

# Private Robust Estimation by Stabilizing Convex Relaxations

**Pravesh K. Kothari**

*Carnegie Mellon University*

PRAVESHK@CS.CMU.EDU

**Pasin Manurangsi**

**Ameya Velingker**

*Google Research*

PASIN@GOOGLE.COM

AMEYAV@GOOGLE.COM

**Editors:** Po-Ling Loh and Maxim Raginsky

## Abstract

We give the first polynomial time and sample  $(\epsilon, \delta)$ -differentially private (DP) algorithm to estimate the mean, covariance and higher moments in the presence of a constant fraction of adversarial outliers. Our algorithm succeeds for families of distributions that satisfy two well-studied properties in prior works on robust estimation: *certifiable subgaussianity* of directional moments and *certifiable hypercontractivity* of degree 2 polynomials. Our recovery guarantees hold in the “right affine-invariant norms”: Mahalanobis distance for mean, multiplicative spectral and relative Frobenius distance guarantees for covariance and injective norms for higher moments. Prior works obtained private robust algorithms for mean estimation of subgaussian distributions with bounded covariance. For covariance estimation, ours is the first efficient algorithm (even in the absence of outliers) that succeeds without any condition-number assumptions.

Our algorithms arise from a new framework that provides a general blueprint for modifying convex relaxations for robust estimation to satisfy strong *worst-case stability* guarantees in the *appropriate parameter norms* whenever the algorithms produce *witnesses of correctness* in their run. We verify such guarantees for a modification of standard sum-of-squares (SoS) semidefinite programming relaxations for robust estimation. Our privacy guarantees are obtained by combining stability guarantees with a new “estimate dependent” noise injection mechanism in which noise scales with the eigenvalues of the estimated covariance. We believe this framework will be useful more generally in obtaining DP counterparts of robust estimators.

Independently of our work, [Ashtiani and Liaw \(2021\)](#) also obtained a polynomial time and sample private robust estimation algorithm for Gaussian distributions.

**Keywords:** private statistics; sum-of-squares; robust statistics

## 1. Introduction

In this work, we consider the problem of efficiently estimating the mean, covariance and, more generally, the higher moments of an unknown high-dimensional probability distribution on  $\mathbb{R}^d$ , given a sample  $y_1, y_2, \dots, y_n \in \mathbb{R}^d$ , under two design constraints: outlier robustness and privacy. The first demands that we build estimators for such basic parameters of probability distributions that tolerate a fixed (dimension-independent) constant fraction of adversarial outliers in the input data. The second demands that our estimators preserve the privacy of individual points  $y_i$ s (that we model as being contributed by different individuals) participating in our input data.

Sans privacy constraints, the problem of robustly estimating the basic parameters of an unknown distribution has been the focus of intense research in algorithmic robust statistics starting with the pioneering works of [Diakonikolas et al. \(2016\)](#); [Lai et al. \(2016\)](#) from 2016. In addition to new (and often, information-theoretically optimal) algorithms for several basic robust estimation

tasks (Kothari and Steurer, 2017; Kothari and Steinhardt, 2017a; Hopkins and Li, 2018; Bakshi and Kothari, 2020a; Diakonikolas et al., 2020b), this line of work has led to a deeper understanding of the properties of the underlying distribution (algorithmic certificates of analytic properties such as subgaussianity (Kothari and Steurer, 2017), hypercontractivity (Bakshi and Kothari, 2020a; Diakonikolas et al., 2020a; Bakshi and Kothari, 2020b) and anti-concentration (Karmalkar et al. (2019); Raghavendra and Yau (2019), resilience (Steinhardt et al., 2018)) that make robust estimation possible along with general frameworks such as outlier filtering and the sum-of-squares (SoS) method for attacking algorithmic problems in robust statistics.

Sans outlier robustness constraints, the task of *private estimation* of the mean and covariance of probability distributions has also seen considerable progress in the recent years. *Differential privacy* (Dwork et al., 2006b) has emerged as a widely-used standard for providing strong individual privacy guarantees. Under this notion, a single sample is not allowed to have too significant of an impact on the output distribution of an algorithm that operates on a dataset. Differential privacy has now been deployed in a number of production systems, including those at Google (Erlingsson et al., 2014; Bittau et al., 2017), Microsoft (Ding et al., 2017), Apple (Greenberg, 2016; Apple Differential Privacy Team, 2017), and the US Census Bureau (Abowd, 2018). While initial approaches to estimating the mean and covariance under differential privacy required *a priori* bounds on the support of the samples, a more recent work (Karwa and Vadhan, 2018) managed to obtain the first private mean estimation algorithm for samples with unbounded support. Subsequent works have built on this progress to obtain differentially private algorithms for mean estimation and covariance estimation (under assumptions on the condition number of the unknown covariance) of Gaussian and heavy-tailed distributions (Kamath et al., 2019a; Bun and Steinke, 2019; Bun et al., 2019; Cai et al., 2019; Biswas et al., 2020; Kamath et al., 2020; Du et al., 2020; Wang et al., 2020; Aden-Ali et al., 2021; Brown et al., 2021).

In this paper, we focus on the task of finding efficient estimation algorithms for mean, covariance and, more generally, higher moments with recovery guarantees in multiplicative spectral distance and relative Frobenius distance. These notions of distance are well-motivated: the multiplicative spectral guarantee is affine invariant and is necessary, for example, to whiten the data or put a set of points in approximate isotropic position. The relative Frobenius guarantee is necessary for obtaining total variation close estimates of an unknown high-dimensional Gaussian. A very recent work of Liu et al. (2021a) found the first private and robust algorithm for mean estimation under natural distributional assumptions with bounded covariance. However, their techniques do not appear to extend to covariance estimation. Informally, this is because in order to obtain privacy guarantees, we need robust estimation algorithms that are *stable*, i.e., whose output suffers from a bounded perturbation when a single data point is changed arbitrarily. When the unknown covariance is bounded, one can effectively assume that the change in a single data point is bounded. However, in general, the covariance of the unknown distribution can be exponentially (in the underlying dimension) varying eigenvalues which precludes such a method (even in the outlier-free regime).

**This work** In this paper, we give the first algorithms for differentially private robust moment estimation with polynomial time and sample complexity. Our algorithms, in fact, provide a general blueprint for transforming any robust estimation algorithm into a differentially private robust estimation algorithm with similar accuracy guarantees as long as the robust estimation algorithm satisfies two key properties: 1) the algorithm is “*witness-producing*,” i.e., the algorithm finds a sequence of “weights” on the input corrupted sample that induce a distribution with a relevant property of the

unknown distribution family (such as certifiable subgaussianity or hypercontractivity) and 2) the algorithm allows for finding weights that minimize a natural strongly convex objective function in polynomial time. Such properties are naturally satisfied by robust estimation algorithms based on sum-of-squares semidefinite programs. Our main technical result is a simple framework that transforms such an algorithm into one that satisfies *worst-case stability* under input perturbation *in the relevant norms on the parameters*. The final ingredient in our framework is a new noise injection mechanism that uses the stability guarantees so obtained to derive privacy guarantees. This mechanism allows obtaining privacy guarantees even though the distribution of the noise being added depends on the unknown quantity being estimated. In particular, such a subroutine allows us to obtain private robust covariance estimation without any assumptions on the condition number of the unknown covariance. We note that even without the robustness constraints, a private covariance estimation algorithm without any assumptions on the condition number was not known prior to our work.

**Robustness implies privacy?** Our blueprint presents an intuitively appealing picture—that robustness, when obtained by estimators that satisfy some additional but generic conditions, implies privacy via a generic transformation. This connection might even appear natural: privacy follows by “adding noise” to the estimates obtained via algorithms that are insensitive or stable with respect to changing any single point in the input, while robustness involves finding estimators that are *insensitive* to the effects of even up to a constant fraction of outliers. Despite this apparent similarity, there are two key differences that prevent such an immediate connection from being true: 1) privacy is a worst-case guarantee while robustness guarantees are only sensible under distributional assumptions, and, 2) privacy guarantees need insensitivity even against “inliers.” Nevertheless, our main result shows that robustness, when obtained via algorithms that satisfy some natural additional conditions, does yield stable/insensitive algorithms that satisfy differential privacy.

In what follows, we describe our results and techniques in more detail.

## 1.1. Our Results

Formally, our results provide differentially private robust estimation algorithms in the strong contamination model, which we define below.

**Definition 1 (Strong Contamination Model)** *Let  $\eta > 0$  be the outlier rate. Given a distribution  $D$  on  $\mathbb{R}^d$  and a parameter  $n \in \mathbb{N}$ , the strong contamination model with outlier rate  $\eta$  gives access to a set  $Y \subseteq \mathbb{R}^d$  of  $n$  points generated as follows: 1) Generate  $X \subseteq \mathbb{R}^d$ , an i.i.d. sample from  $D$  of size  $n$ , 2) Return any (potentially adversarially chosen)  $Y$  such that  $|Y \cap X| \geq (1 - \eta)n$ . In this case, we say that  $Y$  is an  $\eta$ -corruption of  $X$ .*

In the context of analyzing privacy, we will say that two subsets of  $n$  points  $Y, Y' \subseteq \mathbb{R}^d$  (a.k.a. *databases*) are *adjacent* if they differ in exactly one point (i.e.  $|Y \cap Y'| \geq n - 1$ .) We now present our main theorem, which provides a differentially private robust algorithm for moment estimation of an unknown certifiably subgaussian distribution in the strong contamination model.

Our formal guarantees hold for moment estimation of certifiably subgaussian distributions. A distribution  $D$  is *C-subgaussian* if for any direction  $v$  and any  $t \in \mathbb{N}$ ,  $\mathbb{E}_D \langle x - \mu(D), v \rangle^{2t} \leq (Ct)^t (\mathbb{E}_D \langle x - \mu(D), v \rangle^2)^t$  where  $\mu(D)$  is the mean of the distribution  $D$ . Certifiably subgaussianity is a stricter version of such a property that additionally demands that the difference between

the two sides of the inequality be a sum-of-squares (SoS) polynomial in the variable  $v$ . Gaussian distributions, uniform distributions on product domains, all strongly log-concave distributions and, more generally, any distribution that satisfies a Poincaré inequality with a dimension-independent constant (Kothari and Steinhardt, 2017a) are known to satisfy certifiable subgaussianity. See Definition 26 and the preliminaries for a detailed discussion.

Our first result is an algorithm for moment estimation of certifiably subgaussian distributions that runs in polynomial time and has polynomial sample complexity.

**Theorem 2** *Fix  $C_0 > 0$  and  $k \in \mathbb{N}$ . Then, there exists an  $\eta_0 > 0$  such that for any given outlier rate  $0 < \eta \leq \eta_0$  and  $\varepsilon, \delta > 0$ , there exists a randomized algorithm **Alg** that takes an input of  $n \geq n_0 = \tilde{\Omega} \left( \frac{d^{4k}}{\eta^2} \left( 1 + \left( \frac{\ln(1/\delta)}{\varepsilon} \right)^4 + \left( \frac{\ln(1/\delta)}{\varepsilon} \right)^{\frac{2k}{k-1}} \right) \cdot C^{4k} k^{4k+6} \right)$  points  $Y = \{y_1, y_2, \dots, y_n\} \subseteq \mathbb{Q}^d$  (where  $C = 2C_0 + \frac{3 \ln(3/\delta)}{\varepsilon} + \frac{9}{\varepsilon} + 1$ ), runs in time  $(Bn)^{O(k)}$  (where  $B$  is the bit complexity of the entries of  $Y$ ) and outputs either “reject” or estimates  $\hat{\mu} \in \mathbb{Q}^d$ ,  $\hat{\Sigma} \in \mathbb{Q}^d$ , and  $\hat{M}^{(t)} \in \mathbb{Q}^{d \times d \times \dots \times d}$  (for all even  $t < 2k$  such that  $t$  divides  $2k$ ) satisfying the following guarantees:*

1. **Privacy:** *Alg* is  $(\varepsilon, \delta)$ -differentially private with respect to the input  $Y$ , viewed as a  $d$ -dimensional database of  $n$  individuals.
2. **Utility:** *Let  $X = \{x_1, x_2, \dots, x_n\}$  be an i.i.d. sample of size  $n \geq n_0$  from a certifiably  $C_0$ -subgaussian distribution  $\mathcal{D}$  with mean  $\mu_*$ , covariance  $\Sigma_* \succeq 2^{-\text{poly}(d)} I$ , and moment tensors  $M_*^{(t)}$  for  $t \geq 2$ . If  $Y = \{y_1, y_2, \dots, y_n\}$  is an  $\eta$ -corruption of  $X$ , then with probability at least  $9/10$  over the draw of  $X$  and random choices of the algorithm, **Alg** does not reject and outputs estimates  $\hat{\mu} \in \mathbb{Q}^d$ ,  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$ , and  $\hat{M}^{(t)} \in \mathbb{Q}^{d \times d \times \dots \times d}$  (for all  $t < 2k$  such that  $t$  divides  $2k$ ) satisfying the following guarantees:*

$$\forall u \in \mathbb{R}^d, \langle \hat{\mu} - \mu_*, u \rangle \leq O(\sqrt{Ck}) \eta^{1-1/2k} \sqrt{u^\top \Sigma_* u},$$

and,

$$\left( 1 - O((Ck)^{t/2k}) \eta^{1-1/k} \right) \Sigma_* \preceq \hat{\Sigma} \preceq \left( 1 + O((Ck)^{t/2k}) \eta^{1-1/k} \right) \Sigma_*,$$

and, for all even  $t < 2k$  such that  $t$  divides  $2k$ ,

$$\left( 1 - O(Ck) \eta^{1-t/2k} \right) \langle u^{\otimes t}, M_*^{(t)} \rangle \leq \langle u^{\otimes t}, \hat{M}^{(t)} \rangle \leq \left( 1 + O(Ck) \eta^{1-t/2k} \right) \langle u^{\otimes t}, M_*^{(t)} \rangle.$$

In the above and subsequent theorems, we use the  $\tilde{\Omega}$  notation to hide multiplicative logarithmic factors in  $d, C, k, 1/\eta, 1/\varepsilon$ , and  $\ln(1/\delta)$ .

