



Pedro Abdalla · Nikita Zhivotovskiy

Covariance estimation: Optimal dimension-free guarantees for adversarial corruption and heavy tails

Received 27 June 2022; revised 20 February 2023

Abstract. We provide an estimator of the covariance matrix that achieves the optimal rate of convergence (up to constant factors) in the operator norm under two standard notions of data contamination. We allow the adversary to corrupt an η -fraction of the sample arbitrarily, while the distribution of the remaining data points only satisfies that the L_p -marginal moment with some $p \geq 4$ is equivalent to the corresponding L_2 -marginal moment. Despite requiring the existence of only a few moments of the distribution, our estimator achieves the same tail estimates as if the underlying distribution were Gaussian. As a part of our analysis, we prove a non-asymptotic, dimension-free Bai–Yin type theorem in the regime $p > 4$.

Keywords: covariance estimation, heavy tails, adversarial contamination, Bai–Yin theorem.

1. Introduction

Estimation of the covariance matrix is a classic topic. In high-dimensional statistics, the role of sample covariance matrices is central to principal component analysis (PCA) and to linear least squares. Most of the existing work focuses on estimation of covariance matrices under different structural assumptions allowing minimax estimation in the high-dimensional setup. We refer to the line of work [6, 7, 10, 12, 24, 40] and the recent surveys [11, 25].

A line of research on non-asymptotic guarantees for sample covariance matrices was initiated by a question of Kannan, Lovász and Simonovits [34] on the computation of the volume of a convex body. For the class of log-concave measures, the optimal $\sqrt{\frac{d}{N}}$ (d is the dimension and N is the sample size) rate of convergence in the operator norm was first obtained in the renowned work of Adamczak, Litvak, Pajor and Tomczak-Jaegermann [1].

Pedro Abdalla: Department of Mathematics, IFOR, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland; pedro.abdallateixeira@ifor.math.ethz.ch

Nikita Zhivotovskiy: Department of Mathematics, IFOR, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland; Department of Statistics, University of California at Berkeley, 367 Evans Hall, Berkeley, CA 94720-3860, USA; zhivotovskiy@berkeley.edu

Mathematics Subject Classification 2020: 62G35 (primary); 62G15 (secondary).

Since then several authors focused on making less assumptions on the underlying distribution [29,51,52,62]. The best known result in this direction is due to Tikhomirov [65], who proved the optimal rate of convergence $\sqrt{\frac{d}{N}}$ for the sample covariance matrix assuming only the existence of $p > 4$ moments.

However, as discussed by Chen, Gao and Ren [16], if there exists only one outlier in the whole sample, the statistical performance of the sample covariance matrix can be compromised. Thus, one is interested in estimators of the covariance matrix robust to adversarial contamination of the data [17,55]. For a standard perspective on robustness, building on the ideas of contaminated models, influence functions and breakdown points, we refer to the monographs [30,33,61].

On the other hand, there is a growing interest in getting dimension-free guarantees for estimating the covariance matrix. Given a covariance matrix Σ , the *effective rank* (see [71]) of Σ is defined as

$$\mathbf{r}(\Sigma) = \frac{\text{Tr}(\Sigma)}{\|\Sigma\|},$$

where, for the rest of the paper, $\|\cdot\|$ denotes the operator norm of the matrix and the Euclidean norm of the vector. Koltchinskii and Lounici [39] proved the optimal high probability bound for the sample covariance matrix in the Gaussian case that depends on the effective rank rather than the dimension. Since then, their result was recovered and extended multiple times and via different techniques [38,68,72,73].

Finally, we mention the recent interest in getting the so-called *sub-Gaussian estimators* when the data is heavy-tailed. This direction was initiated by Catoni in [13], where the sub-Gaussian estimation of the mean of a random variable is considered. To explain informally, one aims to construct statistical estimators performing as well as the sample mean does for the Gaussian distribution, while making as weak assumptions on the distribution as possible. For a recent survey with focus on multivariate mean estimation, we refer to [44]. The central ideas behind the robust mean estimation found their applications in many related problems such as regression [9,19,31,46,49,56], covariance estimation [14,15,21,53,55,58], and clustering [36]. For related results in the context of covariance estimation for heavy-tailed distributions, we refer to the recent survey [35].

Our goal is to provide an estimator that simultaneously achieves all the properties described above:

- We allow the adversarial contamination and recover the optimal dependence on the contamination level based on the number of moments of the underlying distribution.
- Our bounds do not contain unnecessary logarithmic factors and the convergence rates coincide with the classical asymptotic result of Bai and Yin [4] provided that there are at least four moments of the distribution.
- The convergence rates scale with the effective rank $\mathbf{r}(\Sigma)$ rather than the dimension d .
- We allow the distributions satisfying certain weak norm equivalence assumptions instead of more restrictive Gaussian/log-concave assumptions appearing in the literature. At the same time, we provide the same high probability bounds as if the data were Gaussian.

We begin with the following definition. We say that the distribution of a zero mean random vector X satisfies the L_p - L_2 norm equivalence (hypercontractivity), if for all $v \in \mathbb{R}^d$ and $2 \leq q \leq p$,

$$(\mathbb{E}|\langle X, v \rangle|^q)^{1/q} \leq \kappa(q)(\mathbb{E}|\langle X, v \rangle|^2)^{1/2}, \quad (1)$$

where $\kappa(\cdot)$ is a function of q and $\langle \cdot, \cdot \rangle$ denotes the inner product. Without loss of generality, we assume that κ is a non-decreasing function. We say that the sample $\tilde{X}_1, \dots, \tilde{X}_N$ is η -corrupted if it is obtained from the sample X_1, \dots, X_N of independent copies of X by replacing at most ηN points by arbitrary vectors that might depend on X_1, \dots, X_N . This corruption model is described in detail [47] and captures the standard setups in robust statistics such as the Huber contamination model [32]. This model is sometimes called the model of η -corruption or the *strong contamination model* [22]. We discuss this model in more detail in Section 3.1. We present a simplified version of our main result.

Theorem 1.1 (Simplified). *Assume that X is a zero mean random vector in \mathbb{R}^d with covariance Σ satisfying L_p - L_2 norm equivalence with $p \geq 4$. Fix the corruption level $\eta \in [0, 1]$ and the confidence level $\delta \in (0, 1)$. Assume that $\tilde{X}_1, \dots, \tilde{X}_N$ is an η -corrupted sample. There is an estimator*

$$\hat{\Sigma}_{\eta, \delta} = \hat{\Sigma}_{\eta, \delta}(\tilde{X}_1, \dots, \tilde{X}_N)$$

depending on η, δ such that if

$$N \geq c(p) \left(\mathbf{r}(\Sigma) + \log\left(\frac{1}{\delta}\right) \right),$$

then with probability at least $1 - \delta$, it holds that

$$\|\hat{\Sigma}_{\eta, \delta} - \Sigma\| \leq C \|\Sigma\| \left(\sqrt{\frac{\mathbf{r}(\Sigma) + \log(1/\delta)}{N}} + \kappa(p)^2 \eta^{1-2/p} \right),$$

where $C > 0$ is an absolute constant that depends only on the value $\kappa(4)$, and $c(p)$ depends only on p and $\kappa(p)$. Moreover, under these assumptions, no estimator can perform better (up to multiplicative constant factors).

We show that the term $\kappa(p)^2 \eta^{1-2/p}$ scales as $\sqrt{\eta}$ when $p = 4$ and as $\eta \log(\frac{1}{\eta})$ for sub-Gaussian distributions. We also show that both rates cannot be improved. In particular, in the special case of variance estimation, which corresponds to $d = 1$, our bound matches the recent rates of Comminges, Collier, Ndaoud, and Tsybakov [20, Table 1], whose analysis covers a less general contamination model called the sparse vector model and only in-expectation bounds. A detailed version of Theorem 1.1 with the corresponding explicit estimators is stated in Theorem 3.2 in the regime $p = 4$ and in Theorem 4.3 in the regime $p > 4$, respectively. For a detailed discussion on the optimality of our result, we refer to Section 5. Since our estimator is robust to both heavy tails and adversarial corruption, it should necessarily differ from the sample covariance matrix. However, it still

has a relatively simple form. The estimator depends on a specifically tuned scalar $\lambda > 0$, which in turn depends on some parameters including δ and η (the details are postponed to Theorems 3.2 and 4.3). We define the truncation function

$$\psi(x) = \begin{cases} x & \text{for } x \in [-1, 1], \\ \text{sign}(x) & \text{for } |x| > 1. \end{cases} \tag{2}$$

Then, when $p = 4$, our estimator has the form

$$\widehat{\Sigma}_{\eta, \delta} = \operatorname{argmin}_{\widehat{\Sigma} \in \mathbb{S}_+^d} \sup_{v \in S^{d-1}} \left| \frac{1}{\lambda N} \sum_{i=1}^N \psi(\lambda \langle \widetilde{X}_i, v \rangle^2) - v^\top \widehat{\Sigma} v \right|, \tag{3}$$

where \mathbb{S}_+^d is the set of $d \times d$ positive semi-definite matrices and S^{d-1} denotes the unit sphere in \mathbb{R}^d . A simple observation is that, if $\lambda \rightarrow 0$, our estimator coincides with the sample covariance matrix $\frac{1}{N} \sum_{i=1}^N \widetilde{X}_i \otimes \widetilde{X}_i$. For a fixed value of λ , our estimator can be seen as the second-order extension of the classical trimmed mean estimator [5, 47, 63, 67]. We are aiming to remove a fraction of extreme observations and then average over the remaining sample. The only difference is that, in the matrix case, we need to take into account all possible directions in the unit sphere. The result of Theorem 1.1 has multiple advantages over the best known results in the literature:

- *Adversarial corruption*: One of the well-known results in the adversarial corruption setup is due to Chen, Gao and Ren [16]. They analyze a version of Tukey’s median under the restrictive assumption of elliptical distributions satisfying certain growth conditions. Moreover, their results are dimension-dependent. See [55] for some recent extensions.
- *Heavy tails*: The best known result in the literature in the heavy-tailed setup (the only assumption is L_4 – L_2 hypercontractivity) is due to Mendelson and the second-named author of this paper [53]. However, their results can be further improved. Firstly, there is an additional logarithmic factor $\log \mathbf{r}(\Sigma)$ due to the application of the non-commutative Bernstein inequality in the analysis.¹ Secondly, since their estimator is based on median-of-means, it is currently not known how to significantly improve the dependence on η beyond $\sqrt{\eta}$ in the bound of Theorem 1.1.

¹We note that under the L_4 – L_2 norm equivalence the logarithmic factor $\log \mathbf{r}(\Sigma)$ can be removed as implicitly follows from previous arguments. In particular, without adversarial contamination the technique of Catoni and Giulini [15, Proposition 4.1] can be adapted to recover the desired rate (see also the discussion in [27] and the bounds in [28] where at least a $\log \log N$ term appears); their estimator involves some explicit Gaussian integrals in the parameter space and will depend on additional parameters of the distribution such as $\operatorname{Tr}(\Sigma)$ and $\|\Sigma\|$.

Shortly after the first version of this paper was made public, the authors were notified that Rico and Oliveira achieved the same bound (only for $p = 4$) without the logarithmic factor in the setup without any contamination. Their result appeared later in the Ph.D. thesis of Rico [60] and in the preprint [57].

The most involved step is to get the correct dependence on the corruption level η when $p > 4$ in (1). One of the technical results used in this paper is a dimension-free version of the classical Bai–Yin theorem [4]. In a nutshell, the result of Bai and Yin implies that if X_1, \dots, X_N are independent copies of a zero mean random vector X in \mathbb{R}^d with unit covariance and independent identically distributed coordinates having four bounded moments, that is, $p = 4$, then almost surely as $d, N \rightarrow \infty$ so that $\frac{d}{N} \rightarrow \beta \in (0, 1]$,

$$\left\| \frac{1}{N} \sum_{i=1}^N X_i \otimes X_i - I_d \right\| \rightarrow 2\sqrt{\frac{d}{N}} + \frac{d}{N}.$$

Moreover, if the distribution only has $p < 4$ moments, no convergence at such a rate is possible. Our result is non-asymptotic and is somewhat stronger, since we do not require that the coordinates of X_i are independent. Moreover, our focus is on the dimension-free bound. At the same time, we require that $p > 4$ and the constant in the bound depends on how well $p - 4$ is separated from zero.

Theorem 1.2 (A non-asymptotic, dimension-free Bai–Yin type theorem). *Assume that Y is a zero mean random vector in \mathbb{R}^d with covariance Σ satisfying L_p – L_2 norm equivalence with $p > 4$. Let Y_1, \dots, Y_N be a sample of independent copies of Y . Consider the truncated vectors $X_i = Y_i \mathbb{1}_{\{\|Y_i\| \leq (N \operatorname{Tr}(\Sigma)\|\Sigma\|)^{1/4}\}}$ for $i = 1, \dots, N$. If $N \geq c(p)\mathbf{r}(\Sigma)$, then it holds that*

$$\mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N (X_i \otimes X_i - \mathbb{E} X_i \otimes X_i) \right\| \leq C(p)\|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}},$$

where $c(\cdot)$ and $C(\cdot)$ are non-increasing and both satisfy $C(p), c(p) \rightarrow \infty$ as $p \rightarrow 4$.

A natural attempt to prove this bound would be to apply the matrix Bernstein inequality (see the survey [66] for more details on matrix concentration inequalities) as in [53]. However, in this case, we would get an additional multiplicative $\log \mathbf{r}(\Sigma)$ -term that does not allow us to recover the optimal Bai–Yin rate of convergence in Theorem 1.1 for $p > 4$. Alternatively, one can try the generic chaining machinery. We refer to the monograph of Talagrand [64] for a detailed exposition of this method. For example, the analysis based on the generic chaining is known to recover the optimal dependence on the effective dimension in the sub-Gaussian case [39]. The problem with using the generic chaining in our case, is that we only assume the existence of $p > 4$ moments, which does not allow to control the increments of empirical processes with the required accuracy. The proof of Theorem 1.2 is based on the combination of the (dimension-dependent) approach developed by Tikhomirov [65] with the variational approach applied by the second-named author of this paper in [73].

Remark 1.3. Whether the bound of Theorem 1.2 holds for $p = 4$ is a well-known open problem (see [70] and the survey [71]) even in the isotropic case where $\Sigma = I_d$. The analysis of Theorem 1.1 bypasses the bound of Theorem 1.2 in the regime where $p = 4$. This makes our analysis different from the approach used in [53].

Remark 1.4. Theorem 1.2 can be used to remove the logarithm factor in the median-of-means estimator proposed in [53] for $p > 4$.

On the breakdown point of our estimator. To explain informally, the breakdown point of an estimator is defined as the largest proportion of outliers in the data for which the estimator gives a non-vacuous result [30]. The form of the bounds we are interested in is a high probability bound of the form $\|\widehat{\Sigma} - \Sigma\| \leq \varepsilon \|\Sigma\|$, where $\widehat{\Sigma}$ is our estimator and $\varepsilon > 0$ is precision parameter. Since $\widehat{\Sigma} = 0$ satisfies this inequality with $\varepsilon = 1$, in Theorem 1.1 we are focusing on $\eta \in [0, c]$, where c is a small enough constant so that for large enough sample size we guarantee $\varepsilon < 1$.

Practical considerations. Our main focus is on the statistical properties of the estimator. In fact, we are aiming to achieve the best possible statistical performance and we left computational questions aside. While there is some algorithmic progress in the case of adversarial contamination (see the survey [22]), in a recent paper [18], the authors provide evidence that achieving the same for heavy-tailed distributions is computationally hard at least when median-of-means estimators are used. For some practical estimators in the context of heavy-tailed covariance estimation, we refer to [35, 55, 59].

Technical overview. Our analysis connects three separate arguments in the literature. First, we prove a dimension-free version of the Bai–Yin theorem (Theorem 1.2) that allows us to eliminate unnecessary logarithmic factors in our bounds. The proof of this result consists of two parts. The analysis of the so-called *peaky part* is based on the arguments of Tikhomirov [65], which in itself improves the line of research [1, 8, 29, 52]. The dimension-free analysis of the *spread part* is based on the variational inequality techniques applied by the second-named author of this paper in [73]. The latter approach traces back to the works of Catoni and co-authors on robust mean estimation [3, 14, 15, 28]. Extending the results in [73], we prove Theorem 1.1 in the case where $p = 4$. Proposition 3.4 provides an optimal estimator for the largest eigenvalue of the covariance matrix and is based on Catoni’s estimator [13, 14]. Finally, the proof of Theorem 1.1 in the regime $p > 4$ is based on combining Theorem 1.2 with the second-order version of the multivariate trimmed mean estimator of Lugosi and Mendelson [47]. Our lower bounds are based on reducing to corresponding lower bounds in the multivariate mean estimation setup.

Structure of the paper. The rest of the paper is organized as follows. In Section 2, we present a proof of Theorem 1.2. In Section 3, we provide a proof Theorem 1.1 in the regime $p = 4$. In Section 3.3, we provide an auxiliary result on estimating the largest eigenvalue of the covariance matrix. Then, using Theorem 1.2 in Section 4, we give a proof Theorem 1.1 in the regime $p > 4$. We conclude with a detailed discussion of the lower bounds, showing the optimality of our results in Section 5.

