

---

# Robust Testing in High-Dimensional Sparse Models

Supplementary (full version)

---

**Anand Jerry George**

School of Computer and Communication Sciences  
École Polytechnique Fédérale de Lausanne (EPFL)  
anand.george@epfl.ch

**Clément L. Canonne**

School of Computer Science  
The University of Sydney  
clement.canonne@sydney.edu.au

## Abstract

We consider the problem of robustly testing the norm of a high-dimensional sparse signal vector under two different observation models. In the first model, we are given  $n$  i.i.d. samples from the distribution  $\mathcal{N}(\theta, I_d)$  (with unknown  $\theta$ ), of which a small fraction has been arbitrarily corrupted. Under the promise that  $\|\theta\|_0 \leq s$ , we want to correctly distinguish whether  $\|\theta\|_2 = 0$  or  $\|\theta\|_2 > \gamma$ , for some input parameter  $\gamma > 0$ . We show that any algorithm for this task requires  $n = \Omega(s \log \frac{ed}{s})$  samples, which is tight up to logarithmic factors. We also extend our results to other common notions of sparsity, namely,  $\|\theta\|_q \leq s$  for any  $0 < q < 2$ . In the second observation model that we consider, the data is generated according to a sparse linear regression model, where the covariates are i.i.d. Gaussian and the regression coefficient (signal) is known to be  $s$ -sparse. Here too we assume that an  $\varepsilon$ -fraction of the data is arbitrarily corrupted. We show that any algorithm that reliably tests the norm of the regression coefficient requires at least  $n = \Omega(\min(s \log d, 1/\gamma^4))$  samples. Our results show that the complexity of testing in these two settings significantly increases under robustness constraints. This is in line with the recent observations made in robust mean testing and robust covariance testing.

## 1 Introduction

Hypothesis testing is a fundamental task in statistics and a staple of the scientific method, in which we seek to test the validity of a pre-specified hypothesis based on empirical observations. In this work, we are concerned with the problem of testing whether a given high-dimensional sparse signal vector is zero under two common and well-studied observation models: 1) the Gaussian location model and 2) the Gaussian linear regression model. Specifically, in the Gaussian location model, we observe i.i.d. samples from a  $d$ -dimensional spherical Gaussian distribution with an unknown sparse mean vector, and seek to detect whether its  $\ell_2$  norm is large (equivalently, we observe a set of measurements subject to white noise, and seek to determine whether there exists an underlying (sparse) signal). Similarly, in the Gaussian linear regression model we seek to detect whether the  $\ell_2$  norm of the sparse regression coefficient is large. We further assume that our samples are imperfect or even corrupted, allowing an adversary to arbitrarily tamper with up to an  $\varepsilon$ -fraction of the observations. Our objective is to characterize the minimum number of samples required to perform these testing tasks, and, crucially, to understand the effect that requiring robustness to this adversarial corruption has on the complexity of the problems.

It is known [DKS17, DK21] that, for a variety of high-dimensional tasks, robust testing becomes as costly (in terms of sample complexity) as the corresponding *estimation* task. This is in contrast to the non-robust version, where testing is typically much more efficient – often by a quadratic factor in the dimension. However, it is unclear how sparsity enters the picture, and for instance if robustness only starts becoming “costly” when the signal vector is sufficiently dense – i.e., whether the problem

exhibits a phase transition. This is particularly relevant, as the non-robust versions of the problems we consider are known to present such a phase transition at sparsity  $s \approx \sqrt{d}$ .

*How does robustness affect the sample and computational complexities of testing norm of the signal vector in high-dimensional sparse models? Does testing remain easier than learning?*

This type of question, framed in a minimax setting, sits at the intersection of theoretical computer science (where it is captured under the framework of distribution testing) and robust statistics. Specifically, for  $\Theta \subseteq \mathbb{R}^d$ , let  $\{\mathbf{p}_\theta\}_{\theta \in \Theta}$  be a family of distributions and  $\mathbf{p}_0$  be a reference distribution (simple null hypothesis—in our case the standard Gaussian). Then, we say that an algorithm  $T$  reliably tests the  $\ell_2$ -norm of  $\theta$  if it satisfies the condition

$$\max \left\{ \Pr_{X \sim \mathbf{p}_0} [T(X) = \text{reject}], \sup_{\substack{\theta \in \Theta \\ \|\theta\|_2 \geq \gamma}} \Pr_{X \sim \mathbf{p}_\theta} [T(X) = \text{accept}] \right\} \leq \delta \quad (1)$$

where  $\delta \in (0, 1]$  is the failure probability, which following the literature we will hereafter set to  $1/3$ .<sup>1</sup> The quantity of interest here is the *sample complexity*, that is the minimum number of samples  $n$  required by any algorithm to solve the problem.

The above task, however, assumes access to “perfect” samples from the unknown distribution  $\mathbf{p}_\theta$ . This is often an unrealistic assumption, as a fraction of the  $n$  samples could be imperfect or corrupted. This motivates the setting of *robust testing*. The problem is then similar to the formulation in (1), with a crucial difference: the algorithm  $T$  does not have access to the i.i.d. samples  $X = (X_1, \dots, X_n) \in \mathbb{R}^{d \times n}$ , but instead to a “contaminated” version  $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_n) \in \mathbb{R}^{d \times n}$  obtained by arbitrarily modifying up to  $\varepsilon n$  of the  $X_i$ ’s (i.e., an  $\varepsilon$ -fraction). We will refer to this as the  *$\varepsilon$ -corruption model*.

In this work we consider two instances of the general testing task (1) in the robust testing setting. First, let us introduce some notation common to the problems. For  $q \in (0, 2)$  let

$$\mathcal{B}_{s,q} := \{\theta \in \mathbb{R}^d : \|\theta\|_q \leq s\} \quad (2)$$

be the  $\ell_q$ -ball of radius  $s$  in  $\mathbb{R}^d$ . For  $q = 0$ , we get the usual notion of sparsity, and will simply write  $\mathcal{B}_s$  for  $0 \leq s \leq d$ . Let  $\mathcal{N}(\theta, I_d)$  denote the  $d$ -dimensional Gaussian with mean  $\theta \in \mathbb{R}^d$  and identity covariance.

The first problem that we consider is the *sparse Gaussian mean testing*, in which, given an  $\varepsilon$ -corrupted dataset of  $n$  samples from  $\mathcal{N}(\theta, I_d)$ , where  $\theta \in \mathcal{B}_{s,q}$  is unknown, our goal is to robustly distinguish between (1)  $\|\theta\|_2 = 0$ , and (2)  $\|\theta\|_2 \geq \gamma$  (equivalently, the total variation distance between  $\mathcal{N}(\theta, I_d)$  and the standard Gaussian  $\mathcal{N}(0, I_d)$  is  $\Omega(\gamma)$ ), for some input parameter  $\gamma \in (0, 1]$ .

The second problem that we consider is *testing in the sparse linear regression model*. In the sparse linear regression model the data is generated according to the following process: Let  $X_1, X_2, \dots, X_n$  be i.i.d. samples from  $\mathcal{N}(0, I_d)$ . Let  $\theta \in \mathcal{B}_s$  be unknown and let  $\xi_i$ ’s be i.i.d. samples from  $\mathcal{N}(0, 1)$  (and independent from the  $X_i$ ’s), for  $1 \leq i \leq n$ . Then, the  $y_i$ ’s are generated as follows:

$$y_i = \langle X_i, \theta \rangle + \xi_i \quad \text{for all } 1 \leq i \leq n. \quad (3)$$

Note that for a given  $\theta$ , the joint distribution of  $(X_i, y_i)$  is  $\mathcal{N}(0, \Sigma_\theta)$ , where  $\Sigma_\theta = \begin{bmatrix} I_d & \theta \\ \theta^T & 1 + \|\theta\|^2 \end{bmatrix}$ .

Our aim is to robustly distinguish between (1)  $\|\theta\|_2 = 0$ , and (2)  $\|\theta\|_2 \geq \gamma$ , given an  $\varepsilon$ -corrupted version of the observations  $(X_1, y_1), (X_2, y_2), \dots, (X_n, y_n)$ .

Note that for both the problems, one can restrict themselves to the case  $\gamma \geq \varepsilon$ , as otherwise the problem becomes trivially information-theoretically impossible.

## 1.1 Our Contributions

Our main contribution are the characterization of the sample complexity of robust sparse Gaussian mean testing for a range of notions of sparsity, and a lower bound on the sample complexity of robust

<sup>1</sup>The choice of the value  $1/3$  here is arbitrary, and any fixed value greater than  $1/2$  would suffice, as one can amplify the success probability to  $1 - \delta$ , for any  $\delta > 0$ , using a standard majority vote.

testing in the sparse linear regression model. Together, these results fully answer the above question, and provide more evidence to the belief that “robustness requirements make the testing tasks as hard as the corresponding estimation tasks.” To establish our lower bounds, we draw upon and combine a variety of methods from the literature, in order to upper bound the  $\chi^2$ -divergence between a point and a mixture distribution before concluding by Le Cam’s two-point method. We elaborate further on those aspects below.

### 1.1.1 Sparse Gaussian Mean Testing

It is known [CCT17] that, in the *non-robust* setting described in (1), the sample complexity of sparse Gaussian mean testing is

$$n_0(s, d, \gamma) = \begin{cases} \Theta\left(\frac{s}{\gamma^2} \log\left(1 + \frac{d}{s^2}\right)\right) & \text{if } s < \sqrt{d} \\ \Theta\left(\frac{\sqrt{d}}{\gamma^2}\right) & \text{if } s \geq \sqrt{d} \end{cases}$$

for  $q = 0$ , and, for  $q \in (0, 2)$ ,

$$n_q(s, d, \gamma) = \begin{cases} \Theta\left(\frac{m}{\gamma^2} \log\left(1 + \frac{d}{m^2}\right)\right) & \text{if } m < \sqrt{d} \\ \Theta\left(\frac{\sqrt{d}}{\gamma^2}\right) & \text{if } m \geq \sqrt{d} \end{cases},$$

where  $m := \max\{u \in [d] : \gamma^2 u^{\frac{2}{q}-1} \leq s^2\}$  is the *effective sparsity*. In particular, both sample complexities present a phase transition at  $\sqrt{d}$ , after which the sparsity no longer helps decreasing the sample complexity of the problem, which defaults to the “folklore” non-sparse bound of  $\Theta\left(\frac{\sqrt{d}}{\gamma^2}\right)$ .

Our main result in this setting is a lower bound on *robust* sparse mean testing, which shows a significantly different landscape:

**Theorem 1** (Informal; see Theorems 4 and 5). *For every constant  $\varepsilon, \gamma$ , the sample complexity of robust sparse Gaussian mean testing in the  $\varepsilon$ -corruption model is*

$$\Omega\left(s \log \frac{ed}{s}\right)$$

for  $q = 0$ , and  $\Omega\left(m \log \frac{ed}{m}\right)$  for  $q \in (0, 2)$ , where  $m = \max\{u \in [d] : \gamma^2 u^{\frac{2}{q}-1} \leq s^2\}$ .

Moreover, our bound for standard  $s$ -sparsity is tight, in view of the known  $O\left(\frac{s}{\gamma^2} \log \frac{ed}{s}\right)$  sample complexity bound for robust sparse mean *estimation* [Li17, DK19] (which implies the same bound for robust testing). This not only shows that the robust testing problem is much harder than its non-robust counterpart especially in the dense regime (and actually as hard as the robust *estimation* problem), but also that the robust setting no longer presents any threshold phenomenon. If we set  $s = d$ , we further recover the result in [DKS17] that robustness requirement increases the sample complexity of Gaussian mean testing. Figure 1 illustrates the sample complexity of robust and non-robust sparse Gaussian mean testing as a function of the sparsity.

The tightness of our bound, however, only follows from previous work in the case  $q = 0$  (standard sparsity). We provide a (near) matching upper bound for the case  $q \in (0, 2)$ , which essentially resolves the question: showing that the aforementioned hardness and disappearance of a phase transition apply to all types of sparsity.

**Theorem 2** (Informal; see Theorem 6). *For every  $\varepsilon, \gamma$ , the sample complexity of robust sparse Gaussian mean testing in the  $\varepsilon$ -corruption model is*

$$O\left(\frac{m}{\varepsilon^2} \log \frac{ed}{\gamma}\right)$$

for  $q \in (0, 2)$ , where  $m = \max\{u \in [d] : \gamma^2 u^{\frac{2}{q}-1} \leq s^2\}$ .

