_\omega f \right], \\ v &= \frac{1}{\lambda + \hat{v}}, & \hat{m} &= \alpha \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 \partial_\mu Z_0 \left( y, \frac{\sigma m}{\sqrt{q}} \zeta, \sigma^2 \beta_*^2 - \frac{\sigma^2 m^2}{q} \right) f \right], \end{aligned} \quad (64)$$

where as usual  $f := \arg \min_u \left[ \frac{\sigma^2 v u^2}{2} + \ell(y, \omega + \sigma^2 v u) \right]$  and  $\omega = \sigma \sqrt{q} \zeta$ . Let us now perform the rescaling  $v \mapsto \alpha v$ ,  $\hat{q} \mapsto \alpha \hat{q}$ ,  $\hat{m} \mapsto \alpha \hat{m}$ , and  $\hat{v} \mapsto \alpha \hat{v}$ , where  $v = O(1)$ ,  $\hat{v} = O(1)$ ,  $\hat{m} = O(1)$ ,  $\hat{q} = O(1)$ . Then, under these assumptions, in the  $\alpha \rightarrow +\infty$  limit

$$\begin{aligned} q &= \frac{\beta_*^2 \hat{m}^2}{\hat{v}^2} & \hat{q} &= \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0 \left( y, \frac{\sigma m}{\sqrt{q}} \zeta, \sigma^2 \beta_*^2 - \frac{\sigma^2 m^2}{q} \right) f^2 \right], \\ m &= \frac{\beta_*^2 \hat{m}}{\hat{v}} & \hat{v} &= - \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 Z_0 \left( y, \frac{\sigma m}{\sqrt{q}} \zeta, \sigma^2 \beta_*^2 - \frac{\sigma^2 m^2}{q} \right) \partial_\omega f \right], \\ v &= \frac{1}{\hat{v}}, & \hat{m} &= \int_{\mathbf{y}} d\mathbf{y} \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 \partial_\mu Z_0 \left( y, \frac{\sigma m}{\sqrt{q}} \zeta, \sigma^2 \beta_*^2 - \frac{\sigma^2 m^2}{q} \right) f \right], \end{aligned} \quad (65)$$

Moreover, in the large  $\alpha$  limit,  $f = -\partial_\eta \ell(y, \eta)|_{\eta = \sigma \sqrt{q} \zeta}$ . It follows that, independently from the adopted loss, the angle  $\pi^{-1} \arccos \frac{m}{\beta_* \sqrt{q}}$  between the estimator  $\hat{\boldsymbol{\beta}}_\lambda$  and the teacher  $\boldsymbol{\beta}_*$  goes to zero as  $\alpha \rightarrow +\infty$ . In this limit, it is easily found that  $\varepsilon_{\text{est}} \rightarrow \beta_*^{-2} (m - \beta_*^2)^2$ , hence the estimator is consistent if  $m \rightarrow \beta_*^2$ .

**Uncorrelated teachers: universality** To study the universality properties in the ridge setting, let us introduce two possible new assumptions.

**Assumption B.1.** For all  $c \in [K]$ ,  $\lim_{d \rightarrow +\infty} t_c^0 = 0$ .

This assumption holds, for example, by assuming the centroids  $\boldsymbol{\mu}_c \sim \mathcal{N}(\mathbf{0}, 1/d \mathbf{I}_d)$ . It expresses in general the fact that the teacher  $\boldsymbol{\beta}_*$  is completely uncorrelated from the different centroids  $\boldsymbol{\mu}_c$ .

**Assumption B.2.** *The following symmetry properties hold*

$$P_0(y|\tau) = P_0(-y|-\tau), \quad \ell(y, \eta) = \ell(-y, -\eta). \quad (66)$$

Under Assumption B.1 and Assumption B.2,  $t_c = \hat{t}_c = 0 \forall c$  is a saddle-point solution of the equations 63. Indeed, if  $\hat{t}_c = 0$  the prior equation implies  $t_c = 0$ . On the other hand, if  $t_c = 0$ ,  $\hat{t}_c = 0$  for parity reason (Pesce et al., 2023). We recover therefore in our setting the *mean universality* discussed by Pesce et al. (2023) in the Gaussian setting: the learning task is mean-independent and equivalent to one on  $c$  clouds all centered in the origin, i.e., a problem obtained assuming  $\mathbf{x} \sim \sum_c p_c \mathbb{E}_{\sigma_c}[\mathcal{N}(\mathbf{0}, \sigma_c/d\mathbf{I}_d)]$ . Note that in the Gaussian setting (i.e., assuming  $\rho_c(\sigma) = \delta(\sigma - \bar{\sigma}_c)$  for some fixed  $\bar{\sigma}_c$  for each  $c \in [C]$ ) Pesce et al. (2023) also observed that in the limit  $\lambda \rightarrow 0^+$ , *covariance universality* holds in the Gaussian case:  $\varepsilon_g$  and  $\varepsilon_\ell$  are independent on the covariance of the clouds. This fact does not extend to the case in which the distribution of  $\sigma_c$  is not atomic (not even in the case in which  $\sigma_c$  are all identically distributed), as it has been verified by Adomaityte et al. (2023).