Notation. Throughout the proofs, $C(p)$ and $c(p)$ will denote the constants depending only on p and (possibly) on $\kappa(p)$, where $\kappa(\cdot)$ is given by (1). The exact values of $C(p)$ and $c(p)$ may change from line to line. For an integer N , we set $[N] = \{1, \dots, N\}$. Let $\mathbb{1}\{A\}$ denote the indicator of an event A . The symbol \mathbb{R}_+ denotes the set of positive

reals. For any two functions (or random variables) f, g defined in some common domain, the notation $f \lesssim g$ means that there is an absolute constant c such that $f \leq cg$ and $f \sim g$ means that $f \lesssim g$ and $g \lesssim f$. For $(x_i)_{i=1}^m \in \mathbb{R}^m$, the sequence $(x_i^*)_{i=1}^m \in \mathbb{R}^m$ is a non-increasing rearrangement of $(|x_i|)_{i=1}^m$. For a set $I \subseteq [N]$, let $I^c = [N] \setminus I$. Let \mathbb{S}_+^d denote the set of $d \times d$ positive-definite matrices. The symbol $\|\cdot\|$ denotes the operator norm of a matrix or the Euclidean norm of the vector depending on the context. The symbol $\|a\|_0$ corresponds to the number of non-zero components of the vector a . For a random variable X , let $\|X\|_\infty$ denote its essential supremum. Let $\mathcal{KL}(\rho, \mu) = \int \log\left(\frac{d\rho}{d\mu}\right) d\rho$ denote the Kullback–Leibler divergence between a pair of measures ρ and μ .

2. Proof of Theorem 1.2

We start by proving the dimension-free version of the Bai–Yin theorem. Fix $\lambda > 0$. In the notation of Theorem 1.2, let us write the following decomposition:

$$\begin{aligned} \sup_{v \in S^{d-1}} \left| \frac{1}{N} \sum_{i=1}^N \langle X_i, v \rangle^2 - \mathbb{E} \langle X, v \rangle^2 \right| &\leq \underbrace{\sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, v \rangle^2 \mathbb{1}\{\lambda \langle X_i, v \rangle^2 > 1\}}_{\text{peaky part}} \\ &+ \underbrace{\sup_{v \in S^{d-1}} \left| \frac{1}{N\lambda} \sum_{i=1}^N \psi(\lambda \langle X_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right|}_{\text{spread part}}. \end{aligned}$$

The terminology comes from [64, Chapter 9.4]. Our analysis will consist of two steps, where we first analyze the *peaky part* and then analyze the *spread part*.

2.1. Dimension-free upper bound on the peaky part

To analyze this term, we follow the strategy of Tikhomirov in [65], which in itself improves a line of research [1, 8, 29, 52, 64]. Although this part of our proof follows their steps, we need several modifications to avoid any explicit dependence on the dimension. Following [65], given two sets $C, I \subset [N]$ and a positive integer k , we define the quantities $f(k, C)$, $g(k, C, I)$ and $W_{v,i}$ as

$$\begin{aligned} f(k, C) &= \sup_{\substack{\|y\|_2=1 \\ \|y\|_0 \leq k \\ \text{supp } y \subset C}} \left\| \sum_{i=1}^N y_i X_i \right\|^2, \\ g(k, C, I) &= \sup_{\substack{\|y\|_2=1 \\ \|y\|_0 \leq k}} \sup_{\substack{\|z\|_2=1 \\ \|z\|_0 \leq k}} \left\langle \sum_{i \in I \cap C} y_i X_i, \sum_{j \in I^c \cap C} z_j X_j \right\rangle, \\ W_{v,i} &= \left\langle X_i, \sum_{j=1}^N v_j X_j \right\rangle. \end{aligned} \tag{4}$$

In particular, as in [64, (9.140) and (9.142)], we have

$$\sup_{v \in S^{d-1}} \sup_{|I| \leq k} \sum_{i \in I} \langle v, X_i \rangle^2 = \sup_{|I| \leq k} \sup_{\sum_{i \in I} a_i^2 \leq 1} \left\| \sum_{i \in I} a_i X_i \right\|^2 = f(k, [N]). \tag{5}$$

We show below that, to control the peaky part, it is sufficient to upper bound the quantity $f(k, [N])$. The role of $g(K, C, I)$ goes as follows: Using a standard decoupling argument (see the derivation in [65, p. 11]), we obtain that

$$f(k, C) \leq \max_{i \leq N} \|X_i\|^2 + 2^{-N+2} \sum_{I \subset [N]} g(k, C, I).$$

Since $\max_{i \leq N} \|X_i\|^2 \leq (N \operatorname{Tr}(\Sigma) \|\Sigma\|)^{1/2}$ by assumption, we limit our attention on bounding $g(k, C, I)$. We use the following auxiliary bound. Its application should be compared with the application of [65, Lemma 2.6]. Due to the dimension-free nature of our bound, we do not capture the best possible dependence on p in the bound below and only exploit the case of $p = 4$. Despite that, this still allows us to prove the desired result.

Lemma 2.1. *Let $Y \in \mathbb{R}^d$ be a zero mean random vector with covariance Σ satisfying the L_p - L_2 norm equivalence for $p > 4$ with $\kappa(\cdot)$ in (1). Denote $X = Y \mathbb{1}\{\|Y\| \leq R\}$, where $R > 0$, and let $\kappa = \kappa(4)$. It holds that*

$$\mathbb{P}(\|X\|_2 \geq t) \leq \kappa^4 \frac{\operatorname{Tr}(\Sigma)^2}{t^4}.$$

Proof. We have

$$\mathbb{E}\|X\|_2^4 = \mathbb{E}\left(\sum_{i=1}^d \langle X_i, e_i \rangle^2\right)^2 = \mathbb{E}\sum_{i=1}^d \langle X, e_i \rangle^4 + \mathbb{E}\sum_{k \neq l} \langle X, e_k \rangle^2 \langle X, e_l \rangle^2.$$

We use that, for all $i \in [d]$,

$$(\mathbb{E}\langle X, e_i \rangle^4)^{1/4} \leq (\mathbb{E}\langle Y, e_i \rangle^4)^{1/4} \leq \kappa (\mathbb{E}\langle Y, e_i \rangle^2)^{1/2} = \kappa \Sigma_{ii}^{1/2}.$$

Now, by the above display and Hölder’s inequality we have

$$\begin{aligned} \kappa^4 \sum_{i=1}^d (\mathbb{E}\langle X, e_i \rangle^2)^2 + \kappa^4 \sum_{k \neq l} \mathbb{E}\langle X, e_k \rangle^2 \mathbb{E}\langle X, e_l \rangle^2 &= \kappa^4 \left(\sum_{j=1}^d \Sigma_{jj}^2 + \sum_{k \neq l} \Sigma_{kk} \Sigma_{ll} \right) \\ &= \kappa^4 \operatorname{Tr}(\Sigma)^2. \end{aligned}$$

The Markov inequality concludes the proof. ■

Following [65], the next step is to split the vectors X_1, \dots, X_N in several groups so that the vectors in each group are almost orthogonal. Fix $H > 0$. Consider the random graph $\mathcal{G}_H = ([N], E)$, where the set of edges is

$$E = \{(i, j) : 1 \leq i < j \leq N, |\langle X_i, X_j \rangle| > H \max_{k \leq N} \|X_k\|\}.$$

If we color this graph so that no two adjacent vertices share the same color, we have that the vertices of the same color correspond to a pair of vectors having a small inner product. Let $\chi(\mathcal{G}_H)$ denote the chromatic number of \mathcal{G}_H . We now induce the sets $\{\mathcal{C}_m^H\}_{m=1}^N$ performing the partition of $[N]$ such that $\mathcal{C}_m^H = \emptyset$ for $m > \chi(\mathcal{G}_H)$ and each \mathcal{C}_m^H contains the vertices of the same color. Following the proof of the dimension-dependent result, we provide a dimension-free bound on the chromatic number of this random graph.

Proposition 2.2. *Let $Y \in \mathbb{R}^d$ be a zero mean random vector with covariance Σ satisfying the L_p - L_2 norm equivalence for $p > 4$ with $\kappa(\cdot)$ as in (1). Define, for an arbitrary positive R , $X = Y \mathbb{1}\{\|Y\| \leq R\}$, and let $\kappa = \kappa(4)$. Then, for any $H > 0$ and any integer $m > 1$, with probability at least $1 - (\kappa(p)^p \|\Sigma\|^{p/2} N H^{-p})^{m-1} N \kappa(p)^4 \text{Tr}(\Sigma)^2 H^{-4}$,*

$$\chi(\mathcal{G}_H) \leq m.$$

Proof. We consider a greedy coloring process. Let $Y(1), Y(2), \dots$ be an auxiliary random process such that $Y(1) = 1$ and

$$Y(i) = \min\{r \in \mathbb{N} : \forall j < i \text{ with } Y(j) = r \text{ we have } |\langle X_i, X_j \rangle| \leq H \|X_j\|\}.$$

Note that, by definition, $\chi(\mathcal{G}_H) \leq \max_{i \in [N]} Y(i)$. Next, for each $i > 1$ and $m \geq 1$, we have

$$\begin{aligned} \mathbb{P}(Y(i) = m + 1) &\leq \mathbb{P}(\exists l \leq i - 1 \text{ such that } |\langle X_i, X_j \rangle| > H \|X_j\| \text{ and } Y(l) = m) \\ &\leq \sum_{l=1}^{i-1} \mathbb{P}(|\langle X_i, X_j \rangle| > H \|X_j\| \text{ and } Y(l) = m) \\ &= \sum_{l=1}^{i-1} \mathbb{P}(|\langle X_i, X_j \rangle| > H \|X_j\| : Y(l) = m) \mathbb{P}(Y(l) = m) \\ &\leq \kappa(p)^p \frac{\|\Sigma\|^{p/2}}{H^p} \sum_{l=1}^{i-1} \mathbb{P}(Y(l) = m) \\ &\leq \kappa(p)^p \frac{\|\Sigma\|^{p/2}}{H^p} \mathbb{E}\{j \leq N : Y(j) = m\}, \end{aligned}$$

where we used that since X_i is independent of $Y(l)$, and by Markov's inequality

$$\begin{aligned} &\mathbb{P}(|\langle X_i, X_j \rangle| > H \|X_j\| : Y(l) = m) \\ &= \mathbb{P}(|\langle X_i, X_j \rangle| > H \|X_j\| \text{ and } \|X_j\| \neq 0 : Y(l) = m) \\ &\leq \frac{\kappa(p)^p \|\Sigma\|^{p/2}}{H^p}. \end{aligned}$$

We obtain the following recursion:

$$\mathbb{E}\{j \leq N : Y(j) = m + 1\} \leq \kappa(p)^p \|\Sigma\|^{p/2} N \frac{1}{H^p} \mathbb{E}\{j \leq N : Y(j) = m\}.$$

Now we apply Lemma 2.1 and use the monotonicity of $\kappa(\cdot)$ to obtain

$$\mathbb{E}|\{j \leq N : Y(j) = 2\}| \leq N\mathbb{P}(\|X\| > H) \leq N\kappa(p)^4 \text{Tr}(\Sigma)^2 H^{-4}.$$

Combining the estimates above, we get

$$\mathbb{E}|\{j \leq N : Y(j) = m + 1\}| \leq (\kappa(p)^p \|\Sigma\|^{p/2} NH^{-p})^{m-1} N\kappa(p)^4 \text{Tr}(\Sigma)^2 H^{-4}.$$

Finally, we obtain

$$\begin{aligned} \mathbb{P}(\chi(\mathcal{G}_H) \geq m + 1) &\leq \mathbb{P}(\exists j \leq N : Y(j) = m + 1) \\ &\leq \mathbb{E}|\{j \leq N : Y(j) = m + 1\}| \\ &\leq (\kappa(p)^p \|\Sigma\|^{p/2} NH^{-p})^{m-1} N\kappa(p)^4 \text{Tr}(\Sigma)^2 H^{-4}. \quad \blacksquare \end{aligned}$$

We need the following result. It can be seen as one of the main ingredients of the proof and follows from the so-called *sparsifying* lemma [65, Lemma 4.1]. Our key observation here is that this lemma does not involve explicitly the dimension d . Recall that for $(x_i)_{i=1}^m \in \mathbb{R}^m$, the sequence $(x_i^*)_{i=1}^m \in \mathbb{R}^m$ is a non-increasing rearrangement of $(|x_i|)_{i=1}^m$ and v is an s -sparse vector if $\|v\|_0 \leq s$. We remark that the statement below is deterministic and holds for any realization of X_1, \dots, X_N .

Proposition 2.3 ([65, Proposition 4.4]). *There exists an absolute constant $C > 0$ such that the following holds. Fix $I \subseteq [N]$ and let $\gamma \in (0, \frac{1}{3})$ with $k \geq \frac{24}{\gamma^2}$ and $N \geq 128C\gamma^{-2}k$. Set $t = \lfloor \log_2 \frac{\gamma^2 k}{24} \rfloor$ and $k_j = \lfloor \frac{k}{2^j} \rfloor$ for $0 \leq j \leq t$. There are subsets $\mathcal{N}_j, \mathcal{N}'_j$ for $0 \leq j \leq t - 1$ supported on I and I^c , respectively, and consisting of unit γk_j -sparse vectors such that*

$$\max\{|\mathcal{N}_j, \mathcal{N}'_j|\} \leq \left(\frac{CN}{\gamma k_j}\right)^{2\gamma k_j},$$

and for any $\mathcal{C} \subseteq [N]$,

$$\begin{aligned} g(k, \mathcal{C}, I) &\leq \frac{C}{\gamma^2} \left(\log k \max_{i \neq j \in \mathcal{C}} |\langle X_i, X_j \rangle| \right. \\ &\quad \left. + \sum_{j=0}^{t-1} \sqrt{k_j} \left(\sup_{u \in \mathcal{N}_j} (A_u)_{\lfloor (k_{j+1})/16 \rfloor}^* + \sup_{v \in \mathcal{N}'_j} (B_v)_{\lfloor (k_{j+1})/16 \rfloor}^* \right) \right), \end{aligned}$$

where the random vectors $A_u \in \mathbb{R}^{|I^c|}$ and $B_v \in \mathbb{R}^{|I|}$ are given by $A_u = (|W_{u,i}|)_{i \in I^c}$ and $B_v = (|W_{v,i}|)_{i \in I}$, and $W_{u,i}, W_{v,i}$ are given by (4).

The next result is a dimension-free analog of [65, Proposition 5.1] in the regime $p > 4$.

Proposition 2.4. *There exists an absolute constant $C > 0$ such that if $I \subset [N]$ with $|I| \leq s$ is fixed and $\log \frac{N}{s} \geq C$, then simultaneously for all $\mathcal{C} \subset [N]$, with probability at least $1 - \frac{1}{N^3}$, $g(s, \mathcal{C}, I)$ is at most*

$$C \left(\log^2 \frac{N}{s} \log s \max_{i \neq j \in \mathcal{C}} |\langle X_i, X_j \rangle| + p\kappa(p) \|\Sigma\|^{1/2} \log^2 \frac{N}{s} \sqrt{s} \left(\frac{N}{s}\right)^{1/p} \sqrt{f(s, [N])} \right).$$

Proof. We provide the required changes of the original proof. First, we notice that if $s < C \log^2 \frac{N}{s}$, then

$$g(s, \mathcal{C}, I) \leq s \max_{i \neq j \in \mathcal{C}} |\langle X_i, X_j \rangle| < C \log^2 \frac{N}{s} \max_{i \neq j \in \mathcal{C}} |\langle X_i, X_j \rangle|,$$

and the claim trivially follows.

For the rest of the proof, we assume $s \geq C \log^2 \frac{N}{s}$. We define $\gamma = \frac{1}{\log(N/s)}$, $t = \lfloor \log_2 \frac{\gamma^2 s}{C} \rfloor$ and $k_j = \lfloor \frac{s}{2^j} \rfloor$. We choose $C > 24$ to be able to apply Proposition 2.3 with this choice of γ , t and k_j . For a fixed $0 \leq j < t$, consider $u \in \mathcal{N}_j$, where \mathcal{N}_j is given by Proposition 2.3. By its definition, u is supported on I . Therefore, for any $l \in I^c$ recalling (4), we have $W_{u,l} = \langle X_l, \sum_{i \in I} u_i X_i \rangle$. Observe that conditionally on the realization of $X_i, i \in I$, we have that $W_{u,l}$ are independent random variables for all $l \in I^c$. Therefore, by the norm equivalence assumption (1), we have

$$\mathbb{E}[\|W_{u,l}\|^p : X_i, i \in I] \leq \kappa(p)^p \|\Sigma\|^{p/2} \left\| \sum_{i \in I} u_i X_i \right\|^p \leq \kappa(p)^p \|\Sigma\|^{p/2} f(s, [N])^{p/2}.$$

For the same vector u and the same set I , consider the random vector

$$A_u = (\|W_{u,l}\|)_{l \in I^c}.$$

We apply the standard bound [65, Lemma 2.5] to control the coordinates of A_u with

$$\tau_j = (32e)^{1/p} \|\Sigma\|^{1/2} \kappa(p) \sqrt{f(s, [N])} \left(\frac{N}{k_{j+1}}\right)^{p^{-1}(1+256\gamma)}$$

to obtain

$$\mathbb{P}((A_u)_{\lfloor k_{j+1}/16 \rfloor}^* \geq \tau_j) \leq \left(\frac{e\kappa^p \|\Sigma\|^{p/2} f(s, [N])^{p/2} N}{\tau_j^p \lfloor k_{j+1}/16 \rfloor}\right)^{\lfloor k_{j+1}/16 \rfloor} \leq \left(\frac{k_{j+1}}{N}\right)^{4\gamma k_j}.$$

We union bound over the net \mathcal{N}_j , whose size is bounded in Proposition 2.3, to get for some absolute constant $C_1 > 0$,

$$\begin{aligned} \mathbb{P}\left(\sup_{u \in \mathcal{N}_j} (A_u)_{\lfloor k_{j+1}/16 \rfloor}^* \geq \tau_j\right) &\leq \left(\frac{k_{j+1}}{N}\right)^{4\gamma k_j} |\mathcal{N}_j| \leq \left(\frac{C_1 k_j}{\gamma N}\right)^{2\gamma k_j} \leq \left(\frac{k_j}{N}\right)^{\gamma k_j} \\ &\leq \left(\frac{k_t}{N}\right)^{\gamma k_t} \leq \frac{1}{N^4}, \end{aligned}$$

where we assumed $N \geq C_1^2 \gamma^{-2} k_j$. Repeating the same arguments for \mathcal{N}'_j and summing over all j in Proposition 2.3, we obtain, with probability at least $1 - \frac{1}{N^3}$,

$$g(s, \mathcal{C}, I) \lesssim \gamma^{-2} \left(\log s \max_{i \neq j \in \mathcal{C}} |\langle X_i, X_j \rangle| + \sum_{j=0}^{t-1} \tau_j \sqrt{k_j} \right).$$

The sum $\sum_{j=0}^{t-1} \tau_j \sqrt{k_j}$ can be upper bounded exactly the same way as in the end of the proof of [65, Proposition 5.1]. Using $p > 4$, we conclude the proof. ■

The next result applies Proposition 2.4 to upper bound $f(s, \mathcal{C}_m^H)$, where \mathcal{C}_m^H is a class of coloring of the sample.