**Discussion** Our algorithm above achieves an error guarantee in the “right” affine-invariant norms similar to the robust moment estimation algorithm of Kothari and Steurer (2017). In particular, the error in the mean in any direction scales proportional to the variance of the unknown distribution providing recovery error bounds in the strong “Mahalanobis error.” Similarly, the error in the covariance is multiplicative in the Löwner ordering. Our algorithm succeeds in the standard word RAM model of computation. In particular, the lower bound assumption on the eigenvalue of the unknown covariance in the statement above is entirely an artifact of numerical issues. Such an assumption can be removed (and in particular, we can deal with rank deficient covariances) if we assume that

the unknown covariance  $\Sigma_*$  has rational entries with polynomial bit complexity. We choose to make an assumption on the smallest eigenvalue of  $\Sigma_*$  for the sake of simpler exposition.

Our algorithm above is obtained by applying a general blueprint that applies to any robust estimation algorithms that use “one-shot rounding” to produce a differentially private version. We explain our general blueprint in more detail in Section 2.

**Applications** Our differentially private moment estimation algorithm immediately allows us to obtain a differentially private mechanism to implement an outlier-robust *method of moments*. This allows us to learn parameters of statistical models that rely on the method of moments, such as mixtures of spherical Gaussians with linearly independent means (Hsu and Kakade, 2013) (that rely on decomposing 3rd moments) as well as independent component analysis (De Lathauwer et al., 2007) (that relies on decomposing fourth moments). We direct the reader to the work on robust moment estimation that details such applications (Kothari and Steurer, 2017).

**Covariance estimation in relative Frobenius error** The above theorem provides a multiplicative spectral guarantee. Such a guarantee, however, only yields a dimension-dependent bound on the Frobenius norm of the error. While this is provably unavoidable for the class of certifiably subgaussian distributions, recent work (Bakshi and Kothari, 2020d) showed that for distributions that satisfy the stronger property of having certifiably hypercontractive degree 2 polynomials (informally speaking, this is the analog of certifiable subgaussianity for moments of degree 2 polynomials instead of linear polynomials  $\langle x, v \rangle$  of the random variable  $x$ ), one can obtain a *dimension-independent* bound on the Frobenius estimation error that vanishes as the fraction of outliers tends to zero. Their algorithm relies on rounding an SoS relaxation with a slightly different constraint system. By working with their constraint system and applying our blueprint for obtaining a “stable” version, we obtain a version of the above theorem with the stronger Frobenius estimation guarantee (see Theorem 64).

By combining our privacy analysis above with the recent work that shows that the algorithm in Bakshi and Kothari (2020d) gives optimal estimation error when analyzed for corrupted samples from a Gaussian distribution, we obtain the following stronger guarantees for private mean and covariance estimation for Gaussian distributions.

**Theorem 3 (Mean and Covariance Estimation for Gaussian Distributions)** Fix  $\varepsilon, \delta > 0$ . Then, there exists an absolute constant  $\eta_0 > 0$  such that for any given outlier rate  $0 < \eta \leq \eta_0$ , there exists a randomized algorithm **Alg** that takes an input of  $n \geq n_0 = \tilde{\Omega}\left(\frac{d^8}{\eta^4}\left(1 + \frac{\ln(1/\delta)}{\varepsilon}\right)^4\right)$  points  $Y \subseteq \mathbb{Q}^d$ , runs in time  $(Bn)^{O(1)}$  (where  $B$  is the bit complexity of the entries of  $Y$ ) and outputs either “reject” or estimates  $\hat{\mu} \in \mathbb{Q}^d$  and  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$  with the following guarantees:

1. **Privacy:** **Alg** is  $(\varepsilon, \delta)$ -differentially private with respect to the input  $Y$ , viewed as a  $d$ -dimensional database of  $n$  individuals.
2. **Utility:** Let  $X = \{x_1, x_2, \dots, x_n\}$  be an i.i.d. sample of size  $n \geq n_0$  from a Gaussian distribution with mean  $\mu_*$  and covariance  $\Sigma_* \succeq 2^{-\text{poly}(d)}I$  such that  $Y$  is an  $\eta$ -corruption of  $X$ . Then, with probability at least  $9/10$  over the random choices of the algorithm, **Alg** outputs estimates  $\hat{\mu} \in \mathbb{Q}^d$  and  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$  satisfying the following guarantees:

$$\forall u \in \mathbb{R}^d, \langle \hat{\mu} - \mu_*, u \rangle \leq \tilde{O}\left(\eta \cdot \frac{\log(1/\delta)}{\varepsilon}\right) \sqrt{u^\top \Sigma_* u},$$

and,

$$\left\| \Sigma_*^{-1/2} \hat{\Sigma} \Sigma_*^{-1/2} - I \right\|_F \preceq \tilde{O} \left( \eta \cdot \sqrt{\frac{\log(1/\delta)}{\varepsilon}} \right),$$

where the  $\tilde{O}$  hides multiplicative logarithmic factors in  $1/\eta$ . In particular,  $d_{\text{TV}}(\mathcal{N}(\hat{\mu}, \hat{\Sigma}), \mathcal{N}(\mu_*, \Sigma_*)) < \tilde{O}(\eta \log(1/\delta)/\varepsilon)$ .

## 1.2. Related Work

Since the works of [Diakonikolas et al. \(2016\)](#); [Lai et al. \(2016\)](#), there has been a spate of works designing robust estimation algorithms for a wide variety of problems, e.g., mean and covariance estimation ([Diakonikolas et al., 2017a,b](#); [Cheng et al., 2019](#); [Dong et al., 2019](#); [Hopkins et al., 2020](#); [Hopkins, 2020](#); [Li and Ye, 2020](#)), mixture models ([Hopkins and Li, 2018](#); [Kothari et al., 2018](#); [Bakshi and Kothari, 2020c](#); [Diakonikolas et al., 2020b](#); [Bakshi et al., 2020](#)), principal component analysis (PCA) ([Kong et al., 2020](#); [Jambulapati et al., 2020](#)), etc. (see [Diakonikolas and Kane \(2019\)](#) for details on recent advances in robust statistics). Moreover, the criterion of *resilience* formulated in [Steinhardt et al. \(2018\)](#) as a sufficient condition for robustly learning a property of a dataset was later generalized in [Zhu et al. \(2019\)](#) in order to deal with a more general class of perturbations.

In the setting of high-dimensional parameter estimation, release of statistics can often reveal significant information about individual data points, which can be problematic in a number of applications in which it is desirable to protect the privacy of individuals while still providing useful aggregate information (e.g., medical data or census data). Attacks exploiting such properties have been investigated in a long line of works ([Dinur and Nissim, 2003](#); [Bun et al., 2014](#); [Dwork et al., 2015](#); [Steinke and Ullman, 2015](#); [Dwork et al., 2017](#); [Shokri et al., 2017](#)). There has thus been much interest in designing statistical algorithms that protect the privacy of individual samples in a dataset.

In the area of differentially privacy, various works have explored private estimation pertaining to Gaussian mixtures ([Nissim et al., 2007](#); [Kamath et al., 2019b](#)), identity testing ([Canonne et al., 2020](#)), Markov random fields ([Zhang et al., 2020](#)), etc.

**Concurrent and subsequent related works** The problem of private robust mean and covariance estimation has been the subject of great interest, resulting in a number of concurrent and independent related works appearing around the time of the initial posting of this work. [Kamath et al. \(2021\)](#) gives an efficient differentially private algorithm for mean and covariance estimation of Gaussian distributions without making condition number assumptions on the covariance. The work of [Liu et al. \(2021b\)](#) gives a statistical feasibility of private robust estimation with optimal sample complexity via a computationally *inefficient* algorithm. Moreover, [Hopkins et al. \(2021\)](#) also uses the sum-of-squares semidefinite programs to obtain a private robust mean estimation algorithm for bounded covariance distributions with *pure differential privacy*. Our result is most directly related to the work of [Ashtiani and Liaw \(2021\)](#), which also obtains efficient private and robust mean and covariance estimation for Gaussian distributions. Finally, we note that [Tsfadia et al. \(2021\)](#) established a framework similar to the one of [Ashtiani and Liaw \(2021\)](#) and applied it to mean estimation and clustering tasks. Subsequently, the authors released an update showing that the framework is also able to produce an efficient private algorithm for (non-robust) covariance estimation of unbounded Gaussians.

## 2. Technical Overview

We now present a high-level overview of our general blueprint for obtaining private versions of robust estimation algorithms. As a running example, we will focus on the problem of obtaining private and robust mean and covariance estimators. Specifically, our goal is to design an algorithm that takes input consisting of  $n$  points, say  $Y \subseteq \mathbb{R}^d$ , along with an *outlier rate*  $\eta$  and returns mean and covariance estimates. We wish the algorithm to be  $(\varepsilon, \delta)$ -differentially private for every  $Y$  (i.e., a “worst-case” guarantee), viewed as a database in which each  $d$ -dimensional point in  $Y$  is contributed by an individual. We would like the outputs of the algorithm to provide faithful estimates whenever  $Y$  is an  $\eta$ -corruption of a i.i.d. sample from a distribution that has  $C$ -subgaussian fourth moments.

For the purpose of the first part of this overview, we recommend the reader to ignore the distinction between certifiable subgaussianity and “vanilla” subgaussianity. Recall that a distribution  $D$  on  $\mathbb{R}^d$  has  $C$ -subgaussian fourth moments if for every  $v \in \mathbb{R}^d$ ,  $\mathbb{E}_{x \sim D} \langle x - \mu(D), v \rangle^4 \leq 4C(\mathbb{E}_{x \sim D} \langle x - \mu(D), v \rangle^2)^2$ . It turns out that the uniform distribution on a  $O(d^2)$  size i.i.d. sample  $X$  from a  $C$ -subgaussian distribution has  $2C$ -subgaussian fourth moments.

**Stable robust estimation algorithms** In order to design differentially private algorithms, we need to find robust moment estimation algorithms that are *stable*. Specifically, a robust moment estimation algorithm **Alg** is stable if the outputs of **Alg** on any pair of adjacent inputs  $Y, Y'$  (i.e., inputs that differ in at most one point but arbitrarily so) are close. Such a guarantee must hold over *worst-case* pairs  $Y, Y'$ —in particular,  $Y$  may not be obtained by taking an  $\eta$ -corruption of an i.i.d. sample from a distribution following our assumptions. This presents a problem at the outset as robust moment estimation algorithms are typically analyzed under *distributional assumptions*. The work of [Liu et al. \(2021a\)](#) addresses this issue by “opening up” an iterative filter based algorithm for robust moment estimation and effectively making every step of the algorithm stable.