Note that, for constant  $\gamma$ , our upper bound only differs from the lower bound of Theorem 1 by a logarithmic dependence on the effective sparsity  $m$ . Finally, we note that while the upper bounds from [Li17] and Theorem 2 are achieved by computationally inefficient algorithms (time complexity exponential in  $s$ ), this is actually inherent; indeed, [BB20] recently proved that any computationally efficient algorithm for robust sparse mean *estimation* must have much higher sample complexity, namely  $\Omega(s^2)$ . A simple inspection of their proof shows that this result extends to the robust *testing* problem.

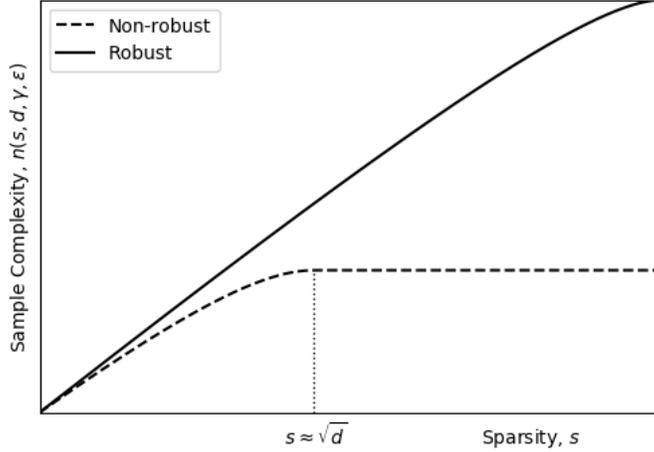


Figure 1: Sample complexity as a function of sparsity for non-robust and robust sparse Gaussian mean testing (the behavior applies to testing in the sparse regression model as well).

### 1.1.2 Testing in Sparse Linear Regression Model

From the results in [CCC<sup>+</sup>19] we can deduce that the sample complexity of *non-robust* testing in the sparse linear regression model is given by

$$n_0(s, d, \gamma) = \begin{cases} \Theta\left(\min\left(\frac{s}{\gamma^2} \log\left(1 + \frac{d}{s^2}\right), \frac{1}{\gamma^4}\right)\right) & \text{if } s < \sqrt{d} \\ \Theta\left(\min\left(\frac{\sqrt{d}}{\gamma^2}, \frac{1}{\gamma^4}\right)\right) & \text{if } s \geq \sqrt{d}. \end{cases}$$

This expression looks very similar to the sample complexity of sparse Gaussian mean testing, except for an additional  $1/\gamma^4$  term. This term is essentially due to the fact that we can ignore the observations  $X_i$ 's and estimate  $\gamma$  just by using  $y_i$ 's, since their variance is  $1 + \gamma^2$ . This would require  $O(1/\gamma^4)$  samples. Nevertheless, the testing in sparse linear regression still exhibits a phase transition at  $s \approx \sqrt{d}$  as was observed in the sparse Gaussian mean testing. It is thus natural to wonder whether the parallels between these two problems extend to the robust setting as well. We show that this is indeed the case: the sample complexity of testing in sparse linear regression significantly increases when introducing the robustness condition.

**Theorem 3** (Informal; see Theorem 7). *For any sufficiently small  $\gamma > 0$ ,  $\epsilon = \frac{\gamma}{C}$  for a sufficiently large  $C$ , and  $s = d^{1-\delta}$  for any  $\delta \in (0, 1)$ , the sample complexity of testing in sparse linear regression under  $\epsilon$ -corruption model is*

$$\Omega\left(\min\left(s \log d, \frac{1}{\gamma^4}\right)\right).$$

The tightness of this bound follows from the results in [LSLC20, Theorem 2.1] which states that any algorithm for robust sparse mean estimation can be used for robust sparse linear regression with a  $\text{polylog}(1/\gamma)$  increase in the sample complexity. Hence the agnostic hypothesis selection via tournaments algorithm in [Li17], which is a statistically optimal algorithm for robust sparse mean estimation works in this case as well.

## 1.2 Our Techniques

To establish the lower bounds in Theorem 1, we combine a range of tools from information theory and the literature on robust estimation. First, we argue that it is enough to consider the standard sparsity case ( $q = 0$ ), as this will imply the analogous result for  $q \in (0, 2)$ . Indeed, the definition of effective sparsity will enable us to deduce that any  $x \in \mathbb{R}^d$  with  $\|x\|_0 = m$  and  $\|x\|_2 = \gamma$  belong to the set  $\mathcal{B}_{s,q}$ , letting us establish the lower bound given in Theorem 1 for general  $q$  from the  $q = 0$  case. In this discussion, we therefore focus on standard sparsity, and assume that our observations are from a distribution for which the unknown parameter  $\theta$  is in the set  $\mathcal{B}_s$ .

To simplify further, we note that in order to establish the desired lower bound it suffices to consider the weaker  $\varepsilon$ -Huber contamination model instead of the (more general)  $\varepsilon$ -corruption model [DK19]. In the  $\varepsilon$ -Huber contamination model,  $n$  i.i.d. samples from a distribution  $\mathbf{p}$  are, after contamination, modeled as a set of  $n$  i.i.d. samples from a distribution  $(1 - \varepsilon)\mathbf{p} + \varepsilon N$ , where  $N$  is an arbitrary and unknown probability distribution (that is, the adversary is “oblivious:” limited to choosing, ahead of time, a “bad” mixture component  $N$  to fool the algorithm). We thus can restrict ourselves, for our lower bound, to the setting where the  $n$  i.i.d. samples are from some distribution in

$$\{(1 - \varepsilon)\mathcal{N}(\theta, I_d) + \varepsilon N_{\varepsilon, \theta} : \|\theta\|_0 \leq s\}$$

where we get to design the distributions  $N_{\varepsilon, \theta}$ . Let  $\tilde{\mathcal{C}}_s := \{(1 - \varepsilon)\mathcal{N}(\theta, I_d) + \varepsilon N_{\varepsilon, \theta} : \|\theta\|_0 \leq s, \|\theta\|_2 \geq \gamma\}$  denote the set of  $\varepsilon$ -Huber contaminated versions of Gaussian distributions whose mean has  $\ell_2$  norm greater than  $\gamma$ . Our goal can be rephrased as choosing a suitable set of parameters  $\Theta$  and a corresponding ensemble of distributions  $\{\mathbf{q}_\theta\}_{\theta \in \Theta} \subseteq \tilde{\mathcal{C}}_s$ , and argue that no algorithm can distinguish  $n$  samples drawn from a (randomly chosen) element  $\mathbf{q}_\theta$  from  $n$  samples drawn from  $\mathcal{N}(0, I_d)$ .

To do so, define the mixture  $\mathbf{q} := \frac{1}{|\Theta|} \sum_{\theta} \mathbf{q}_\theta^n$ . (equivalently  $\mathbf{q} = \mathbb{E}_{\theta \sim \mathbf{u}_\Theta}[\mathbf{q}_\theta^n]$ ). Le Cam’s two-point method allows us to reduce the above indistinguishability problem to showing that  $d_{\text{TV}}(\mathbf{q}, \mathcal{N}(0, I_d)^n) = o(1)$ , for which, in view of the standard inequality,  $d_{\text{TV}}(\mathbf{q}, \mathbf{p})^2 \leq \frac{1}{4} \chi^2(\mathbf{q} \parallel \mathbf{p})$ , it suffices to show that  $\chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_d)^n) = o(1)$ . By the Ingster–Suslina method, this  $\chi^2$ -divergence between a point and a mixture distribution can then be simplified as

$$1 + \chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_d)^n) = 1 + \chi_{\mathcal{N}(0, I_d)^n}(\mathbb{E}_\theta[\mathbf{q}_\theta^n], \mathbb{E}_{\theta'}[\mathbf{q}_{\theta'}^n]) = \mathbb{E}_{\theta, \theta'}[(1 + \chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'}))^n], \quad (4)$$

where  $\theta, \theta' \sim \mathbf{u}_\Theta$ , and  $\chi_{\mathbf{p}}(\mathbf{q}, \mathbf{q}') := \int \frac{d\mathbf{q}d\mathbf{q}'}{d\mathbf{p}} - 1$  is the “ $\chi^2$ -correlation” between  $\mathbf{q}$  and  $\mathbf{q}'$  with respect to  $\mathbf{p}$ . Thus, the challenge is to design an ensemble  $\{\mathbf{q}_\theta\}$  that yields a good upper bound for the r.h.s. of (4), while still being simple enough to analyze.

Building upon previous works [DKS17, BB20], we define our ensemble  $\{\mathbf{q}_\theta\}$  as follows: let

$$\Theta := \left\{ \theta \in \left\{ -\frac{1}{\sqrt{s}}, 0, \frac{1}{\sqrt{s}} \right\}^d : \|\theta\|_0 = s \right\} \quad (5)$$

and, for  $\theta \in \Theta$ ,  $\mathbf{q}_\theta := (1 - \varepsilon)\mathcal{N}(\gamma\theta, I_d) + \varepsilon\mathcal{N}(\mu_\theta, I_d)$ , where  $\mu_\theta$  is chosen such that  $(1 - \varepsilon)\gamma\theta + \varepsilon\mu_\theta = 0$ . Note that indeed, for all  $\theta \in \Theta$ ,  $\mathbf{q}_\theta \in \tilde{\mathcal{C}}_s$  and  $\|\theta\|_2 = 1$ . This choice of  $\mathbf{q}_\theta$  combined with the above outline will enable us to prove Theorem 1, by carefully upper bounding  $\mathbb{E}_{\theta, \theta'}[(1 + \chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'}))^n]$  as a function of  $n, d$ , and  $s$ .

One can attempt to prove Theorem 3 (the lower bound for testing in the sparse linear regression model) in a similar vein. By restricting ourselves to the  $\varepsilon$ -Huber contamination model (which only strengthens the resulting lower bound by constraining the adversary), we get the following set of distributions in the sparse linear regression problem:

$$\{(1 - \varepsilon)\mathcal{N}(0, \Sigma_\theta) + \varepsilon N_{\varepsilon, \theta} : \|\theta\|_0 \leq s\}, \quad (6)$$

where  $\Sigma_\theta := \begin{bmatrix} I_d & \theta \\ \theta^T & 1 + \|\theta\|^2 \end{bmatrix}$ . Further defining  $\tilde{\mathcal{D}}_s := \{(1 - \varepsilon)\mathcal{N}(0, \Sigma_\theta) + \varepsilon N_{\varepsilon, \theta} : \|\theta\|_0 \leq s, \|\theta\|_2 \geq \gamma\}$ , as earlier our problem boils down to finding an ensemble  $\{\mathbf{q}_\theta\} \subseteq \tilde{\mathcal{D}}_s$  such that  $\chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_{d+1})^n) = o(1)$  for  $\mathbf{q} := \frac{1}{|\Theta|} \sum_{\theta} \mathbf{q}_\theta^n$ .

Notice that in the sparse linear regression model,  $y_i$  has marginal distribution  $\mathcal{N}(0, 1 + \|\theta\|^2)$  and conditioned on  $y_i$  the distribution of  $X_i$  is given by  $\mathcal{N}\left(\frac{y_i \theta}{1 + \|\theta\|^2}, I - \frac{\theta\theta^T}{1 + \|\theta\|^2}\right)$ . This observation, along with the indistinguishability result established while proving Theorem 1 encourage us to choose  $\Theta$  as in (5) and  $\mathbf{q}_\theta$  in the following manner:

$$\begin{aligned} \mathbf{q}_\theta(y_i) &= \mathcal{N}(0, 1 + \gamma^2) \\ \mathbf{q}_\theta(X_i \mid y_i) &= (1 - \varepsilon)\mathcal{N}\left(\frac{y_i \gamma \theta}{1 + \gamma^2}, I - \frac{\theta\theta^T}{1 + \gamma^2}\right) + \varepsilon\mathcal{N}\left(\frac{y_i \mu_\theta}{1 + \gamma^2}, I - \frac{\theta\theta^T}{1 + \gamma^2}\right), \end{aligned}$$

where  $\mu_\theta$  is chosen such that  $(1 - \varepsilon)\gamma\theta + \varepsilon\mu_\theta = 0$ . Again, note that  $\mathbf{q}_\theta \in \tilde{\mathcal{D}}$ . Although this setup looks promising in giving us the right lower bound, it turns out that certain tail events can cause  $\chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_{d+1})^n)$  to blow up to infinity. To circumvent this, one of the remedies available in the literature is to use the *conditional second moment method* [RXZ19, WX18]. In this method, one carefully conditions out the rare events that preclude us from evaluating the  $\chi^2$ -divergence. Specifically, we define an event  $\mathcal{E}$  generated by the random variables  $(y_1, y_2, \dots, y_n)$  such that  $\mathbf{q}(\mathcal{E}^C) = o(1)$ . We then define  $\mathbf{q}^\mathcal{E}$  as the distribution  $\mathbf{q}$  conditioned on the “good” event  $\mathcal{E}$ . That is,

$$\mathbf{q}^\mathcal{E}(X, Y) = \mathbf{q}(X, Y \mid \mathcal{E}) = \frac{\mathbb{E}_\theta[\mathbf{q}_\theta(X, Y)\mathbb{1}_\mathcal{E}\{Y\}]}{\mathbf{q}(\mathcal{E})}. \quad (7)$$