Lemma 2.5. *Let $s \leq N$ be fixed and assume that $\log \frac{N}{s} \geq C_1$, where $C_1 > 0$ is a large enough absolute constant. Let \mathcal{C}_m^H be the class from the coloring of the sample X_1, \dots, X_N with threshold $H > 0$. Then there exists an absolute constant $C > 0$ such that, with probability at least $1 - \frac{1}{N^2}$, $f(s, \mathcal{C}_m^H)$ is at most*

$$\max_{i \leq N} \|X_i\|^2 + C \log^2 \frac{N}{s} \left(H \log s \max_{i \leq N} \|X_i\| + p\kappa \|\Sigma\|^{1/2} \sqrt{s} \left(\frac{N}{s} \right)^{1/p} \sqrt{f(s, [N])} \right).$$

Proof. The proof follows exactly the same steps (an application of Proposition 2.4 in our case) as the proof of [65, Lemma 5.2] by replacing d (their notation for the dimension is n) by s . ■

Our final step is to upper bound the desired function $f(s, [N])$ using the upper bound on $f(s, \mathcal{C}_m^H)$ for different coloring classes. Our new observation is that since we are only interested in the regime $p > 4$, we can worsen the dependence on some logarithmic factors appearing in the corresponding result in [65]. This allows us to build our analysis on a somewhat weaker Lemma 2.1 (compare it with a dimension-dependent result in [65, Lemma 2.6]).

Lemma 2.6. *Let s, N satisfy $\log \frac{N}{s} \geq C_1$, where $C_1 > 0$ is a large enough absolute constant. Then there exists $\chi = \chi(p)$ that depends only on p such that $f(s, [N])$ is at most*

$$\chi^2 \max_{i \leq N} \|X_i\|^2 + \kappa^2 \|\Sigma\| s \left(\frac{N}{s} \right)^{4/(4+p)} \log^4 \frac{N}{s} + \chi^2 p^2 \kappa(p)^2 \|\Sigma\| s \log^4 \frac{N}{s} \left(\frac{N}{s} \right)^{2/p},$$

with probability at least

$$1 - N^{(4-p)(\chi-2)/(p+4)} s^{-(p^2(\chi-1)+4p)/(2p+8)} (\log s)^{p(\chi-1)+4} \mathbf{r}^2(\Sigma) - \frac{\chi}{N^2}.$$

Proof. We define $\chi(\mathcal{G}_H)$ to be the chromatic number of the random graph induced by the partition of the sample X_1, \dots, X_N into the classes \mathcal{C}_m^H . Fix

$$H = \kappa \|\Sigma\|^{1/2} \left(\frac{N}{s} \right)^{2/(4+p)} \frac{\sqrt{s}}{\log s},$$

where $\kappa = \kappa(p)$, and apply Lemma 2.5 together with union bound to obtain that, with probability at least $1 - \frac{\chi}{N^2}$,

$$\begin{aligned} \frac{1}{\chi} \sum_{m=1}^{\chi} f(s, \mathcal{C}_m^H) &\leq \max_{i \leq N} \|X_i\|^2 + C \log^2 \frac{N}{s} \left(H \log s \max_{i \leq N} \|X_i\| \right. \\ &\quad \left. + p\kappa \|\Sigma\|^{1/2} \sqrt{s} \left(\frac{N}{s} \right)^{1/p} \sqrt{f(s, [N])} \right). \end{aligned}$$

We now apply the dimension-free version of the coloring Proposition 2.2 to obtain that the chromatic number $\chi(\mathcal{E}_H)$ is at most χ , with probability at least

$$1 - (\kappa^p \|\Sigma\|^{p/2} N H^{-p})^{\chi-1} N \kappa^4 \text{Tr}(\Sigma)^2 H^{-4}.$$

Now, we estimate the latter probability. We have

$$\begin{aligned} & (\kappa^p \|\Sigma\|^{p/2} N H^{-p})^{\chi-1} N \kappa^4 \text{Tr}(\Sigma)^2 H^{-4} \\ &= (N^{1-2p/(p+4)} s^{2p/(p+4)-p/2})^{\chi-1} N^{1-8/(4+p)} s^{8/(4+p)-2} (\log s)^{p(\chi-1)+4} \mathbf{r}^2(\Sigma) \\ &= N^{(4-p)(\chi-2)/(p+4)} s^{-(p^2(\chi-1)+4p)/(2p+8)} (\log s)^{p(\chi-1)+4} \mathbf{r}^2(\Sigma). \end{aligned}$$

Following exactly the lines of the proof of [65, Proposition 5.3], we obtain

$$\begin{aligned} \chi^{-1}(\mathcal{E}_{\mathcal{H}}) f(s, [N]) &\leq \chi^{-1}(\mathcal{E}_{\mathcal{H}}) \sum_{m=1}^{\chi(\mathcal{E}_H)} f(s, \mathcal{C}_m^H) \\ &\leq \max_{i \leq N} \|X_i\|^2 + C \kappa \|\Sigma\|^{1/2} \left(\frac{N}{s}\right)^{2/(4+p)} \sqrt{s} \log^2 \frac{N}{s} \max_{i \leq N} \|X_i\| \\ &\quad + C p \kappa \|\Sigma\| \log^2 \frac{N}{s} \sqrt{s} \left(\frac{N}{s}\right)^{1/p} \sqrt{f(s, [N])}, \end{aligned}$$

with the probability in the statement. Now, we solve the inequality using

$$2ab \leq \gamma a^2 + \frac{b^2}{\gamma}$$

for $a, b, \gamma \geq 0$ twice and solving with respect to $f(s, [N])$. ■

The following bound is the main result of this section.

Theorem 2.7. *Assume that for some large enough absolute constant $c > 0$, it holds that $N \geq c\mathbf{r}(\mathbf{E})$. For large enough sample size N , simultaneously for all integers s satisfying $\mathbf{r}(\Sigma) \leq s \leq \frac{N}{c}$, with probability at least $1 - \frac{c(p)}{N}$, it holds that*

$$f(s, [N]) \leq C(p) \left(\max_{i \leq N} \|X_i\|^2 + \|\Sigma\| s \left(\frac{N}{s}\right)^{4/(4+p)} \log^4 \frac{N}{s} \right), \quad (6)$$

where $C(p)$ and $c(p)$ depend only on p and $\kappa(p)$.

Proof. The proof is based on the application of Lemma 2.6. Let $s \geq \mathbf{r}(\Sigma)$ and fix $\chi = \max(10, \frac{4p}{p-4})$. Consider two cases:

- If $\mathbf{r}(\Sigma) \leq 100$, we have

$$N^{(4-p)(\chi-2)/(p+4)} s^{(p^2(\chi-1)+4p)/(2p+8)} (\log s)^{p(\chi-1)+4} \mathbf{r}^2(\Sigma) \leq c(p) N^{-2}.$$

- Otherwise, if $\mathbf{r}(\Sigma) > 100$, then we have $s \geq 100$ and $\log s \leq s^{1/3}$. Therefore,

$$\begin{aligned} & N^{(4-p)(\chi-2)/(p+4)} s^{-(p^2(\chi-1)+4p)/(2p+8)+(p(\chi-1)+4)/3} \mathbf{r}^2(\Sigma) \\ &\leq s^{-2} \mathbf{r}^2(\Sigma) N^{-2} \leq N^{-2}. \end{aligned}$$

Therefore, for a fixed integer s satisfying $\mathbf{r}(\Sigma) \leq s \leq \frac{N}{c}$, we have that $\log \frac{N}{s}$ is sufficiently large so that, by Lemma 2.6 and the above calculations, with probability at least $1 - \frac{c(p)}{N^2}$,

$$f(s, [N]) \leq C(p) \left(\max_{i \leq N} \|X_i\|^2 + \|\Sigma\|_s \left(\frac{N}{s} \right)^{4/(4+p)} \log^4 \frac{N}{s} \right).$$

By taking the union bound with respect to the value of s , we conclude the proof. ■

2.2. Analysis of the spread part

The main tool of this section is the following lemma sometimes called the PAC-Bayesian inequality (see [15, Proposition 2.1] or [2] for a detailed proof taking care of measurability questions). In our proofs, we adapt several computations appearing in [73]. We also refer to [73] for a detailed exposition of related results.

Lemma 2.8. *Assume that $X_i, i = 1, \dots, N$, are i.i.d. random variables defined on some measurable space. Assume also that Θ (called the parameter space) is a subset of \mathbb{R}^d for some $d \geq 1$. Let μ be a distribution (called prior) on Θ and let ρ be any distribution (called posterior) on Θ such that $\rho \ll \mu$. Then, simultaneously for any such ρ we have, with probability at least $1 - \exp(-t)$,*

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}_\rho f(X_i, \theta) \leq \mathbb{E}_\rho \log(\mathbb{E}_X \exp(f(X, \theta))) + \frac{\mathcal{KL}(\rho, \mu) + t}{N}.$$

Here θ is distributed according to ρ . Moreover,

$$\mathbb{E} \sup_{\rho} \left(\sum_{i=1}^N \mathbb{E}_\rho f(X_i, \theta) - N \mathbb{E}_\rho \log(\mathbb{E}_X \exp(f(X, \theta))) - \mathcal{KL}(\rho, \mu) \right) \leq 0.$$

Our analysis will also exploit the following elementary relation.

Lemma 2.9 ([73, Lemma 4]). *Let the truncation function ψ be given by (2) and let Z be a square integrable random variable. We have*

$$\psi(\mathbb{E}Z) \leq \mathbb{E} \log(1 + Z + Z^2) + \min \left\{ 1, \frac{\mathbb{E}Z^2}{6} \right\}.$$

Moreover, for any $a > 0$, it holds that

$$\begin{aligned} & \mathbb{E} \log(1 + Z + Z^2) + a \mathbb{E} \min \left\{ 1, \frac{Z^2}{6} \right\} \\ & \leq \mathbb{E} \log \left(1 + Z + \left(1 + \frac{(7 + \sqrt{6})(\exp(a) - 1)}{6} \right) Z^2 \right). \end{aligned}$$

The main result of this section is the following.

Proposition 2.10. *Assume that Y is a zero mean random vector with covariance Σ satisfying the L_p - L_2 norm equivalence with $p \geq 4$. Let Y_1, \dots, Y_N be a sample of independent copies of Y . Consider the truncated vectors $X_i = Y_i \mathbb{1}\{\|Y_i\| \leq R\}$ for $i = 1, \dots, N$ and some $R > 0$. For a fixed truncation level $\lambda > 0$, it holds that*

$$\mathbb{E} \sup_{v \in S^{d-1}} \left| \frac{1}{N\lambda} \sum_{i=1}^N \psi(\lambda \langle X_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| \lesssim \frac{\mathbf{r}(\Sigma)}{\lambda N} + \lambda \kappa^4 \|\Sigma\|^2,$$

where $\kappa = \kappa(4)$. In particular, when $\lambda = \frac{1}{\kappa^2 \|\Sigma\|} \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}$, we have

$$\mathbb{E} \sup_{v \in S^{d-1}} \left| \frac{1}{N\lambda} \sum_{i=1}^N \psi(\lambda \langle X_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| \lesssim \kappa^2 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}.$$

Proof. Our aim is to choose the distributions μ and ρ in Lemma 2.8. Let $\Theta = (\mathbb{R}^d)^2$. We choose μ to be a product of two zero mean multivariate Gaussians with covariance $\beta^{-1} I_d$. For $v \in S^{d-1}$, let ρ_v be a product of two multivariate Gaussian distribution with mean v and covariance $\beta^{-1} I_d$. Because of this, if (θ, v) is distributed according to ρ_v , we have $\mathbb{E}_{\rho_v}(\theta, v) = (v, v)$. By the additivity of \mathcal{KL} -divergence for product measures and the standard formula, we have

$$\mathcal{KL}(\rho_v, \mu) = \beta.$$

By the first part of Lemma 2.9, we have

$$\begin{aligned} \psi(\lambda \langle X, v \rangle^2) &= \psi(\lambda \mathbb{E}_{\rho_v} \langle X, \theta \rangle \langle X, v \rangle) \\ &\leq \mathbb{E}_{\rho_v} \log(1 + \lambda \langle X, \theta \rangle \langle X, v \rangle + \lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2) \\ &\quad + \min\left\{1, \frac{\lambda^2 \mathbb{E}_{\rho_v} (\langle X, \theta \rangle \langle X, v \rangle)^2}{6}\right\}. \end{aligned}$$

Observe that

$$\mathbb{E}_{\rho_v} (\langle X, \theta \rangle \langle X, v \rangle)^2 = (\langle X, v \rangle^2 + \beta^{-1} \|X\|^2)^2 \leq 2\langle X, v \rangle^4 + 2\beta^{-2} \|X\|^4, \quad (7)$$

and we can write

$$\min\left\{1, \frac{\lambda^2 \mathbb{E}_{\rho_v} (\langle X, \theta \rangle \langle X, v \rangle)^2}{6}\right\} \leq \min\left\{1, \frac{2\lambda^2 \langle X, v \rangle^4}{6}\right\} + \min\left\{1, \frac{2\lambda^2 \beta^{-2} \|X\|^4}{6}\right\}.$$

Conditionally on X the distribution of $\langle X, \theta \rangle$ is Gaussian with mean $\langle X, v \rangle$. Since it is symmetric, we have that

$$\Pr_{\rho_v}(\langle X, \theta \rangle^2 \geq \langle X, v \rangle^2) \geq \frac{1}{4},$$

and this holds trivially when $X = 0$. Therefore,

$$\min\left\{1, \frac{2\lambda^2 \langle X, v \rangle^4}{6}\right\} \leq 8 \mathbb{E}_{\rho_v} \min\left\{1, \frac{\lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2}{6}\right\}.$$

By the second part of Lemma 2.9, we have for some absolute constant $c > 0$,

$$\begin{aligned} \mathbb{E}_{\rho_v} \log(1 + \lambda \langle X, \theta \rangle \langle X, v \rangle + \lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2) + 8 \mathbb{E}_{\rho_v} \min \left\{ 1, \frac{\lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2}{6} \right\} \\ \leq \mathbb{E}_{\rho_v} \log(1 + \lambda \langle X, \theta \rangle \langle X, v \rangle + c \lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2). \end{aligned}$$

Following the proof of Lemma 2.1, we get

$$\mathbb{E} \|X\|^4 \leq \mathbb{E} \|Y\|^4 \leq \kappa^4 (\text{Tr}(\Sigma))^2 \quad \text{and} \quad \mathbb{E} \langle X, v \rangle^2 \leq \mathbb{E} \langle Y, v \rangle^2 \leq \|\Sigma\|, \quad (8)$$

where $v \in S^{d-1}$. Using $\log(1 + y) \leq y$ for $y \geq -1$ and Fubini's theorem, we have

$$\begin{aligned} \mathbb{E}_{\rho_v} \log \mathbb{E}(1 + \lambda \langle X, \theta \rangle \langle X, v \rangle + c \lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2) \\ \leq \lambda \mathbb{E}_{\rho_v} \mathbb{E} \langle X, \theta \rangle \langle X, v \rangle + c \lambda^2 \mathbb{E} \mathbb{E}_{\rho_v} \langle X, \theta \rangle^2 \langle X, v \rangle^2 \\ \leq \lambda \mathbb{E} \langle X, v \rangle^2 + 2c \lambda^2 (\mathbb{E} \langle X, v \rangle^4 + \mathbb{E} \beta^{-2} \|X\|^4) \end{aligned} \quad (\text{by (7)})$$

$$\leq \lambda \mathbb{E} \langle X, v \rangle^2 + 2c \lambda^2 \kappa^4 ((\mathbb{E} \langle X, v \rangle^2)^2 + \beta^{-2} (\text{Tr}(\Sigma))^2) \quad (\text{by (8)})$$

$$\leq \lambda \mathbb{E} \langle X, v \rangle^2 + 2c \lambda^2 \kappa^4 (\|\Sigma\|^2 + \beta^{-2} (\text{Tr}(\Sigma))^2).$$

We plug

$$f(X, \theta, v) = \log(1 + \lambda \langle X, \theta \rangle \langle X, v \rangle + c \lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2)$$

in the second part of Lemma 2.8. Choosing $\beta = \mathbf{r}(\Sigma)$ and dividing both sides by N , we have

$$\mathbb{E} \sup_{v \in S^{d-1} \cup \{0\}} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\rho_v} f(X_i, \theta, v) - \lambda \mathbb{E} \langle X, v \rangle^2 \right) \leq \frac{\mathbf{r}(\Sigma)}{N} + 4c \lambda^2 \kappa^4 \|\Sigma\|^2, \quad (9)$$

where we added the 0 vector by considering μ as a posterior distribution and observing that $0 = \mathcal{KL}(\mu, \mu) \leq \beta$. By adding the 0 vector, we guarantee that the supremum in (9) is always non-negative. Adding the term

$$\frac{1}{N} \mathbb{E} \sum_{i=1}^N \min \left\{ 1, \frac{2\lambda^2 (\mathbf{r}(\Sigma))^{-2} \|X_i\|^4}{6} \right\} \leq \lambda^2 \frac{\kappa^4 \|\Sigma\|^2}{3}$$

to inequality (9) and using the derivations at the beginning of the proof, we have

$$\mathbb{E} \sup_{v \in S^{d-1} \cup \{0\}} \left(\frac{1}{N} \sum_{i=1}^N \psi(\lambda \langle X, v \rangle^2) - \lambda \mathbb{E} \langle X, v \rangle^2 \right) \leq \frac{\mathbf{r}(\Sigma)}{N} + \left(4c + \frac{1}{3} \right) \lambda^2 \kappa^4 \|\Sigma\|^2.$$

The same argument works for $-\lambda$ instead. This leads to

$$\mathbb{E} \sup_{v \in S^{d-1}} \left| \frac{1}{\lambda N} \sum_{i=1}^N \psi(\lambda \langle X, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| \leq 2 \frac{\mathbf{r}(\Sigma)}{\lambda N} + 2 \left(4c + \frac{1}{3} \right) \lambda \kappa^4 \|\Sigma\|^2.$$

Our choice of λ concludes the proof. ■

2.3. Combining peaky and spread parts

We conclude the proof of Theorem 1.2.