### B.1 SQUARE LOSS

If we consider a square loss  $\ell(y, \eta) = \frac{1}{2}(y - \eta)^2$ , then an explicit formula for the proximal can be found, namely

$$f_c = \frac{y - \omega_c}{1 + v_c \sigma_c^2}, \quad \omega_c = t_c + \sigma_c \sqrt{q} \zeta, \quad (67)$$

so that the second set of saddle-point equations Eq. 50 can be made even more explicit. Let us assume, that labels are generated according to the linear model in Eq. 1a, where the noise term  $\eta_i$  has  $\mathbb{E}[\eta_i] = 0$  and  $\hat{\sigma}_0^2 := \mathbb{E}[\eta_i^2] < +\infty$ . In this setting, the channel equations can be written as

$$\hat{q}_c = \alpha p_c \hat{\sigma}_0^2 \mathbb{E}_{\sigma_c} \left[ \frac{\sigma_c^2}{(1 + v_c \sigma_c^2)^2} \right] + \alpha p_c \mathbb{E}_{\sigma_c} \left[ \left( \frac{\sigma_c^2}{1 + v_c \sigma_c^2} \right)^2 \right] (\beta_\star^2 - 2m + q + (t_c^0 - t_c)^2), \quad (68a)$$

$$\hat{v} = \alpha \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \frac{\sigma_c^2}{1 + v_c \sigma_c^2} \right], \quad (68b)$$

$$\hat{m}_c = \alpha p_c \mathbb{E}_{\sigma_c} \left[ \frac{\sigma_c^2}{1 + v_c \sigma_c^2} \right] \quad (68c)$$

$$\hat{t}_c = \alpha p_c (t_c^0 - t_c) \mathbb{E}_{\sigma_c} \left[ \frac{1}{1 + v_c \sigma_c^2} \right]. \quad (68d)$$

In this setting the generalisation error becomes

$$\varepsilon_g := \lim_{d \rightarrow +\infty} \mathbb{E} \left[ \left( y - \hat{\beta}_\lambda^\top \mathbf{x} \right)^2 \right] = \hat{\sigma}_0^2 + \sum_c p_c (t_c^0 - t_c)^2 + (\beta_\star^2 - 2m + q) \sum_c p_c \mathbb{E}_{\sigma_c} [\sigma_c^2], \quad (69)$$

which is finite if and only if  $\mathbb{E}_{\sigma_c} [\sigma_c^2] < +\infty$  for all  $c$ . Note that the dependence on  $\hat{v}$  is through its second moment only. Observe that the possible power-law behavior of the noise on the label *does not* influence the generalisation performances, that only depends on the noise variance  $\hat{\sigma}_0^2$ . The training loss, on the other hand, is

$$\begin{aligned} \varepsilon_\ell := \lim_{d \rightarrow +\infty} \frac{1}{2n} \sum_{i=1}^n \left( y_i - \hat{\beta}_\lambda^\top \mathbf{x}_i \right)^2 &\xrightarrow{d \rightarrow +\infty} \frac{\hat{\sigma}_0^2}{2} \sum_c p_c \mathbb{E}_{\sigma_c} \left[ \frac{1}{(1 + v_c \sigma_c^2)^2} \right] + \\ &+ \sum_c \frac{p_c}{2} \left[ \mathbb{E}_{\sigma_c} \left[ \frac{1}{(1 + v_c \sigma_c^2)^2} \right] (t_c^0 - t_c)^2 + \mathbb{E}_{\sigma_c} \left[ \frac{\sigma_c^2}{(1 + v_c \sigma_c^2)^2} \right] (\beta_\star^2 - 2m + q) \right]. \end{aligned} \quad (70)$$

**Strong universality of  $\varepsilon_\ell$  for  $\lambda \rightarrow 0^+$**  We will show now that, under the Assumption B.1 (Assumption B.2 is satisfied in the setting under consideration), the strong universality of the training loss observed by Pesce et al. (2023) is preserved. Let us put ourselves in the

case of a single cluster centered in the origin (an assumption that is not restrictive, as shown above). In this case, let us introduce

$$S_v := \mathbb{E}_\sigma \left[ \frac{1}{1 + v\sigma^2} \right] \quad (71)$$

which can be interpreted in terms of the Stieljes transform of the random variable  $\sigma^2$ . The saddle-point equations are

$$\begin{aligned} \hat{q} &= -\alpha \hat{\sigma}_0^2 \partial_v S_v + \alpha (1 - S_v + v \partial_v S_v) \frac{\beta_\star^2 - 2m + q}{v^2}, & q &= \frac{\hat{m}^2 \beta_\star^2 + \hat{q}}{(\lambda + \hat{v})^2} \\ \hat{v} &= \alpha \frac{1 - S_v}{v}, & m &= \frac{\beta_\star^2 \hat{m}}{\lambda + \hat{v}} \\ \hat{m} &= \alpha \frac{1 - S_v}{v}, & v &= \frac{1}{\lambda + \hat{v}}. \end{aligned} \quad (72)$$