Note that we have the relation  $\mathbf{q} = \mathbf{q}(\mathcal{E})\mathbf{q}^\mathcal{E} + \mathbf{q}(\mathcal{E}^C)\mathbf{q}^{\mathcal{E}^C}$ . The convexity of total variation distance consequently gives

$$\begin{aligned} d_{\text{TV}}(\mathbf{q}, \mathcal{N}(0, I_{d+1})^n) &\leq \mathbf{q}(\mathcal{E}) d_{\text{TV}}(\mathbf{q}^\mathcal{E}, \mathcal{N}(0, I_{d+1})^n) + \mathbf{q}(\mathcal{E}^C) d_{\text{TV}}(\mathbf{q}^{\mathcal{E}^C}, \mathcal{N}(0, I_{d+1})^n) \\ &\leq d_{\text{TV}}(\mathbf{q}^\mathcal{E}, \mathcal{N}(0, I_{d+1})^n) + \mathbf{q}(\mathcal{E}^C) \end{aligned}$$

Thus, in view of the fact that  $\mathbf{q}(\mathcal{E}^C) = o(1)$ , deriving a lower bound reduces to showing that  $d_{\text{TV}}(\mathbf{q}^\mathcal{E}, \mathcal{N}(0, I_{d+1})^n)$  is small, for which we saw earlier that it was sufficient (and more convenient) to show that  $\chi^2(\mathbf{q}^\mathcal{E} \parallel \mathcal{N}(0, I_{d+1})^n) = o(1)$ . This is the roadmap we will follow – first, defining a suitable event  $\mathcal{E}$ , before showing that  $\chi^2(\mathbf{q}^\mathcal{E} \parallel \mathcal{N}(0, I_{d+1})^n) = o(1)$ .

The above outlines our approach to proving Theorems 1 and 3. To establish the upper bound in Theorem 2, we show that the “Agnostic hypothesis selection via Tournaments” algorithm [DKK<sup>+</sup>16] achieves the stated sample complexity. The argument, in turn, is very similar to the proof of the upper bound for robust sparse mean estimation given in [Li17].

### 1.3 Related Work

The study of high-dimensional signals with a sparse underlying structure has enjoyed a significant amount of attention in statistics and signal processing for the past few decades. This line of research has led to the discovery of surprising phenomena such as phase transitions and computational hardness in problems involving sparse signals. The main motivation to study sparse signals is that the sample complexity of statistical tasks involving sparsity is typically significantly smaller than that of dense signals, leading to much more data-efficient algorithms. This yields significant savings whenever the problem is expected to exhibit such a sparse structure, e.g., for physical or biological reasons, or due to a specific design choice when engineering a system. Yet, practical scenarios seldom involve noise-free or perfect signals, which effectively destroys the sparsity of the signal one would have capitalized on. This led to the study of these questions under various noise models, in order to understand if one could still see the same type of sample size savings in these settings.

There is a large body of work on the estimation and detection of signals with a sparse structure under noise (e.g., [DJ04, Bar02, Ver12]). Most recently, [CCT17] gave the tight characterization of minimax rates of estimating linear and quadratic functionals of sparse signals under Gaussian noise. The authors also derived the minimum detection level required for reliably testing the  $\ell_2$  norm of sparse signals under Gaussian noise. Another prominent sparse signal model studied in the literature is that of sparse linear regression [ITV10, Ver12, RXZ19]. In the non-asymptotic setting, [CCC<sup>+</sup>19] established the tight sample complexity of testing in the sparse linear regression model. However, these line of works focused on *random* noise, and not the more challenging types of noise allowing for *adversarial* corruptions – what is commonly known as seeking *robust* algorithms.

The systematic study of the robustness of statistical procedures was initiated in the foundational works of Huber [Hub64] and Tukey [Tuk60]. Several statistically optimal procedures were found to break down even under slight model misspecification or sample contamination. Although there was substantial progress in the field of robust statistics, surprisingly, *computationally efficient* procedures remained elusive until recently, even for simple tasks such as high-dimensional mean estimation.

In this regard, the past few years witnessed an incredible progress in algorithmic robust statistics. A line of work initiated by [DKK<sup>+</sup>16] and [LRV16] provided computationally efficient optimal robust estimators for various estimation tasks in high dimension [DKK<sup>+</sup>19]. The surprising upshot from these papers is that even with robustness requirements, the sample complexity of many high-dimensional estimation tasks remains essentially the same, at no extra computational cost: i.e.,

that robustness and computational efficiency are not at odds for those estimation tasks. And still, a subsequent result of [DKS17] shows that, surprisingly, imposing robustness constraints *does* significantly increase the sample complexity of the Gaussian mean *testing* problem, and makes the sample complexity as large as that of Gaussian mean estimation: that is, for testing, robustness comes at a very high cost, and negates the usual savings that testing allows over estimation. Extending their results on mean testing, [DK21] shows the analogue in the case of Gaussian covariance testing under the Frobenius norm. Yet, those striking results focus on testing *dense* parameters; our work seeks to combine the two lines of work – inference under sparsity guarantees, and robustness – to understand if an analogous jump in sample complexity occurs for testing in high-dimensional *sparse* models.

We further note that several works have shown evidence for the existence of statistical-computation gaps in estimation problems with sparse signal structure [BR13, CW20]. [DKS17] gave the first evidence for the presence of an  $s$  to  $s^2$  statistical-computation gap in robust sparse mean estimation, in the form of a Statistical Query (SQ) lower bound (i.e., for a restricted type of algorithms). This was complemented by [Li17] and [BDLS17], which gave computationally efficient algorithms for robust sparse estimation achieving  $O(s^2)$  sample complexity. Finally, [BB20] recently used average-case reductions to prove the algorithmic hardness of robust sparse mean estimation and robust sparse linear regression.

## 2 Preliminaries and Notation

Given a probability distribution  $\mathbf{p}$  and integer  $n \geq 1$ , we denote by  $\mathbf{p}^n$  the  $n$ -fold product distribution with marginals  $\mathbf{p}$ , and given a set of distributions  $\mathcal{D}$  write  $\mathcal{D}^n = \{ \mathbf{p}^n : \mathbf{p} \in \mathcal{D} \}$ . For  $0 < q \leq 2$  and  $r > 0$ , we let  $\mathcal{B}_{q,r} = \{ \theta \in \mathbb{R}^d : \|\theta\|_q \leq r \}$  the  $\ell_q$  ball of radius  $r$ . We say a vector  $\theta \in \mathbb{R}^d$  is  $s$ -sparse if  $\|\theta\|_0 \leq s$ ; for  $q \in (0, 2]$ , we will accordingly say that  $\theta$  is  $s$ -sparse in  $\ell_q$  norm whenever  $\|\theta\|_q \leq s$  (note that  $s$  need not be an integer). We denote the uniform distribution over a set  $S$  by  $\mathbf{u}_S$ . The delta measure (point mass) at an element  $x$  is denoted by  $\delta_x$ . The total variation distance between two distributions  $\nu$  and  $\mu$  is defined as:

$$d_{\text{TV}}(\nu, \mu) = \frac{1}{2} \int |d\nu - d\mu|,$$

which is to be interpreted as  $\frac{1}{2} \int \left| \frac{d\nu}{d\lambda} - \frac{d\mu}{d\lambda} \right| d\lambda$ , where  $\lambda$  is any distribution dominating both  $\nu$  and  $\mu$ ; equivalently,  $d_{\text{TV}}(\nu, \mu) = \sup_S (\nu(S) - \mu(S))$  where the supremum is over all measurable sets  $S$ . The  $\chi^2$ -divergence between  $\nu$  and  $\mu$  is given by

$$\chi^2(\nu \parallel \mu) = \int \frac{d\nu d\nu}{d\mu} - 1,$$

and satisfies  $d_{\text{TV}}(\nu, \mu) \leq \frac{1}{2} \sqrt{\chi^2(\nu \parallel \mu)}$ . Finally, given a reference probability measure  $\mu$ , the  $\chi^2$ -correlation between  $\nu$  and  $\lambda$  with respect to  $\mu$  is defined as

$$\chi_\mu(\nu, \lambda) = \int \frac{d\nu d\lambda}{d\mu} - 1;$$

note that  $\chi^2(\nu \parallel \mu) = \chi_\mu(\nu, \nu)$ .

We will also rely in several occasions on the following technical lemmas: the first will be useful to bound the  $\chi^2$  correlation between Gaussians.

**Lemma 1.** *Let  $\phi_{\mu, \Sigma}$  denote the density function of  $\mathcal{N}(\mu, \Sigma)$ . Then,*

$$\mathbb{E}_{X \sim \phi_{0, I}} \left[ \frac{\phi_{\mu_1, \Sigma_1}(X) \phi_{\mu_2, \Sigma_2}(X)}{\phi_{0, I}^2(X)} \right] = \frac{\exp\left(\frac{1}{2} (\mu'^T A^{-1} \mu' - \mu_1^T \Sigma_1^{-1} \mu_1 - \mu_2^T \Sigma_2^{-1} \mu_2)\right)}{\det(A)^{\frac{1}{2}} \det(\Sigma_1 \Sigma_2)^{\frac{1}{2}}},$$

where  $A := \Sigma_1^{-1} + \Sigma_2^{-1} - I$  and  $\mu' := \Sigma_1^{-1} \mu_1 + \Sigma_2^{-1} \mu_2$ .

*Proof.* The l.h.s. can be written explicitly as

$$\frac{1}{(2\pi)^{d/2} \det(\Sigma_1 \Sigma_2)^{1/2}} \int_{\mathbb{R}^d} \frac{\exp\left(-\frac{1}{2} ((x - \mu_1)^T \Sigma_1^{-1} (x - \mu_1) + (x - \mu_2)^T \Sigma_2^{-1} (x - \mu_2))\right)}{\exp\left(-\frac{1}{2} x^T x\right)} dx.$$

Integrating this quantity by completing the square gives the result.  $\square$

The second, due to Cai, Ma, and Wu, will let us bound the moment generating function of the square of a random sum of Rademacher random variables, which will arise when we consider sparse priors in our lower bounds.

**Lemma 2** (Lemma 1 in [CMW15]). *Fix  $d \in \mathbb{N}$  and  $s \in [d]$ . Let  $H \sim \text{Hypergeometric}(d, s, s)$ , and let  $\xi_1, \xi_2, \dots, \xi_s$  be i.i.d. Rademacher. Define the random variable  $Y$  as*

$$Y := \sum_{i=1}^H \xi_i.$$

*Then there exists a function  $\tau: (0, \frac{1}{36}) \rightarrow (1, \infty)$  with  $\tau(0^+) = 1$ , such that for any  $0 < b < \frac{1}{36}$ ,*

$$\mathbb{E}[\exp(\lambda Y^2)] \leq \tau(b),$$

*where  $\lambda := \frac{b}{s} \log \frac{ed}{s}$ .*

Finally, we will require the two facts below on Gaussians and Hypergeometric distributions.