Proof Theorem 1.2. Using our decomposition, we have

$$\begin{aligned} \mathbb{E} \sup_{v \in \mathcal{S}^{d-1}} \left| \frac{1}{N} \sum_{i=1}^N \langle X_i, v \rangle^2 - \mathbb{E} \langle X, v \rangle^2 \right| &\leq \mathbb{E} \sup_{v \in \mathcal{S}^{d-1}} \left| \frac{1}{N\lambda} \sum_{i=1}^N \psi(\lambda \langle X_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| \\ &\quad + \mathbb{E} \sup_{v \in \mathcal{S}^{d-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, v \rangle^2 \mathbb{1}\{\lambda \langle X_i, v \rangle^2 > 1\}. \end{aligned}$$

We choose $\lambda = \frac{1}{\kappa(4)^2 \|\Sigma\|} \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}$. By Proposition 2.10, we have

$$\mathbb{E} \sup_{v \in \mathcal{S}^{d-1}} \left| \frac{1}{N\lambda} \sum_{i=1}^N \psi(\lambda \langle X_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| \lesssim \kappa(4)^2 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}.$$

We proceed with the remaining term. Define the random set

$$I_v = \{i \in [N] : \langle X_i, v \rangle^2 > \lambda^{-1}\}.$$

Let $m = \sup_{v \in \mathcal{S}^{d-1}} |I_v|$. By (5), we have

$$\frac{m}{N\lambda} \leq \sup_{v \in \mathcal{S}^{d-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, v \rangle^2 \mathbb{1}\{\lambda \langle X_i, v \rangle^2 > 1\} \leq \frac{1}{N} f(m, [N]). \tag{10}$$

Observe that for any $p > 4$, there is $C(p) > 0$ such that for all $x \geq 1$,

$$x^{1-4/(4+p)} \log^4 x \leq C(p) \sqrt{x}.$$

We now want to apply Theorem 2.7 with the (random) value m . First, we assume that $\mathbf{r}(\Sigma) \leq m \leq \frac{N}{c}$. In this case, by Theorem 2.7, with probability at least $1 - \frac{c(p)}{N}$, it holds that

$$\frac{1}{N} f(m, [N]) \leq C(p) \left(\frac{1}{N} \max_{i \leq N} \|X_i\|^2 + \|\Sigma\| \left(\frac{m}{N} \right)^{1/2} \right).$$

By (10) and since $\|X_i\| \leq (N \operatorname{Tr}(\Sigma) \|\Sigma\|)^{1/4}$, we have on the same event

$$\begin{aligned} m &\leq C(p) \lambda \left(\max_{i \leq N} \|X_i\|^2 + \sqrt{mN} \|\Sigma\| \right) \\ &\leq C(p) \frac{1}{\|\Sigma\|} \sqrt{\frac{\mathbf{r}(\Sigma)}{N}} (\sqrt{mN} \|\Sigma\| + \sqrt{\operatorname{Tr}(\Sigma) \|\Sigma\| N}). \end{aligned}$$

Solving the inequality above with respect to m , we have on the same event

$$m \leq C(p) \mathbf{r}(\Sigma).$$

We consider the case where m does not satisfy $\mathbf{r}(\Sigma) \leq m \leq \frac{N}{c}$. If $m < \mathbf{r}(\Sigma)$, then we recover the same upper bound as above. If $m > \frac{N}{c}$, then on the event of Theorem 2.7 for $k = \lfloor \frac{N}{c} - 1 \rfloor$ by (10) and monotonicity, we have

$$\frac{k}{\lambda} \leq f(k, [N]) \leq C(p)(\sqrt{N \operatorname{Tr}(\Sigma)\|\Sigma\|} + \|\Sigma\|(kN)^{1/2}).$$

For our choice of λ , this bound leads to contradiction if $N \geq c(p)\mathbf{r}(\Sigma)$ for large enough $c(p)$. We are now ready to plug our bound

$$m \leq C(p)\mathbf{r}(\Sigma)$$

in (10) to obtain that, with probability at least $1 - \frac{c(p)}{N}$,

$$\frac{1}{N} f(C(p)\mathbf{r}(\Sigma), [N]) \leq C(p)\|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}.$$

Finally, for any $v \in S^{d-1}$, we have

$$|\langle X_i, v \rangle| \leq \|X_i\| \leq (N \operatorname{Tr}(\Sigma)\|\Sigma\|)^{1/4}.$$

Therefore,

$$\begin{aligned} \mathbb{E} \sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N \langle X_i, v \rangle^2 \mathbb{1}\{\lambda \langle X_i, v \rangle^2 > 1\} &\leq C(p)\|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}} + \frac{c(p)}{N} \sqrt{N \operatorname{Tr}(\Sigma)\|\Sigma\|} \\ &= (C(p) + c(p))\|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}. \end{aligned}$$

The claim of Theorem 1.2 follows. ■

3. Proof of Theorem 1.1 in the regime $p = 4$

First, we provide a high probability version of Proposition 2.10. In the definition of our estimator, we assume that both $\|\Sigma\|$ and $\operatorname{Tr}(\Sigma)$ are known, so that our estimator depends on such quantities. In Section 3.1, we provide optimal guarantees for estimating these parameters. The following result can be seen as a special case of [73, Lemma 5]. We provide a detailed sketch of the proof for the sake of completeness.

Proposition 3.1. *Assume that X is a zero mean random vector with covariance Σ satisfying the L_4 – L_2 norm equivalence. Let X_1, \dots, X_N be a sample of independent copies of X . For a fixed truncation level $\lambda > 0$, with probability at least $1 - \delta$, it holds that*

$$\sup_{v \in S^{d-1}} \left| \frac{1}{N\lambda} \sum_{i=1}^N \psi(\lambda \langle X_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| \lesssim \frac{\mathbf{r}(\Sigma) + \log(1/\delta)}{\lambda N} + \lambda \kappa^4 \|\Sigma\|^2,$$

where $\kappa = \kappa(4)$.

Proof. We repeat the lines of the proof of Proposition 2.10. However, instead we plug

$$f(X, \theta, v) = \log(1 + \lambda \langle X, \theta \rangle \langle X, v \rangle + c \lambda^2 (\langle X, \theta \rangle \langle X, v \rangle)^2)$$

in the first part of Lemma 2.8. This implies that, with probability at least $1 - \delta$,

$$\sup_{v \in S^{d-1}} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\rho_v} f(X_i, \theta, v) - \lambda \mathbb{E} \langle X, v \rangle^2 \right) \leq \frac{\mathbf{r}(\Sigma) + \log(1/\delta)}{N} + 4c \lambda^2 \kappa^4 \|\Sigma\|^2, \quad (11)$$

where $c > 0$ is the same constant as in the proof of Proposition 2.10. Furthermore, by the Bernstein inequality we have

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \min \left\{ 1, \frac{\lambda^2 (\mathbf{r}(\Sigma))^{-2} \|X_i\|^4}{3} \right\} &\leq \mathbb{E} \min \left\{ 1, \frac{2\lambda^2 (\mathbf{r}(\Sigma))^{-2} \|X\|^4}{6} \right\} \\ &\quad + \sqrt{\frac{2 \log(1/\delta)}{N} \mathbb{E} \min \left\{ 1, \frac{2\lambda^2 (\mathbf{r}(\Sigma))^{-2} \|X\|^4}{6} \right\}} \\ &\quad + \frac{2 \log(1/\delta)}{3N}, \end{aligned}$$

where we used that each summand belongs to the interval $[0, 1]$ and therefore the variance of each summand is bounded by its expectation. Following the proof of Lemma 2.1, we have $\mathbb{E} \|X\|^4 \leq \kappa^4 (\text{Tr}(\Sigma))^2$. This implies that, with probability at least $1 - \delta$,

$$\frac{1}{N} \sum_{i=1}^N \min \left\{ 1, \frac{2\lambda^2 (\mathbf{r}(\Sigma))^{-2} \|X_i\|^4}{6} \right\} \leq \frac{2\lambda^2 \kappa^4 \|\Sigma\|^2}{3} + \frac{3 \log(1/\delta)}{N}. \quad (12)$$

Using this line and the union bound together with (11) and the derivations in the proof of Proposition 2.10, we obtain the one-sided version of our claim. Repeating the same lines by replacing λ by $-\lambda$ and using the union bound again, we finish the proof. ■

Our next result concludes the proof of Theorem 1.1 when $p = 4$. Recall that κ denotes $\kappa(4)$.

Theorem 3.2. *There is an absolute constant $C > 0$ such that the following holds. Assume that X is a zero mean random vector with covariance Σ satisfying the L_4 - L_2 norm equivalence. Fix the corruption level $\eta \in [0, 1]$ and the confidence level $\delta \in (0, 1)$. Assume that $\tilde{X}_1, \dots, \tilde{X}_N$ is an η -corrupted sample. Then there exists an estimator $\hat{\Sigma}_{\eta, \delta}$ such that, with probability at least $1 - \delta$,*

$$\|\hat{\Sigma}_{\eta, \delta} - \Sigma\| \leq C \kappa^2 \|\Sigma\| \left(\sqrt{\frac{\mathbf{r}(\Sigma) + \log(1/\delta)}{N}} + \sqrt{\eta} \right).$$

We begin with presenting the estimator. As we mentioned, we assume the knowledge of η , $\|\Sigma\|$, $\text{Tr}(\Sigma)$ and the constant $\kappa = \kappa(4)$. In Section 3.1, we discuss how to estimate $\|\Sigma\|$ and $\text{Tr}(\Sigma)$ up to multiplicative constant factors. The dependence on κ in our estimator can also be ignored, though in this case we obtain a slightly weaker dependence on this parameter in the final bound.

The estimator in Theorem 3.2:

(1) Given δ , $\text{Tr}(\Sigma)$, $\|\Sigma\|$, κ , and an η -corrupted sample $\tilde{X}_1, \dots, \tilde{X}_N$, we set

$$\lambda = \frac{1}{\kappa^2 \|\Sigma\|} \sqrt{\frac{\mathbf{r}(\Sigma) + \log(1/\delta) + \eta N}{N}}.$$

(2) Define

$$\Gamma = \bigcap_{v \in \mathcal{S}^{d-1}} \left\{ A \in \mathbb{S}_+^d : \left| \frac{1}{\lambda N} \sum_{i=1}^N \psi(\lambda \langle \tilde{X}_i, v \rangle^2) - v^\top A v \right| \leq \frac{C \lambda \kappa^4 \|\Sigma\|^2}{2} \right\}.$$

(3) Let $\hat{\Sigma}_{\eta, \delta}$ be any matrix in the set Γ . If the set is empty, we output $\hat{\Sigma}_{\eta, \delta} = 0$.

Proof of Theorem 3.2. Since the truncation function $\psi(\cdot)$ is bounded, we have

$$\begin{aligned} & \sup_{v \in \mathcal{S}^{d-1}} \left| \frac{1}{N \lambda} \sum_{i=1}^N \psi(\lambda \langle \tilde{X}_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| \\ & \leq \sup_{v \in \mathcal{S}^{d-1}} \left| \frac{1}{N \lambda} \sum_{i=1}^N \psi(\lambda \langle X_i, v \rangle^2) - \mathbb{E} \langle X, v \rangle^2 \right| + \frac{\eta}{\lambda}. \end{aligned}$$

Combining this with Proposition 3.1, we have for some $C > 0$, with probability at least $1 - \delta$,

$$\sup_{v \in \mathcal{S}^{d-1}} \left| \frac{1}{N \lambda} \sum_{i=1}^N \psi(\lambda \langle \tilde{X}_i, v \rangle^2) - v^\top \Sigma v \right| \leq \frac{C}{4} \left(\frac{\mathbf{r}(\Sigma) + \log(1/\delta) + \eta N}{\lambda N} + \lambda \kappa^4 \|\Sigma\|^2 \right).$$

Using the triangle inequality, the definition of $\hat{\Sigma}_{\eta, \delta}$ and the line above, we have on the same event

$$\begin{aligned} \|\hat{\Sigma}_{\eta, \delta} - \Sigma\| & \leq \sup_{v \in \mathcal{S}^{d-1}} \left(\left| \frac{1}{N \lambda} \sum_{i=1}^N \psi(\lambda \langle \tilde{X}_i, v \rangle^2) - v^\top \hat{\Sigma}_{\eta, \delta} v \right| \right. \\ & \quad \left. + \left| \frac{1}{N \lambda} \sum_{i=1}^N \psi(\lambda \langle \tilde{X}_i, v \rangle^2) - v^\top \Sigma v \right| \right) \\ & \leq C \kappa^2 \|\Sigma\| \left(\sqrt{\frac{\mathbf{r}(\Sigma) + \log(1/\delta)}{N}} + \sqrt{\eta} \right). \end{aligned}$$

Our choice of λ concludes the proof. ■

Remark 3.3. It is straightforward to verify that the matrix $\hat{\Sigma}_{\eta, \delta}$ defined by (3) satisfies the assumption of Theorem 3.2 and can be used as an estimator of the covariance matrix.

Finally, we show how to estimate the value of the truncation parameter λ using the corrupted observations.

3.1. Estimating the truncation level λ in the regime $p = 4$

We begin this section with a brief discussion of our contamination model. To simplify the technical aspects of our analysis, we impose a slight restriction on the strong contamination model described in the introduction. In our analysis, we require the estimation of multiple parameters, which can be accomplished by dividing the sample into several equal parts and estimating these parameters using independent sub-samples. This approach is standard and has been analyzed, for example, in [53] in the same context. However, in the strong contamination model, the adversary has the ability to replace some observations within each sub-sample in a manner that depends on the entire sample. This can result in dependent sub-samples and a more complicated stability type analysis appearing in the literature (see [23, Section 6]). We therefore restrict the adversary from making the sub-samples dependent on each other. We note that this assumption is standard and has been made (implicitly) in some recent works [47, 57]. Furthermore, it is automatically satisfied in standard contamination models, including the Huber model.

An important aspect of the proof of Theorem 3.2 is that we only need to know $\|\Sigma\|$ and $\text{Tr}(\Sigma)$ up to multiplicative constant factors. That is, we want to find two estimators $\hat{\varphi}_1$ and $\hat{\varphi}_2$ such that for some $c > 1$, with high probability with respect to the realization of the corrupted sample $\tilde{X}_1, \dots, \tilde{X}_N$,

$$\frac{1}{c} \text{Tr}(\Sigma) \leq \hat{\varphi}_1 \leq c \text{Tr}(\Sigma) \quad \text{and} \quad \frac{1}{c} \|\Sigma\| \leq \hat{\varphi}_2 \leq c \|\Sigma\|.$$

If these estimators are available, we can assume, without the loss of generality, that our initial sample is of size $3N$. We use the first $2N$ elements to estimate $\text{Tr}(\Sigma)$ and $\|\Sigma\|$ and then plug them into the truncation level λ in Theorem 3.2. A similar sample splitting strategy is used in [53]. Our contribution is that we provide estimators $\hat{\varphi}_1$ and $\hat{\varphi}_2$ such that they take the adversarial corruption into account as well as do not contain an unnecessary logarithmic factor. For the rest of the section, we fix $\kappa = \kappa(4)$, where $\kappa(\cdot)$ is defined in (1).