The training loss can be written as

$$\varepsilon_\ell = -\frac{\hat{\sigma}_0^2 \partial_v S_v}{2} - \frac{(\beta_\star^2 - 2m + q) \partial_v S_v}{2}. \quad (73)$$

In the limit  $\lambda \rightarrow 0$ ,

$$x := \frac{\beta_\star^2 - 2m + q}{v} = \frac{(1 - S_v + v \partial_v S_v)x - v \partial_v S_v \hat{\sigma}_0^2}{1 - S_v} \Rightarrow x = \hat{\sigma}_0^2 \quad (74)$$

so that  $\varepsilon_\ell = \frac{1}{2} S_v \hat{\sigma}_0^2$ . The quantity  $S_v$  can be extracted from the equation for  $v$ , as it has to satisfy, in the zero regularisation limit,  $\alpha(1 - S_v) = 1 \Rightarrow S_v = 1 - \frac{1}{\alpha}$  which is a valid solution for  $\alpha > 1$  only. As a result, we obtain a *universal* formula for the training loss

$$\varepsilon_\ell = \frac{\hat{\sigma}_0^2}{2} \left( 1 - \frac{1}{\alpha} \right)_+, \quad \text{where } (x)_+ := x\theta(x). \quad (75)$$

Note that the formula above is valid for *any* distribution of  $\sigma$ , including distributions with no second moment.

**Generalisation error rate** — We conclude this section by extracting the generalisation error rate for  $n \rightarrow +\infty$  and large but fixed  $d$ , i.e., for  $\alpha \rightarrow +\infty$ . For simplicity, let us focus, once again, on the case  $K = 1$  and  $\boldsymbol{\mu}_1 = \mathbf{0}$ , corresponding to the fixed-point equations given in Eq. 72. Let us assume that  $\hat{\sigma}_0^2 < +\infty$  and that  $\sigma_0^2 := \mathbb{E}[\sigma^2] < +\infty$ . From Eq. 72  $v$  satisfies the equation  $\alpha(1 - S_v) = 1 - \lambda v$ : as  $S_v \in [0, 1]$  and  $v > 0$ , then for  $\alpha \rightarrow +\infty$  we must have  $S_v \rightarrow 1$  and  $v \rightarrow 0$ , so that for  $\alpha \rightarrow +\infty$ ,  $S_v = 1 - \frac{1}{\alpha} + O(\alpha^{-1})$ . In this limit, therefore, by direct inspection of the fixed-point equations,  $q \rightarrow \beta_\star^2$  and  $m \rightarrow \beta_\star^2$  so that  $\varepsilon_{\text{est}} \rightarrow 0$  and the estimator  $\hat{\beta}_\lambda$  is unbiased.

In the hypothesis that  $\sigma_0^2$  is finite (i.e.,  $\varrho(\sigma) \sim \sigma^{-2a-1}$  with  $a > 1$  for  $\sigma \gg 1$ ), then, for small  $v$ , as  $S_v \simeq 1 - v\sigma_0^2 + o(v)$ , it is found that

$$q = \beta_\star^2 + \frac{\hat{\sigma}_0^2 - 2\beta_\star^2 \lambda}{\sigma_0^2} \frac{1}{\alpha} + o\left(\frac{1}{\alpha}\right), \quad m = \beta_\star^2 - \frac{\lambda \beta_\star^2}{\sigma_0^2} \frac{1}{\alpha} + o\left(\frac{1}{\alpha}\right), \quad v = \frac{1}{\sigma_0^2 \alpha} + o\left(\frac{1}{\alpha}\right), \quad (76)$$

which, together with our general formulas for  $\varepsilon_{\text{est}}$ , imply  $\varepsilon_{\text{est}} \sim \alpha^{-1}$  for large  $\alpha$ .