### 2.1. A Prototypical Robust Estimator to Privatize

To understand our ideas, it is helpful to work with a “prototypical” but inefficient robust estimation algorithm that we can eventually swap with an efficient one. Let us thus start with a simple (but inefficient) robust estimation algorithm that we call **Alg** in the discussion below.

Input:  $Y = \{y_1, y_2, \dots, y_n\} \subseteq \mathbb{R}^d$  and outlier rate  $\eta > 0$ .

Output: Estimates  $\hat{\mu}, \hat{\Sigma}$  of mean and covariance or “reject.”

Operation:

- (a) Find a *witness* set of  $n$  points  $X' \subseteq \mathbb{R}^d$  such that the uniform distribution on  $X'$  has subgaussian fourth moments and  $|Y \cap X'| \geq (1 - \eta)n$ . Reject if no such  $X'$  exists.
- (b) Return the mean and covariance of  $X'$ .

Observe that the property of having subgaussian fourth moments requires verifying an inequality for every  $v \in \mathbb{R}^d$ , and, in general, there is no efficient (or even sub-exponential time) algorithm known (or expected, modulo the small-set expansion hypothesis) for this problem. Nevertheless, in [Kothari and Steurer \(2017\)](#) (see Section 2), the authors prove that a variant of the above program (which we discuss this at the end of this overview) produces estimates that are guaranteed to be

close to the mean and covariance of  $D$  if  $Y$  is an  $\eta$ -corruption of an i.i.d. sample  $X$  from  $D$ . Note that, though inefficient, such a result is sufficient to establish statistical identifiability of mean and covariance of  $D$  from  $O_\eta(d^2)$  samples. The closeness guarantees in [Kothari and Steurer \(2017\)](#) hold from a more general and basic result that is useful to us in this exposition, which we note below:

**Fact 4 (See Section 2 of [Kothari and Steurer \(2017\)](#), Parameter Closeness from TV Closeness)**

*Suppose  $D, D'$  are two distributions such that 1) both have subgaussian fourth moments and 2) the total variation distance between  $D, D'$  is at most  $\eta$ . Then, for every  $v \in \mathbb{R}^d$ ,  $\langle \mu(D') - \mu(D), v \rangle \leq O(\eta^{3/4})\sqrt{v^\top(\Sigma(D) + \Sigma(D'))v}$  and  $v^\top(\Sigma(D') - \Sigma(D))v \leq O(\sqrt{\eta})v^\top(\Sigma(D) + \Sigma(D'))v$ . We will say that the means (covariances, respectively) of  $D, D'$  are close to within  $O(\eta^{3/4})$  ( $O(\eta^{1/2})$ , respectively) in Mahalanobis distance, to summarize such a guarantee.*

This fact effectively says that if two distributions both have bounded fourth moments and are close in total variation distance, then their parameters (mean and covariance) must be close. In fact, the closeness is in strong *affine-invariant* norms—often called the *Mahalanobis distance*.

**2.2. Robustness Implies Weak Stability of Alg with a Randomized Outlier Rate**

Consider the lack of stability of the above inefficient algorithm: as written, there must be two adjacent  $Y, Y'$  such that **Alg** rejects on  $Y$  but not on  $Y'$ . We therefore introduce our first simple idea to patch the algorithm to prevent it from displaying such “drastic” change in its behavior.

**Randomizing the outlier rate** The following is a simple but useful observation: If **Alg** does not reject on input  $Y$  with outlier rate  $\eta$ , then, **Alg** must also not reject on  $Y'$  outlier rate  $\eta + 1/n$ . To see why, let  $X$  be the set of points with subgaussian fourth moments that intersects  $Y$  in  $(1 - \eta)n$  points. Then, since  $Y$  and  $Y'$  differ in at most one point,  $Y'$  must intersect  $X$  in at least  $(1 - \eta)n - 1 = (1 - (\eta + 1/n))n$  points. Thus, if, instead of a fixed outlier rate  $\eta$ , we ran **Alg** above with an appropriately “randomized” outlier rate, we might expect the rejection probabilities of **Alg** on  $Y, Y'$  to be similar. We can formalize this with a simple truncated Laplace noise injection procedure.

**Robustness implies weak stability in Mahalanobis norms** We now consider whether the estimates computed on  $Y, Y'$  (assuming **Alg** does not reject on  $Y, Y'$ ) are close. First note that the outlier robustness of **Alg** already ensures a *weak stability* property. Specifically, suppose  $X, X'$  are the sets of size  $n$  generated by **Alg** when run on inputs  $Y, Y'$ . Then, since  $Y \cap Y'$  is of size  $n - 1$ ,  $|X \cap X'| \geq (1 - 2\eta)n - 1$ . Next, observe that intersection bound above is equivalent to the uniform distributions on  $X, X'$  having a total variation distance of at most  $2\eta + 1/n$ . Thus, from [Fact 4](#), we know that the parameters of  $X, X'$  are  $O(\eta^{O(1)})$  close in the relative Mahalanobis distance defined above. Thus we obtain stability properties in the *right norms* directly! However, this is a weak stability guarantee since it only provides a fixed constant distance guarantee instead of  $o_n(1)$  that one may expect given that  $Y$  and  $Y'$  differ in at most 1 out of  $n$  points. Nevertheless, our discussion shows that *robustness, via the inefficient algorithm above, immediately implies weak stability*.

**2.3. A Simple Private Robust Mean Estimator from Weak Stability**

Can we derive private algorithms from the weak stability guarantees? If the unknown covariance happens to be *spherical* (i.e., has all of its eigenvalues equal to each other), then the Mahalanobis distance guarantees are in fact equivalent (up to constant factor scaling) to Euclidean distance guarantees. As a result, simply adding Gaussian noise calibrated to the sensitivity bounds yields a private



robust mean estimation algorithm! Indeed, 1) randomizing the outlier rate, 2) working with the SoS relaxation of the above program and 3) adding Gaussian noise to the resulting estimate, immediately yields a simple, straightforward private robust mean estimator that gives essentially optimal sample complexity guarantees (i.e., matching those of the known non-private robust estimators).

**Weak stability is not enough for covariance estimation** The challenge in using weak stability to obtain private robust covariance estimators arises when the covariance is non-spherical (e.g., is rank-deficient or has eigenvalues of vastly different scales), in which case our Mahalanobis or multiplicative spectral stability guarantee does not translate into Euclidean/spectral norm distance guarantees. In particular, if we were to add Gaussian noise, we would end up *scrambling all small eigenvalues up* and end up with no non-trivial recovery guarantee. Indeed, this challenge necessitates a rethink of noise injection mechanisms for covariance estimation in general—standard noise addition mechanisms do not appear meaningful in faithfully preserving eigenvalues of different scales. Prior works (Kamath et al., 2019a) deal with this by iteratively computing some approximate preconditioning matrices. We instead explore *one-shot, blackbox* noise injection mechanisms that still provide us the right guarantees for covariance estimation.

#### 2.4. Noise Injection in Estimate-Dependent Norms

If we wanted to faithfully preserve all eigenvalues (of varying scales) of the unknown covariance, a natural mechanism would be to add noise *linearly transformed with respect to the computed estimate*. For example, if  $\hat{\Sigma}$  is the computed estimate, we would like to consider the mechanism that returns  $\hat{\Sigma} + \hat{\Sigma}^{1/2} Z \hat{\Sigma}^{1/2}$  where  $Z$  is a matrix of random Gaussians. The upshot of such a mechanism is that it adds noise that is scaled relative to the eigenvalues of the estimate  $\hat{\Sigma}$ —directions where  $v^\top \hat{\Sigma} v$  is small get a smaller additive noise than directions where the same quadratic form is large.

However, the distribution of the added noise in this mechanism *depends on the non-privately estimated quantity itself*. Thus, *a priori*, it provides no useful privacy guarantee!

**Key Observation:** Our main idea to rescue this plan is to note that the mechanism above does indeed provide meaningful privacy guarantees (by standard computations from the celebrated Gaussian mechanism) if we are able to guarantee that on any adjacent inputs  $Y, Y'$ , the non-privately computed estimates are  $o_n(1)$  close in relative Frobenius distance (this follows from elementary arguments and is presented in Lemmas 56 and 57)! This observation crucially needs the distance between covariances (in relative Frobenius norm) to tend to 0 as  $n \rightarrow \infty$ ; in fact, we need an inverse polynomial rate to achieve polynomial sample complexity. However, our weak stability guarantee ensures only a weak  $O(\eta^{1/2})$  bound on multiplicative spectral distance, implying a relative Frobenius bound of  $O(\eta^{1/2} \sqrt{d})$ —not only does this not tend to 0 as  $n \rightarrow \infty$  but it, in fact, explodes as  $d \rightarrow \infty$ . Thus, using the above mechanism for covariance estimation requires coming up with significantly stronger (and asymptotically vanishing) stability guarantees. We investigate this next.

#### 2.5. Strong Stability for Robust Estimation Algorithms

**Lack of stability because of multiple differing solutions** There is an important barrier that prevents Alg from offering the strong stability guarantees we need in the covariance estimation mechanism above. Consider the case when  $Y$  is an i.i.d. sample from a one-dimensional standard Gaussian distribution with mean 0 and variance 1 without any outliers added to it. Then,  $\mathcal{N}(0, 1 \pm c\eta)$  for a small enough constant  $c$  is  $\eta$ -close in total variation distance to  $\mathcal{N}(0, 1)$ . By a standard argument,

this means we can choose  $X'$  to be an i.i.d. sample of size  $n$  from  $\mathcal{N}(0, 1 \pm c\eta)$ —if  $n$  is large enough,  $X'$  will have subgaussian fourth moments and intersect  $Y$  in  $(1 - \eta)n$  points. The two difference distributions (and corresponding samples  $X'$ ) however, have variances differing by an additive  $O(\eta)$ —a constant independent of the sample size  $n$ . This shows that even in one dimension, **Alg** has feasible solutions with variance both  $1 - O(\eta)$  and  $1 + O(\eta)$ . Observe that this issue concerns the output of **Alg** itself, which can belong to a range significantly larger than what we can tolerate—we have not yet broached the issue of what happens when we change  $Y$  to an adjacent  $Y'$ .

**Convexification and entropy surrogates** In order to modify **Alg** to output a canonical solution (and with an eye for satisfying the stronger stability property), we wish to make the feasible solution space of **Alg** belong to a convex set (instead of the discrete set of solutions  $X'$  that intersect with  $Y$  in  $(1 - \eta)n$  points). With no fear of computational complexity, this is easy to do in a canonical way: we search instead for a *probability distribution* over  $X'$  that satisfy the constraints that **Alg** imposes. Unlike  $X'$ , distributions on  $X'$  that satisfy the constraints are easily seen to form a convex set.

Given such a convex set, we can resolve our difficulty of not having canonical solutions for any given  $Y$  by simply finding a solution (i.e., a probability distribution  $\zeta$  over  $X'$ ) that minimizes an appropriate strongly convex objective function. Specifically, for any  $X'$ , let  $w_i$  be the 0-1 indicator of those indices  $i$  where  $x_i = y_i$ . Then, the constraints in **Alg** force  $\sum_i w_i \geq (1 - \eta)n$ , and the distribution  $\zeta$  can be thought to be over  $(X', w)$  in a natural way.

To ensure that **Alg** finds a canonical solution, a natural idea is to search over distributions  $\zeta$  over  $(X', w)$  while minimizing some strongly convex function. We choose the simplest:  $\|\mathbb{E}_\zeta[w]\|_2^2$ <sup>1</sup>. We view this objective as a surrogate for finding “maximum entropy solutions” as, when viewing  $\mathbb{E}_\zeta[w_i]$  as defining a probability distribution over  $y_i$ , minimizing the  $\ell_2$  norm favors “spread-out” or high entropy solutions. Since  $\|\mathbb{E}_\zeta[w]\|_2^2$  is a convex function being minimized over convex set of expectations with respect to  $\zeta$ , we expect that the minimizing solution  $\mathbb{E}_{\zeta^*}[w]$  should be unique.

This is not immediately true, however, as our **Alg** as stated outputs the mean of  $X'$  (there could be “multiple”  $X'$  with the same intersection with  $Y$ , in principle).