3.2. Estimating $\text{Tr}(\Sigma)$

One can estimate $\text{Tr}(\Sigma)$ using, for example, the trimmed mean estimator analyzed in [47]. Let e_1, \dots, e_d denote the canonical basis in \mathbb{R}^d . Observe that

$$\text{Tr}(\Sigma) = \sum_{i=1}^d \mathbb{E} \langle X, e_i \rangle^2.$$

Using the L_4 - L_2 norm equivalence and the proof of Lemma 2.1, we estimate the variance $\text{Var}(\sum_{i=1}^d \langle X, e_i \rangle^2) \leq \kappa^4 (\text{Tr}(\Sigma))^2$. Given an η -corrupted sample of size $2N$, the trimmed mean estimator $\hat{\varphi}_1$ applied to the random variable $\sum_{i=1}^d \langle X, e_i \rangle^2$ implies that, for any $\delta \in (\frac{\exp(-N)}{4}, 1)$, with probability at least $1 - 4 \exp(-\frac{N}{12})$,

$$|\hat{\varphi}_1 - \text{Tr}(\Sigma)| \leq 10\sqrt{\varepsilon}\kappa^2 \text{Tr}(\Sigma), \tag{13}$$

where $\varepsilon = 8\eta + 12 \frac{\log(4/\delta)}{N}$. This bound is presented in [47, Theorem 1]. For the sake of brevity, we do not provide the details of the multivariate trimmed mean estimator. Observe that for this choice of ε , it holds that $1 - 4 \exp(-\varepsilon \frac{N}{12}) \geq 1 - \delta$. Provided that $10\sqrt{\varepsilon}\kappa^2 \leq \frac{1}{2}$ in (13), we obtain that, with probability at least $1 - \delta$,

$$\frac{1}{2} \text{Tr}(\Sigma) \leq \hat{\varphi}_1 \leq 2 \text{Tr}(\Sigma).$$

It is only left to check that it is sufficient to have $\eta \in [0, \frac{1}{40^2\kappa^2}]$ and $N \geq 12 \cdot 40^2\kappa^4 \log(\frac{4}{\delta})$.

3.3. Estimating $\|\Sigma\|$

A simple way of estimating $\|\Sigma\|$ is to first estimate the full covariance matrix Σ as in Theorem 3.2 and then compute the operator norm of the estimator. However, the difficulty of estimating $\|\Sigma\|$ is that we are not allowed to choose the truncation level λ depending on $\|\Sigma\|$. In particular, all previous results we are aware of lead to an additional logarithmic factor in the assumption on the sample size (see [15, 53]). Instead, we provide an adaptive estimator similar in spirit to the original robust estimator of Catoni [13] for estimating the mean of a random variable (see also [14] for related computations). To simplify the proof, we assume that the distribution of X satisfies $\Pr(X = 0) = 0$. This assumption is mild and can always be satisfied if we add a small Gaussian perturbation to X without changing the covariance matrix of X too much.

Proposition 3.4. *Assume that X is a zero mean random vector with covariance Σ satisfying the L_4 - L_2 norm equivalence. Assume additionally that $\Pr(X = 0) = 0$. Let $c \geq 1$ be a large enough absolute constant. Fix the corruption level $\eta \in [0, \frac{1}{300c\kappa^4}]$ and the confidence level $\delta \in (0, \frac{1}{4})$. Assume that $\tilde{X}_1, \dots, \tilde{X}_N$ is an η -corrupted sample. Then there is a unique value $\hat{\alpha} > 0$ satisfying*

$$\frac{1}{N} \sup_{v \in S^{d-1}} \sum_{i=1}^N \psi(\hat{\alpha}^2 \langle \tilde{X}_i, v \rangle^2) = \frac{1}{20c\kappa^4} + \eta.$$

If $N \geq 100c\kappa^4 \mathbf{r}(\Sigma) + 400c\kappa^4 \log(\frac{1}{\delta})$, then with probability at least $1 - 4\delta$, it holds that

$$\frac{1}{4} \|\Sigma\| \leq \frac{1}{24c\kappa^4 \hat{\alpha}^2} \leq 4 \|\Sigma\|.$$

Proof. We begin with the analysis of an uncorrupted sample X_1, \dots, X_N . Our first aim is to choose the distributions μ and ρ in Lemma 2.8. Let $\Theta = (\mathbb{R}^d)^2$ and choose μ to be a product of two zero mean multivariate Gaussians with covariance $\beta^{-1}I_d$. For $v \in S^{d-1}$ and $\alpha > 0$, let $\rho_{\alpha,v}$ be a product of two multivariate Gaussian distribution with mean αv and covariance $\beta^{-1}I_d$. Because of this, if (θ, v) is distributed according to $\rho_{\alpha,v}$, we have $\mathbb{E}_{\rho_v}(\theta, v) = (\alpha v, \alpha v)$. By the additivity of \mathcal{KL} -divergence for product measures and the standard formula, we have

$$\mathcal{KL}(\rho_{\alpha,v}, \mu) = \alpha^2 \beta.$$

For the rest of the proof, we sometimes write ρ instead of $\rho_{\alpha,v}$. By the first part of Lemma 2.9, we have

$$\begin{aligned} \psi(\alpha^2 \langle X, v \rangle^2) &= \psi(\mathbb{E}_\rho \langle X, \theta \rangle \langle X, v \rangle) \leq \mathbb{E}_\rho \log(1 + \langle X, \theta \rangle \langle X, v \rangle + (\langle X, \theta \rangle \langle X, v \rangle)^2) \\ &\quad + \min\left\{1, \frac{\mathbb{E}_\rho(\langle X, \theta \rangle \langle X, v \rangle)^2}{6}\right\}. \end{aligned}$$

Observe that

$$\mathbb{E}_\rho(\langle X, \theta \rangle \langle X, v \rangle)^2 = (\alpha^2 \langle X, v \rangle^2 + \beta^{-1} \|X\|^2)^2 \leq 2\alpha^4 \langle X, v \rangle^4 + 2\beta^{-2} \|X\|^4,$$

and we can write

$$\min\left\{1, \frac{\mathbb{E}_\rho(\langle X, \theta \rangle \langle X, v \rangle)^2}{6}\right\} \leq \min\left\{1, \frac{2\alpha^4 \langle X, v \rangle^4}{6}\right\} + \min\left\{1, \frac{2\beta^{-2} \|X\|^4}{6}\right\}.$$

Repeating the lines of the proof of Proposition 2.10, we get

$$\begin{aligned} \mathbb{E}_\rho \log \mathbb{E}(1 + \langle X, \theta \rangle \langle X, v \rangle + c(\langle X, \theta \rangle \langle X, v \rangle)^2) \\ \leq \alpha^2 v^\top \Sigma v + 2c\kappa^4(\alpha^4 \|\Sigma\|^2 + \beta^{-2}(\text{Tr}(\Sigma))^2), \end{aligned}$$

where $c \geq 1$ is an absolute constant. Using the first part of Lemma 2.8, we obtain that, with probability at least $1 - \delta$, simultaneously for all $v \in S^{d-1}$ and $\alpha > 0$,

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \psi(\alpha^2 \langle X_i, v \rangle^2) &\leq \alpha^2 v^\top \Sigma v + 2c\kappa^4(\alpha^4 \|\Sigma\|^2 + \beta^{-2}(\text{Tr}(\Sigma))^2) \\ &\quad + \frac{\alpha^2 \beta + \log(1/\delta)}{N} + \frac{1}{N} \sum_{i=1}^N \min\left\{1, \frac{2\beta^{-2} \|X_i\|^4}{6}\right\}. \end{aligned}$$

Using the Bernstein inequality as in (12), we have, with probability at least $1 - \delta$,

$$\frac{1}{N} \sum_{i=1}^N \min\left\{1, \frac{2\beta^{-2} \|X_i\|^4}{6}\right\} \leq \frac{2\kappa^4 \beta^{-2} (\text{Tr}(\Sigma))^2}{3} + \frac{3 \log(1/\delta)}{N}.$$

We choose $\beta = 10c\kappa^4 \text{Tr}(\Sigma)$ and using the union bound, we obtain that, with probability at least $1 - 2\delta$, simultaneously for all $v \in S^{d-1}$ and $\alpha > 0$,

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \psi(\alpha^2 \langle X_i, v \rangle^2) &\leq \alpha^2 v^\top \Sigma v + 2c\kappa^4 \alpha^4 \|\Sigma\|^2 + \frac{2}{75c\kappa^4} + \frac{10c\kappa^4 \alpha^2 \text{Tr}(\Sigma)}{N} \\ &\quad + \frac{4 \log(1/\delta)}{N}. \end{aligned}$$

Since $N \geq 100c\kappa^4 \mathbf{r}(\Sigma) + 400c\kappa^4 \log(\frac{1}{\delta})$, we have on the same event

$$\frac{1}{N} \sum_{i=1}^N \psi(\alpha^2 \langle X_i, v \rangle^2) \leq \alpha^2 v^\top \Sigma v + \frac{1}{10} \alpha^2 \|\Sigma\| + 2c\kappa^4 \alpha^4 \|\Sigma\|^2 + \frac{11}{300c\kappa^4}.$$

Returning to the η -corrupted sample $\tilde{X}_1, \dots, \tilde{X}_N$ and since $|\psi(\cdot)| \leq 1$,

$$\frac{1}{N} \sum_{i=1}^N \psi(\alpha^2 \langle \tilde{X}_i, v \rangle^2) \leq \alpha^2 v^\top \Sigma v + \frac{1}{10} \alpha^2 \|\Sigma\| + 2c\kappa^4 \alpha^4 \|\Sigma\|^2 + \frac{11}{300c\kappa^4} + \eta. \quad (14)$$

Observe that the function $\varphi: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ given by

$$\varphi(\alpha) = \frac{1}{N} \sup_{v \in S^{d-1}} \sum_{i=1}^N \psi(\alpha^2 \langle \tilde{X}_i, v \rangle^2)$$

is continuous (it is easy to check that $\varphi(\cdot)$ is a supremum over equi-Lipschitz functions) and non-decreasing. Furthermore, there is $w \in S^{d-1}$ such that $\langle X_i, w \rangle \neq 0$ for all $i \in [N]$. Indeed, to construct such a vector, we pick a random vector U distributed as a zero mean multivariate Gaussian random vector with unit covariance. Since $X_i \neq 0$ almost surely for all $i \in [N]$, we have that $\langle U, X_i \rangle$ is a zero mean Gaussian (conditionally on X_i) with non-zero covariance. Thus, taking $w = \frac{U}{\|U\|}$ we have almost surely $\langle X_i, w \rangle \neq 0$ for all $i \in [N]$. We have $\varphi(0) = 0$ and $\varphi(\frac{1}{\min_{i \in [N]} |\langle X_i, w \rangle|}) \geq 1 - \eta$ since at most ηN vectors \tilde{X}_i can be orthogonal to w . Therefore, since $\kappa, c \geq 1$ and $\eta \leq \frac{1}{300c\kappa^4}$, there is $\hat{\alpha} > 0$ such that

$$\varphi(\hat{\alpha}) = \frac{1}{20c\kappa^4} + \eta.$$

Using (14), we get on the same event

$$2c\kappa^4 \hat{\alpha}^4 \|\Sigma\|^2 + \frac{11}{10} \hat{\alpha}^2 \|\Sigma\| - \frac{1}{75c\kappa^4} \geq 0.$$

Solving this as a quadratic equation (only the positive root plays a role) with respect to $\hat{\alpha}^2 \|\Sigma\|$, we get

$$\frac{-11/10 + \sqrt{(11/10)^2 + 8/75}}{4c\kappa^4 \hat{\alpha}^2} \leq \|\Sigma\|.$$

This provides a lower bound on $\|\Sigma\|$ in terms of $\hat{\alpha}^2$.

We proceed with an upper bound. For $v \in S^{d-1}$ and $\alpha > 0$, let $\rho_{\alpha, -\alpha, v}$ be a product of two multivariate Gaussian distributions with means αv and $-\alpha v$, respectively, and the same covariance matrix $\beta^{-1} I_d$. Repeating the proof, we obtain the following analog of inequality (14). With probability at least $1 - 2\delta$, for all $v \in S^{d-1}$ and $\alpha > 0$,

$$\begin{aligned} -\frac{1}{N} \sum_{i=1}^N \psi(\alpha^2 \langle \tilde{X}_i, v \rangle^2) &\leq -\alpha^2 v^\top \Sigma v + \frac{1}{10} \alpha^2 \|\Sigma\| + 2c\kappa^4 \alpha^4 \|\Sigma\|^2 \\ &\quad + \frac{11}{300c\kappa^4} + \eta. \end{aligned} \quad (15)$$

Let v_1 be a maximizer of $v^\top \Sigma v$ in S^{d-1} . We have $v_1^\top \Sigma v_1 = \|\Sigma\|$. Furthermore, we have

$$-\sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N \psi(\alpha^2 \langle \tilde{X}_i, v \rangle^2) \leq -\frac{1}{N} \sum_{i=1}^N \psi(\alpha^2 \langle \tilde{X}_i, v_1 \rangle^2).$$

Since $-\varphi(\cdot)$ is non-increasing, this implies that simultaneously for all $\alpha \in [0, \hat{\alpha}]$ on the event where (15) holds, we have

$$-\frac{1}{20c\kappa^4} - \eta \leq -\frac{9}{10}\alpha^2\|\Sigma\| + 2c\kappa^4\alpha^4\|\Sigma\|^2 + \frac{11}{300c\kappa^4} + \eta.$$

Since $\eta \leq \frac{1}{300c\kappa^4}$, we reduce this equation to

$$2c\kappa^4\alpha^4\|\Sigma\|^2 - \frac{9}{10}\alpha^2\|\Sigma\| + \frac{7}{75c\kappa^4} \geq 0. \quad (16)$$

As a function of $\alpha^2\|\Sigma\|$, this quadratic equation has two non-negative roots:

$$x_1 = \frac{9/10 - \sqrt{(9/10)^2 - 56/75}}{4c\kappa^4} \quad \text{and} \quad x_2 = \frac{9/10 + \sqrt{(9/10)^2 - 56/75}}{4c\kappa^4}.$$

We show that since (16) holds for all $\alpha \in [0, \hat{\alpha}]$, we should have that $\hat{\alpha}^2\|\Sigma\| \leq x_1$. Indeed, assume that instead $\hat{\alpha}^2\|\Sigma\| \geq x_2$. In this case $0 \leq \frac{x_1+x_2}{2} \leq \hat{\alpha}^2\|\Sigma\|$ and thus (16) should be satisfied for $\alpha^2\|\Sigma\| = \frac{x_1+x_2}{2}$. The obtained contradiction proves that $\hat{\alpha}^2\|\Sigma\| \leq x_1$. Finally, by union bound we have, with probability at least $1 - 4\delta$,

$$\frac{-11/10 + \sqrt{(11/10)^2 + 8/75}}{4c\kappa^4\hat{\alpha}^2} \leq \|\Sigma\| \leq \frac{9/10 - \sqrt{(9/10)^2 - 56/75}}{4c\kappa^4\hat{\alpha}^2}.$$

Algebraic computations conclude the proof. \blacksquare

We end up by observing that the estimator proposed in this section admits a straightforward generalization for (uncontaminated) higher-order tensors. In this setup, one is aiming to estimate $\mathbb{E}\langle X, v \rangle^s$ uniformly over $v \in S^{d-1}$ for some integer $s \geq 2$, as in the recent work of Mendelson [50]. In order to provide this extension, we apply [73, Lemma 5] together with the estimators of the trace and the operator norm provided above. We note that getting the optimal dependence on the contamination level η seems more subtle when $s > 2$.

4. Proof of Theorem 1.1 in the regime $p > 4$

In this section, we prove Theorem 1.1 in the regime $p > 4$. We begin with the following symmetrized version of Theorem 1.2.

Corollary 4.1. *Assume that Y is a zero mean random vector with covariance Σ satisfying the L_p - L_2 norm equivalence with $p > 4$. Let Y_1, \dots, Y_N be a sample of independent copies of Y . Consider the truncated vectors $X_i = Y_i \mathbb{1}\{\|Y_i\| \leq (N \text{Tr}(\Sigma)\|\Sigma\|)^{1/4}\}$ for $i = 1, \dots, N$. Let $\varepsilon_1, \dots, \varepsilon_N$ be independent Rademacher random signs. If $N \geq c(p)\mathbf{r}(\Sigma)$, then it holds that*

$$\mathbb{E} \sup_{v \in S^{d-1}} \left| \frac{1}{N} \sum_{i=1}^N \varepsilon_i \langle X_i, v \rangle^2 \right| \leq C(p)\|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}},$$

where $c(p)$ and $C(p)$ depend only on p and $\kappa(p)$ and the expectation is taken with respect to both X_i and ε_i , $i = 1, \dots, N$.

Remark 4.2. We highlight that the results of Theorem 1.2 and Corollary 4.1 admit extensions to higher-order power in the spirit of [69]. For the related derivations in the light-tailed case, see the proof of [73, Theorem 2].

Proof of Corollary 4.1. We use the standard desymmetrization argument [37, Section 2.1]. Assume that X'_1, \dots, X'_N are independent copies of X , and let \mathbb{E}' denote the expectation with respect to these random vectors. Using the triangle inequality, Jensen’s inequality and the contraction principle [41, Theorem 4.4], we get

$$\begin{aligned} & \mathbb{E} \sup_{v \in S^{d-1}} \left| \sum_{i=1}^N \varepsilon_i \langle X_i, v \rangle^2 \right| \\ & \leq \mathbb{E} \sup_{v \in S^{d-1}} \left| \sum_{i=1}^N \varepsilon_i (\langle X_i, v \rangle^2 - \mathbb{E}' \langle X'_i, v \rangle^2) \right| + \mathbb{E} \sup_{v \in S^{d-1}} \left| \sum_{i=1}^N \varepsilon_i \mathbb{E} \langle X_i, v \rangle^2 \right| \\ & \leq \mathbb{E} \mathbb{E}' \sup_{v \in S^{d-1}} \left| \sum_{i=1}^N \langle X_i, v \rangle^2 - \langle X'_i, v \rangle^2 \right| + \|\Sigma\| \sqrt{N} \\ & \leq 2 \mathbb{E} \sup_{v \in S^{d-1}} \left| \sum_{i=1}^N \langle X_i, v \rangle^2 - \mathbb{E} \langle X_i, v \rangle^2 \right| + \|\Sigma\| \sqrt{N}. \end{aligned}$$

The bound of Theorem 1.2 concludes the proof. ■

In the regime $p > 4$, we show how to adapt the proposed estimator to obtain the optimal dependence on the contamination level η . The main insight is to interpret our estimator as a second-order version of the trimmed mean estimator of Lugosi and Mendelson [47]. Our proof mainly follows their steps, however two important modifications are required. First, we need to exploit Theorem 1.2 when controlling the expected supremum of quadratic processes, while in [47] the control of the corresponding process follows from a simple application of the Cauchy–Schwarz inequality. To do so, we also need to truncate the norms of our observations as explained after Proposition 4.5. Second, since the quadratic forms of interest are non-negative, we consider only a one-sided truncation. The main result of this section is the following.