**Lemma 3.** *Let  $P$  be the density function of  $\mathcal{N}(0, 1 + \gamma^2)$  for some  $0 \leq \gamma \leq 1/\sqrt{3}$ , and  $Q$  be the density function of  $\mathcal{N}(0, 1)$ . Then*

$$\mathbb{E}_{y \sim Q} \left[ \left( \frac{P(y)}{Q(y)} \right)^2 y^{2\ell} \right] \leq \frac{c \cdot 4^\ell \ell^\ell}{\sqrt{1 - \gamma^4}},$$

*where  $c > 0$  is a universal constant.*

*Proof.* We have

$$\left( \frac{P(y)}{Q(y)} \right)^2 = \frac{1}{1 + \gamma^2} \exp\left(\frac{\gamma^2 y^2}{1 + \gamma^2}\right).$$

Therefore,

$$\begin{aligned} \mathbb{E}_{y \sim Q} \left[ \left( \frac{P(y)}{Q(y)} \right)^2 y^{2\ell} \right] &= \frac{1}{\sqrt{2\pi}} \frac{1}{1 + \gamma^2} \int_{\mathbb{R}} y^{2\ell} \exp\left(\frac{-y^2}{2} \left(1 - \frac{2\gamma^2}{1 + \gamma^2}\right)\right) \\ &= \frac{1}{\sqrt{1 - \gamma^4}} \left(\frac{1 + \gamma^2}{1 - \gamma^2}\right)^\ell \mathbb{E}_{y \sim Q} [y^{2\ell}] \\ &\stackrel{(a)}{\leq} \frac{c}{\sqrt{1 - \gamma^4}} 4^\ell \ell^\ell, \end{aligned}$$

where in (a) we used the facts that, for all  $\ell \geq 0$ ,  $\mathbb{E}_{y \sim \mathcal{N}(0,1)} [y^{2\ell}] \leq c(2\ell)^\ell$  for some absolute constant  $c > 0$ ; and that  $\gamma \in [0, 1/\sqrt{3}]$ .  $\square$

**Lemma 4.** *Let  $H \sim \text{Hypergeometric}(d, s, s)$ . Then,*

$$\Pr[H = h] \leq \left( \frac{es^2}{h(d - s + 1)} \right)^h.$$

*Proof.* This directly follows from bounds on binomial coefficients:

$$\Pr[H = h] = \binom{s}{h} \frac{\binom{d-s}{s-h}}{\binom{d}{s}} \leq \left(\frac{es}{h}\right)^h \left(\frac{s}{d-s+1}\right)^h = \left(\frac{es^2}{h(d-s+1)}\right)^h.$$

$\square$

### 3 Main Results and Proofs

In this section, we formally state and give proofs for the theorems outlined informally in Section 1.1. Recall that a lower bound on the hardness of hypothesis testing between two distributions  $\mathbf{p}$  and  $\mathbf{q}$  can be characterized by the total variation distance between them. Indeed, by the Pearson-Neyman lemma, if there exists a test which successfully distinguishes between two distributions  $\mathbf{p}_n$  and  $\mathbf{q}_n$  (which in our case will correspond to distributions over  $n$  tuples of i.i.d. samples) with probability at least  $2/3$ , then one must have  $d_{\text{TV}}(\mathbf{p}_n, \mathbf{q}_n) \geq 1/3$ . Hence, to prove indistinguishability for a given  $n$ , it suffices to show  $d_{\text{TV}}(\mathbf{p}_n, \mathbf{q}_n) = o(1)$ ; since  $d_{\text{TV}}(\mathbf{p}_n, \mathbf{q}_n)^2 \leq \frac{1}{4}\chi^2(\mathbf{q}_n \parallel \mathbf{p}_n)$ , one can then focus on showing  $\chi^2(\mathbf{q}_n \parallel \mathbf{p}_n) = o(1)$ .

In our problems, we formulate  $\mathbf{p}_n$  as a product distribution (product of high-dimensional Gaussians) and  $\mathbf{q}_n$  as a mixture distribution. In such cases, the following tensorization property of  $\chi^2$ -divergence helps in upper bounding it: Let  $\mathbf{q} := \int \mathbf{q}_\theta^n d\theta$  be a mixture distribution. Then,

$$\begin{aligned} 1 + \chi^2(\mathbf{q} \parallel \mathbf{p}^n) &= \int_{\mathbb{R}^n} \frac{\mathbf{q}(x)^2}{\mathbf{p}^n(x)} dx = \int_{\mathbb{R}^n} \frac{\int \mathbf{q}_\theta^n(x) d\theta \int \mathbf{q}_{\theta'}^n(x) d\theta'}{\mathbf{p}^n(x)} dx \\ &= \int_{\theta} \int_{\theta'} \int_{\mathbb{R}^n} \frac{\mathbf{q}_\theta^n(x) \mathbf{q}_{\theta'}^n(x)}{\mathbf{p}^n(x)} dx d\theta d\theta' = \int_{\theta} \int_{\theta'} \left( \int_{\mathbb{R}} \frac{\mathbf{q}_\theta(x) \mathbf{q}_{\theta'}(x)}{\mathbf{p}(x)} dx \right)^n d\theta d\theta' \\ &= \mathbb{E}_{\theta, \theta'} [(1 + \chi_{\mathbf{p}}(\mathbf{q}_\theta, \mathbf{q}_{\theta'}))^n]. \end{aligned}$$

This approach is widely known as the Ingster–Suslina method [IS03], and is the starting point of many minimax lower bounds.

#### 3.1 Sparse Gaussian Mean Testing

In this section, we state and prove our results related to sparse Gaussian mean testing. First, we derive the lower bounds in Theorem 1 using the techniques outlined in Section 1.2 and then show a matching upper bound (up to logarithmic factors) for  $q > 0$  case.

**Theorem 4.** *Let  $\varepsilon, \gamma > 0$  be fixed. Let  $X_1, X_2, \dots, X_n$  be i.i.d. samples from an unknown distribution  $\mathbf{p}$ . Moreover, suppose an  $\varepsilon$ -fraction of these  $n$  samples are arbitrarily corrupted. Then, if there exists an algorithm that distinguishes between the cases  $\mathbf{p} = \mathcal{N}(0, I_d)$  and  $\mathbf{p} \in \{\mathcal{N}(\theta, I_d) : \|\theta\|_2 \geq \gamma, \|\theta\|_0 \leq s\}$  with probability greater than  $2/3$ , we must have  $n = \Omega\left(s \log \frac{ed}{s}\right)$ .*

To prove this theorem, we will require the following lemma due to Diaconikolas, Kane, and Stewart [DKS17]; we provide below an alternative proof, which we believe is simpler than the original.

**Lemma 5** ([DKS17, Lemma 6.9]). *Fix  $\theta, \theta' \in \mathbb{R}^d$ ,  $\varepsilon \in (0, 1/3]$ , and let  $\mathbf{q}_\theta$  be defined as*

$$\mathbf{q}_\theta := (1 - \varepsilon)\mathcal{N}(\theta, I_d) + \varepsilon\mathcal{N}\left(-\frac{(1 - \varepsilon)}{\varepsilon}\theta, I_d\right).$$

Then,

$$1 + |\chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'})| \leq \exp\left(\frac{\langle \theta, \theta' \rangle^2}{\varepsilon^4}\right).$$

*Proof.* We here provide a simple proof. Let  $p := \frac{1 - \varepsilon}{\varepsilon}$ , which is at least 2 since  $\varepsilon \leq 1/3$ . The distribution  $\mathbf{q}_\theta$  can be written as follows:

$$\mathbf{q}_\theta = \mathbb{E}_b[\mathcal{N}(b\theta, I_d)],$$

where  $b \sim (1 - \varepsilon)\delta_1 + \varepsilon\delta_{-p}$ . Let  $\phi_{\mu, \Sigma}$  denote the density function of  $\mathcal{N}(\mu, \Sigma)$ . Then,

$$\begin{aligned} \chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'}) &= \mathbb{E}_{X \sim \mathcal{N}(0, I_d)} \left[ \frac{\mathbb{E}_b[\phi_{b\theta, I_d}(X)] \mathbb{E}_{b'}[\phi_{b'\theta', I_d}(X)]}{\phi_{0, I_d}^2(X)} \right] - 1 \\ &= \mathbb{E}_{bb'} \left[ \mathbb{E}_{X \sim \mathcal{N}(0, I_d)} \left[ \frac{\phi_{b\theta, I_d}(X) \phi_{b'\theta', I_d}(X)}{\phi_{0, I_d}^2(X)} \right] \right] - 1. \end{aligned}$$

By Lemma 1, we have

$$\mathbb{E}_{X \sim \mathcal{N}(0, I_d)} \left[ \frac{\phi_{b\theta, I_d}(X) \phi_{b'\theta', I_d}(X)}{\phi_{0, I_d}^2(X)} \right] = \exp(bb' \langle \theta, \theta' \rangle),$$

from which

$$\chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'}) = \mathbb{E}_{bb'}[\exp(bb' \langle \theta, \theta' \rangle)] - 1 = \sum_{\ell=1}^{\infty} \frac{\mathbb{E}_b[b^\ell]^2 \langle \theta, \theta' \rangle^\ell}{\ell!}.$$

Note that  $\mathbb{E}[b] = 0$ , and that  $|\mathbb{E}_b[b^\ell]| = |(1 - \varepsilon) + (-1)^\ell p^\ell \varepsilon| = (1 - \varepsilon) |1 + (-1)^\ell p^{\ell-1}| \leq 2p^{\ell-1}$ . Thus, we have

$$\begin{aligned} |\chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'})| &\leq \frac{4}{p^2} \sum_{\ell=2}^{\infty} \frac{p^{2\ell} |\langle \theta, \theta' \rangle|^\ell}{\ell!} \\ &= \frac{4}{p^2} (\exp(p^2 |\langle \theta, \theta' \rangle|) - p^2 |\langle \theta, \theta' \rangle| - 1) \\ &\stackrel{(a)}{\leq} \frac{4}{p^2} (\exp(p^4 \langle \theta, \theta' \rangle^2) - 1) \\ &\stackrel{(b)}{\leq} \exp\left(\frac{\langle \theta, \theta' \rangle^2}{\varepsilon^4}\right) - 1, \end{aligned}$$

where (a) is due to the fact that  $e^x - x \leq e^{x^2}$  for all  $x \geq 0$  and (b) follows from  $p \geq 2$ .  $\square$

With this in hand, we are now able to establish Theorem 4.

*Proof of Theorem 4.* For the purpose of deriving a lower bound, it is enough to consider the weaker  $\varepsilon$ -Huber model for the corruption of samples, as then a lower bound on the sample complexity of the hypothesis testing problem in this setting will also be a lower bound for robust sparse Gaussian mean testing in the adversarial one.

$$\begin{aligned} \mathcal{H}_0 &: X_i \sim \mathcal{N}(0, I_d) \\ \mathcal{H}_1 &: X_i \sim (1 - \varepsilon)\mathcal{N}(\theta, I_d) + \varepsilon N_{\varepsilon, \theta} \quad \text{s. t.} \quad \|\theta\|_2 \geq \gamma \text{ and } \|\theta\|_0 \leq s, \end{aligned}$$

for  $1 \leq i \leq n$ . Here the distributions  $N_{\varepsilon, \theta}$  need to be chosen appropriately.

Let  $B = \{\beta \in \{-1, 0, 1\}^d : \|\beta\|_0 = s\}$  and  $\beta \sim \mathbf{u}_B$ . We can think of  $\beta$  as being generated according to the following process: First pick an element uniformly from the set  $\{b \in \{0, 1\}^d : \|b\|_0 = s\}$  and then set the non-zero elements to be i.i.d. Rademacher random variables.

Define  $\Theta := \frac{1}{\sqrt{s}}B$  and  $\theta := \frac{1}{\sqrt{s}}\beta$ , so that  $\|\theta\|_2 = 1$ . Define  $\mathbf{q}_\theta$  as

$$\mathbf{q}_\theta = (1 - \varepsilon)\mathcal{N}(\gamma\theta, I_d) + \varepsilon\mathcal{N}\left(-\frac{(1 - \varepsilon)}{\varepsilon}\gamma\theta, I_d\right).$$

It is immediate to see that  $\mathbf{q}_\theta \in \mathcal{H}_1$  for all  $\theta \in \Theta$ . Let  $\mathbf{q} := \mathbb{E}_{\theta \sim \mathbf{u}_\Theta}[\mathbf{q}_\theta^n]$ . Then by (4),

$$1 + \chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_d)^n) = \mathbb{E}_{\theta, \theta'}[(1 + \chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'}))^n].$$

By Lemma 5, we get

$$1 + \chi_{\mathcal{N}(0, I_d)}(\mathbf{q}_\theta, \mathbf{q}_{\theta'}) \leq \exp\left(\frac{\gamma^4 \langle \theta, \theta' \rangle^2}{\varepsilon^4}\right) = \exp\left(\frac{\gamma^4 \langle \beta, \beta' \rangle^2}{\varepsilon^4 s^2}\right).$$

Hence,

$$1 + \chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_d)^n) \leq \mathbb{E}_{\beta, \beta'} \left[ \exp\left(\frac{n\gamma^4 \langle \beta, \beta' \rangle^2}{\varepsilon^4 s^2}\right) \right]. \quad (8)$$

Now by symmetry of the problem, it suffices to evaluate the expectation for any fixed  $\beta'$ , say to  $\beta_0 := \underbrace{(1, 1, \dots, 1)}_s, 0, 0, \dots, 0$ . We can characterize  $\langle \beta, \beta_0 \rangle$  as follows: let  $H$  be a random

variable with distribution Hypergeometric( $d, s, s$ ) and let  $\xi_1, \xi_2, \dots, \xi_s$  be i.i.d. Rademacher random variables independent of  $H$ . Then, we have  $Y := \langle \beta, \beta_0 \rangle = \sum_{i=1}^H \xi_i$ , where  $Y$  can be thought of as a symmetric random walk with Hypergeometric stopping time. We can then rewrite (8) as

$$1 + \chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_d)^n) \leq \mathbb{E} \left[ \exp \left( \frac{n\gamma^4 Y^2}{\varepsilon^4 s^2} \right) \right].$$

By Lemma 2, if  $n \leq c \frac{\varepsilon^4}{\gamma^4} \cdot s \log \frac{ed}{s}$  for a sufficiently small  $c > 0$  (e.g.,  $c < 1/36$  suffices),

$$\chi^2(\mathbf{q} \parallel \mathcal{N}(0, I_d)^n) \leq \mathbb{E} \left[ \exp \left( \frac{n\gamma^4 Y^2}{\varepsilon^4 s^2} \right) \right] - 1 \leq \tau(c) - 1.$$

Since  $\tau$  is continuous at  $0^+$ , we can choose  $c > 0$  to make  $\tau(c) - 1$  arbitrarily small. This implies that it is impossible to distinguish between  $\mathcal{H}_0$  and  $\mathcal{H}_1$  with high probability if  $n = o(s \log \frac{ed}{s})$ , establishing the theorem.  $\square$

Next, we extend this lower bound to other sparsity notions, namely, with respect to the  $\ell_q$ -norms.