**Modifying the output of Alg** In order to fit our framework better, we modify the above blueprint in **Alg** to output the weighted average of points in  $Y$  instead of  $X'$ . While such a procedure is not directly analyzed in [Kothari and Steurer \(2017\)](#), the methods there can be naturally adapted to our setting. As a result we obtain the following modified version of **Alg** that we can now work with:

Input:  $Y = \{y_1, y_2, \dots, y_n\} \subseteq \mathbb{R}^d$  and an outlier rate  $\eta > 0$ .

Output: Estimates  $\hat{\mu}, \hat{\Sigma}$  of mean and covariance or “reject.”

Operation:

(a) Find a probability distribution  $\zeta$  over a *witness* set of  $n$  points  $X' \subseteq \mathbb{R}^d$  and intersection indicator  $w \in \{0, 1\}^n$  that minimizes  $\|\mathbb{E}_\zeta[w]\|_2^2$  and is supported on  $(X', w)$  such that 1) the uniform distribution on  $X'$  has subgaussian fourth moments and 2)  $\sum_i w_i \geq (1 - \eta)n$ . Reject if no such  $\zeta$  exists.

1. The exponent of the polynomials appearing in our sample complexity bounds improve if we use a strongly convex function with respect to 1-norm such as  $\|x\|_q^2$  for  $q = 1 + 1/\log d$ . Our interest is in presenting a general “privatizing” blueprint so we continue with the simpler choice above in this work.

(b) Return  $\hat{\mu} = \frac{1}{Z} \sum_i \mathbb{E}_\zeta[w_i] y_i$ ,  $\hat{\Sigma} = \frac{1}{Z} \sum_i \mathbb{E}_\zeta[w_i] (y_i - \hat{\mu})(y_i - \hat{\mu})^\top$  where  $Z = \sum_i \mathbb{E}_\zeta[w_i]$ .

With this modification, **Alg** outputs a canonical single solution on any given  $Y$  (or rejects).

**Stability of Alg from the stability of the entropy potential** We now return to the issue of stability. What happens if we switch the input  $Y$  of **Alg** above to  $Y'$ ? The strongly convex objective we imposed in the above discussion comes in handy here! Namely, by basic convex analysis (see Proposition 24), it follows that if optimum entropy potential values of **Alg** on  $Y$  and  $Y'$  are say,  $O(1)$ -close, then, the vectors  $\mathbb{E}_\zeta[w](Y)$  and  $\mathbb{E}_\zeta[w](Y')$  are themselves  $O(1)$  close. Recall that each  $\mathbb{E}_\zeta[w_i]$  is a number in  $[0, 1]$  and that these numbers add up to 1. Hence, intuitively speaking,  $O(1)$ -closeness of  $\|\mathbb{E}_\zeta[w]\|_2^2$  corresponds to constant perturbation in a constant number of coordinates.

Thus, working with the strongly convex objective above reduces our stability analysis of **Alg** to understanding how much can our entropy potential change when changing a single point in  $Y$ .

Unfortunately, this change can be large in general.  $\|\mathbb{E}_\zeta[w]\|_2^2$  varies between  $(1 - \eta)n$  and  $(1 - \eta)^2 n$ . The additive difference between these two extremes is  $O(\eta n) \gg O(1)$ .

**Stabilizing the entropy potential: private stable selection** Before describing our key idea, we first make a simple observation: Fix an input  $Y$  and consider the optimum value of the entropy potential of **Alg** when run with outlier rate  $\eta$ . What happens if we change  $\eta$  to  $\eta + 1/n$ ? Clearly, the potential cannot *increase*: any solution  $\zeta$  with outlier rate  $\eta$  is also a solution for outlier rate  $\eta + 1/n$ . The potential can decrease arbitrarily though.

More specifically, we show that to make the entropy potential stable under a change of  $Y$  to an adjacent  $Y'$ , it suffices to run  $Y$  with an outlier rate  $\eta' = O(\eta)$  such that the entropy potential of **Alg** on  $Y$  for any outlier rate in  $[\eta' - L/n, \eta' + L/n]$  is within an additive  $\tilde{O}(L/n)$  of any other.

To see why this claim could be true, informally speaking, observe that if  $Y'$  is obtained from  $Y$  by changing at most a single point, then a solution  $\zeta$  with outlier rate  $\eta'$  can be modified into a solution  $\zeta'$  for  $Y'$  with outlier rate  $\eta' + 1/n$  by simply zeroing out the  $w_i$  for the index  $i$  where  $Y'$  and  $Y$  differ. This allows us to relate the potentials for neighboring outlier rates on  $Y$  and  $Y'$ . Under the above assumption, the potential remains stable in an interval around  $\eta'$  on  $Y$ . This allows us to conclude that the same must be true for  $Y'$  for the interval  $[\eta' - L/n + 1, \eta' + L/n - 1]$ .

The above reasoning allows us to obtain strong stability guarantees if we can 1) show that a stable interval as above exists and 2) find such an interval via a stable process.

**A stable selection procedure via the exponential mechanism** We show that a stable interval as above (for  $L = \tilde{O}_n(1)$ ) exists via a simple Markov-like argument. Using an appropriate scoring rule, we show that the standard exponential mechanism can then be used to produce a stable interval like above via a stable algorithm (see Section A.6.4).

**Putting things together** Altogether, we obtain a version of **Alg** that outputs a sequence of weights (i.e.,  $\mathbb{E}_\zeta[w_i]$ ) that are stable under the modification of a single point in  $Y$ . When viewed as a distribution on  $Y$ , the stability guarantee we obtain corresponds to an  $\ell_1$ -stability of  $\tilde{O}(1/\sqrt{n})$  compared to the  $O(\eta)$  (fixed constant) stability that follows from any naive robust estimation algorithm.

We note that  $\tilde{O}(1/\sqrt{n})$  can be upgraded to  $\tilde{O}(1/n)$  if we work with a more sophisticated potential function  $\|x\|_q^2$  for  $q = 1 + 1/\log n$ .

By applying Fact 4, we immediately get that if **Alg** does not reject on  $Y, Y'$ , then the parameters of the respective inputs must be close in the Mahalanobis distance up to a *polynomially vanish-*

ing function of  $n$ , as desired. This allows us to implement the estimate-dependent noise injection mechanism for covariance estimation!

We note that the discussion above can be formalized into an *information-theoretic private identifiability algorithm* (i.e., an inefficient private robust algorithm). We next discuss how to transform the above blueprint result into an efficient algorithm.

## 2.6. From Ideal Algorithms to Efficient Algorithms

Let us now go back and summarize 1) facts about the idealized inefficient algorithm and 2) our general blueprint for making such an algorithm Alg private.

1. **Witness Production:** We have used that the fact that Alg searches over witnesses  $X'$  that share the relevant property of the distributional model we have chosen (e.g., subgaussianity of fourth moments in the above discussion).
2. **Strongly Convex Entropy Potential:** We have minimized a strongly convex potential function in order to ensure that Alg outputs a canonical solution.
3. **Stable Outlier Rate Selection:** We have implemented a randomized stable selection scheme (via the exponential mechanism) for the outlier rate in order to argue that the optimum entropy potential of Alg is stable under the modification of a single point in the input  $Y$ .

We can apply this scheme to any algorithm that outputs a sequence of weights on the input sample  $Y$ , subject to the constraint that 1) the weights induce the relevant property of the distributional model, and 2) they minimize a strongly convex potential function.

**Witness-producing SoS-based robust estimation algorithms** It turns out that we can ensure all the above properties for *efficient* robust estimation algorithms based on “one-shot rounding” of convex relaxations. We specifically rely on the algorithms for robust estimation based on SoS semidefinite programs in this work. The SoS-based algorithms in [Bakshi and Kothari \(2020a\)](#); [Kothari and Steurer \(2017\)](#) that we use almost fit our requirements except with two technical constraints:

1. The algorithms in these works do not output weights on  $Y$  explicitly, but we are able to show that a natural modification that outputs such weights can be analyzed by the same methods.
2. The algorithms in the above prior works were analyzed under distributional assumptions on  $Y$  without the need to explicitly argue that the weights induce good witnesses (a desirable property). Indeed, arguing that these algorithms produce such witnesses on worst-case datasets  $Y$  (whenever they don’t reject) appears difficult. However, we are able to get by without such a statement by observing that we can adapt the analyses of the algorithms in the prior works to infer the following statement: if the algorithm returns a good witness on  $Y$ , then under a small perturbation of the parameters, it must also return a good witness on an adjacent  $Y'$ .

While verifying the properties makes our transformation not entirely blackbox at the moment, we strongly believe that our blueprint demonstrates a conceptually appealing connection between robust algorithm design and private algorithm design. Concretely, we expect our blueprint to be useful in designing more private (and robust) estimation algorithms. Indeed, we believe our techniques immediately extend to other problems where SoS-based robust estimation algorithms are known,

such as linear regression (Klivans et al., 2018; Bakshi and Prasad, 2020) and clustering spherical and non-spherical mixtures (Diakonikolas et al., 2020b; Bakshi and Kothari, 2020a; Hopkins and Li, 2018; Kothari and Steinhardt, 2017b; Fleming et al., 2019).

The detailed analysis with proofs based on the above technical overview appears in [Appendix B](#), while discussion of the upgraded guarantees for certifiably hypercontractive distributions appears in [Appendix C](#). Preliminaries, including notation and definitions, are presented in [Appendix A](#).

## Acknowledgments

Pravesh Kothari is supported by an NSF CAREER Award #2047933, a Google Research Scholar Award, and an Alfred P. Sloan Fellowship.

## References

- John M. Abowd. The u.s. census bureau adopts differential privacy. KDD '18, page 2867, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355520. doi: 10.1145/3219819.3226070. URL <https://doi.org/10.1145/3219819.3226070>.
- Ishaq Aden-Ali, Hassan Ashtiani, and Gautam Kamath. On the sample complexity of privately learning unbounded high-dimensional gaussians. In Vitaly Feldman, Katrina Ligett, and Sivan Sabato, editors, *Algorithmic Learning Theory, 16-19 March 2021, Virtual Conference, Worldwide*, volume 132 of *Proceedings of Machine Learning Research*, pages 185–216. PMLR, 2021. URL <http://proceedings.mlr.press/v132/aden-ali21a.html>.
- Apple Differential Privacy Team. Learning with privacy at scale. *Apple Machine Learning Journal*, 2017. URL <https://machinelearning.apple.com/docs/learning-with-privacy-at-scale/appledifferentialprivacysystem.pdf>.
- Hassan Ashtiani and Christopher Liaw. Private and polynomial time algorithms for learning gaussians and beyond. *CoRR*, abs/2111.11320, 2021. URL <https://arxiv.org/abs/2111.11320>.
- A. Bakshi and P. Kothari. Outlier-robust clustering of non-spherical mixtures. *CoRR*, abs/2005.02970, 2020a. URL <https://arxiv.org/abs/2005.02970>.
- A. Bakshi and P. Kothari. List-decodable subspace recovery via sum-of-squares. *arXiv preprint arXiv:2002.05139*, 2020b.
- A. Bakshi and A. Prasad. Robust linear regression: Optimal rates in polynomial time. *arXiv preprint arXiv:2007.01394*, 2020.
- A. Bakshi, I. Diakonikolas, S. B. Hopkins, D. Kane, S. Karmalkar, and P. K. Kothari. Outlier-robust clustering of gaussians and other non-spherical mixtures. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, pages 149–159. IEEE, 2020. doi: 10.1109/FOCS46700.2020.00023. URL <https://doi.org/10.1109/FOCS46700.2020.00023>.