On the other hand, let us consider the case in which  $\varrho(\sigma) \sim \sigma^{-2a-1}$  for  $\sigma \gg 1$ , with  $0 < a < 1$ . In this case,  $\sigma_0^2 = +\infty$  and  $S_v$  has an expansion in the form  $S_v = 1 - \tilde{\sigma}_0^2 v^a + O(v)$  for some finite positive quantity  $\tilde{\sigma}_0^2$ . Such asymptotic implies that  $v \simeq (\tilde{\sigma}_0^2 \alpha)^{-1/a}$  for  $\alpha \gg 1$ . By replacing this in the fixed point equations, it is found that  $m = \beta_\star^2 - \lambda \beta_\star^2 (\tilde{\sigma}_0^2 \alpha)^{-1/a} + o(1/\alpha)$  and  $q = \beta_\star^2 + [\hat{\sigma}_0^2 - 2\lambda \beta_\star^2] (\tilde{\sigma}_0^2 \alpha)^{-1/a} + o(1/\alpha)$ , so that in the end  $\varepsilon_{\text{est}} \sim \alpha^{-1/a}$ .

The  $a = 1$  case is marginal, as  $S_v \simeq 1 + \tilde{\sigma}_0^2 v \log v$  for  $\alpha \gg 1$  for some positive constant  $\tilde{\sigma}_0^2$ . Therefore  $v = (\tilde{\sigma}_0^2 \alpha \ln \alpha)^{-1}$ . By consequence, for  $\alpha \gg 1$   $m \simeq \beta_\star^2 - \lambda \beta_\star^2 (\tilde{\sigma}_0^2 \alpha \ln \alpha)^{-1}$  and  $q \simeq \beta_\star^2 + [\hat{\sigma}_0^2 - 2\lambda \beta_\star^2] (\tilde{\sigma}_0^2 \alpha \ln \alpha)^{-1}$ , so that  $\varepsilon_{\text{est}} \sim (\alpha \ln \alpha)^{-1}$ .



## B.2 HUBER LOSS AND ROBUST REGRESSION IN THE PRESENCE OF FAT TAILS

### B.2.1 A MODEL FOR THE STUDY OF ROBUSTNESS

A toy model for the study of robustness has been introduced recently by [Vilucchio et al. \(2023\)](#). Here we will consider a more general setting to include the possibility of having fat tails. We consider the case of one cloud only,  $K = 1$ , so that  $P(\mathbf{x}) = \mathbb{E}_\sigma[\mathcal{N}(\mathbf{x}; \mathbf{0}, \sigma^2 \mathbf{I}_d)]$ , and  $P_0(y|\tau) = \mathbb{E}_{\hat{\sigma}}[\mathcal{N}(y; \hat{\eta}\tau, \hat{\sigma}^2)]$ , where the expectation is taken over the joint distribution  $\hat{\rho}$  for the pair  $(\hat{\eta}, \hat{\sigma})$  of (possibly correlated) random variables. [Vilucchio et al. \(2023\)](#) adopted, in particular, the distribution  $\hat{\rho}(\hat{\eta}, \hat{\sigma}) = \epsilon \delta_{\hat{\eta}, \hat{\eta}_{\text{out}}} \delta_{\hat{\sigma}, \hat{\sigma}_{\text{out}}} + (1 - \epsilon) \delta_{\hat{\eta}, 1} \delta_{\hat{\sigma}, \hat{\sigma}_{\text{in}}}$  for  $\epsilon \in [0, 1]$ , with  $(\hat{\eta}_{\text{out}}, \hat{\sigma}_{\text{out}})$  referring to ‘‘outlier labels’’, and  $(1, \hat{\sigma}_{\text{in}})$  referring to ‘‘inlier labels’’. The general fixed-point equations can be adapted to this case quite easily. We assume, once again, a ridge regularisation. Here we comment on the fact that in this case, it can be interesting to consider, beyond the ERM estimator  $\hat{\beta}_\lambda$  and the Bayes-optimal estimator  $\hat{\beta}_{\text{BO}}$ , the estimator that minimises the (posterior-averaged) mean-squared test error

$$\hat{\beta}_{g, \text{BO}} = \arg \min_{\beta} \mathbb{E}_{\hat{\rho}|\mathcal{D}} \left[ \mathbb{E}_{(y, \mathbf{x})|\hat{\beta}} \left[ (y - \beta^\top \mathbf{x})^2 \right] \right] = \mathbb{E}[\hat{\eta}] \hat{\beta}_{\text{BO}}. \quad (77)$$