Theorem 4.3. *Assume that X is a zero mean random vector with covariance Σ satisfying the L_p – L_2 norm equivalence with $p > 4$. Fix the corruption level $\eta \in [0, 1]$ and the confidence level $\delta \in (0, 1)$. Then there exists an estimator $\widehat{\Sigma}_{\eta, \delta}$ such that, with probability at least $1 - \delta$,*

$$\|\widehat{\Sigma}_{\eta, \delta} - \Sigma\| \leq C(p) \|\Sigma\| \left(\sqrt{\frac{\mathbf{r}(\Sigma) + \log(1/\delta)}{n}} + \kappa(p)^2 \eta^{1-2/p} \right).$$

Here $C(p)$ is a non-increasing function of p that satisfies $C(p) \rightarrow \infty$ as $p \rightarrow 4$.

Let us first recall the form of our estimator in the case where $p = 4$. For some specifically chosen $\lambda > 0$, the estimator in Theorem 3.2 is defined by the following set:

$$\Gamma = \bigcap_{v \in S^{d-1}} \left\{ A \in \mathbb{S}_+^d : \left| \frac{1}{\lambda N} \sum_{i=1}^N \psi(\lambda \langle \tilde{X}_i, v \rangle^2) - v^\top A v \right| \leq \frac{C \lambda \kappa^4 \|\Sigma\|^2}{2} \right\},$$

where $\psi(\cdot)$ is the truncation function given by (2). We defined $\hat{\Sigma}_{\eta, \delta}$ to be any matrix in Γ . In the regime $p > 4$, we introduce a more involved estimator that uses a direction-dependent value of the one-sided threshold. For simplicity, we assume that we observe a sample of size $2N$. For a set of positive semi-definite matrices Γ , we define its diameter as $\Delta(\Gamma) = \sup_{A, B \in \Gamma} \|A - B\|$.

The estimator in Theorem 4.3:

- (1) Given the confidence level δ , the corruption level η and the η -corrupted sample $\tilde{Y}_1, \dots, \tilde{Y}_{2N}$, we split the sample in two parts of equal size: Truncate each vector $\tilde{Z}_i = \tilde{Y}_i \mathbb{1}\{\|\tilde{Y}_i\| \leq R\}$ for $i \in \{1, \dots, N\}$ and $\tilde{X}_i = \tilde{Y}_i \mathbb{1}\{\|\tilde{Y}_i\| \leq R\}$ for $i \in \{N + 1, \dots, 2N\}$, where $R = (N \text{Tr}(\Sigma) \|\Sigma\|)^{1/4}$.
- (2) Set $\varepsilon = \max(20\eta, 560 \frac{\log(2/\delta)}{N})$.
- (3) For the first half of the sample, define $q_v = (\langle \tilde{Z}_i, v \rangle^2)_{N\varepsilon/2}^*$.
- (4) For the second half of the sample $\tilde{X}_1, \dots, \tilde{X}_N$, we proceed as follows: For every $Q > 0, v \in S^{d-1}$, define the trimming level $\lambda_v(Q) = (q_v + Q)^{-1}$ and set

$$\Gamma(Q) = \bigcap_{v \in S^{d-1}} \left\{ A \in \mathbb{S}_+^d : \left| \frac{1}{\lambda_v(Q) N} \sum_{i=1}^N \psi(\lambda_v(Q) \langle \tilde{X}_i, v \rangle^2) - v^\top A v \right| \leq 4\varepsilon Q \right\}.$$

- (5) Let $\hat{\Sigma}_{\eta, \delta}$ be any matrix in the set $\Gamma(2^{i^*})$, where i^* minimizes the diameter $\Delta(\Gamma(2^i))$ over all integers i subject to be non-empty.

Remark 4.4. Due to the bounds in Section 3.1, we explicitly assume the knowledge of the operator norm $\|\Sigma\|$ and the effective rank $\mathbf{r}(\Sigma)$ up to multiplicative constant factors.

Our first technical observation is that, by truncating the norms of the vectors at the level $R = (N \text{Tr}(\Sigma) \|\Sigma\|)^{1/4}$, the covariance matrix does not change too much. At the same time, the truncation allows us to apply Theorem 1.2 in our analysis.

Proposition 4.5. *Assume that Y is a zero mean random vector with covariance Σ satisfying the L_p - L_2 norm equivalence with $p \geq 4$. Consider the truncated vector $X = Y \mathbb{1}\{\|Y\| \leq (N \text{Tr}(\Sigma) \|\Sigma\|)^{1/4}\}$ and define $\tilde{\Sigma} = \mathbb{E} X \otimes X$. It holds that*

$$\|\Sigma - \tilde{\Sigma}\| \leq \kappa(4)^4 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}.$$

Moreover, $\|\tilde{\Sigma}\| \leq \|\Sigma\|$ and $\text{Tr}(\tilde{\Sigma}) \leq \text{Tr}(\Sigma)$.

Proof. The proof follows from the proof of [53, Lemma 2.1]. For the second part of the proof, we observe that $\Sigma - \tilde{\Sigma}$ is positive semi-definite. ■

The result of Proposition 4.5 allows us to focus on estimating the matrix of second moments $\tilde{\Sigma}$. Indeed, by the triangle inequality, the estimator $\hat{\Sigma}_{\eta,\delta}$ satisfies

$$\|\hat{\Sigma}_{\eta,\delta} - \Sigma\| \leq \|\hat{\Sigma}_{\eta,\delta} - \tilde{\Sigma}\| + \kappa(4)^4 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}},$$

and the last term only affects the multiplicative constant factors in our bound. We proceed by defining the following quantity:

$$Q_0 = C_{Q_0} \max \left\{ \frac{1}{\varepsilon} \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma) + \log(2/\delta)}{N}}, \varepsilon^{-2/p} \kappa(p)^2 \|\Sigma\| \right\}.$$

Here $C_{Q_0} = C_{Q_0}(p) > 0$ is to be chosen later, depends only on p and is non-increasing with respect to p . The first term in the definition above is responsible for the rate in the uncontaminated case, and the second term captures the rate of η according to the value of p . We start with an important lemma that is the second-order analog of [47, Lemma 1]. For the rest of this section, $\tilde{\Sigma}$ is the matrix of second moments of the truncated vector

$$X = Y \mathbb{1}\{\|Y\| \leq (N \text{Tr}(\Sigma)\|\Sigma\|)^{1/4}\}$$

and

$$\varepsilon = \max\left(20\eta, 560 \frac{\log(2/\delta)}{N}\right),$$

Lemma 4.6. *Let $Y \in \mathbb{R}^d$ be a mean zero random vector satisfying the L_p - L_2 norm equivalence assumption with $p > 4$. Let Y_1, \dots, Y_N be i.i.d. copies of Y and set*

$$Z_i = Y_i \mathbb{1}\{\|Y_i\| \leq (N \text{Tr}(\Sigma)\|\Sigma\|)^{1/4}\}$$

for every $i \in [N]$, then, with probability at least $1 - \frac{\delta}{2}$,

$$\sup_{v \in S^{d-1}} |\{i \in [N] : \langle Z_i, v \rangle^2 - v^\top \tilde{\Sigma} v \geq Q_0\}| \leq \frac{\varepsilon}{4} N.$$

Proof. For simplicity of notation, we write

$$\bar{Z}_i(v) = \langle Z_i, v \rangle^2 - v^\top \tilde{\Sigma} v.$$

The proof is a standard application of a small ball argument in empirical process theory. Consider a function $\xi: \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$\xi(x) = \begin{cases} 0, & x \leq \frac{Q_0}{2}, \\ \frac{2x}{Q_0} - 1, & x \in \left(\frac{Q_0}{2}, Q_0\right], \\ 1, & x > Q_0. \end{cases}$$

It is clear that $\mathbb{1}\{\bar{Z}_i(v) \geq Q_0\} \leq \xi(\bar{Z}_i(v)) \leq \mathbb{1}\{\bar{Z}_i(v) \geq \frac{Q_0}{2}\}$ and ξ is a Lipschitz function with constant $\frac{2}{Q_0}$. By symmetrization [26] and the contraction lemma for Rademacher processes [41, Theorem 4.4], we have

$$\begin{aligned} \mathbb{E} \sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\bar{Z}_i(v) \geq Q_0\} &\leq \mathbb{E} \sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N \xi(\bar{Z}_i(v)) \\ &\leq 2\mathbb{E} \sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N |\varepsilon_i \xi(\bar{Z}_i(v))| + \sup_{v \in S^{d-1}} \mathbb{E} \xi(\bar{Z}_i(v)) \\ &\leq \frac{4}{Q_0} \mathbb{E} \sup_{v \in S^{d-1}} \frac{1}{N} \left| \sum_{i=1}^N \varepsilon_i \bar{Z}_i(v) \right| + \sup_{v \in S^{d-1}} \mathbb{E} \xi(\bar{Z}_i(v)). \end{aligned}$$

Now, we define $C_{BY} = C_{BY}(p) > 0$ to be the constant in the conclusion of Corollary 4.1. We remark that C_{BY} is at most an absolute constant for any $p \geq c > 4$, where c is another absolute constant. We now apply Corollary 4.1 and the same arguments as in its proof and obtain

$$\begin{aligned} \mathbb{E} \sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N |\varepsilon_i (\langle Z_i, v \rangle^2 - v^\top \tilde{\Sigma} v)| &\leq \mathbb{E} \sup_{v \in S^{d-1}} \frac{1}{N} \left| \sum_{i=1}^n \varepsilon_i \langle Z_i, v \rangle^2 \right| + \frac{\|\Sigma\|}{N} \mathbb{E} \left| \sum_{i=1}^N \varepsilon_i \right| \\ &\leq 2C_{BY} \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}} \leq \frac{Q_0 \varepsilon}{128}. \end{aligned}$$

The last steps follow from the fact that $\mathbf{r}(\Sigma) \geq 1$ and $\|\tilde{\Sigma}\| \leq \|\Sigma\|$. By the definition of Q_0 , we conclude that the first term is at most $\frac{\varepsilon}{32}$ if we choose a sufficiently large constant $C_{Q_0} > 0$. We now proceed to bound the second term. We use Markov’s inequality together with the L_p - L_2 norm equivalence to obtain

$$\begin{aligned} \mathbb{E} \xi(\langle Z, v \rangle^2 - v^\top \tilde{\Sigma} v) &\leq \mathbb{P} \left(\langle Z, v \rangle^2 \geq \frac{Q}{2} + v^\top \tilde{\Sigma} v \right) \\ &\leq \frac{\kappa(p)^p}{(Q/2 + v^\top \tilde{\Sigma} v)^{p/2}} (\mathbb{E} \langle Z, v \rangle^2)^{p/2} \leq \left(\frac{2\kappa(p)^2 \|\Sigma\|}{Q} \right)^{p/2}. \end{aligned}$$

Again, by the definition of Q_0 , we can choose $C_{Q_0} > 0$ such that the right-hand side above is at most $\frac{\varepsilon}{32}$. By Talagrand’s concentration inequality for supremum of empirical process (Massart’s version [48]), with probability at least $1 - \exp(-x)$,

$$\sup_{v \in S^{d-1}} \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\bar{Z}_i(v) \geq Q\} \leq \frac{\varepsilon}{8} + \sqrt{\frac{8x}{N}} \frac{\sqrt{\varepsilon}}{8} + \frac{35x}{N}.$$

We choose $x = \frac{\varepsilon N}{560}$ to conclude the proof. ■

We now describe the first part of the estimation procedure. First, we assume without loss of generality that $\varepsilon < 1$. When the conclusion of the previous lemma holds, we immediately obtain that, for every $v \in S^{d-1}$ and $Q \in (2Q_0, 4Q_0)$,

$$q_v - v^\top \tilde{\Sigma} v + Q \leq 5Q_0 \quad \text{and} \quad q_v - v^\top \tilde{\Sigma} v + Q \geq Q_0.$$

To see why the latter condition holds, suppose by contradiction that $q_v - v^\top \tilde{\Sigma} v + Q \leq Q_0$, then $q_v \leq v^\top \tilde{\Sigma} v - Q_0$. Since $\varepsilon < 1$, by the choice of Q_0 , we have $Q_0 > \|\Sigma\|$ and then we get $0 \leq q_v \leq v^\top \tilde{\Sigma} v - Q_0 < 0$, the latter is clearly a contradiction. Moreover, both inequalities above still hold for an η -corrupted sample, since $\varepsilon \geq 20\eta$ and by the definition of ε , we have that the fraction of points greater than Q_0 is at most $\frac{\varepsilon}{4} + \eta \leq \varepsilon(\frac{1}{20} + \frac{1}{4}) < \frac{\varepsilon}{2}$. Thus, we can use $\tilde{Z}_1, \dots, \tilde{Z}_N$ to estimate q_v , and since they are independent from the second half of the sample, from now on we work conditionally on the event of Lemma 4.6.

The second part of the analysis consists in proving that $\Gamma(v, Q)$ is non-empty. The formal statement is the proposition below.

Proposition 4.7. *Under the notation of Theorem 4.3, fix $Q \in [2Q_0, 4Q_0]$ and assume that the event of Lemma 4.6 holds. There exists an absolute constant $C > 0$ such that if $C_{Q_0} \geq C$, then, with probability at least $1 - \frac{\delta}{2}$,*

$$\sup_{v \in S^{d-1}} \left| \frac{1}{\lambda_v(Q)N} \sum_{i=1}^N \psi(\lambda_v(Q)\langle \tilde{X}_i, v \rangle^2) - v^\top \tilde{\Sigma} v \right| \leq 4\varepsilon Q.$$

It will be convenient to work with two-sided function in the analysis. We define the following function for arbitrary positive $\lambda_1 \geq \lambda_2$:

$$\psi_{\lambda_1, \lambda_2}(x) = \begin{cases} \frac{1}{\lambda_2}, & x > \frac{1}{\lambda_2}, \\ x, & x \in \left[\frac{1}{\lambda_1}, \frac{1}{\lambda_2} \right], \\ \frac{1}{\lambda_1}, & x < \frac{1}{\lambda_1}. \end{cases}$$

Clearly, if $\lambda_1 = \lambda_2 = \lambda$, our definition agrees with (2). That is,

$$\frac{1}{\lambda} \psi(\lambda x) = \psi_{\lambda, \lambda}(x).$$

Proof of Proposition 4.7. First, for every $v \in S^{d-1}$, due to the assumption that the event of Lemma 4.6 holds, the uncorrupted observations satisfy

$$\left| \frac{1}{\lambda_v(Q)N} \sum_{i=1}^N \psi(\lambda_v(Q)\langle \tilde{X}_i, v \rangle^2) - \frac{1}{\lambda_v(Q)N} \sum_{i=1}^N \psi(\lambda_v(Q)\langle X_i, v \rangle^2) \right| \leq 2\eta Q \leq \frac{\varepsilon Q}{10}.$$

We centralize the empirical process to obtain that

$$\frac{1}{\lambda_v(Q)N} \sum_{i=1}^N \psi(\lambda_v(Q)\langle X_i, v \rangle^2) - v^\top \tilde{\Sigma} v = \frac{1}{\lambda_v^c(Q)N} \sum_{i=1}^N \psi(\lambda_v^c(Q)(\langle X_i, v \rangle^2 - v^\top \tilde{\Sigma} v)),$$

where $\lambda_v^c(Q) = (q_v + Q - v^\top \tilde{\Sigma} v)^{-1}$. To see why this identity holds, observe that both sides are equal except when $\langle X_i, v \rangle^2 - v^\top \tilde{\Sigma} v \leq -q_v - Q + v^\top \tilde{\Sigma} v$. In this case, $0 \leq \langle X_i, v \rangle^2 \leq -q_v - Q + 2v^\top \tilde{\Sigma} v$. However, since $Q \geq 2\|\Sigma\|$ if $C_{Q_0} \geq 2$ and $\varepsilon < 1$, the latter

inequality cannot be satisfied. Now, on the event of Lemma 4.6, $Q_0 \leq q_v - v^\top \tilde{\Sigma} v + Q \leq 5Q_0$ and therefore

$$\frac{1}{\lambda_v^\varepsilon(Q)N} \sum_{i=1}^N \psi(\lambda_v^\varepsilon(Q)(\langle X_i, v \rangle^2 - v^\top \tilde{\Sigma} v)) \leq \frac{1}{N} \sum_{i=1}^N \psi_{-1/Q_0, 1/(5Q_0)}(\langle X_i, v \rangle^2 - v^\top \tilde{\Sigma} v).$$