- Ainesh Bakshi and Pravesh Kothari. Outlier-robust clustering of non-spherical mixtures. *CoRR*, abs/2005.02970, 2020c. URL <https://arxiv.org/abs/2005.02970>.
- Ainesh Bakshi and Pravesh Kothari. Outlier-robust clustering of non-spherical mixtures. 2020d.
- B. Barak, J. A. Kelner, and D. Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method [extended abstract]. In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 143–151. ACM, New York, 2015.
- Sourav Biswas, Yihe Dong, Gautam Kamath, and Jonathan R. Ullman. Coinpress: Practical private mean and covariance estimation. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/a684ecee76fc522773286a895bc8436-Abstract.html>.
- Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnés, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, October 28-31, 2017*, pages 441–459. ACM, 2017. doi: 10.1145/3132747.3132769. URL <https://doi.org/10.1145/3132747.3132769>.
- Gavin Brown, Marco Gaboardi, Adam D. Smith, Jonathan R. Ullman, and Lydia Zakyntinou. Covariance-aware private mean estimation without private covariance estimation. *CoRR*, abs/2106.13329, 2021. URL <https://arxiv.org/abs/2106.13329>.
- Mark Bun and Thomas Steinke. Average-case averages: Private algorithms for smooth sensitivity and mean estimation. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 181–191, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/3ef815416f775098fe977004015c6193-Abstract.html>.
- Mark Bun, Jonathan Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *STOC*, pages 1–10. ACM, 2014.
- Mark Bun, Gautam Kamath, Thomas Steinke, and Zhiwei Steven Wu. Private hypothesis selection. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 156–167, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/9778d5d219c5080b9a6a17bef029331c-Abstract.html>.
- T. Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *CoRR*, abs/1902.04495, 2019. URL <http://arxiv.org/abs/1902.04495>.

- Clément L. Canonne, Gautam Kamath, Audra McMillan, Jonathan R. Ullman, and Lydia Zakyntinou. Private identity testing for high-dimensional distributions. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/72b32a1f754ba1c09b3695e0cb6cde7f-Abstract.html>.
- Yu Cheng, Ilias Diakonikolas, Rong Ge, and David P. Woodruff. Faster algorithms for high-dimensional robust covariance estimation. In Alina Beygelzimer and Daniel Hsu, editors, *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pages 727–757. PMLR, 2019. URL <http://proceedings.mlr.press/v99/cheng19a.html>.
- Lieven De Lathauwer, Joséphine Castaing, and Jean-François Cardoso. Fourth-order cumulant-based blind identification of underdetermined mixtures. *IEEE Trans. Signal Process.*, 55(6, part 2):2965–2973, 2007. ISSN 1053-587X. doi: 10.1109/TSP.2007.893943. URL <http://dx.doi.org/10.1109/TSP.2007.893943>.
- I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Robust estimators in high dimensions without the computational intractability. In *Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 655–664, 2016.
- I. Diakonikolas, S. B. Hopkins, D. Kane, and S. Karmalkar. Robustly learning any clusterable mixture of gaussians. *CoRR*, abs/2005.06417, 2020a. URL <https://arxiv.org/abs/2005.06417>.
- Ilias Diakonikolas and Daniel M. Kane. Recent advances in algorithmic high-dimensional robust statistics. *CoRR*, abs/1911.05911, 2019. URL <http://arxiv.org/abs/1911.05911>.
- Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Being robust (in high dimensions) can be practical. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 999–1008. PMLR, 2017a.
- Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robustly learning a gaussian: Getting optimal error, efficiently. *CoRR*, abs/1704.03866, 2017b.
- Ilias Diakonikolas, Samuel B. Hopkins, Daniel Kane, and Sushrut Karmalkar. Robustly learning any clusterable mixture of gaussians. *CoRR*, abs/2005.06417, 2020b. URL <https://arxiv.org/abs/2005.06417>.
- Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 3571–3580, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/253614bbac999b38b5b60cae531c4969-Abstract.html>.

- Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *PODS*, pages 202–210. ACM, 2003.
- Yihe Dong, Samuel B. Hopkins, and Jerry Li. Quantum entropy scoring for fast robust mean estimation and improved outlier detection. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 6065–6075, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/a4d92e2cd541fca87e4620aba658316d-Abstract.html>.
- Wenxin Du, Canyon Foot, Monica Moniot, Andrew Bray, and Adam Groce. Differentially private confidence intervals. *CoRR*, abs/2001.02285, 2020. URL <http://arxiv.org/abs/2001.02285>.
- Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 371–380. ACM, 2009. doi: 10.1145/1536414.1536466. URL <https://doi.org/10.1145/1536414.1536466>.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014. doi: 10.1561/04000000042. URL <https://doi.org/10.1561/04000000042>.
- Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503. Springer, 2006a. doi: 10.1007/11761679\_29. URL [https://doi.org/10.1007/11761679\\_29](https://doi.org/10.1007/11761679_29).
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006b. doi: 10.1007/11681878\_14. URL [https://doi.org/10.1007/11681878\\_14](https://doi.org/10.1007/11681878_14).
- Cynthia Dwork, Adam D. Smith, Thomas Steinke, Jonathan R. Ullman, and Salil P. Vadhan. Robust traceability from trace amounts. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 650–669. IEEE Computer Society, 2015. doi: 10.1109/FOCS.2015.46. URL <https://doi.org/10.1109/FOCS.2015.46>.
- Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. Exposed! a survey of attacks on private data. *Annual Review of Statistics and Its Application*, 4(1):61–84, 2017. doi: 10.1146/annurev-statistics-060116-054123. URL <https://doi.org/10.1146/annurev-statistics-060116-054123>.



- Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *CCS*, pages 1054–1067, 2014.
- Noah Fleming, Pravesh Kothari, and Toniann Pitassi. Semialgebraic proofs and efficient algorithm design. *Foundations and Trends® in Theoretical Computer Science*, 14(1-2):1–221, 2019. ISSN 1551-305X. doi: 10.1561/04000000086. URL <http://dx.doi.org/10.1561/04000000086>.
- Andy Greenberg. Apple’s “differential privacy” is about collecting your data – but not your data. *Wired*, June, 13, 2016.
- S. B. Hopkins and J. Li. Mixture models, robustness, and sum of squares proofs. In *Proc. 50th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1021–1034, 2018.
- Samuel B. Hopkins. Mean estimation with sub-Gaussian rates in polynomial time. *The Annals of Statistics*, 48(2):1193 – 1213, 2020. doi: 10.1214/19-AOS1843. URL <https://doi.org/10.1214/19-AOS1843>.
- Samuel B. Hopkins, Jerry Li, and Fred Zhang. Robust and heavy-tailed mean estimation made simple, via regret minimization. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/8a1276c25f5efe85f0fc4020fbf5b4f8-Abstract.html>.
- Samuel B. Hopkins, Gautam Kamath, and Mahbod Majid. Efficient mean estimation with pure differential privacy via a sum-of-squares exponential mechanism. *CoRR*, abs/2111.12981, 2021. URL <https://arxiv.org/abs/2111.12981>.
- Daniel Hsu and Sham M. Kakade. Learning mixtures of spherical Gaussians: moment methods and spectral decompositions. In *ITCS’13—Proceedings of the 2013 ACM Conference on Innovations in Theoretical Computer Science*, pages 11–19. ACM, New York, 2013.
- Arun Jambulapati, Jerry Li, and Kevin Tian. Robust sub-gaussian principal component analysis and width-independent Schatten packing. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/b58144d7e90b5a43edc1ca9e642882-Abstract.html>.
- Cédric Jozs and Didier Henrion. Strong duality in Lasserre’s hierarchy for polynomial optimization. *Optim. Lett.*, 10(1):3–10, 2016. ISSN 1862-4472. doi: 10.1007/s11590-015-0868-5. URL <http://dx.doi.org/10.1007/s11590-015-0868-5>.
- Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan R. Ullman. Privately learning high-dimensional distributions. In Alina Beygelzimer and Daniel Hsu, editors, *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pages 1853–1902. PMLR, 2019a. URL <http://proceedings.mlr.press/v99/kamath19a.html>.

- Gautam Kamath, Or Sheffet, Vikrant Singhal, and Jonathan R. Ullman. Differentially private algorithms for learning mixtures of separated gaussians. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 168–180, 2019b. URL <https://proceedings.neurips.cc/paper/2019/hash/68d30a9594728bc39aa24be94b319d21-Abstract.html>.
- Gautam Kamath, Vikrant Singhal, and Jonathan R. Ullman. Private mean estimation of heavy-tailed distributions. In Jacob D. Abernethy and Shivani Agarwal, editors, *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pages 2204–2235. PMLR, 2020. URL <http://proceedings.mlr.press/v125/kamath20a.html>.
- Gautam Kamath, Argyris Mouzakis, Vikrant Singhal, Thomas Steinke, and Jonathan R. Ullman. A private and computationally-efficient estimator for unbounded gaussians. *CoRR*, abs/2111.04609, 2021. URL <https://arxiv.org/abs/2111.04609>.
- S. Karmalkar, A. Klivans, and P. Kothari. List-decodable linear regression. In *Advances in Neural Information Processing Systems*, pages 7423–7432, 2019.
- Vishesh Karwa and Salil P. Vadhan. Finite sample differentially private confidence intervals. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pages 44:1–44:9. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi: 10.4230/LIPIcs.ITCS.2018.44. URL <https://doi.org/10.4230/LIPIcs.ITCS.2018.44>.
- Manuel Kauers, Ryan O’Donnell, Li-Yang Tan, and Yuan Zhou. Hypercontractive inequalities via sos, and the frankl-rödl graph. In *SODA*, pages 1644–1658. SIAM, 2014.
- A. Klivans, P. Kothari, and R. Meka. Efficient algorithms for outlier-robust regression. In *Proc. 31st Annual Conference on Learning Theory (COLT)*, pages 1420–1430, 2018.
- Weihaio Kong, Raghav Somani, Sham M. Kakade, and Sewoong Oh. Robust meta-learning for mixed linear regression with small batches. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/3214a6d842cc69597f9edf26df552e43-Abstract.html>.
- P. K. Kothari and J. Steinhardt. Better agnostic clustering via relaxed tensor norms. *CoRR*, abs/1711.07465, 2017a. URL <http://arxiv.org/abs/1711.07465>.
- P. K. Kothari and D. Steurer. Outlier-robust moment-estimation via sum-of-squares. *CoRR*, abs/1711.11581, 2017. URL <http://arxiv.org/abs/1711.11581>.
- P. K. Kothari, J. Steinhardt, and D. Steurer. Robust moment estimation and improved clustering via sum of squares. In *Proc. 50th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1035–1046, 2018.

- Pravesh K. Kothari and Jacob Steinhardt. Better agnostic clustering via relaxed tensor norms. *CoRR*, abs/1711.07465, 2017b. URL <http://arxiv.org/abs/1711.07465>.
- Pravesh K. Kothari, Peter Manohar, and Brian Hu Zhang. Polynomial-time sum-of-squares can robustly estimate mean and covariance of gaussians optimally, 2021.
- K. A. Lai, A. B. Rao, and S. Vempala. Agnostic estimation of mean and covariance. In *Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 665–674, 2016.
- Jean B. Lasserre. New positive semidefinite relaxations for nonconvex quadratic programs. In *Advances in convex analysis and global optimization (Pythagorion, 2000)*, volume 54 of *Nonconvex Optim. Appl.*, pages 319–331. Kluwer Acad. Publ., Dordrecht, 2001. doi: 10.1007/978-1-4613-0279-7\_18. URL [http://dx.doi.org/10.1007/978-1-4613-0279-7\\_18](http://dx.doi.org/10.1007/978-1-4613-0279-7_18).
- Jerry Li and Guanghao Ye. Robust gaussian covariance estimation in nearly-matrix multiplication time. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/9529fbba677729d3206b3b9073d1e9ca-Abstract.html>.
- Xiyang Liu, Weihao Kong, Sham M. Kakade, and Sewoong Oh. Robust and differentially private mean estimation. *CoRR*, abs/2102.09159, 2021a. URL <https://arxiv.org/abs/2102.09159>.
- Xiyang Liu, Weihao Kong, and Sewoong Oh. Differential privacy and robust statistics in high dimensions. *CoRR*, abs/2111.06578, 2021b. URL <https://arxiv.org/abs/2111.06578>.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 94–103. IEEE Computer Society, 2007. doi: 10.1109/FOCS.2007.41. URL <https://doi.org/10.1109/FOCS.2007.41>.
- Yurii Nesterov. Squared functional systems and optimization problems. In *High performance optimization*, volume 33 of *Appl. Optim.*, pages 405–440. Kluwer Acad. Publ., Dordrecht, 2000. doi: 10.1007/978-1-4757-3216-0\_17. URL [http://dx.doi.org/10.1007/978-1-4757-3216-0\\_17](http://dx.doi.org/10.1007/978-1-4757-3216-0_17).
- Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. Smooth sensitivity and sampling in private data analysis. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 75–84. ACM, 2007. doi: 10.1145/1250790.1250803. URL <https://doi.org/10.1145/1250790.1250803>.
- Pablo A Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.
- Prasad Raghavendra and Morris Yau. List decodable learning via sum of squares. In *Manuscript*, 2019.

- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, pages 3–18. IEEE Computer Society, 2017. doi: 10.1109/SP.2017.41. URL <https://doi.org/10.1109/SP.2017.41>.
- N. Z. Shor. Quadratic optimization problems. *Izv. Akad. Nauk SSSR Tekhn. Kibernet.*, (1):128–139, 222, 1987. ISSN 0002-3388.
- J. Steinhardt, M. Charikar, and G. Valiant. Resilience: A criterion for learning in the presence of arbitrary outliers. In *Proc. 9th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 45:1–45:21, 2018.
- Thomas Steinke and Jonathan R. Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In Peter Grünwald, Elad Hazan, and Satyen Kale, editors, *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, volume 40 of *JMLR Workshop and Conference Proceedings*, pages 1588–1628. JMLR.org, 2015. URL <http://proceedings.mlr.press/v40/Steinke15.html>.
- T. Tao. *Topics in Random Matrix Theory*. Graduate studies in mathematics. American Mathematical Society, 2012. ISBN 9780821874301. URL <https://books.google.com/books?id=L51VAwAAQBAJ>.
- Eliad Tsfadia, Edith Cohen, Haim Kaplan, Yishay Mansour, and Uri Stemmer. Friendlycore: Practical differentially private aggregation. *CoRR*, abs/2110.10132, 2021. URL <https://arxiv.org/abs/2110.10132>.
- Salil P. Vadhan. The complexity of differential privacy. In Yehuda Lindell, editor, *Tutorials on the Foundations of Cryptography*, pages 347–450. Springer International Publishing, 2017. doi: 10.1007/978-3-319-57048-8\_7. URL [https://doi.org/10.1007/978-3-319-57048-8\\_7](https://doi.org/10.1007/978-3-319-57048-8_7).
- Di Wang, Hanshen Xiao, Srinivas Devadas, and Jinhui Xu. On differentially private stochastic convex optimization with heavy-tailed data. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 10081–10091. PMLR, 2020. URL <http://proceedings.mlr.press/v119/wang20y.html>.
- Miaoyan Wang, Khanh Dao Duc, Jonathan Fischer, and Yun S. Song. Operator norm inequalities between tensor unfoldings on the partition lattice. *Linear algebra and its applications*, 520: 44–66, 2017.
- Huanyu Zhang, Gautam Kamath, Janardhan Kulkarni, and Zhiwei Steven Wu. Privately learning markov random fields. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 11129–11140. PMLR, 2020. URL <http://proceedings.mlr.press/v119/zhang20l.html>.
- Banghua Zhu, Jiantao Jiao, and Jacob Steinhardt. Generalized resilience and robust statistics. *CoRR*, abs/1909.08755, 2019. URL <http://arxiv.org/abs/1909.08755>.

## Appendix A. Preliminaries

In this work, we will deal with algorithms that operate on numerical inputs. In all such cases, we will rely on the standard word RAM model of computation and assume that all the numbers are rational represented as a pair of integers describing the numerator and the denominator. In order to measure the running time of our algorithms, we will need to account for the length of the numbers that arise during the run of the algorithm. The following definition captures the size of the representations of rational numbers:

**Definition 5 (Bit Complexity)** *The bit complexity of an integer  $p \in \mathbb{Z}$  is  $1 + \lceil \log_2 p \rceil$ . The bit complexity of a rational number  $p/q$  where  $p, q \in \mathbb{Z}$  is the sum of the bit complexities of  $p$  and  $q$ .*

For any finite set  $X$  of points in  $\mathbb{R}^d$ , we will use  $\mu(X), \Sigma(X), M^{(t)}(X)$  to denote the mean, covariance and the  $t$ -th moment tensor of the uniform distribution on  $X$ .

In the following, we include the relevant background for analysis of algorithms using the sum-of-squares method. We refer the reader to the monograph [Fleming et al. \(2019\)](#) for detailed exposition and proofs of the simple statements included in this section.

### A.1. Pseudo-Distributions

Pseudo-distributions are generalizations of probability distributions and form dual objects to sum-of-squares proofs in a precise sense that we will describe below.

**Definition 6 (Pseudo-distribution, Pseudo-expectations, Pseudo-moments)** *A degree- $\ell$  pseudo-distribution is a finitely-supported function  $D : \mathbb{R}^n \rightarrow \mathbb{R}$  such that  $\sum_x D(x) = 1$  and  $\sum_x D(x) f(x)^2 \geq 0$  for every polynomial  $f$  of degree at most  $\ell/2$ . (Here, the summations are over the support of  $\mu$ .)*

*The pseudo-expectation of a function  $f$  on  $\mathbb{R}^d$  with respect to a pseudo-distribution  $D$ , denoted  $\tilde{\mathbb{E}}_{D(x)} f(x)$ , as*

$$\tilde{\mathbb{E}}_{D(x)} f(x) = \sum_x D(x) f(x) . \tag{A.1}$$

*In particular, the mean  $\mu$  of a pseudo-distribution is defined naturally as the pseudo-expectation of  $f(x) = x$ , i.e.,  $\mu \tilde{\mathbb{E}}_{D(x)} x$ .*

*The degree- $\ell$  moment tensor of a pseudo-distribution  $\mu$  is the tensor  $\mathbb{E}_{\mu(x)}(1, x_1, x_2, \dots, x_n)^{\otimes \ell}$ . In particular, the moment tensor has an entry corresponding to the pseudo-expectation of every monomial of degree at most  $\ell$  in  $x$ .*

Observe that if a pseudo-distribution  $\mu$  satisfies, in addition, that  $\mu(x) \geq 0$  for every  $x$ , then it is a mass function of some probability distribution. Further, a straightforward polynomial-interpolation argument shows that every degree- $\infty$  pseudo-distribution satisfies  $\mu \geq 0$  and is thus an actual probability distribution. The set of all degree- $\ell$  moment tensors of probability distribution is a convex set. Similarly, the set of all degree- $\ell$  moment tensors of degree- $d$  pseudo-distributions is also convex.

We now define what it means for  $\tilde{\mathbb{E}}$  to (approximately) satisfy constraints.

**Definition 7 (Satisfying constraints)** For a polynomial  $g$ , we say that a degree- $k$   $\tilde{\mathbb{E}}$  satisfies the constraint  $\{g = 0\}$  exactly if for every polynomial  $p$  of degree  $\leq k - \deg(g)$ ,  $\tilde{\mathbb{E}}[pg] = 0$  and  $\tau$ -approximately if  $|\mathbb{E}[pg_j]| \leq \tau \|p\|_2$ . We say that  $\tilde{\mathbb{E}}$  satisfies the constraint  $\{g \geq 0\}$  exactly if for every polynomial  $p$  of degree  $\leq k/2 - \deg(g)/2$ , it holds that  $\tilde{\mathbb{E}}[p^2g] \geq 0$  and  $\tau$ -approximately if  $\tilde{\mathbb{E}}[p^2g] \geq -\tau \|p\|_2^2$ .

The following fact describes the precise sense in which pseudo-distributions are duals to sum-of-squares proofs. (See Theorem 3.70 in Fleming et al. (2019) for an exposition.)

**Fact 8 (Strong Duality, (Josz and Henrion, 2016))** Let  $p_1, p_2, \dots, p_k$  be real-coefficient polynomials in  $x_1, x_2, \dots, x_n$ . Suppose there is a degree- $d$  sum-of-squares refutation of the system  $\{p_i(x) \geq 0\}_{i \leq k}$ . Then, there is no pseudo-distribution  $\mu$  of degree  $\geq d$  satisfying  $\{p_i(x) \geq 0\}_{i \leq k}$ . On the other hand, suppose that there is a pseudo-distribution  $\mu$  of degree  $d$  consistent with  $\{p_i(x) \geq 0\}_{i \leq k}$ . Suppose further that the set  $\{p_1, p_2, \dots, p_k\}$  contains the quadratic polynomial  $R - \sum_i x_i^2$  for some  $R > 0$ . Then, there is no degree- $d$  sum-of-squares refutation of the system  $\{p_i(x) \geq 0\}_{i \leq k}$ .

**Basic sum-of-squares (SoS) proofs** In the following, we write  $\{f_i \geq 0 \mid i \leq r\} \Big|_k \{g \geq 0\}$  if every degree- $k$  pseudo-expectation that satisfies  $\{f_i \geq 0 \mid i \leq r\}$  also satisfies  $\{g \geq 0\}$ . Furthermore, we will sometimes write  $\Big|_{\substack{a_1, \dots, a_\ell \\ k}}$  instead of just  $\Big|_k$  to emphasize that the indeterminates are  $a_1, \dots, a_\ell$ .

**Fact 9 (Operator norm Bound)** Let  $A$  be a symmetric  $d \times d$  matrix with rational entries with numerators and denominators upper-bounded by  $2^B$  and  $v$  be a vector in  $\mathbb{R}^d$ . Then, for every  $\varepsilon \geq 0$ ,

$$\Big|_{\frac{v}{2}} \left\{ v^\top A v \leq \|A\|_2 \|v\|_2^2 + \varepsilon \right\}$$

The total bit complexity of the proof is  $\text{poly}(B, d, \log 1/\varepsilon)$ .

**Fact 10 (SoS Hölder's Inequality)** Let  $f_i, g_i$  for  $1 \leq i \leq s$  be indeterminates. Let  $p$  be an even positive integer. Then,

$$\Big|_{\frac{f, g}{p^2}} \left\{ \left( \frac{1}{s} \sum_{i=1}^s f_i g_i^{p-1} \right)^p \leq \left( \frac{1}{s} \sum_{i=1}^s f_i^p \right) \left( \frac{1}{s} \sum_{i=1}^s g_i^p \right)^{p-1} \right\}.$$

The total bit complexity of the SoS proof is  $s^{O(p)}$ .

Observe that using  $p = 2$  yields the SoS Cauchy-Schwarz inequality.

**Fact 11 (SoS Almost Triangle Inequality)** Let  $f_1, f_2, \dots, f_r$  be indeterminates. Then,

$$\Big|_{\frac{f_1, f_2, \dots, f_r}{2t}} \left\{ \left( \sum_{i \leq r} f_i \right)^{2t} \leq r^{2t-1} \left( \sum_{i=1}^r f_i^{2t} \right) \right\}.$$

The total bit complexity of the SoS proof is  $r^{O(t)}$ .

**Fact 12 (SoS AM-GM Inequality, see Appendix A of Barak et al. (2015))** Let  $f_1, f_2, \dots, f_m$  be indeterminates. Then,

$$\{f_i \geq 0 \mid i \leq m\} \Big|_{\frac{f_1, f_2, \dots, f_m}{m}} \left\{ \left( \frac{1}{m} \sum_{i=1}^m f_i \right)^m \geq \prod_{i \leq m} f_i \right\}.$$

The total bit complexity of the SoS proof is  $\exp(O(m))$ .

We will also use the following two consequence of the SoS AM-GM inequality:

**Proposition 13** Let  $a, b$  be indeterminates. Then,

$$\Big|_{\frac{a, b}{2t}} \{a^{2j} b^{2t-2j} \leq j a^{2t} + (t-j) b^{2t}\}.$$

The total bit complexity of the SoS proof is  $\exp(O(t))$ .

**Proof** We apply the SoS AM-GM inequality with  $f_i = a^2$  for  $i = 1, \dots, j$  and  $f_i = b^2$  for  $i = j+1, \dots, t$ . We thus obtain:

$$\Big|_{\frac{a, b}{2t}} \{(j/t a^2 + (1-j/t) b^2)^t \geq a^{2j} b^{2t-2j}\}$$

By the SoS Almost Triangle inequality, we have:

$$\Big|_{\frac{a, b}{2t}} \{(j/t a^2 + (1-j/t) b^2)^t \leq (j a^{2t} + (t-j) b^{2t})\}$$

Combining the above two claims completes the proof. The total bit complexity of the SoS proof follows immediately by using the bounds for the two constituent inequalities used in the proof above.  $\blacksquare$

**Proposition 14** Let  $a, b$  be indeterminates. Then, for any positive integers  $i, t$  such that  $i$  is odd and  $2t \geq i$ , we have:

$$\Big|_{\frac{a, b}{2t}} \left\{ a^i b^{2t-i} \leq \frac{1}{2} (a^{i-1} b^{2t-i+1} + a^{i+1} b^{2t-i-1}) \right\}.$$

The total bit complexity of the SoS proof is  $\exp(O(t))$ .