In the expression above,  $\mathbb{E}_{(y, \mathbf{x})|\hat{\beta}}$  expresses the fact that the pair  $(y, \mathbf{x})$  has been generated with a teacher vector  $\hat{\beta}$ , sampled by the posterior. Using the results on the Bayes optimal estimator, it is simple to derive the errors obtained by using  $\hat{\beta}_{g, \text{BO}}$ . Under the assumptions that  $\sigma_0^2 := \mathbb{E}[\sigma^2] < +\infty$  and  $\hat{\sigma}_0^2 := \mathbb{E}[\hat{\sigma}^2]$ ,

$$\varepsilon_{g, \text{BO}} := \mathbb{E}_{(y, \mathbf{x})} \left[ (y - \mathbf{x}^\top \beta_{g, \text{BO}})^2 \right] = \sigma_0^2 (\beta_\star^2 \mathbb{E}[\hat{\eta}^2] - \mathbb{E}[\hat{\eta}]^2 \mathbf{q}) + \hat{\sigma}_0^2,$$

where  $\mathbf{q}$  is provided by Eq. 60. As in the pure Gaussian case, by imposing the ansatz  $\mathbf{q} = \beta_\star^2 - \frac{1}{\alpha} \mathbf{q}_0 + \Theta(\alpha^{-2})$ , and consequently  $\hat{\mathbf{q}} = \alpha \hat{\mathbf{q}}_0 + \Theta(1)$  for large  $\alpha$ , we can obtain

$$\frac{1}{\hat{\mathbf{q}}_0} = \hat{\mathbf{q}}_0 = \int_y dy \mathbb{E}_{\sigma, \zeta} \left[ \sigma^2 P_0(y|\omega) (\partial_\omega \ln P_0(y|\omega))^2 \Big|_{\omega = \sigma \beta_\star \zeta} \right]. \quad (78)$$

In the  $\alpha \rightarrow +\infty$  limit, then,  $\mathbf{q} \rightarrow \rho$  and  $\varepsilon_{\text{est}, \text{BO}} := \lim_{d \rightarrow +\infty} \frac{1}{d} \mathbb{E}_{\mathcal{D}} [\|\hat{\beta}_{g, \text{BO}} - \beta\|_2^2] = \frac{\mathbf{q}_0}{\alpha} + \Theta(\alpha^{-2}) \rightarrow 0$ . On the other hand,  $\varepsilon_{g, \text{BO}} = \hat{\sigma}_0^2 + \sigma_0^2 \beta_\star^2 \text{Var}[\hat{\eta}] - \frac{\sigma_0^2 \mathbf{q}_0}{\alpha} + \Theta(\alpha^{-2})$ .

### B.2.2 HUBER LOSS AND ITS APPLICATION

The Huber loss is a strongly convex loss depending on a tunable parameter  $\delta \geq 0$  and is defined as

$$\ell_\delta(y, \eta) = \begin{cases} \frac{(y - \eta)^2}{2} & \text{if } |y - \eta| < \delta \\ \delta |y - \eta| - \frac{\delta^2}{2} & \text{otherwise.} \end{cases} \quad (79)$$

This loss is widely adopted in robust regression as it is less sensitive to outliers than the most commonly adopted square loss, and is associated with the following expression for the proximal

$$h_c = \omega_c + \frac{(y - \omega_c) v_c \sigma_c^2}{\max(\delta^{-1} |y - \omega_c|, 1 + v_c \sigma_c^2)} \Leftrightarrow f_c = \frac{y - \omega_c}{\max(\delta^{-1} |y - \omega_c|, 1 + v_c \sigma_c^2)}, \quad \omega_c = t_c + \sigma_c \sqrt{q} \zeta. \quad (80)$$

The prior equations are therefore the usual in Eq. 64. The channel equations are instead

$$\hat{m} = \alpha \mathbb{E} \left[ \frac{\sigma^2 \hat{\eta} \text{erf } \chi}{1 + v \sigma^2} \right] \quad (81a)$$

$$\hat{q} = \alpha \mathbb{E} \left[ \frac{\sigma^2 \psi \text{erf } \chi}{(1 + v \sigma^2)^2} + \sigma^2 \delta^2 (1 - \text{erf } \chi) - \sqrt{\frac{2\psi}{\pi}} \frac{\sigma^2 \delta e^{-\chi^2}}{1 + v \sigma^2} \right] \quad (81b)$$

$$\hat{v} = \alpha \mathbb{E} \left[ \frac{\sigma^2 \text{erf } \chi}{1 + v \sigma^2} \right]. \quad (81c)$$

where the expectation is over all random variables involved in the expressions (namely,  $\sigma$ ,  $\hat{\sigma}$ , and  $\hat{\eta}$ ) and we used the short-hand notation

$$\psi := \hat{\sigma}^2 + \sigma^2(\hat{\eta}^2\beta_\star^2 - 2\hat{\eta}m + q), \quad \chi := \frac{\delta(1 + v\sigma^2)}{\sqrt{2\psi}} \quad (82)$$

Note that we recover the expressions obtained for the square loss for  $\delta \rightarrow +\infty$ .