**Theorem 5.** *Let  $\varepsilon, \gamma > 0$  be fixed, and  $q \in (0, 2)$ . Let  $X_1, X_2, \dots, X_n$  be i.i.d. samples from an unknown distribution  $\mathbf{p}$ . Moreover, suppose an  $\varepsilon$ -fraction of these  $n$  samples are arbitrarily corrupted. Then, if there exists an algorithm that distinguishes between the cases  $\mathbf{p} = \mathcal{N}(0, I_d)$  and  $\mathbf{p} \in \{\mathcal{N}(\theta, I_d) : \|\theta\|_2 \geq \gamma, \|\theta\|_q \leq s\}$  with probability greater than  $2/3$ , we must have  $n = \Omega(m \log \frac{ed}{m})$ , where  $m$  is the effective sparsity defined as:*

$$m := \max\{u \in [d] : \gamma^2 u^{\frac{2}{q}-1} \leq s^2\}.$$

*Proof.* The parameter set of the alternative hypothesis in this problem is

$$\Theta = \{\theta \in \mathbb{R}^d : \|\theta\|_q \leq s \text{ and } \|\theta\|_2 \geq \gamma\}.$$

To conclude, it suffices to show that every  $m$ -sparse (in  $\ell_0$  sense) vector with  $\ell_2$  norm equal to  $\gamma$  belongs to  $\Theta$ , and then appeal to the proof of Theorem 4 (whose hard instances had mean with magnitude exactly  $\gamma$ ). To do so, let  $x \in \mathbb{R}^d$  be  $m$ -sparse such that  $\|x\|_2 = \gamma$ . Then,

$$\|x\|_q^2 = \left( \sum_{i=1}^d |x_i|^q \right)^{\frac{2}{q}} \stackrel{\text{(convexity)}}{\leq} m^{\frac{2}{q}-1} \|x\|_2^2 = m^{\frac{2}{q}-1} \gamma^2 \stackrel{\text{(a)}}{\leq} s^2,$$

where (a) is due to the definition of  $m$ , and the convexity step follows from  $\frac{q}{2} \leq 1$ . Hence  $x \in \Theta$ . Thus the lower bound in Theorem 4 holds here with  $m$  replacing  $s$ .  $\square$

We will now proceed to show that the lower bounds in Theorems 4 and 5 are tight up to a logarithmic factor. First, we note that the lower bound given in Theorem 4 is optimal due to the result that the sample complexity of robust sparse Gaussian mean estimation is upper bounded by the same value [DK19, Li17], and the folklore fact that testing is no harder than estimation. So we need to prove the upper bound only for the cases where  $q > 0$ . We will show that the sample complexity lower bound given in Theorem 5 is tight by proving that the algorithm for robust sparse Gaussian mean estimation given in [Li17] works in the  $\ell_q$ -norm constrained case as well (and so, again, the upper bound for testing will follow from the upper bound for estimation), in the regime  $\gamma = \Theta(\varepsilon)$ . Intuitively, this is because all but  $m$  coordinates of a vector  $\theta$  would be small if  $\|\theta\|_q \leq s$  and  $\|\theta\|_2 = \gamma$ . Our proof is similar to the proof of [Li17, Fact A.1] with a minor modification.

**Theorem 6.** *Let  $\varepsilon, \delta > 0$  and,  $q \in (0, 2)$  be fixed. Let  $X_1, X_2, \dots, X_n$  be i.i.d. samples from an unknown distribution  $\mathcal{N}(\mu, I_d)$ , where  $\|\mu\|_q \leq s$ . Moreover, suppose an  $\varepsilon$ -fraction of these  $n$  samples are arbitrarily corrupted. Then, there exists an algorithm that, for  $n = O(\frac{1}{\varepsilon^2} (m \log \frac{d}{\varepsilon} + \log \frac{1}{\delta}))$ , upon being given the  $X_1, \dots, X_n$  outputs a  $\mu'$  such that  $\|\mu' - \mu\|_2 \leq O(\varepsilon)$  with probability  $1 - \delta$ . Here  $m$  is the effective sparsity defined as:*

$$m = \max\{u \in [d] : \varepsilon^2 u^{\frac{2}{q}-1} \leq s^2\}.$$

We will use the following lemma, originally stated in [DKK<sup>+</sup>16], to prove Theorem 6.

**Lemma 6** ([Li17, Lemma A.5]). *Let  $\mathcal{C}$  be a class of probability distributions. Suppose that for some fixed  $N, \varepsilon, \gamma, \delta > 0$  there exists an algorithm that given  $N$  independent samples from some  $D \in \mathcal{C}$ , of which up to an  $\varepsilon$ -fraction is corrupted, returns a list of  $M$  distributions such that, with probability  $1 - \delta/3$ , there exists a  $D'$  in the list with  $d_{\text{TV}}(D', D) < \gamma$ . Suppose furthermore that with probability  $1 - \delta/3$ , the distributions returned by this algorithm are all in some fixed set  $\mathcal{M}$ . Then there exists another algorithm, which given  $O\left(N + \frac{1}{\varepsilon^2} \left(\log(|\mathcal{M}|) + \log \frac{1}{\delta}\right)\right)$  samples from  $D$ , an  $\varepsilon$ -fraction of which have been arbitrarily corrupted, returns a single distribution  $D'$  so that with  $1 - \delta$  probability  $d_{\text{TV}}(D', D) < O(\gamma + \varepsilon)$ .*

*Proof of Theorem 6.* Let  $\mathcal{M}_A$  be the set of distributions defined as follows:

$$\mathcal{M}_A = \left\{ \mathcal{N}(\mu', I_d) : \|\mu'\|_0 \leq m, \|\mu' - \mu\|_2 \leq A \text{ and, for all } i, \mu'_i = k_i \frac{\varepsilon}{\sqrt{d}} \text{ for some } k_i \in \mathbb{Z} \right\}.$$

First, we will show that there exists a  $\mathcal{N}(\mu', I_d) \in \mathcal{M}_A$  such that  $\|\mu' - \mu\|_2 \leq O(\varepsilon)$ . Let  $J \subseteq [d]$  be the set of indices of  $m$  coordinates of  $\mu$  with largest magnitude. Define  $\mu'$  as a vector with coordinates  $\mu'_j = \frac{\varepsilon}{\sqrt{d}} \left\lceil \frac{\sqrt{d}\mu_j}{\varepsilon} \right\rceil \mathbb{1}_J$ , where  $\lceil \cdot \rceil$  is the rounding operator. Then we have

$$\|\mu' - \mu\|_2^2 = \sum_{j \in J} (\mu_j - \mu'_j)^2 + \sum_{j \in J^c} \mu_j^2 \leq \frac{\varepsilon^2 m}{d} + \sum_{j \in J^c} \mu_j^2. \quad (9)$$

The first term of the r.h.s. is at most  $\varepsilon^2$ ; as for the second, it can be upper bounded as follows:

$$\begin{aligned} \sum_{j \in J^c} \mu_j^2 &= \sum_{j \in J^c} |\mu_j|^{2-q} |\mu_j|^q \leq \max_{u \in J^c} |\mu_u|^{2-q} \sum_{j \in J^c} |\mu_j|^q \\ &\stackrel{(a)}{\leq} \frac{\|\mu\|_q^{2-q} \|\mu\|_q^q}{(m+1)^{\frac{2-q}{q}}} = \frac{\|\mu\|_q^2}{(m+1)^{2/q-1}} \\ &\leq s^2 (m+1)^{1-2/q} \stackrel{(b)}{\leq} \varepsilon^2, \end{aligned}$$

where (a) uses the fact that  $\max_{u \in J^c} |\mu_u|$  is the  $(m+1)^{\text{th}}$  largest coordinate of  $\mu$  and (b) is by the definition of  $m$ . Hence, we get  $\|\mu' - \mu\|_2 \leq \sqrt{2}\varepsilon$  (and, in particular,  $\mu' \in \mathcal{M}_A$  for every  $A \geq \sqrt{2}\varepsilon$ ).

It is easy to see that, for all  $A > 0$ ,  $|\mathcal{M}_A| \leq \binom{d}{m} (A\sqrt{d}/\varepsilon)^m$ : there are  $\binom{d}{m}$  different ways to select the non-zero coordinates, and for each chosen set of non-zero coordinates there are at most  $(A\sqrt{d}/\varepsilon)^m$  elements in  $\mathcal{M}_A$ .

Now, consider the following algorithm: First, use a naive pruning algorithm with  $N$  samples as input, to output an approximation of  $\mu$ , denoted by  $\mu_0$ . Such a pruning is given and analyzed in [DKK<sup>+</sup>16], which outputs  $\mu_0$  such that  $\|\mu_0 - \mu\|_2 \leq B := O\left(\sqrt{d \log(N/\delta)}\right)$  with probability at least  $1 - \delta$ . Next, round each coordinate of  $\mu_0$  to its nearest integer multiple of  $\frac{\varepsilon}{\sqrt{d}}$ , and output the set of distributions

$$\mathcal{M}' = \left\{ \mathcal{N}(\mu'', I) : \|\mu''\|_0 \leq m, \|\mu'' - \mu_0\|_2 \leq B \text{ and, for all } i, \mu''_i = k_i \frac{\varepsilon}{\sqrt{d}} \text{ for some } k_i \in \mathbb{Z} \right\}.$$

By the triangle inequality, with probability at least  $1 - \delta$  we have  $\mathcal{M}' \subseteq \mathcal{M}_{2B}$ . Hence, by Lemma 6 there exists an algorithm that, with probability  $1 - \delta$ , outputs a  $\mu'$  such that  $\|\mu' - \mu\|_2 \leq O(\varepsilon)$  and the number of samples required is

$$\begin{aligned} O\left(N + \frac{\log |\mathcal{M}_{2B}| + \log \frac{1}{\delta}}{\varepsilon^2}\right) &= O\left(\frac{\log \binom{d}{m} + m \log \frac{d}{\varepsilon} + \log \frac{1}{\delta}}{\varepsilon^2}\right) \\ &= O\left(\frac{m \log \frac{ed}{m} + m \log \frac{d}{\varepsilon} + \log \frac{1}{\delta}}{\varepsilon^2}\right). \end{aligned}$$

This proves Theorem 6. □

### 3.2 Testing in Sparse Linear Regression Model

In this section, we state the formal version of Theorem 3, and provide its proof.