We define $\bar{U}_Q(v)$ to be the right-hand side in the inequality above. We first study how the empirical process of interest concentrates around the mean. To simplify the analysis, we consider a simpler quantity, $\bar{W}_Q(v) = \frac{1}{N} \sum_{i=1}^N \psi_{-1/(3Q), 1/(3Q)}(\langle X_i, v \rangle^2 - v^\top \tilde{\Sigma} v)$. Our starting point is

$$\begin{aligned} \sup_{v \in S^{d-1}} (\bar{U}_Q(v) - \mathbb{E} \bar{U}_Q(v)) &\leq \sup_{v \in S^{d-1}} (\bar{U}_Q(v) - \bar{W}_Q(v)) + \sup_{v \in S^{d-1}} (\bar{W}_Q(v) - \mathbb{E} \bar{W}_Q(v)) \\ &\quad + \sup_{v \in S^{d-1}} (\mathbb{E} \bar{W}_Q(v) - \mathbb{E} \bar{U}_Q(v)) \\ &= (a) + (b) + (c). \end{aligned}$$

For the first term, observe that it is non-zero if and only if the argument of the function is above $5Q_0$ or below $-Q_0$. The latter cannot happen because of the same contradiction argument as above. The magnitude of the difference is at most $3Q$. By Lemma 4.6, there are at most $\frac{\varepsilon N}{4}$ points in the range $(5Q_0, \infty)$, therefore, we obtain that the first term (a) is at most $\frac{3\varepsilon Q}{4}$. Similarly, we apply the Markov inequality together with the L_p - L_2 norm equivalence assumption to obtain that

$$\begin{aligned} \sup_{v \in S^{d-1}} \mathbb{E}(\bar{W}_Q(v) - \bar{U}_Q(v)) &\leq 3Q \mathbb{P}(\langle X, v \rangle^2 - v^\top \tilde{\Sigma} v > 5Q_0) \\ &\leq 3Q \left(\frac{\kappa(p)^2 \|\Sigma\|}{5Q_0} \right)^{p/2} \leq \frac{3\varepsilon Q}{32}. \end{aligned}$$

This concludes the bound for (c). The second term (b) is handled via Talagrand’s concentration inequality for empirical processes (Massart’s version [48]). Observe that, for every $v \in S^{d-1}$,

$$|\psi_{-1/(3Q), 1/(3Q)}(\langle X, v \rangle^2 - v^\top \tilde{\Sigma} v)| \leq 3Q,$$

and

$$\mathbb{E} |\psi_{-1/(3Q), 1/(3Q)}(\langle X, v \rangle^2)|^2 \leq \mathbb{E} \langle X, v \rangle^4 \leq \kappa(4)^4 \|\Sigma\|^2.$$

Moreover, for every $Q > 0$, $\psi(\cdot)_{-1/Q, 1/Q}$ is a 1-Lipschitz function that passes through the origin. By the same symmetrization and contraction arguments as in the proof of Lemma 4.6, we obtain that

$$\mathbb{E} \sup_{v \in S^{d-1}} |\bar{W}_Q(v) - \mathbb{E} \bar{W}_Q(v)| \leq 2\mathbb{E} \sup_{v \in S^{d-1}} \left| \sum_{i=1}^N \varepsilon_i \langle X_i, v \rangle^2 \right| \leq 2C_{BY} \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}}.$$

Therefore, by Talagrand’s concentration inequality for empirical processes, with probability at least $1 - 2e^{-x}$, it holds that

$$\sup_{v \in S^{d-1}} |W_Q(v) - \mathbb{E} W_Q(v)| \leq 4C_{BY} \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{N}} + \kappa(4)^2 \|\Sigma\| \sqrt{\frac{8x}{N}} + 105Q \frac{x}{N}.$$

We choose $x = 2 \log(\frac{2}{\delta})$ to obtain that the right-hand side above is less than $\frac{\varepsilon Q}{2}$. Now, the problem boils down to control uniformly the typical magnitude of $\bar{U}_Q(v)$. For simplicity, we define $\bar{X}_i(v) = \langle X_i, v \rangle^2 - v^\top \tilde{\Sigma} v$ and write

$$\begin{aligned} \sup_{v \in S^{d-1}} \mathbb{E} \bar{U}_Q(v) &= \sup_{v \in S^{d-1}} \mathbb{E} [\psi_{-1/2Q_0, 1/(5Q_0)}(\langle X, v \rangle^2 - v^\top \tilde{\Sigma} v) \\ &\quad \times (\mathbb{1}\{\bar{X}_i(v) \leq 5Q_0\} + \mathbb{1}\{\bar{X}_i(v) > 5Q_0\})] \\ &\leq \sup_{v \in S^{d-1}} 5Q_0 \mathbb{E} \mathbb{1}\{\bar{X}(v) \geq 5Q_0\} \leq 5Q \frac{\varepsilon}{64}. \end{aligned}$$

The final bound for the first part becomes

$$\varepsilon Q \left(\frac{1}{10} + \frac{3}{4} + \frac{3}{32} + \frac{5}{64} + \frac{1}{2} \right) \leq 2\varepsilon Q.$$

For the other side bound, observe that

$$\begin{aligned} v^\top \tilde{\Sigma} v - \frac{1}{\lambda_v(Q)N} \sum_{i=1}^N \psi(\lambda_v(Q) \langle X_i, v \rangle^2) &= \frac{1}{\lambda_v^c(Q)N} \sum_{i=1}^N \psi(\lambda_v^c(Q) v^\top \tilde{\Sigma} v - \langle X_i, v \rangle^2) \\ &\leq \frac{1}{N} \sum_{i=1}^N \psi_{-\frac{1}{Q_0}, \frac{1}{5Q_0}}(v^\top \tilde{\Sigma} v - \langle X_i, v \rangle^2). \end{aligned}$$

The same analysis follows and concludes the proof. ■

Proof of Theorem 4.3. The proof of Theorem 4.3 follows directly from Lemma 4.6 and Proposition 4.7. Indeed, consider the event \mathcal{E} where the events of both Lemma 4.6 and Proposition 4.7 hold. It occurs with probability at least $1 - \delta$ as desired. Our next argument is standard in the literature [47] and is similar to Lepskii’s method [42]. Let i_0 be the smallest integer satisfying $2^{i_0} \in (2Q_0, 4Q_0)$. By Lemma 4.6, we know that at most $\frac{\varepsilon N}{4}$ samples are outside of the range $q_v + 2^{i_0}$. We conclude that the difference between $U_Q(v)$ and $U_{2Q}(v)$ is at most $\frac{\varepsilon Q}{2}$. By induction, this implies that $\Gamma(2^{i_0}) \subset \Gamma(2^{i_0+1})$ for every $i \geq i_0$. This implies that corresponding sets are nested. The diameter of the set $\bigcap_{i \geq i^*} \Gamma(2^i)$ is at most $8\varepsilon Q$ and the matrix of second moments $\tilde{\Sigma}$ belongs to an even smaller set. Thus, $\|\hat{\Sigma} - \tilde{\Sigma}\| \leq 8\varepsilon Q$. We conclude by applying Proposition 4.5. ■

5. Optimality

It is natural to ask about the optimality of our statistical guarantees of Theorem 1.1. It is well known that the term $C \|\Sigma\| \sqrt{\frac{r(\Sigma)}{N}}$ appears in the minimax lower bound for covariance estimation [43, Theorem 2] as well as in the lower bound for the performance of the sample covariance matrix in the Gaussian case [39, Theorem 4]. Clearly, this term is necessary because the adversary can always leave the sample uncorrupted and the multivariate Gaussian distribution satisfies the norm equivalence assumption (1). Further, we show that the term $C \|\Sigma\| \sqrt{\frac{\log(1/\delta)}{N}}$ cannot be improved in general. As shown

in [53] (following the lower bound in [45]), the term scaling as $R\sqrt{\frac{\log(1/\delta)}{N}}$ appears in the lower bound for any estimator of the covariance matrix. Here R^2 is the so-called *weak variance* defined by

$$R^2 = \sup_{v \in \mathcal{S}^{d-1}} \mathbb{E}(v^\top (X \otimes X - \mathbb{E}X \otimes X)v)^2.$$

In the multivariate Gaussian case, using the standard relation between moments, we get

$$R^2 \geq \sup_{v \in \mathcal{S}^{d-1}} \mathbb{E}\langle v, X \rangle^4 - \|\Sigma\|^2 = 3 \sup_{v \in \mathcal{S}^{d-1}} (\mathbb{E}\langle v, X \rangle)^2 - \|\Sigma\|^2 = 2\|\Sigma\|^2.$$

This shows the necessity of the term $C\|\Sigma\|\sqrt{\frac{\log(1/\delta)}{N}}$ in Theorem 1.1.

Finally, we can restrict the discussion to the convergence rate with respect to the fraction of corrupted samples η . Some closely related lower bounds with matching rates of convergence appear in the analysis of the sparse vector model [20, Theorems 5 and 6]. Their model can be seen as a special case of our setup when $d = 1$. Instead of exploiting their techniques for the lower bounds, we provide a separate analysis using explicit examples. Our key argument to derive the minimax optimality with respect to η is to reduce our problem to a mean estimation problem. For the mean estimation, a simple computation shows a lower bound with respect to η . To describe this result, we first consider a basic definition.

Definition 5.1. For a random variable X , we define the quantile

$$Q_q(X) = \sup\{M \in \mathbb{R} : \mathbb{P}(X \geq M) \geq 1 - q\}.$$

Formally, the following simple lower bound holds.

Proposition 5.2 ([47, (2.3)]). *Let X be a random variable with mean μ , variance σ_X , and with an absolutely continuous distribution. Suppose that $\tilde{X}_1, \dots, \tilde{X}_N$ is an η -corrupted sample that sampled according to the distribution of X . Let $\bar{X} = X - \mathbb{E}X$ and define $\epsilon(\bar{X}, \eta)$ as follows:*

$$\begin{aligned} \epsilon(\bar{X}, \eta) = \max\{ & \mathbb{E}|\bar{X} - Q_{\eta/2}(\bar{X})| \mathbb{1}\{\bar{X} \leq Q_{\eta/2}(\bar{X})\}, \\ & \mathbb{E}|\bar{X} - Q_{1-\eta/2}(\bar{X})| \mathbb{1}\{\bar{X} \geq Q_{1-\eta/2}(\bar{X})\}\}. \end{aligned}$$

Then, no estimator $\hat{\mu} = \hat{\mu}(\tilde{X}_1, \dots, \tilde{X}_N)$ of the mean μ can perform better than

$$|\hat{\mu} - \mu| \leq \epsilon(\bar{X}, \eta).$$

To obtain the minimax rates with respect to η , it is enough to restrict our attention to the one-dimensional case. Assume that the distribution of a zero mean random variable X satisfies the L_4 - L_2 norm equivalence (1). That is, for some $\kappa \geq 1$, it holds that

$$(\mathbb{E}X^4)^{1/4} \leq \kappa(\mathbb{E}X^2)^{1/2}.$$

Denote $Y = X^2$ and observe that Y is non-negative. In this case, the estimation of the variance of X can be seen as the estimation of the mean of Y . The norm equivalence assumption can be rewritten as $(\mathbb{E}Y^2)^{1/2} \leq \kappa^2\mathbb{E}Y$.

Example 5.3 (Optimality of the $\sqrt{\eta}$ -term when $p = 4$). First, a few intermediate random variables are considered before the definition of Y . Define

$$Y_1 = \begin{cases} -\frac{1}{\sqrt{\eta}} & \text{with probability } \frac{\eta}{2}, \\ -1 & \text{with probability } \frac{1-\eta}{2}, \\ 1 & \text{with probability } \frac{1-\eta}{2}, \\ \frac{1}{\sqrt{\eta}} & \text{with probability } \frac{\eta}{2}. \end{cases}$$

Clearly, $\mathbb{E}Y_1 = 0$. Let σ_{Y_1} denote the standard deviation of Y_1 . It holds that

$$\sigma_{Y_1} = (\mathbb{E}Y_1^2)^{1/2} = \sqrt{2-\eta} \leq \sqrt{2}.$$

Assume that $\eta \leq \frac{1}{4}$. We have that $Q_{\eta/2}(Y_1) = -1$. Indeed, $\mathbb{P}(Y_1 \geq -1) = 1 - \frac{\eta}{2}$ and $\mathbb{P}(Y_1 \geq 1) = \frac{1}{2} < 1 - \frac{\eta}{2}$. We conclude that

$$\epsilon(Y_1, \eta) \geq \mathbb{E}|Y_1 + 1| \mathbb{1}\{Y_1 \leq -1\} = \left(\frac{1}{\sqrt{\eta}} - 1\right) \frac{\eta}{2} = \frac{\sqrt{\eta}}{2} - \frac{\eta}{2} \geq \frac{1}{4} \sqrt{\eta}.$$

We now consider the random variable

$$Y_2 = \frac{\sigma^2}{\sqrt{2-\eta}} Y_1,$$

where $\sigma > 0$ is a positive real number. It still holds that $\mathbb{E}Y_2 = 0$, but now, $(\mathbb{E}Y_2^2)^{1/2} = \sigma^2$. By homogeneity, we have

$$\epsilon(Y_2, \eta) \geq \frac{\sigma^2 \sqrt{\eta}}{4\sqrt{2-\eta}} \geq \frac{\sigma^2 \sqrt{\eta}}{4\sqrt{2}}.$$

Finally, observe that $Y = Y_2 + \|Y_2\|_\infty$ is a non-negative random variable and $\epsilon(Y, \eta) = \epsilon(Y_2, \eta)$ since the difference between Y and Y_2 is a constant. We conclude by observing that $\mathbb{E}Y = \|Y_2\|_\infty$ and $\mathbb{E}Y^2 = \mathbb{E}Y_2^2 + \|Y_2\|_\infty^2 \leq 2\|Y_2\|_\infty^2$. Thus, the norm equivalence assumption $(\mathbb{E}Y^2)^{1/2} \leq \kappa^2 \mathbb{E}Y$ holds with $\kappa^2 = \sqrt{2}$.

We now present an example that attains the sub-exponential minimax rate in the mean estimation and therefore confirms the optimality of our covariance estimation results in the sub-Gaussian case. Indeed, observe that if X is a sub-Gaussian random variable, then $Y = X^2$ is a sub-exponential random variable (see [72, Lemma 2.7.6]). Therefore, when making a reduction from covariance estimation to a mean estimation problem as in Example 5.3, we have to analyze the mean estimation problem in the sub-exponential regime. We note that a similar sub-Gaussian construction appears in [47, Remark before Section 2.1].²

²To make their construction work in the Gaussian case, we changed the definition of Q , so that it becomes an $\frac{\eta}{4}$ -quantile instead of an $\frac{\eta}{2}$ -quantile claimed in [47].

Example 5.4 (Optimality of $\eta \log(\frac{1}{\eta})$ -term in sub-Gaussian covariance estimation). Let

$$Y = \min\{1, |g|^2\} \mathbb{1}\{|g|^2 < Q\} + |g|^2 \mathbb{1}\{|g|^2 \geq Q\},$$

where g is a standard Gaussian random variable and Q is such that $\mathbb{P}(|g|^2 \geq Q) = \frac{\eta}{4}$. Consequently, we have $\mathbb{P}(|g| \geq \sqrt{Q}) = \frac{\eta}{4}$. It is easy to see that for η sufficiently small, it holds that $Q > 1$. Let us now prove that $Q_{1-\eta/2}(Y) = 1$. Indeed, since $\Pr(Y \geq 1) = \Pr(|g| \geq 1) > \frac{\eta}{2}$ for small enough η , then $Q_{1-\eta/2}(Y) \geq 1$. It is clear that if $Q_{1-\eta/2}(Y) > 1$, then $Q_{1-\eta/2}(Y) \geq Q$, which contradicts $\mathbb{P}(|g|^2 \geq Q) = \frac{\eta}{4} < \frac{\eta}{2}$. Thus, we have

$$\varepsilon(\bar{Y}, \eta) \geq \mathbb{E}|\bar{Y} - Q_{1-\eta/2}(\bar{Y})| \mathbb{1}\{\bar{Y} \geq Q_{1-\eta/2}(\bar{Y})\} = \mathbb{E}(Y - 1) \mathbb{1}\{Y \geq 1\}.$$

This leads to

$$\begin{aligned} \mathbb{E}(Y - 1) \mathbb{1}\{Y \geq 1\} &= \mathbb{E}Y \mathbb{1}\{1 \leq Y < Q\} + \mathbb{E}Y \mathbb{1}\{Y \geq Q\} - \Pr(Y \geq 1) \\ &= \Pr(1 \leq |g|^2 < Q) + \mathbb{E}|g|^2 \mathbb{1}\{|g|^2 \geq Q\} - \Pr(|g|^2 \geq 1) \\ &= \mathbb{E}|g|^2 \mathbb{1}\{|g|^2 \geq Q\} - \Pr(|g|^2 \geq Q). \end{aligned}$$

Using the standard computation (see [72, Exercise 2.1.4]), we obtain

$$\mathbb{E}|g|^2 \mathbb{1}\{|g|^2 \geq Q\} = \frac{2}{\sqrt{2\pi}} \sqrt{Q} \exp\left(-\frac{Q}{2}\right) + \Pr(|g| \geq \sqrt{Q}).$$

Therefore, since $\frac{\eta}{4} = \mathbb{P}(|g| \geq \sqrt{Q})$, and by the standard Gaussian integration [72, Proposition 2.1.2], we have

$$\varepsilon(\bar{Y}, \eta) \geq \frac{2}{\sqrt{2\pi}} \sqrt{Q} \exp\left(-\frac{Q}{2}\right) \geq Q \Pr(|g| \geq \sqrt{Q}) = \frac{Q\eta}{4}.$$

Using the same Gaussian integration formula, we conclude that, in order to obtain $\Pr(|g| \geq \sqrt{Q}) = \frac{\eta}{4}$, we need to choose $Q \sim \log(\frac{1}{\eta})$. This implies that

$$\varepsilon(\bar{Y}, \eta) \geq \frac{Q\eta}{4} \gtrsim \eta \log\left(\frac{1}{\eta}\right).$$

We claim that the rate $\eta \log(\frac{1}{\eta})$ is sharp. To do so, we compute the upper bound with respect to η in Theorem 1.1. If the original random variable X is a sub-Gaussian, then $\kappa(p)^2 \sim p$ for every $p \geq 1$ (see [72, Proposition 2.5.2]) and the upper bound scales as $\frac{\eta p}{\eta^{2/p}}$. We choose $p \sim \log(\frac{1}{\eta})$ to optimize the fraction and obtain the claimed rate.