With the usual notation convention  $\hat{\sigma}_0^2 := \mathbb{E}[\hat{\sigma}^2]$  and  $\sigma_0^2 := \mathbb{E}[\sigma^2]$ , the estimation error is given by the general formula in Eq. 57, whereas the generalisation error is

$$\varepsilon_g := \mathbb{E} \left[ (y - \hat{\beta}_\lambda^\top \mathbf{x})^2 \right] = \hat{\sigma}_0^2 + (\beta_\star^2 \mathbb{E}[\hat{\eta}^2] - 2\mathbb{E}[\hat{\eta}]m + q)\sigma_0^2, \quad (83)$$

$\varepsilon_g$  being finite if  $\sigma^2 < +\infty$ ,  $\hat{\sigma}_0^2 < +\infty$  and  $\mathbb{E}[\hat{\eta}] < +\infty$ . We aim now at extrapolating the large- $\alpha$  behavior of such errors and at studying the consistency of  $\hat{\beta}_\lambda$  with respect to the Bayes optimal estimators discussed in Section A.2. To do so, we rescale  $\hat{m} \mapsto \alpha\hat{m}$ ,  $\hat{v} \mapsto \alpha\hat{v}$ ,  $v \mapsto \alpha^{-1}v$  and  $\hat{q} \mapsto \alpha\hat{q}$ . We also assume that  $\lambda \mapsto \lambda + \alpha\lambda'$  (the role of  $\lambda' \neq 0$  will be clear in the following). The set of fixed point equations become, for  $\alpha \rightarrow +\infty$

$$\begin{aligned} \hat{m} &= \mathbb{E} \left[ \sigma^2 \hat{\eta} \operatorname{erf} \bar{\chi} \right] & q &= \frac{\beta_\star^2 \hat{m}^2}{(\lambda' + \hat{v})^2} \\ \hat{q} &= \mathbb{E} \left[ \sigma^2 \psi \operatorname{erf} \bar{\chi} + \sigma^2 \delta^2 (1 - \operatorname{erf} \bar{\chi}) - \sqrt{\frac{2\psi}{\pi}} \sigma^2 \delta e^{-\bar{\chi}^2} \right] & m &= \frac{\beta_\star^2 \hat{m}}{\lambda' + \hat{v}}, \quad \bar{\chi} := \frac{\delta}{\sqrt{2\psi}}. \\ \hat{v} &= \mathbb{E} \left[ \sigma^2 \operatorname{erf} \bar{\chi} \right]. & v &= \frac{1}{\lambda' + \hat{v}} \end{aligned} \quad (84)$$

In this limit, as  $\beta_\star^2 q = m^2$ ,  $\psi = \hat{\sigma}_0^2 + \frac{\sigma_0^2}{\beta_\star^2} (m - \beta_\star^2 \hat{\eta})^2$ , so that

$$\varepsilon_{\text{est}} = \frac{(m - \beta_\star^2)^2}{\beta_\star^2}, \quad \varepsilon_g = \hat{\sigma}_0^2 + \sigma^2 \frac{\mathbb{E}[(m - \hat{\eta}\beta_\star^2)^2]}{\beta_\star^2}. \quad (85)$$

It is possible to choose  $\lambda'$  so that  $\lim_{\alpha \rightarrow +\infty} \varepsilon_g = \lim_{\alpha \rightarrow +\infty} \varepsilon_g^{\text{BO}}$ , i.e.,

$$\hat{\sigma}_0^2 + \sigma_0^2 \frac{\mathbb{E}[(m - \hat{\eta}\beta_\star^2)^2]}{\beta_\star^2} = \hat{\sigma}_0^2 + \sigma_0^2 \beta_\star^2 \operatorname{Var}[\hat{\eta}] \Rightarrow m = \beta_\star^2 \mathbb{E}[\hat{\eta}]. \quad (86)$$

We can try to satisfy this condition by tuning properly  $\lambda'$ , under the constraint that  $\lambda' \geq 0$ . We can write in particular

$$\lambda' = \frac{\hat{m}}{\mathbb{E}[\hat{\eta}]} - \hat{v} = \frac{\mathbb{E}[\sigma^2 \hat{\eta} \operatorname{erf} \bar{\chi}] - \mathbb{E}[\hat{\eta}] \mathbb{E}[\sigma^2 \operatorname{erf} \bar{\chi}]}{\mathbb{E}[\hat{\eta}]} \geq 0 \Rightarrow \mathbb{E}[\sigma^2 (\hat{\eta} - \mathbb{E}[\hat{\eta}]) \operatorname{erf} \bar{\chi}] \geq 0 \quad (87)$$

to be computed with

$$\bar{\chi} \equiv \frac{\delta}{\sqrt{2\psi_g}}, \quad \psi_g = \hat{\sigma}_0^2 + \sigma_0^2 \beta_\star^2 (\mathbb{E}[\hat{\eta}] - \hat{\eta})^2. \quad (88)$$

Note that the condition is always satisfied in the case of the square loss (i.e., for  $\delta \rightarrow +\infty \Leftrightarrow \bar{\chi} \rightarrow 1$ ).