**Theorem 7.** *Let  $\gamma > 0$  be sufficiently small,  $\varepsilon = \frac{\gamma}{C}$  for a sufficiently large  $C$ , and  $s = d^{1-\delta}$  for some  $\delta \in (0, 1)$ . Let  $(X_1, y_1), (X_2, y_2), \dots, (X_n, y_n)$  be i.i.d. samples obtained from the sparse linear regression model described in (3). Moreover, suppose an  $\varepsilon$ -fraction of these  $n$  samples are arbitrarily corrupted. Then, if there exists an algorithm that distinguishes between the cases  $\|\theta\|_2 = 0$  and  $\|\theta\|_2 \geq \gamma$  with probability greater than  $2/3$ , we must have  $n = \Omega\left(\min\left(s \log d, \frac{1}{\gamma^4}\right)\right)$ .*

*Proof.* Given  $(X_1, y_1), (X_2, y_2), \dots, (X_n, y_n)$  as in the statement, let  $X := (X_1, X_2, \dots, X_n)$  and  $Y := (y_1, y_2, \dots, y_n)$ . As noted earlier, while deriving a lower bound, it is enough to consider the weaker  $\varepsilon$ -Huber model for the corruption of samples. We define a hypothesis testing task compliant with the  $\varepsilon$ -Huber contamination model: a lower bound on the sample complexity of this hypothesis testing task will then constitute a lower bound for the robust testing in sparse linear regression model. Let  $\Theta := \{\theta \in \{-\frac{1}{\sqrt{s}}, 0, \frac{1}{\sqrt{s}}\}^d : \|\theta\|_0 = s\}$  and  $p := \frac{1-\varepsilon}{\varepsilon}$ . Further, denote the probability distribution under the hypotheses  $\mathcal{H}_0$  and  $\mathcal{H}_1$  by  $\mathbf{p}$  and  $\mathbf{q}$ , respectively. We define  $\mathcal{H}_0$  and  $\mathcal{H}_1$  as follows:

$$\begin{aligned} \mathcal{H}_0 : (X_i, y_i) &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_{d+1}) \\ \mathcal{H}_1 : \theta &\sim \mathbf{u}_\Theta \\ y_i &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1 + \gamma^2) \\ \mathbf{q}(X_i | y_i, \theta) = \mathbf{q}_\theta(X_i | y_i) &= (1 - \varepsilon)\mathcal{N}\left(\frac{y_i \gamma \theta}{1 + \gamma^2}, I - \frac{\gamma^2 \theta \theta^T}{1 + \gamma^2}\right) \\ &\quad + \varepsilon \mathcal{N}\left(-\frac{(1 - \varepsilon)}{\varepsilon} \frac{y_i \gamma \theta}{(1 + \gamma^2)}, I - \frac{\gamma^2 \theta \theta^T}{1 + \gamma^2}\right) \end{aligned} \quad (10)$$

Note that we can rewrite

$$\mathbf{q}_\theta(X_i | y_i) = \mathbb{E}_b \left[ \mathcal{N}\left(\frac{b y_i \gamma \theta}{1 + \gamma^2}, I - \frac{\gamma^2 \theta \theta^T}{1 + \gamma^2}\right) \right], \quad \text{where } b \sim (1 - \varepsilon)\delta_1 + \varepsilon\delta_{-p};$$

in particular,  $\mathbf{q}_\theta$  is an  $\varepsilon$ -Huber contaminated version of  $\mathcal{N}(0, \Sigma_{\gamma\theta})$ , where  $\Sigma_{\gamma\theta} = \begin{bmatrix} I_d & \gamma\theta \\ \gamma\theta^T & 1 + \gamma^2 \end{bmatrix}$  and thus a valid distribution to consider in this problem. A natural approach to prove the impossibility of detection between  $\mathcal{H}_0$  and  $\mathcal{H}_1$  would be to try and show that  $\chi^2(\mathbf{q} \parallel \mathbf{p}) = o(1)$ . However, as previously discussed, this method does not yield a useful lower bound for this problem, due to low-probability events which cause  $\chi^2(\mathbf{q} \parallel \mathbf{p})$  to blow up.

To circumvent this issue, we take recourse in an alternative approach, the *conditional second moment method* [RXZ19, WX18]. In this method, we first define a high-probability event  $\mathcal{E}$  and then evaluate  $\chi^2(\mathbf{q}^\mathcal{E} \parallel \mathbf{p})$ , where  $\mathbf{q}^\mathcal{E}$  is the distribution  $\mathbf{q}$  conditioned on the event  $\mathcal{E}$ . The idea is that, by this conditioning, we can rule out the rare events that cause  $\chi^2(\mathbf{q} \parallel \mathbf{p})$  to go to infinity. Indeed, suppose the event  $\mathcal{E}$  is chosen so that  $\mathbf{q}(\mathcal{E}) = o(1)$ , and that we are able to show that  $\chi^2(\mathbf{q}^\mathcal{E} \parallel \mathbf{p}) = o(1)$ . This latter statement implies that  $d_{\text{TV}}(\mathbf{q}^\mathcal{E}, \mathbf{p}) = o(1)$ , and so, from the relation  $\mathbf{q} = \mathbf{q}(\mathcal{E})\mathbf{q}^\mathcal{E} + \mathbf{q}(\mathcal{E}^C)\mathbf{q}^{\mathcal{E}^C}$  and the convexity of total variation distance, we have

$$d_{\text{TV}}(\mathbf{q}, \mathbf{p}) \leq \mathbf{q}(\mathcal{E}) d_{\text{TV}}(\mathbf{q}^\mathcal{E}, \mathbf{p}) + \mathbf{q}(\mathcal{E}^C) d_{\text{TV}}(\mathbf{q}^{\mathcal{E}^C}, \mathbf{p}) \leq d_{\text{TV}}(\mathbf{q}^\mathcal{E}, \mathbf{p}) + \mathbf{q}(\mathcal{E}^C) = o(1),$$

which proves the impossibility result.

We now implement this roadmap. For  $\nu > 0$ , we define the event  $\mathcal{E}$  as follows:

$$\mathcal{E} = \left\{ Y \in \mathbb{R}^d : \frac{\|Y\|_2^2}{n} \leq 2 + \nu \right\}, \quad (11)$$

so that the conditional probability distribution  $\mathbf{q}^\mathcal{E}$  is given by

$$\mathbf{q}^\mathcal{E} = \mathbf{q}(X, Y | \mathcal{E}) = \frac{\mathbb{E}_\theta[\mathbf{q}_\theta(X, Y) \mathbf{1}_\mathcal{E}\{Y\}]}{\mathbf{q}(\mathcal{E})}. \quad (12)$$

We will later choose  $\nu$  such that  $\mathcal{E}$  is asymptotically a high-probability event, and for now focus on establishing that  $\chi^2(\mathbf{q}^\mathcal{E} \parallel \mathbf{p}) = o(1)$ . Note that  $\chi^2(\mathbf{q}^\mathcal{E} \parallel \mathbf{p}) = \mathbb{E}_{\mathbf{p}} \left[ \left( \frac{\mathbf{q}^\mathcal{E}}{\mathbf{p}} \right)^2 \right] - 1$ . Now, since  $\frac{\mathbf{q}^\mathcal{E}}{\mathbf{p}} = \frac{\mathbb{E}_\theta[\mathbf{q}_\theta(X, Y)\mathbb{1}_{\mathcal{E}}\{Y\}]}{\mathbf{q}(\mathcal{E})\mathbf{p}(X, Y)}$ , we have

$$\left( \frac{\mathbf{q}^\mathcal{E}}{\mathbf{p}} \right)^2 = \frac{1}{\mathbf{q}(\mathcal{E})^2} \mathbb{E}_{\theta, \theta'} \left[ \frac{\mathbf{q}_\theta(X, Y)\mathbf{q}_{\theta'}(X, Y)}{\mathbf{p}(X, Y)^2} \mathbb{1}_{\mathcal{E}}\{Y\} \right],$$

and so

$$\mathbb{E}_{\mathbf{p}} \left[ \left( \frac{\mathbf{q}^\mathcal{E}}{\mathbf{p}} \right)^2 \right] = \frac{1}{\mathbf{q}(\mathcal{E})^2} \mathbb{E}_{\theta, \theta'} \left[ \mathbb{E}_Y \left[ \frac{\mathbf{q}(Y)^2}{\mathbf{p}(Y)^2} \mathbb{E}_X \left[ \frac{\mathbf{q}_\theta(X | Y)\mathbf{q}_{\theta'}(X | Y)}{\mathbf{p}(X)^2} \right] \mathbb{1}_{\mathcal{E}}\{Y\} \right] \right],$$

where  $Y \sim \mathcal{N}(0, 1)^n$  and  $X \sim \mathcal{N}(0, I_d)^n$ . Since  $\mathcal{E}$  is a high-probability event, we have  $\mathbf{q}(\mathcal{E}) = 1 - o(1)$ . Thus,

$$\mathbb{E}_{\mathbf{p}} \left[ \left( \frac{\mathbf{q}^\mathcal{E}}{\mathbf{p}} \right)^2 \right] = (1 + o(1)) \mathbb{E}_{\theta, \theta'} \left[ \underbrace{\mathbb{E}_Y \left[ \frac{\mathbf{q}(Y)^2}{\mathbf{p}(Y)^2} \mathbb{E}_X \left[ \frac{\mathbf{q}_\theta(X | Y)\mathbf{q}_{\theta'}(X | Y)}{\mathbf{p}(X)^2} \right] \mathbb{1}_{\mathcal{E}}\{Y\} \right]}_{:=Z} \right]. \quad (13)$$

We will evaluate  $\mathbb{E}_{\theta, \theta'}[Z]$  by splitting it into two cases depending on the value of  $|\langle \theta, \theta' \rangle|$ . Let  $\tau := \frac{\varepsilon^2}{4e\gamma^2 \log d}$  (the choice of value for this threshold will become clear in the course of the argument).

**Case 1:**  $|\langle \theta, \theta' \rangle| \leq \tau$ . In this case we drop the  $\mathbb{1}_{\mathcal{E}}\{Y\}$  term in the expression for  $Z$ , simply upper bounding it by 1. Note that in the absence of  $\mathbb{1}_{\mathcal{E}}\{Y\}$  term,  $Z$  becomes a product of expectations, enabling the following simplification:

$$\begin{aligned} & \mathbb{E}_{\theta, \theta'} [Z \mathbb{1}\{|\langle \theta, \theta' \rangle| \leq \tau\}] \\ & \leq \mathbb{E}_{\theta, \theta'} \left[ \left( \mathbb{E}_y \left[ \frac{\mathbf{q}(y)^2}{\mathbf{p}(y)^2} \mathbb{E}_x \left[ \frac{\mathbf{q}_\theta(x | y)\mathbf{q}_{\theta'}(x | y)}{\mathbf{p}(x)^2} \right] \right] \right)^n \mathbb{1}\{|\langle \theta, \theta' \rangle| \leq \tau\} \right], \quad (14) \end{aligned}$$

where  $y \sim \mathcal{N}(0, 1)$  and  $x \sim \mathcal{N}(0, I_d)$ . By substituting for  $\mathbf{q}_\theta$  from (10) and using Corollary 7, we get

$$\mathbb{E}_x \left[ \frac{\mathbf{q}_\theta(x | y)\mathbf{q}_{\theta'}(x | y)}{\mathbf{p}(x)^2} \right] \leq \frac{1}{\sqrt{1 - \gamma^4 \langle \theta, \theta' \rangle^2}} \mathbb{E}_{b, b'} [\exp(\gamma^2 y^2 b b' \langle \theta, \theta' \rangle)]. \quad (15)$$

Note that  $\mathbb{E}[b] = 0$  and, for  $\varepsilon \in (0, 1/2)$ ,  $|\mathbb{E}_b[b^\ell]| = |(1 - \varepsilon) + (-1)^\ell p^\ell \varepsilon| = (1 - \varepsilon) |1 + (-1)^\ell p^{\ell-1}| \leq 2p^{\ell-1}$ . Therefore,

$$\begin{aligned} \mathbb{E}_{b, b'} [\exp(\gamma^2 y^2 b b' \langle \theta, \theta' \rangle)] &= \mathbb{E}_{b, b'} \left[ 1 + \sum_{\ell=1}^{\infty} \frac{\gamma^{2\ell} y^{2\ell} b^\ell b'^\ell \langle \theta, \theta' \rangle^\ell}{\ell!} \right] \\ &\leq 1 + \sum_{\ell=2}^{\infty} \frac{\gamma^{2\ell} y^{2\ell} \mathbb{E}_b[b^\ell]^2 |\langle \theta, \theta' \rangle|^\ell}{\ell!} \\ &\leq 1 + \frac{4}{p^2} \sum_{\ell=2}^{\infty} \frac{\gamma^{2\ell} y^{2\ell} p^{2\ell} |\langle \theta, \theta' \rangle|^\ell}{\ell!}. \end{aligned}$$

By monotone convergence theorem, it then follows that

$$\begin{aligned} & \mathbb{E}_y \left[ \frac{\mathbf{q}(y)^2}{\mathbf{p}(y)^2} \mathbb{E}_{b, b'} [\exp(\gamma^2 y^2 b b' \langle \theta, \theta' \rangle)] \right] \\ & \leq \mathbb{E}_y \left[ \frac{\mathbf{q}(y)^2}{\mathbf{p}(y)^2} \right] + \frac{4}{p^2} \sum_{\ell=2}^{\infty} \frac{\gamma^{2\ell} p^{2\ell} |\langle \theta, \theta' \rangle|^\ell}{\ell!} \mathbb{E}_y \left[ \frac{\mathbf{q}(y)^2}{\mathbf{p}(y)^2} y^{2\ell} \right]. \end{aligned}$$