**Proof** Write  $i = 2r - 1$  for some  $r \geq 1$ . Then, we have:  $a^i b^{2t-i} = a^r b^{t-r} a^{r-1} b^{t-r+1}$ . By the SoS AM-GM inequality with  $f_1 = a^r b^{t-r}$  and  $f_2 = a^{r-1} b^{t-r+1}$ , we thus have:

$$\Big|_{\frac{a, b}{2t}} \left\{ a^i b^{2t-i} = a^r b^{t-r} a^{r-1} b^{t-r+1} \leq \frac{1}{2} (a^{i-1} b^{2t-i+1} + a^{i+1} b^{2t-i-1}) \right\}.$$

**Fact 15 (Cancellation within SoS, Constant RHS (Bakshi and Kothari, 2020d))** Suppose  $a$  is indeterminate and  $t \geq 1$ . Then,

$$\{a^{2t} \leq 1\} \Big|_{\frac{a}{2t}} \{a^2 \leq 1\}$$

Further, the total bit complexity of the SoS proof is at most  $2^{O(t)}$ .

**Lemma 16 (Cancellation within SoS (Bakshi and Kothari, 2020d))** Suppose  $a$  and  $c$  are indeterminates and  $t \geq 1$ . Then,

$$\{a \geq 0 \cup a^t \leq c a^{t-1}\} \Big|_{\frac{a, c}{2t}} \{a^{2t} \leq c^{2t}\}.$$

Further, the total bit complexity of the SoS proof is at most  $2^{O(t)}$ .

### A.2. Algorithms and Numerical Accuracy

The following fact follows by using the ellipsoid algorithm for semidefinite programming. The resulting algorithm to compute pseudo-distributions approximately satisfying a given set of polynomial constraints is called the *sum-of-squares algorithm*.

**Fact 17 (Sum-of-squares algorithm (Shor, 1987; Parrilo, 2000; Nesterov, 2000; Lasserre, 2001))**

*There is an algorithm with the following properties: The algorithm takes input  $B \in \mathbb{N}$ ,  $\tau > 0$ , and polynomials  $p_1, p_2, \dots, p_k$  of degree  $\ell$  with rational coefficients of bit complexity  $B$ . If there is a pseudo-distribution of degree  $d$  consistent with the constraints  $\{p_i(x) \geq 0\}_{i \leq k}$ , the algorithm in time  $(Bn)^{O(d)}$  poly  $\log(1/\tau)$  outputs a pseudo-distribution  $\mu$  of degree  $d$  that  $\tau$ -approximately satisfies  $\{p_i(x) \geq 0\}_{i \leq k}$ .*

### A.3. Tensors

Since we will deal with higher moments of distributions, which are naturally represented as tensors, we will need to define some related notation and conventions for the sake of clarity in our exposition.

Let  $[n] = \{1, 2, \dots, n\}$  for any natural number  $n$ . We define the following.

**Definition 18** *Suppose we have an  $m \times n$  matrix  $M$  and an  $m' \times n'$  matrix  $N$ . We define  $M \otimes N$  to be the standard  $mm' \times nn'$  matrix given by the Kronecker product of  $M$  and  $N$ .*

*Moreover, for an  $m \times n$  matrix  $M$ , we denote by  $M^{\otimes t}$  the  $t$ -fold Kronecker product  $\underbrace{M \otimes M \otimes \dots \otimes M}_{t \text{ times}}$  (of dimension  $tm \times tn$ ).*

Given an  $m \times n$  matrix  $M$ , we will also find it convenient to index  $M^{\otimes t}$  as follows: for any  $1 \leq i_1, i_2, \dots, i_t \leq m$  and  $1 \leq j_1, j_2, \dots, j_t \leq n$ , we can refer to the term  $M^{\otimes t}_{(i_1, i_2, \dots, i_t), (j_1, j_2, \dots, j_t)} = \prod_{k=1}^t M_{i_k, j_k}$ .

We also define a useful *flattening* operation on tensors:

**Definition 19** *Given an  $m_1 \times m_2 \times \dots \times m_t$  tensor  $M$ , we define the flattening, or vectorization, of  $M$  to be the  $(m_1 m_2 \dots m_t)$ -dimensional vector, denoted  $\text{vec}(M)$ , whose entries are precisely the entries of  $M$  appearing in the natural lexicographic order on  $[m_1] \times [m_2] \times \dots \times [m_t]$ . In other words, the entry  $M_{i_1, i_2, \dots, i_t}$  appears before  $M_{j_1, j_2, \dots, j_t}$  (where  $i_k, j_k \in [m_k]$  for  $k = 1, 2, \dots, t$ ) in  $\text{vec}(M)$  if and only if there exists some  $1 \leq k \leq t$  such that  $i_k < j_k$  and  $i_l = j_l$  for all  $l < k$ .*

**Definition 20** *Given an  $n$ -dimensional vector  $u$  and an  $\underbrace{n \times n \times \dots \times n}_{t \text{ times}}$ -dimensional tensor  $M$ , we define  $\langle u^{\otimes t}, M \rangle$  to be  $\langle \text{vec}(u^{\otimes t}), \text{vec}(M) \rangle_{\mathbb{R}^d}$ , i.e., the value of the standard inner product (on  $n^t$ -dimensional vectors) between the flattenings of  $u^{\otimes t}$  and  $M$ .*

A convenient fact we will use is a so-called ‘‘mixed product’’ property for matrices.

**Fact 21** *Given an  $m \times n$  matrix  $A$ ,  $m' \times n'$  matrix  $B$ , and  $n \times n'$  matrix  $V$ , we have that*

$$AVB^T = (A \otimes B) \text{vec}(V),$$

where the above is expressed as matrix-vector product.



Finally, we define the moment tensor for a probability distribution.

**Definition 22** Given a probability distribution  $\mathcal{D}$  on  $\mathbb{R}^d$  and an integer  $t > 1$ , we define the  $t^{\text{th}}$  moment tensor  $M$  to be a  $\underbrace{d \times d \times \cdots \times d}_{t \text{ times}}$  tensor whose entries are given by  $M_{i_1, i_2, \dots, i_t} = \mathbb{E}_{X \sim \mathcal{D}}[X_{i_1} X_{i_2} \cdots X_{i_t}]$  for  $i_1, i_2, \dots, i_t \in [d]$ .

#### A.4. Basic Convexity

We will use the following basic propositions about convexity in our analysis.

**Proposition 23 (Neighborhoods of minimizers of convex functions)** Let  $K$  be a closed convex subset of  $\mathbb{R}^N$ . Let  $f$  be a smooth convex function on  $\mathbb{R}^N$ . Let  $x$  be a minimizer of  $f$  on  $K$ . Then, for every  $y \in K$ ,  $\langle y - x, \nabla f(x) \rangle \geq 0$ .

**Proof** If not, then for a small enough positive  $\lambda$ ,  $f(x + \lambda(y - x)) < f(x)$ . But,  $x + \lambda(y - x) = (1 - \lambda)x + \lambda y \in K$ . ■

**Proposition 24 (Pythagorean theorem from strong convexity w.r.t 2 norm)** Let  $K$  be a convex subset of  $\mathbb{R}^d$  for  $d \in \mathbb{N}$ . Let  $x$  be a minimizer of the convex function  $f(x) = \|x\|_2^2$  on  $K$ . Let  $y \in K$ . Then,  $f(y) - f(x) \geq \|y - x\|_2^2$ .

**Proof** We have:  $\|y\|_2^2 = \|y - x\|_2^2 + \|x\|_2^2 + 2\langle y - x, x \rangle$ . The proposition follows by applying Proposition 23 to observe that  $\langle y - x, x \rangle \geq 0$ . ■

We will also need the following basic bound:

**Lemma 25** Suppose  $x, y \in [0, 1]^n$  such that  $\sum_i x_i, \sum_i y_i \geq n/2$  and  $\|x - y\|_1 \leq \beta n$  for  $\beta \leq 1/10$ . Let  $\bar{x} = \frac{x}{\|x\|_1}$  and  $\bar{y} = \frac{y}{\|y\|_1}$  be the normalized versions of  $x, y$ . Then,

$$\|\bar{x} - \bar{y}\|_1 \leq 6\beta.$$

**Proof** Suppose, without loss of generality, that  $\|x\|_1 = c_1 n \geq c_2 n = \|y\|_1$  for  $c_1, c_2 \geq 1/2$ . Then, we know that  $\|y\|_1 = c_2 n \geq (c_1 - \beta)n$ . Thus,  $\|\bar{x} - \bar{y}\|_1 \leq \frac{1}{c_1 c_2 n^2} (\|x\|_1 \|y\|_1 - \|y - x\|_1 \|x\|_1) \leq \frac{1}{c_1 c_2 n^2} (c_1 n \|x - y\|_1 + \beta n^2) \leq 6\beta$ . ■

#### A.5. Certifiable Subgaussianity

**Definition 26 (Certifiable Subgaussianity)** A distribution  $D$  on  $\mathbb{R}^d$  with mean  $\mu_*$  is said to be  $2k$ -certifiably  $C$ -subgaussian if there is a degree  $2k$  sum-of-squares proof of the following polynomial inequality in  $d$ -dimensional vector-valued indeterminate  $v$ :

$$\mathbb{E}_{x \sim D} \langle x - \mu_*, v \rangle^{2k} \leq (Ck)^k \left( \mathbb{E}_{x \sim D} \langle x - \mu_*, v \rangle^2 \right)^k.$$

Furthermore, we say that  $D$  is certifiable  $C$ -subgaussian if it is  $2k$ -certifiably  $C$ -subgaussian for every  $k \in \mathbb{N}$ .

A finite set  $X \subseteq \mathbb{R}^d$  is said to be  $2k$ -certifiably  $C$ -subgaussian if the uniform distribution on  $X$  is  $2k$ -certifiably  $C$ -subgaussian.

**Fact 27 (Consequence of Theorem 1.2 in Kothari and Steurer (2017))** *Let  $Y$  be a collection of  $n$  points in  $\mathbb{R}^d$ . Let  $p, p' \in [0, 1]^n$  be weight vectors satisfying  $\|p\|_1, \|p'\|_1 = 1$ , and  $\|p - p'\|_1 = \tau$ . Suppose that the distributions on  $Y$  where the probability of  $i$  is  $p_i$  ( $p'_i$ , respectively) is  $2k$ -certifiably  $C_1$  ( $C_2$ , respectively) subgaussian. Let  $\mu_p = \sum_i p_i y_i$ ,  $\Sigma_p = \sum_i p_i (y_i - \mu_p)(y_i - \mu_p)^\top$ , and  $M_p^{(t)} = \sum_i p_i y_i^{\otimes t}$  for every  $t \in \mathbb{N}$  be the mean, covariance and  $t$ -th moment tensor of distribution defined  $p$ . Define  $\mu_{p'}, \Sigma_{p'}, M_{p'}^{(t)}$  similarly for the distribution corresponding to  $p'$ .*

*Then, for every  $\tau \leq \eta_0$  for some absolute constant  $\eta_0$ , for every  $u \in \mathbb{R}^d$ ,  $C' = C_1 + C_2$  and  $t \leq k$ :*

$$\begin{aligned} \langle \mu_p - \mu_{p'}, u \rangle &\leq \tau^{1-1/2k} \cdot O(\sqrt{C'k}) \sqrt{u^\top \Sigma_p u}, \\ (1 - O(C'k)\tau^{1-1/k})\Sigma_p &\preceq \Sigma_{p'} \preceq (1 + O(C'k)\tau^{1-1/k})\Sigma_{p'}, \\ (1 - O(C'^{t/2}k^{t/2})\tau^{1-t/2k})\langle u^{\otimes t}, M_p^{(t)} \rangle &\leq \langle u^{\otimes t}, \hat{M}_{p'}^{(t)} \rangle \leq \langle u^{\otimes t}, M_{p'}^{(t)} \rangle, \end{aligned}$$

## A.6. Differential Privacy

In this section, we state a few tools from differential privacy (DP) literature that will be used in our algorithms. We start by recalling the definition of DP:

**Definition 28 (Differential Privacy (Dwork et al., 2006b))** *An algorithm  $\mathcal{M} : \mathcal{Y} \rightarrow \mathcal{O}$  is said to be  $(\varepsilon, \delta)$ -differentially private (or  $(\varepsilon, \delta)$ -DP) for  $\varepsilon, \delta > 0$  iff, for every  $S \subseteq \mathcal{O}$  and every neighboring datasets  $Y, Y'$ , we have*

$$\mathbb{P}[\mathcal{M}(Y) \in S] \leq e^\varepsilon \cdot \mathbb{P}[\mathcal{M}(Y') \in S] + \delta.$$

Throughout this work, our set  $Y$  will consist of  $y_1, \dots, y_n \in \mathbb{R}^d$ .  $Y = (y_1, \dots, y_n)$  and  $Y' = (y'_1, \dots, y'_n)$  are neighbors iff they differ on a single data point, i.e.,  $y'_j = y_j$  for all  $j \neq i$ . Note that this is the so-called *substitution* variant of DP; another popular variant is the *add/remove* DP where a neighboring  $Y'$  results from adding or removing an example from  $Y$ . We remark that it is not hard to extend our algorithm to the add/remove DP setting, by first computing a DP estimate  $\hat{n}$  of  $n$  and either throwing away random elements or adding zero vectors to arrive at an  $n$ -size dataset on which our algorithm can be applied.