**Consistency of the estimator.** The consistency of the estimator can be imposed by properly tuning  $\lambda$ , by requiring that  $\lim_{\alpha \rightarrow +\infty} \varepsilon_{\text{est}} = 0$ , i.e.,  $m = \beta_\star^2$  in this limit. In the same spirit as above, this implies a condition on  $\lambda'$  given by

$$\lambda' = \hat{m} - \hat{v} = \mathbb{E}[\sigma^2 \hat{\eta} \operatorname{erf} \bar{\chi}] - \mathbb{E}[\sigma^2 \operatorname{erf} \bar{\chi}] \geq 0 \Rightarrow \mathbb{E}[\sigma^2 (\hat{\eta} - 1) \operatorname{erf} \bar{\chi}] \geq 0 \quad (89)$$

to be computed with

$$\bar{\chi} \equiv \frac{\delta}{\sqrt{2\psi_{\text{est}}}}, \quad \psi_{\text{est}} = \hat{\sigma}_0^2 + \sigma_0^2 \beta_\star^2 (1 - \hat{\eta})^2. \quad (90)$$

When imposing the equality, the conditions above provide the values of  $\delta$  (if any) for a consistent estimator if  $\lambda = \Theta(1)$  in the  $\alpha \rightarrow +\infty$  limit.

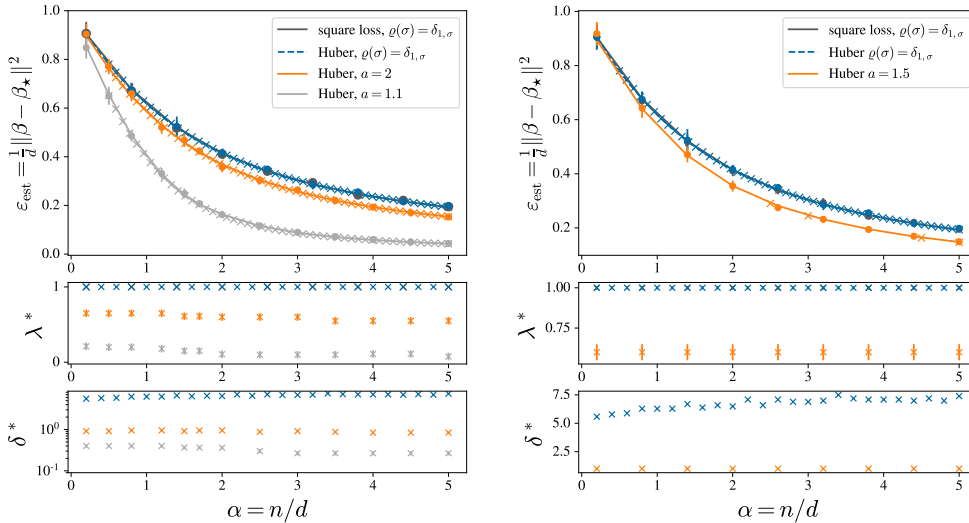


Figure 5: Gaussian covariates fully-contaminated by heavy-tailed noise with distribution  $p_\eta(\eta) = \mathbb{E}_\sigma[\mathcal{N}(x; 0, \sigma^2)]$ , where parameters of various  $\varrho(\sigma)$  are varied: inverse-gamma (see Table. 1, **left**) and Pareto (**right**). (**Top**) Estimation error  $\varepsilon_{\text{est}}$  as a function of the sample complexity  $\alpha = n/d$  for optimally regularised ridge regression (black), Huber with optimal location parameter and optimal regularisation (orange) and Bayes-optimal performance (crosses). (**Center.**) Value of the optimal regularisation parameter  $\lambda^*$  for the Huber loss. (**Bottom.**) Value of the optimal location parameter  $\delta^*$  for the Huber loss. Both optimal values are displayed by varying the scale parameter  $a$  controlling the tails of the noise distribution. Dots indicate numerical simulations averaged over 20 seeds with  $d = 10^3$ .

## C FURTHER NUMERICAL RESULTS

### C.1 FURTHER RESULTS FOR THE CASE OF HEAVY-TAILED NOISE

In this Appendix, we add some details about the case of  $\epsilon$ -contamination in the labels, as in Eq. 17, for different  $\varrho_0$  generating the contaminating noise. Figure 5 compares the performance of various losses for different fully-contaminated ( $\epsilon_n = 1$ ) label noise distributions, obtained picking for  $\varrho_0(\sigma)$  taken to be inverse-Gamma as in Table 1 with  $a = b + 1 > 1$  (left) or Pareto (right). In all cases, the chosen parametrisations enforce unit variance for the noise,  $\mathbb{E}[\eta^2] = 1$ . Taking the limit  $a \rightarrow \infty$  in the inverse-Gamma or in the Pareto distributions results in recovering the Gaussian distribution for label noise. In our experiments, we observe the same phenomenology as in Fig. 2 (bottom) for all these densities which generate different noise label distributions. As long as the label noise variance is kept the same, optimally tuned regularised Huber loss performs the task better with heavier tails, in terms of the estimation Eq. 10. As in the cases discussed in the main text, optimally regularised Huber achieves Bayes-optimal performance.