Finally, we note that for the Gaussian case, it is possible to achieve a slightly improved dependence on η by eliminating the logarithmic factor $\log(\frac{1}{\eta})$ (as shown in the follow-up work of Minasyan and the second-named author of this paper [54]), owing to the rotational invariance property of the Gaussian distribution. As discussed in [54], the estimator needs to be tailored to the Gaussian distribution and therefore it cannot be applied to the more general setup considered in our Theorem 1.1.

Acknowledgments. The authors would like to thank Afonso Bandeira for several valuable discussions, and Alexander Tsybakov for pointing out the connection with [20]. The authors would also like to thank anonymous referees for valuable comments that improved the manuscript.

Funding. Nikita Zhivotovskiy is funded in part by ETH Foundations of Data Science (ETH-FDS).

References

- [1] Adamczak, R., Litvak, A. E., Pajor, A., Tomczak-Jaegermann, N.: [Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles](#). *J. Amer. Math. Soc.* **23**, 535–561 (2010) Zbl [1206.60006](#) MR [2601042](#)
- [2] Appert, G., Catoni, O.: [New bounds for \$k\$ -means and information \$k\$ -means](#). arXiv:[2101.05728](#) (2021)
- [3] Audibert, J.-Y., Catoni, O.: [Robust linear least squares regression](#). *Ann. Statist.* **39**, 2766–2794 (2011) Zbl [1231.62126](#) MR [2906886](#)
- [4] Bai, Z. D., Yin, Y. Q.: [Limit of the smallest eigenvalue of a large dimensional sample covariance matrix](#). *Ann. Probab.* **21**, 1275–1294 (1993) Zbl [0779.60026](#) MR [1235416](#)
- [5] Bickel, P. J.: [On some robust estimates of location](#). *Ann. Math. Statist.* **36**, 847–858 (1965) Zbl [0192.25802](#) MR [177484](#)
- [6] Bickel, P. J., Levina, E.: [Covariance regularization by thresholding](#). *Ann. Statist.* **36**, 2577–2604 (2008) Zbl [1196.62062](#) MR [2485008](#)
- [7] Bickel, P. J., Levina, E.: [Regularized estimation of large covariance matrices](#). *Ann. Statist.* **36**, 199–227 (2008) Zbl [1132.62040](#) MR [2387969](#)
- [8] Bourgain, J.: [Random points in isotropic convex sets](#). In: *Convex geometric analysis* (Berkeley, CA, 1996), Math. Sci. Res. Inst. Publ. 34, Cambridge University Press, Cambridge, 53–58 (1999) Zbl [0941.52003](#) MR [1665576](#)
- [9] Brownlees, C., Joly, E., Lugosi, G.: [Empirical risk minimization for heavy-tailed losses](#). *Ann. Statist.* **43**, 2507–2536 (2015) Zbl [1326.62066](#) MR [3405602](#)
- [10] Cai, T. T., Ren, Z., Zhou, H. H.: [Optimal rates of convergence for estimating Toeplitz covariance matrices](#). *Probab. Theory Related Fields* **156**, 101–143 (2013) Zbl [06176807](#) MR [3055254](#)
- [11] Cai, T. T., Ren, Z., Zhou, H. H.: [Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation](#). *Electron. J. Stat.* **10**, 1–59 (2016) Zbl [1331.62272](#) MR [3466172](#)
- [12] Cai, T. T., Zhang, C.-H., Zhou, H. H.: [Optimal rates of convergence for covariance matrix estimation](#). *Ann. Statist.* **38**, 2118–2144 (2010) Zbl [1202.62073](#) MR [2676885](#)
- [13] Catoni, O.: [Challenging the empirical mean and empirical variance: A deviation study](#). *Ann. Inst. Henri Poincaré Probab. Stat.* **48**, 1148–1185 (2012) Zbl [1282.62070](#) MR [3052407](#)
- [14] Catoni, O.: [PAC-Bayesian bounds for the Gram matrix and least squares regression with a random design](#). arXiv:[1603.05229](#) (2016)
- [15] Catoni, O., Giulini, I.: [Dimension-free PAC-Bayesian bounds for matrices, vectors, and linear least squares regression](#). arXiv:[1712.02747](#) (2017)
- [16] Chen, M., Gao, C., Ren, Z.: [Robust covariance and scatter matrix estimation under Huber’s contamination model](#). *Ann. Statist.* **46**, 1932–1960 (2018) Zbl [1408.62104](#) MR [3845006](#)
- [17] Cheng, Y., Diakonikolas, I., Ge, R., Woodruff, D. P.: [Faster algorithms for high-dimensional robust covariance estimation](#). In: *Proceedings of the Thirty-Second Conference on Learning Theory, Proceedings of Machine Learning Research 99*, ML Research Press, 727–757 (2019)

- [18] Cherapanamjeri, Y., Hopkins, S. B., Kathuria, T., Raghavendra, P., Tripuraneni, N.: [Algorithms for heavy-tailed statistics: regression, covariance estimation, and beyond](#). In: STOC '20—Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, Association for Computing Machinery, New York, 601–609 (2020) Zbl [07298273](#) MR [4141785](#)
- [19] Chinot, G., Lecué, G., Lerasle, M.: [Robust statistical learning with Lipschitz and convex loss functions](#). Probab. Theory Related Fields **176**, 897–940 (2020) Zbl [1436.62178](#) MR [4087486](#)
- [20] Comminges, L., Collier, O., Ndaoud, M., Tsybakov, A. B.: [Adaptive robust estimation in sparse vector model](#). Ann. Statist. **49**, 1347–1377 (2021) Zbl [1476.62063](#) MR [4298867](#)
- [21] Depersin, J.: [Robust subgaussian estimation with VC-dimension](#). Ann. Inst. Henri Poincaré Probab. Stat. **60**, 971–989 (2024) MR [4757514](#)
- [22] Diakonikolas, I., Kane, D. M.: Recent advances in algorithmic high-dimensional robust statistics. arXiv:[1911.05911](#) (2019)
- [23] Diakonikolas, I., Kane, D. M., Lee, J. C. H., Pensia, A.: Outlier-robust sparse mean estimation for heavy-tailed distributions. In: Advances in neural information processing systems, NeurIPS Proceedings 35, Curran Associates, 5164–5177 (2022)
- [24] El Karoui, N.: [Operator norm consistent estimation of large-dimensional sparse covariance matrices](#). Ann. Statist. **36**, 2717–2756 (2008) Zbl [1196.62064](#) MR [2485011](#)
- [25] El Karoui, N.: [Random matrices and high-dimensional statistics: beyond covariance matrices](#). In: Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures, World Scientific, Hackensack, NJ, 2857–2876 (2018) Zbl [1451.62054](#) MR [3966514](#)
- [26] Giné, E., Zinn, J.: [Some limit theorems for empirical processes](#). Ann. Probab. **12**, 929–998 (1984) Zbl [0553.60037](#) MR [757767](#)
- [27] Giulini, I.: [Robust PCA and pairs of projections in a Hilbert space](#). Electron. J. Stat. **11**, 3903–3926 (2017) Zbl [1384.62185](#) MR [3714302](#)
- [28] Giulini, I.: [Robust dimension-free Gram operator estimates](#). Bernoulli **24**, 3864–3923 (2018) Zbl [1415.62012](#) MR [3788191](#)
- [29] Guédon, O., Litvak, A. E., Pajor, A., Tomczak-Jaegermann, N.: [On the interval of fluctuation of the singular values of random matrices](#). J. Eur. Math. Soc. (JEMS) **19**, 1469–1505 (2017) Zbl [1405.60011](#) MR [3635358](#)
- [30] Hampel, F. R.: [A general qualitative definition of robustness](#). Ann. Math. Statist. **42**, 1887–1896 (1971) Zbl [0229.62041](#) MR [301858](#)
- [31] Hsu, D., Sabato, S.: Loss minimization and parameter estimation with heavy tails. J. Mach. Learn. Res. **17**, article no. 18 (2016) Zbl [1360.62380](#) MR [3491112](#)
- [32] Huber, P. J.: [Robust estimation of a location parameter](#). Ann. Math. Statist. **35**, 73–101 (1964) Zbl [0136.39805](#) MR [161415](#)
- [33] Huber, P. J.: [Robust statistics](#). Wiley Ser. Probab. Math. Stat., John Wiley & Sons, New York (1981) Zbl [0536.62025](#) MR [606374](#)
- [34] Kannan, R., Lovász, L., Simonovits, M.: [Random walks and an \$O^*\(n^5\)\$ volume algorithm for convex bodies](#). Random Structures Algorithms **11**, 1–50 (1997) Zbl [0895.60075](#) MR [1608200](#)
- [35] Ke, Y., Minsker, S., Ren, Z., Sun, Q., Zhou, W.-X.: [User-friendly covariance estimation for heavy-tailed distributions](#). Statist. Sci. **34**, 454–471 (2019) Zbl [1429.62312](#) MR [4017523](#)
- [36] Klochov, Y., Kroshnin, A., Zhivotovskiy, N.: [Robust \$k\$ -means clustering for distributions with two moments](#). Ann. Statist. **49**, 2206–2230 (2021) Zbl [1487.62070](#) MR [4319247](#)
- [37] Koltchinskii, V.: [Oracle inequalities in empirical risk minimization and sparse recovery problems](#). Lecture Notes in Math. 2033, Springer, Heidelberg (2011) Zbl [1223.91002](#) MR [2829871](#)

- [38] Koltchinskii, V.: [Asymptotically efficient estimation of smooth functionals of covariance operators](#). *J. Eur. Math. Soc. (JEMS)* **23**, 765–843 (2021) Zbl 1461.62068 MR 4210724
- [39] Koltchinskii, V., Lounici, K.: [Concentration inequalities and moment bounds for sample covariance operators](#). *Bernoulli* **23**, 110–133 (2017) Zbl 1366.60057 MR 3556768
- [40] Lam, C., Fan, J.: [Sparsistency and rates of convergence in large covariance matrix estimation](#). *Ann. Statist.* **37**, 4254–4278 (2009) Zbl 1191.62101 MR 2572459
- [41] Ledoux, M., Talagrand, M.: [Probability in Banach spaces. Isoperimetry and processes](#). *Classics in Math.*, Springer, Berlin (2011) Zbl 1226.60003 MR 2814399
- [42] Lepskii, O. V.: [Asymptotically minimax adaptive estimation. I: Upper bounds. Optimally adaptive estimates](#). *Theory Probab. Appl.* **36**, 682–697 (1991) Zbl 0776.62039 MR 1147167
- [43] Lounici, K.: [High-dimensional covariance matrix estimation with missing observations](#). *Bernoulli* **20**, 1029–1058 (2014) Zbl 1320.62124 MR 3217437
- [44] Lugosi, G., Mendelson, S.: [Mean estimation and regression under heavy-tailed distributions: A survey](#). *Found. Comput. Math.* **19**, 1145–1190 (2019) Zbl 1431.62123 MR 4017683
- [45] Lugosi, G., Mendelson, S.: [Near-optimal mean estimators with respect to general norms](#). *Probab. Theory Related Fields* **175**, 957–973 (2019) Zbl 1431.62234 MR 4026610
- [46] Lugosi, G., Mendelson, S.: [Risk minimization by median-of-means tournaments](#). *J. Eur. Math. Soc. (JEMS)* **22**, 925–965 (2020) Zbl 1436.62312 MR 4055993
- [47] Lugosi, G., Mendelson, S.: [Robust multivariate mean estimation: The optimality of trimmed mean](#). *Ann. Statist.* **49**, 393–410 (2021) Zbl 1461.62069 MR 4206683
- [48] Massart, P.: [About the constants in Talagrand’s concentration inequalities for empirical processes](#). *Ann. Probab.* **28**, 863–884 (2000) Zbl 1140.60310 MR 1782276
- [49] Mendelson, S.: [An unrestricted learning procedure](#). *J. ACM* **66**, article no. 42 (2019) Zbl 1473.68156 MR 4040282
- [50] Mendelson, S.: [Approximating \$L_p\$ unit balls via random sampling](#). *Adv. Math.* **386**, article no. 107829 (2021) Zbl 1469.46012 MR 4270522
- [51] Mendelson, S., Paouris, G.: [On generic chaining and the smallest singular value of random matrices with heavy tails](#). *J. Funct. Anal.* **262**, 3775–3811 (2012) Zbl 1242.60008 MR 2899978
- [52] Mendelson, S., Paouris, G.: [On the singular values of random matrices](#). *J. Eur. Math. Soc. (JEMS)* **16**, 823–834 (2014) Zbl 1290.60006 MR 3191978
- [53] Mendelson, S., Zhivotovskiy, N.: [Robust covariance estimation under \$L_4\$ – \$L_2\$ norm equivalence](#). *Ann. Statist.* **48**, 1648–1664 (2020) Zbl 1451.62084 MR 4124338
- [54] Minasyan, A., Zhivotovskiy, N.: [Statistically optimal robust mean and covariance estimation for anisotropic Gaussians](#). *arXiv:2301.09024* (2023)
- [55] Minsker, S., Wang, L.: [Robust estimation of covariance matrices: Adversarial contamination and beyond](#). *arXiv:2203.02880* (2022)
- [56] Mourtada, J., Vaškevičius, T., Zhivotovskiy, N.: [Distribution-free robust linear regression](#). *Math. Stat. Learn.* **4**, 253–292 (2021) Zbl 1493.62429 MR 4383735
- [57] Oliveira, R. I., Rico, Z. F.: [Improved covariance estimation: optimal robustness and sub-Gaussian guarantees under heavy tails](#). *arXiv:2209.13485* (2024)
- [58] Ostrovskii, D. M., Rudi, A.: [Affine invariant covariance estimation for heavy-tailed distributions](#). *Proceedings of Machine Learning Research* 99, ML Research Press, 2531–2550 (2019)
- [59] Petukhina, A., Klochkov, Y., Härdle, W. K., Zhivotovskiy, N.: [Robustifying Markowitz](#). *J. Econometrics* **239**, article no. 105387 (2024) Zbl 07814007 MR 4708614
- [60] Rico, Z. F.: [Optimal statistical estimation: sub-Gaussian properties, heavy-tailed data, and robustness](#). Ph.D. thesis, Instituto de Matemática Pura e Aplicada (2022)
- [61] Rousseeuw, P. J., Leroy, A. M.: [Robust regression and outlier detection](#). *Wiley Ser. Probab. Math. Stat.*, John Wiley & Sons, New York (1987) Zbl 0711.62030 MR 914792

- [62] Srivastava, N., Vershynin, R.: [Covariance estimation for distributions with \$2 + \varepsilon\$ moments](#). *Ann. Probab.* **41**, 3081–3111 (2013) Zbl [1293.62121](#) MR [3127875](#)
- [63] Stigler, S. M.: [The asymptotic distribution of the trimmed mean](#). *Ann. Statist.* **1**, 472–477 (1973) Zbl [0261.62016](#) MR [359134](#)
- [64] Talagrand, M.: [Upper and lower bounds for stochastic processes. Modern methods and classical problems](#). *Ergeb. Math. Grenzgeb. (3)* 60, Springer, Heidelberg (2014) Zbl [1293.60001](#) MR [3184689](#)
- [65] Tikhomirov, K.: [Sample covariance matrices of heavy-tailed distributions](#). *Int. Math. Res. Not. IMRN* **2018**, 6254–6289 (2018) Zbl [1462.62330](#) MR [3872323](#)
- [66] Tropp, J. A.: [An introduction to matrix concentration inequalities](#). *Found. Trends Mach. Learn.* **8**, 1–230 (2015) Zbl [1391.15071](#)
- [67] Tukey, J. W., McLaughlin, D. H.: [Less vulnerable confidence and significance procedures for location based on a single sample: Trimming/Winsorization. I](#). *Sankhyā Ser. A* **25**, 331–352 (1963) Zbl [0116.10904](#) MR [169354](#)
- [68] van Handel, R.: [Structured random matrices](#). In: *Convexity and concentration, IMA Vol. Math. Appl.* 161, Springer, New York, 107–156 (2017) Zbl [1376.15029](#) MR [3837269](#)
- [69] Vershynin, R.: [Approximating the moments of marginals of high-dimensional distributions](#). *Ann. Probab.* **39**, 1591–1606 (2011) Zbl [1271.62122](#) MR [2857251](#)
- [70] Vershynin, R.: [How close is the sample covariance matrix to the actual covariance matrix?](#) *J. Theoret. Probab.* **25**, 655–686 (2012) Zbl [1365.62208](#) MR [2956207](#)
- [71] Vershynin, R.: [Introduction to the non-asymptotic analysis of random matrices](#). In: *Compressed sensing*, Cambridge University Press, Cambridge, 210–268 (2012) MR [2963170](#)
- [72] Vershynin, R.: [High-dimensional probability](#). *Camb. Ser. Stat. Probab. Math.* 47, Cambridge University Press, Cambridge (2018) Zbl [1430.60005](#) MR [3837109](#)
- [73] Zhivotovskiy, N.: [Dimension-free bounds for sums of independent matrices and simple tensors via the variational principle](#). *Electron. J. Probab.* **29**, article no. 13 (2024) Zbl [1531.60019](#) MR [4693860](#)