Now, we use Lemma 3 and the fact that  $\ell! \geq \left(\frac{\ell}{e}\right)^\ell$ . Recalling that we chose  $\tau = \frac{\varepsilon^2}{4e\gamma^2 \log d}$ , we get

$$\begin{aligned} \mathbb{E}_y \left[ \frac{\mathbf{q}(y)^2}{\mathbf{p}(y)^2} \mathbb{E}_{b,b'} [\exp(\gamma^2 y^2 b b' \langle \theta, \theta' \rangle)] \right] &\leq \frac{1}{\sqrt{1-\gamma^4}} \left( 1 + \frac{4c}{p^2} \sum_{\ell=2}^{\infty} \gamma^{2\ell} 4^\ell e^\ell p^{2\ell} |\langle \theta, \theta' \rangle|^\ell \right) \\ &= \frac{1}{\sqrt{1-\gamma^4}} \left( 1 + \frac{64c}{p^2} \frac{\gamma^4 p^4 e^2 |\langle \theta, \theta' \rangle|^2}{(1-\gamma^2 p^2 4e |\langle \theta, \theta' \rangle|)} \right) \\ &\leq \frac{1}{\sqrt{1-\gamma^4}} \left( 1 + \frac{64c\gamma^4 p^2 e^2 |\langle \theta, \theta' \rangle|^2}{1 - \frac{1}{\log d}} \right). \end{aligned} \quad (16)$$

where  $c > 0$  is the constant from Lemma 3. By (15), and (16) and for large enough  $d$  we have

$$\begin{aligned} \mathbb{E}_y \left[ \frac{\mathbf{q}(y)^2}{\mathbf{p}(y)^2} \mathbb{E}_x \left[ \frac{\mathbf{q}_\theta(x|y) \mathbf{q}_{\theta'}(x|y)}{\mathbf{p}(x)^2} \right] \right] &\leq \frac{1}{\sqrt{1-\gamma^4}} \frac{1}{\sqrt{1-\gamma^4 \langle \theta, \theta' \rangle^2}} \left( 1 + \frac{64c\gamma^4 p^2 e^2 |\langle \theta, \theta' \rangle|^2}{1 - \frac{1}{\log d}} \right) \\ &\leq \exp(\gamma^4) \exp(\gamma^4 \langle \theta, \theta' \rangle^2) \exp\left(\frac{128e^2 c \cdot \gamma^4 \langle \theta, \theta' \rangle^2}{\varepsilon^2}\right). \end{aligned}$$

Substituting in (14), we get, for some absolute constant  $C > 0$  (which one can take to be  $C := 256e^2 c$ ),

$$\begin{aligned} \mathbb{E}_{\theta, \theta'} [Z \mathbb{1}\{|\langle \theta, \theta' \rangle| \leq \tau\}] &\leq \exp(n\gamma^4) \mathbb{E}_{\theta, \theta'} \left[ \exp\left(C \frac{n\gamma^4 \langle \theta, \theta' \rangle^2}{\varepsilon^2}\right) \mathbb{1}\{|\langle \theta, \theta' \rangle| \leq \tau\} \right] \\ &\leq \exp(n\gamma^4) \mathbb{E}_{\theta, \theta'} \left[ \exp\left(C \frac{n\gamma^4 \langle \theta, \theta' \rangle^2}{\varepsilon^2}\right) \right]. \end{aligned}$$

Let  $\beta := \sqrt{s}\theta$  and  $\beta' := \sqrt{s}\theta'$ . Then, we can rewrite the above inequality as

$$\mathbb{E}_{\theta, \theta'} [Z \mathbb{1}\{|\langle \theta, \theta' \rangle| \leq \tau\}] \leq \exp(n\gamma^4) \mathbb{E}_{\beta, \beta'} \left[ \exp\left(C \frac{n\gamma^4 \langle \beta, \beta' \rangle^2}{\varepsilon^2 s^2}\right) \right].$$

From there, we can use an argument similar to that in the proof of Theorem 4. It suffices to evaluate the expectation by fixing  $\beta'$ , say to  $\beta_0$ , where  $\beta_0 = (\underbrace{1, 1, \dots, 1}_s, 0, 0, \dots, 0)$ . The random variable  $G := \langle \beta, \beta_0 \rangle$  is then a symmetric

random walk with Hypergeometric( $d, s, s$ ) stopping time. By Lemma 2, for  $n = o\left(\min\left(\frac{\varepsilon^2 s}{\gamma^4} \log \frac{ed}{s}, \frac{1}{\gamma^4}\right)\right)$ , we have

$$\mathbb{E}_{\theta, \theta'} [Z \mathbb{1}\{|\langle \theta, \theta' \rangle| \leq \tau\}] \leq \exp(n\gamma^4) \mathbb{E}_{\theta, \theta'} \left[ \exp\left(C \frac{n\gamma^4 G^2}{\varepsilon^2 s^2}\right) \right] = 1 + o(1), \quad (17)$$

which concludes the analysis of this case.

**Case 2:**  $|\langle \theta, \theta' \rangle| > \tau$ . From (13) we can rewrite

$$\begin{aligned} &\mathbb{E}_{\theta, \theta'} [Z \mathbb{1}\{|\langle \theta, \theta' \rangle| > \tau\}] \\ &= \mathbb{E}_{\theta, \theta'} \left[ \mathbb{E}_Y \left[ \frac{\mathbf{q}(Y)^2}{\mathbf{p}(Y)^2} \prod_{i=1}^n \mathbb{E}_{x_i} \left[ \frac{\mathbf{q}_\theta(x_i|y_i) \mathbf{q}_{\theta'}(x_i|y_i)}{\mathbf{p}(x_i)^2} \right] \mathbb{1}_{\mathcal{E}}\{Y\} \right] \mathbb{1}\{|\langle \theta, \theta' \rangle| > \tau\} \right]. \end{aligned} \quad (18)$$

To proceed further, we will rely on the following technical lemma, a corollary of Lemma 1 whose proof we defer to the end of the section:

**Lemma 7.** Let  $\phi_{\mu, \Sigma}$  denote the density function of  $\mathcal{N}(\mu, \Sigma)$ . Then, for  $\mu_1 = \frac{\gamma b \theta y}{1 + \gamma^2}$ ,  $\mu_2 = \frac{\gamma b' \theta' y}{1 + \gamma^2}$ ,  $\Sigma_1 = I - \frac{\gamma^2 \theta \theta^T}{1 + \gamma^2}$ , and  $\Sigma_2 = I - \frac{\gamma^2 \theta' \theta'^T}{1 + \gamma^2}$ , we have

$$\mathbb{E}_{X \sim \phi_{0, I}} \left[ \frac{\phi_{\mu_1, \Sigma_1}(X) \phi_{\mu_2, \Sigma_2}(X)}{\phi_{0, I}^2(X)} \right] \leq \frac{1}{\sqrt{1 - \gamma^4 \langle \theta, \theta' \rangle^2}} \exp(\gamma^2 y^2 b b' \langle \theta, \theta' \rangle).$$

Invoking Lemma 7, we then can bound the inner expectations as

$$\mathbb{E}_{x_i} \left[ \frac{\mathbf{q}_\theta(x_i | y_i) \mathbf{q}_{\theta'}(x_i | y_i)}{\mathbf{p}(x_i)^2} \right] \leq \frac{1}{\sqrt{1 - \gamma^4 \langle \theta, \theta' \rangle^2}} \mathbb{E}_{b_i, b'_i} [\exp(\gamma^2 y_i^2 b_i b'_i \langle \theta, \theta' \rangle)]. \quad (19)$$

Let  $b = (b_1, b_2, \dots, b_n)$  and  $b' = (b'_1, b'_2, \dots, b'_n)$  be i.i.d. random variables, where  $b_i, b'_i \sim (1 - \varepsilon)\delta_1 + \varepsilon\delta_{-p}$ . Then

$$\begin{aligned} \prod_{i=1}^n \mathbb{E}_{b_i, b'_i} [\exp(\gamma^2 y_i^2 b_i b'_i \langle \theta, \theta' \rangle)] &= \mathbb{E}_{b, b'} \left[ \exp \left( \gamma^2 \left( \sum_{i=1}^n b_i b'_i y_i^2 \right) \langle \theta, \theta' \rangle \right) \right] \\ &\leq \mathbb{E}_{b, b'} \left[ \exp \left( \gamma^2 \left| \sum_{i=1}^n b_i b'_i y_i^2 \right| |\langle \theta, \theta' \rangle| \right) \right] \\ &\leq \exp \left( \gamma^2 p^2 \|Y\|_2^2 |\langle \theta, \theta' \rangle| \right). \end{aligned} \quad (20)$$

Using (19) and (20) we get

$$\begin{aligned} \mathbb{E}_Y \left[ \frac{\mathbf{q}(Y)^2}{\mathbf{p}(Y)^2} \prod_{i=1}^n \mathbb{E}_{X_i} \left[ \frac{\mathbf{q}_\theta(X_i | y_i) \mathbf{q}_{\theta'}(X_i | y_i)}{\mathbf{p}(X_i)^2} \right] \mathbb{1}_{\mathcal{E}}\{Y\} \right] \\ \leq \frac{1}{(1 - \gamma^4)^{n/2}} \mathbb{E}_Y \left[ \frac{\mathbf{q}(Y)^2}{\mathbf{p}(Y)^2} \exp \left( \gamma^2 p^2 \|Y\|_2^2 |\langle \theta, \theta' \rangle| \right) \mathbb{1}_{\mathcal{E}}\{Y\} \right], \end{aligned}$$

which by using the definition of  $\mathcal{E}$  in (11) can be bounded as

$$\begin{aligned} \mathbb{E}_Y \left[ \frac{\mathbf{q}(Y)^2}{\mathbf{p}(Y)^2} \prod_{i=1}^n \mathbb{E}_{X_i} \left[ \frac{\mathbf{q}_\theta(X_i | y_i) \mathbf{q}_{\theta'}(X_i | y_i)}{\mathbf{p}(X_i)^2} \right] \mathbb{1}_{\mathcal{E}}\{Y\} \right] \\ \leq \frac{\exp(\gamma^2 p^2 n(2 + \nu) |\langle \theta, \theta' \rangle|)}{(1 - \gamma^4)^{n/2}} \mathbb{E}_Y \left[ \frac{\mathbf{q}(Y)^2}{\mathbf{p}(Y)^2} \right] \\ = \frac{\exp(\gamma^2 p^2 n(2 + \nu) |\langle \theta, \theta' \rangle|)}{(1 - \gamma^4)^n}. \end{aligned}$$

Recall that  $|\langle \theta, \theta' \rangle| \leq \frac{H}{s}$ , where  $H \sim \text{Hypergeometric}(d, s, s)$ . Therefore, plugging the above equation in (18) yields, letting  $\lambda := \frac{\gamma^2 p^2 n(2 + \nu)}{s}$ ,

$$\mathbb{E}_{\theta, \theta'} [Z \mathbb{1}\{|\langle \theta, \theta' \rangle| > \tau\}] \quad (21)$$

$$\begin{aligned} &\leq \frac{1}{(1 - \gamma^4)^n} \mathbb{E}_{\theta, \theta'} \left[ \exp \left( \frac{\gamma^2 p^2 n(2 + \nu)}{s} s |\langle \theta, \theta' \rangle| \right) \mathbb{1}\{|\langle \theta, \theta' \rangle| > \tau\} \right] \\ &\leq \frac{1}{(1 - \gamma^4)^n} \mathbb{E}_H [\exp(\lambda H) \mathbb{1}\{H > s\tau\}] \\ &\leq \exp(2n\gamma^4) \mathbb{E}_H [\exp(\lambda H) \mathbb{1}\{H > s\tau\}]. \end{aligned} \quad (22)$$

We use Lemma 4 to evaluate the r.h.s. in the above equation as follows:

$$\begin{aligned}
\mathbb{E}_H[\exp(\lambda H)\mathbf{1}\{H > s\tau\}] &= \sum_{h=\lceil s\tau \rceil}^s \exp(\lambda h) \Pr[H = h] \\
&\leq \sum_{h=\lceil s\tau \rceil}^s \exp(\lambda h) \left( \frac{e s^2}{h(d-s+1)} \right)^h \\
&= \sum_{h=\lceil s\tau \rceil}^s \exp\left(\lambda h - h \log \frac{h(d-s+1)}{e s^2}\right) \\
&\leq \sum_{h=\lceil s\tau \rceil}^{\infty} \exp\left(-h \left(\log \frac{\tau(d-s+1)}{e s} - \lambda\right)\right) \\
&\quad \exp\left(-s\tau \underbrace{\left(\log \frac{\tau(d-s+1)}{e s} - \lambda\right)}_{:=T}\right) \\
&= \frac{\exp\left(-s\tau \left(\log \frac{\tau(d-s+1)}{e s} - \lambda\right)\right)}{1 - \exp\left(-\left(\log \frac{\tau(d-s+1)}{e s} - \lambda\right)\right)}. \tag{23}
\end{aligned}$$

Substituting for  $\lambda$  in  $T$ , we get

$$T = \log \frac{\tau(d-s+1)}{e s} - \frac{\gamma^2 n(2+\nu)(1-\varepsilon)^2}{\varepsilon^2 s}.$$

Recall that we already have fixed  $\tau = \frac{\varepsilon^2}{4e\gamma^2 \log d}$ , but kept  $\nu$  (which appears in the definition of  $\mathcal{E}$ , in (11)) as a free parameter. Set  $\nu := \frac{\log \log d}{n}$ . From the theorem statement, we have  $\varepsilon = \frac{\gamma}{C}$  and  $s = d^{1-\delta}$ , and thus,

$$T = \log\left(\frac{(d-s+1)\varepsilon^2}{4e^2\gamma^2 s \log d}\right) - \frac{\gamma^2(2n + \log \log d)(1-\varepsilon)^2}{\varepsilon^2 s} = \Theta\left(\log d - \log \log d - \frac{n}{s}\right).$$

If  $n = o(s \log d)$ , this implies  $T = \omega(1)$ . Since  $s\tau$  is an increasing function of dimension, (23) then further gives that

$$\mathbb{E}_H[\exp(\lambda H)\mathbf{1}\{H > s\tau\}] = o(1)$$

Therefore, by (21), if  $n = o\left(\min\left(s \log d, \frac{1}{\gamma^4}\right)\right)$ ,

$$\mathbb{E}_{\theta, \theta'}[Z\mathbf{1}\{|\langle \theta, \theta' \rangle| > \tau\}] = o(1).$$

which concludes the analysis of the second case.