## D REBUTTAL PLOTS

This Appendix provides additional experiments requested by the reviewers.

### D.1 LEAST ABSOLUTE DEVIATION ESTIMATOR

We consider the LAD estimator corresponding to taking the  $\ell_1$  loss  $\ell_1(y, \eta) = |y - \eta|$  in Eq. 1b. Note this also corresponding to taking the location parameter  $\delta$  of the Huber loss in Eq. 2 to zero.

Fig. 6 compares the optimally regularised LAD estimator with with the following estimators:




Figure 6: Response contamination including results for the LAD estimator, exactly on the same setting as Fig. 2. **(Left)** Infinite variance contamination. **(Right)** Finite variance contamination. Dots are experiments with dimensionality  $d = 10^3$ , averaged over 20 runs.

- Optimally regularised ridge regression.
- Unregularised Huber with optimally chosen location parameter  $\delta$ .
- Optimally regularised Huber with optimally chosen location parameter  $\delta$

We consider three contamination cases  $\epsilon \in \{0, 0.5, 1\}$ , with both infinite and finite covariate variance. As it can be seen in Fig. 6, the LAD estimator performs similarly or worst to the other estimators in all cases considered.

## D.2 REVIEWER FBSX’S ALGORITHM

In this Appendix, we numerically compare the performance the performance of Reviewer Fbsx’s algorithm to standard ridge and Huber regression. The suggested algorithm consists of the following steps.

1. Given the training data  $\mathcal{D} = \{(\mathbf{x}_i, y_i)_{i \in [n]} \in \mathbb{R}^{d+1}\}$ , we uniformly sample a subset of size  $n/2$ , inducing a random partition of the training data  $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2$ .
2. From subset  $\mathcal{D}_1$  (wlog), we empirically estimate the median norm  $\eta$  of the covariates  $\mathbf{x}_i$ .
3. From subset  $\mathcal{D}_2$ , we select only covariates which have norm smaller than  $\eta$ , i.e., we extract  $\mathcal{D}_3 = \{(\mathbf{x}_i, y_i) \in \mathcal{D}_2: \|\mathbf{x}\|_2 \leq \eta\}$ .
4. Finally, we perform empirical risk minimisation Eq. 1b with either squared or Huber loss on  $\mathcal{D}_3$ .

Figure 7 numerically compares the performance of Reviewer Fbsx’s algorithm to standard ridge and Huber regression on the full data set. We consider covariate contamination in the presence of heavy-tailed response noise. The suggested algorithm performs worse than all measures of performance for both the contaminated  $\epsilon_c = 0.5$  (green) and fully-contaminated  $\epsilon_c = 1$  (black) cases. This is true for both the square and Huber losses, using the optimal parameters obtained from our analysis on the full samples.

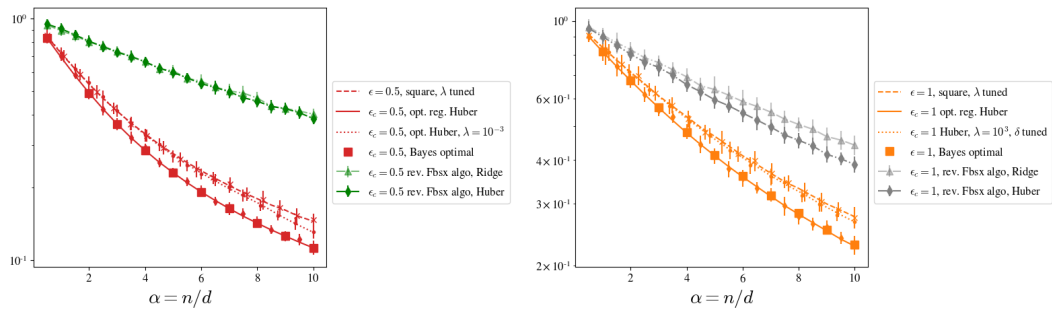


Figure 7: Comparison of the estimation error for vanilla empirical risk minimisation and Reviewer **Fbsx**'s thresholding algorithm. Covariate contamination with data whose variance is distributed as inverse-gamma variance with shape  $a = 1.1$  and scale  $b = 0.1$ , in the presence of heavy-tailed response noise with variance distribution as inverse-gamma with parameters  $a = 2, b = 1$ .