Plugging the bounds derived in Case 1 and Case 2 in (13) imply that  $\mathbb{E}_{\mathbf{p}}\left[\left(\frac{\mathbf{q}^\varepsilon}{\mathbf{p}}\right)^2\right] = 1 + o(1)$  whenever  $n = o\left(\min\left(s \log d, 1/\gamma^4\right)\right)$ , and so

$$\chi^2(\mathbf{q}^\varepsilon \parallel \mathbf{p}) = o(1).$$

It remains to show that  $\mathcal{E}$  is a high probability event for the chosen  $\nu$ , i.e.,  $\nu = \frac{\log \log d}{n}$ . Let  $Q_{\chi_n^2}$  denote the  $Q$ -function of  $\chi^2$  distribution with  $n$  degrees of freedom. Then, by the definition of  $\mathcal{E}$ ,

$$\begin{aligned}
\mathbf{q}(\mathcal{E}^C) &= Q_{\chi_n^2}(n(2+\nu)) \\
&\stackrel{(a)}{\leq} Q_{\chi_n^2}(n(1 + \sqrt{\nu} + \nu/2)) \\
&\stackrel{(b)}{\leq} \exp\left(-\frac{1}{4}n\nu\right) = \exp\left(-\frac{1}{4}\log \log d\right) = o(1),
\end{aligned}$$

where (a) is due to the fact that  $1 + u \geq \sqrt{u} + u/2$  for all  $u \geq 0$  and (b) is due to standard concentration inequalities for  $\chi_n^2$  random variables.

This was the last piece missing: by the argument outlined at the beginning of this proof, we conclude that

$$d_{\text{TV}}(\mathbf{p}, \mathbf{q}) = o(1)$$

that is, that it is impossible to distinguish between  $\mathcal{H}_0$  and  $\mathcal{H}_1$  with constant probability, as long as  $n = o\left(\min\left(s \log d, \frac{1}{\gamma^4}\right)\right)$ . This concludes the proof of Theorem 7.  $\square$

To conclude, it only remains to prove the technical lemma we invoked in the above proof, Lemma 7.

*Proof of Lemma 7.* Plugging in the values of  $\mu_1, \mu_2, \Sigma_1, \Sigma_2$  in Lemma 1, we can expand

$$\mathbb{E}_{X \sim \phi_{0,I}} \left[ \frac{\phi_{\mu_1, \Sigma_1}(X) \phi_{\mu_2, \Sigma_2}(X)}{\phi_{0,I}^2(X)} \right] = \frac{1 + \gamma^2}{\det(A)^{1/2}} \exp\left(\frac{1}{2}U\right), \quad (24)$$

where  $A = I + \gamma^2(\theta\theta^T + \theta'\theta'^T)$  and

$$U = \gamma^2 y^2 (b\theta + b'\theta')^T A^{-1} (b\theta + b'\theta') - b^2 \gamma^2 y^2 + \frac{b^2 \gamma^4 y^2}{1 + \gamma^2} - b'^2 \gamma^2 y^2 + \frac{b'^2 \gamma^4 y^2}{1 + \gamma^2}.$$

Using the matrix inversion identity, we then have

$$\begin{aligned} A^{-1} &= (I + \gamma^2 [\theta, \theta'] [\theta, \theta']^T)^{-1} \\ &= I - \gamma^2 [\theta, \theta'] \begin{bmatrix} 1 + \gamma^2 & \gamma^2 \langle \theta, \theta' \rangle \\ \gamma^2 \langle \theta, \theta' \rangle & 1 + \gamma^2 \end{bmatrix}^{-1} [\theta, \theta']^T \\ &= I - \frac{\gamma^2}{(1 + \gamma^2)^2 - \gamma^4 \langle \theta, \theta' \rangle^2} \underbrace{\left( (1 + \gamma^2)(\theta\theta^T + \theta'\theta'^T) - \gamma^2 \langle \theta, \theta' \rangle (\theta'\theta'^T + \theta\theta^T) \right)}_{:=V}. \end{aligned}$$

Thus,

$$U = \gamma^2 y^2 \|b\theta + b'\theta'\|^2 - b^2 \gamma^2 y^2 - b'^2 \gamma^2 y^2 - \frac{\gamma^4 y^2 ((b\theta + b'\theta')^T V (b\theta + b'\theta'))}{(1 + \gamma^2)^2 - \gamma^4 \langle \theta, \theta' \rangle^2} + \frac{b^2 \gamma^4 y^2}{1 + \gamma^2} + \frac{b'^2 \gamma^4 y^2}{1 + \gamma^2}.$$

For sufficiently small  $\gamma$ , i.e., smaller than some absolute constant  $\gamma_0 > 0$  (recalling that  $\theta, \theta'$  are unit vectors, and so  $V$  is bounded), we get

$$U \leq 2\gamma^2 y^2 b b' \langle \theta, \theta' \rangle. \quad (25)$$

Using the matrix determinant identity, we have

$$\det(A) = \det(I + \gamma^2 [\theta, \theta'] [\theta, \theta']^T) = \det\left(I + \gamma^2 \begin{bmatrix} 1 & \langle \theta, \theta' \rangle \\ \langle \theta, \theta' \rangle & 1 \end{bmatrix}\right) = (1 + \gamma^2)^2 - \gamma^4 \langle \theta, \theta' \rangle^2.$$

Thus,

$$\frac{1 + \gamma^2}{\det(A)^{1/2}} = \frac{1}{\sqrt{1 - \frac{\gamma^4 \langle \theta, \theta' \rangle^2}{(1 + \gamma^2)^2}}} \leq \frac{1}{\sqrt{1 - \gamma^4 \langle \theta, \theta' \rangle^2}}. \quad (26)$$

Eqs. (24), (25) and (26) together prove the lemma.  $\square$

## 4 Discussion and Future Work

In this section, we discuss some of the limitations of our results, which we believe are ground for possible future work. Firstly, we suspect that our lower bounds are not tight with respect to the parameters  $\gamma$  and  $\varepsilon$ . The dependence on  $\gamma$  and  $\varepsilon$  in Theorems 4 and 5, is  $\Omega(\varepsilon^4/\gamma^4)$ . While the exact dependence on these parameters is still an open problem even for robust Gaussian mean testing (non-sparse), we conjecture that the actual dependence on these parameters should scale as  $\varepsilon^2/\gamma^4$ , and hence believe that our dependence on these parameters is suboptimal. In the current proof, the bottleneck appears while obtaining an upper bound for the chi-squared correlation between two

Huber-contaminated distributions, and we suspect that more advanced techniques might be required to get the right dependence on  $\gamma$  and  $\varepsilon$ .

Furthermore, we restricted ourselves to the case when the covariance of the Gaussian distributions under consideration are identity matrices (“spherical Gaussians”); of course, handling the case of arbitrary covariances is a natural and important question. We believe that deriving the sample complexity in the case of non-identity (and unknown) covariance matrices would require significant additional effort, as well as new techniques and ideas. It is worth pointing out that we are not aware of any work addressing the sample complexity results for non-identity covariances even in the non-robust (but sparse) setting.

## References

- [Bar02] Yannick Baraud. Non-asymptotic minimax rates of testing in signal detection. *Bernoulli*, 8(5):577–606, 2002.
- [BB20] Matthew Brennan and Guy Bresler. Reducibility and statistical-computational gaps from secret leakage. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 648–847. PMLR, 09–12 Jul 2020.
- [BDLS17] Sivaraman Balakrishnan, Simon S. Du, Jerry Li, and Aarti Singh. Computationally efficient robust sparse estimation in high dimensions. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 169–212. PMLR, 07–10 Jul 2017.
- [BR13] Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *Ann. Statist.*, 41(4):1780–1815, 2013.
- [CCC<sup>+</sup>19] A. Carpentier, O. Collier, L. Comminges, A. B. Tsybakov, and Yu. Wang. Minimax rate of testing in sparse linear regression. *Autom. Remote Control*, 80(10):1817–1834, oct 2019.
- [CCT17] Olivier Collier, Laëtitia Comminges, and Alexandre B. Tsybakov. Minimax estimation of linear and quadratic functionals on sparsity classes. *Ann. Statist.*, 45(3):923–958, 2017.
- [CMW15] Tony Cai, Zongming Ma, and Yihong Wu. Optimal estimation and rank detection for sparse spiked covariance matrices. *Probab. Theory Related Fields*, 161(3-4):781–815, 2015.
- [CW20] T. Tony Cai and Yihong Wu. Statistical and computational limits for sparse matrix detection. *Ann. Statist.*, 48(3):1593–1614, 2020.
- [DJ04] David Donoho and Jiashun Jin. Higher criticism for detecting sparse heterogeneous mixtures. *Ann. Statist.*, 32(3):962–994, 2004.
- [DK19] Ilias Diakonikolas and Daniel M. Kane. Recent advances in algorithmic high-dimensional robust statistics, 2019, arXiv: 1911.05911.
- [DK21] Ilias Diakonikolas and Daniel M. Kane. The sample complexity of robust covariance testing. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 1511–1521. PMLR, 15–19 Aug 2021.
- [DKK<sup>+</sup>16] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016*, pages 655–664. IEEE Computer Soc., Los Alamitos, CA, 2016.
- [DKK<sup>+</sup>19] Ilias Diakonikolas, Daniel Kane, Sushrut Karmalkar, Eric Price, and Alistair Stewart. Outlier-robust high-dimensional sparse estimation via iterative filtering. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [DKS17] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional Gaussians and Gaussian mixtures (extended

- abstract). In *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*, pages 73–84. IEEE Computer Soc., Los Alamitos, CA, 2017.
- [Hub64] Peter J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73 – 101, 1964.
- [IS03] Yu. I. Ingster and Irina A. Suslina. *Nonparametric Goodness-of-Fit Testing Under Gaussian Models*, volume 169. Springer New York, 2003.
- [ITV10] Yuri I. Ingster, Alexandre B. Tsybakov, and Nicolas Verzelen. Detection boundary in sparse regression. *Electron. J. Stat.*, 4:1476–1526, 2010.
- [Li17] Jerry Li. Robust sparse estimation tasks in high dimensions, 2017, arXiv: 1702.05860.
- [LRV16] Kevin A. Lai, Anup B. Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016*, pages 665–674. IEEE Computer Soc., Los Alamitos, CA, 2016.
- [LSLC20] Liu Liu, Yanyao Shen, Tianyang Li, and Constantine Caramanis. High dimensional robust sparse regression. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 411–421. PMLR, 26–28 Aug 2020.
- [RXZ19] Galen Reeves, Jiaming Xu, and Ilias Zadik. The all-or-nothing phenomenon in sparse linear regression. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2652–2663. PMLR, 25–28 Jun 2019.
- [Tuk60] John W. Tukey. A survey of sampling from contaminated distributions. In *Contributions to probability and statistics*, pages 448–485. Stanford Univ. Press, Stanford, Calif., 1960.
- [Ver12] Nicolas Verzelen. Minimax risks for sparse regressions: ultra-high dimensional phenomena. *Electron. J. Stat.*, 6:38–90, 2012.
- [WX18] Yihong Wu and Jiaming Xu. Statistical problems with planted structures: Information-theoretical and computational limits, 2018, arXiv: 1806.00118.