### A.6.1. HOCKEY-STICK DIVERGENCE

It will be convenient in our analysis to use an equivalent definition of DP based on the *hockey-stick divergence*. For ease of notation, let  $[a]_+ = \max\{a, 0\}$  for all  $a \in \mathbb{R}$ .

**Definition 29 (Hockey-Stick Divergence)** *Let  $p(x), q(x)$  be probability density functions on  $\mathbb{R}^d$ , and  $\alpha$  a non-negative real number. The Hockey-stick divergence  $D_\alpha(p, q)$  between  $p, q$  is defined as:*

$$D_{e^\varepsilon}(p, q) = \int_{x \in \mathbb{R}^d} [p(x) - \alpha \cdot q(x)]_+ dx.$$

The following fact is simple to derive from the definition of DP and is often used in literature.

**Fact 30 ( $(\varepsilon, \delta)$ -DP from Hockey-Stick Divergence Bounds)** *Let  $\mathcal{M} : \mathcal{Y} \rightarrow \mathbb{R}^d$  be a randomized algorithm.  $\mathcal{M}$  is  $(\varepsilon, \delta)$ -DP under condition  $\Psi$  iff for any neighboring pair of databases  $Y, Y'$  both satisfying  $\Psi$ , we have  $D_{e^\varepsilon}(\mathcal{M}(Y), \mathcal{M}(Y')) \leq \delta$ .*

We will need to bound the hockey-stick divergence between two distributions in terms of the hockey-stick divergences to a third distribution. Unfortunately, the hockey-stick divergence does not define a metric and, therefore, does not admit the usual triangle inequality. However, it is possible to prove a looser inequality, which we will find useful:

**Lemma 31** *Suppose  $p(x), q(x), r(x)$  are probability density functions on  $\mathbb{R}^d$ . Then,*

$$D_{e^\varepsilon}(p, r) \leq D_{e^{\varepsilon/2}}(p, q) + e^{\varepsilon/2} \cdot D_{e^{\varepsilon/2}}(q, r).$$

We remark that such a bound is already implicit in the so-called *group differential privacy* (see e.g. (Vadhan, 2017, Lemma 2.2)). Nonetheless, we provide a short proof below.

**Proof** Then, note that

$$\begin{aligned} D_{e^\varepsilon}(p, r) &= \int_{x \in \mathbb{R}^d} [p(x) - e^\varepsilon r(x)]_+ dx \\ &= \int_{x \in \mathbb{R}^d} [(p(x) - e^{\varepsilon/2} q(x)) + (e^{\varepsilon/2} q(x) - e^\varepsilon r(x))]_+ dx \\ &\leq \int_{x \in \mathbb{R}^d} [(p(x) - e^{\varepsilon/2} q(x))]_+ dx + \int_{x \in \mathbb{R}^d} [e^{\varepsilon/2} q(x) - e^\varepsilon r(x)]_+ dx \\ &= \int_{x \in \mathbb{R}^d} [(p(x) - e^{\varepsilon/2} q(x))]_+ dx + e^{\varepsilon/2} \int_{x \in \mathbb{R}^d} [q(x) - e^{\varepsilon/2} r(x)]_+ dx \\ &= D_{e^{\varepsilon/2}}(p, q) + e^{\varepsilon/2} \cdot D_{e^{\varepsilon/2}}(q, r), \end{aligned}$$

as desired. ■

#### A.6.2. LAPLACE MECHANISM AND ITS VARIANTS

The Laplace mechanism (Dwork et al., 2006b) is among the most widely used mechanisms in differential privacy. It works by adding a noise drawn from the Laplace distribution (defined below) to the output of the function one wants to privatize.

**Definition 32 (Laplace Distribution)** *The Laplace distribution with mean  $\mu$  and parameter  $b$  on  $\mathbb{R}$ , denoted by  $\text{Lap}(\mu, b)$ , has the PDF  $\frac{1}{2b} e^{-|x-\mu|/b}$ .*

We will also use the “truncated” version of the Laplace mechanism where the noise distribution is shifted and truncated to be non-negative. The precise definition of the noise distribution and its guarantee is given below.

**Definition 33 (Truncated Laplace Distribution)** *The (negatively) truncated Laplace distribution with mean  $\mu$  and parameter  $b$ , denoted by  $\text{tLap}(\mu, b)$  is defined as  $\text{Lap}(\mu, b)$  conditioned on the value being negative.*

We also state a bound on the tail probability of the truncated Laplace distribution which will be useful in our subsequent analysis.

**Lemma 34** *Suppose  $\mu < 0$  and  $b > 0$ . Let  $X \sim \text{tLap}(\mu, b)$ . Then, for  $y < \mu$ , we have that*

$$\mathbb{P}[X < y] = \frac{e^{(y-\mu)/b}}{2 - e^{\mu/b}}.$$

Finally, we prove the DP guarantee of the truncated Laplace mechanism. Note that this is well known but we provide the proof for completeness.

**Lemma 35 (Truncated Laplace Mechanism)** *Let  $f : \mathcal{Y} \rightarrow \mathbb{R}$  be any function with sensitivity at most  $\Delta$ . Then the algorithm that adds  $\text{tLap}\left(-\Delta\left(1 + \frac{\ln(1/\delta)}{\varepsilon}\right), \Delta/\varepsilon\right)$  to  $f$  satisfies  $(\varepsilon, \delta)$ -DP.*

**Proof** Consider any neighboring datasets  $Y, Y'$  and let  $\mathcal{M}$  denote the truncated Laplace mechanism (with parameter as specified). Let  $p, q$  denote the probability density functions of  $\mathcal{M}(Y), \mathcal{M}(Y')$ . Observe that  $p(x) \leq e^\varepsilon \cdot q(x)$  for all  $x < \min\{f(Y), f(Y')\}$ . Thus, we have

$$\begin{aligned} D_{e^\varepsilon}(p, q) &= \int_{x \in \mathbb{R}} [p(x) - e^\varepsilon q(x)]_+ dx \\ &= \int_{x \geq \min\{f(Y), f(Y')\}} [p(x) - e^\varepsilon q(x)]_+ dx \\ &\leq \int_{x \geq \min\{f(Y), f(Y')\}} p(x) dx \\ &\text{(Since sensitivity of } f \text{ is at most } \Delta) \leq \int_{x \geq f(Y) - \Delta} p(x) dx \\ &\text{(Lemma 34)} \leq \delta, \end{aligned}$$

which means that the truncated Laplace mechanism is indeed  $(\varepsilon, \delta)$ -DP. ■

### A.6.3. COMPOSITION THEOREM

It will be convenient to also consider DP algorithms whose privacy guarantee holds only against subsets of inputs. Specifically, we define:

**Definition 36 (Differential Privacy Under Condition)** *An algorithm  $\mathcal{M} : \mathcal{Y} \rightarrow \mathcal{O}$  is said to be  $(\varepsilon, \delta)$ -differentially private under condition  $\Psi$  (or  $(\varepsilon, \delta)$ -DP under condition  $\Psi$ ) for  $\varepsilon, \delta > 0$  iff, for every  $S \subseteq \mathcal{O}$  and every neighboring datasets  $Y, Y'$  both satisfying  $\Psi$ , we have*

$$\mathbb{P}[\mathcal{M}(Y) \in S] \leq e^\varepsilon \cdot \mathbb{P}[\mathcal{M}(Y') \in S] + \delta.$$

It is not hard to see that an analogue of the basic composition theorem still holds in this setting, which we formalize below. We remark that this is similar to the composition theorem derived in (Dwork and Lei, 2009, Section 5).

**Lemma 37 (Composition for Algorithm with Halting)** *Let  $\mathcal{M}_1 : \mathcal{Y} \rightarrow \mathcal{O}_1 \cup \{\perp\}, \mathcal{M}_2 : \mathcal{O}_1 \times \mathcal{Y} \rightarrow \mathcal{O}_2 \cup \{\perp\}, \dots, \mathcal{M}_k : \mathcal{O}_{k-1} \times \mathcal{Y} \rightarrow \mathcal{O}_k \cup \{\perp\}$  be algorithms. Furthermore, let  $\mathcal{M}$  denote the algorithm that proceeds as follows (with  $o_0$  being empty): For  $i = 1, \dots, k$ , compute  $o_i = \mathcal{M}_i(o_{i-1}, Y)$  and, if  $o_i = \perp$ , halt and output  $\perp$ . Finally, if the algorithm has not halted, then output  $o_k$ .*

*Suppose that:*

- For any  $1 \leq i < k$ , we say that  $Y$  satisfies the condition  $\Psi_i$  if running the algorithm on  $Y$  does not result in halting after applying  $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_i$ .

- $\mathcal{M}_1$  is  $(\varepsilon_1, \delta_1)$ -DP.
- $\mathcal{M}_i$  is  $(\varepsilon_i, \delta_i)$ -DP (with respect to neighboring datasets in the second argument) under condition  $\Psi_{i-1}$  for all  $i = \{2, \dots, k\}$ .

Then,  $\mathcal{M}$  is  $(\sum_{i \in [k]} \varepsilon_i, \sum_{i \in [k]} \delta_i)$ -DP.

The proof of the composition lemma follows from that of the standard adaptive composition of approximate DP proof (Dwork and Lei, 2009, Theorem 16). For the sake of completeness, we provide the proof below. Note that we use the notation  $[x]_+$  to denote  $\max\{x, 0\}$  and  $x \wedge y$  to denote  $\min\{x, y\}$ .

**Proof** It suffices to prove the theorem for  $k = 2$  as we may then apply induction to arrive at the statement for any positive integer  $k$ . To prove the case  $k = 2$ , consider any  $S \subseteq \mathcal{O}_2 \cup \{\perp\}$  and any pair of neighboring datasets  $Y, Y'$ .

For any  $S_1 \subseteq \mathcal{O}_1 \cup \{\perp\}$ , we define the measure  $\mu(S_1) := [\mathbb{P}[\mathcal{M}_1(Y) \in S_1] - e^{\varepsilon_1} \mathbb{P}[\mathcal{M}_1(Y') \in S_1]]_+$ . Note that we have  $\mu(\mathcal{O}_1) \leq \delta_1$  due to our assumption that  $\mathcal{M}_1$  is  $(\varepsilon_1, \delta_1)$ -DP.

Now consider four cases:

- **Both  $Y, Y'$  satisfy  $\Psi_1$ .** In this case, we may appeal to  $(\varepsilon_2, \delta_2)$ -DP under  $\Psi_1$  of  $\mathcal{M}_2$  which implies

$$\mathbb{P}[\mathcal{M}_2(o_1, Y) \in S] \leq (e^{\varepsilon_2} \mathbb{P}[\mathcal{M}_2(o_1, Y') \in S] \wedge 1) + \delta_2. \quad (\text{A.2})$$

For ease of notation, let  $p_Y : \mathcal{O}_1 \rightarrow \mathbb{R}^+$  denote the measure obtained by restricting the probability density function of  $\mathcal{M}_1(Y)$  to  $\mathcal{O}_1$  (note that  $\int_{\mathcal{O}_1} p_Y(o_1) do_1 = 1 - \mathbb{P}[\mathcal{M}_1(Y) = \perp]$ ). Then, observe that

$$\begin{aligned} & \mathbb{P}[\mathcal{M}(Y) \in S] \\ &= \mathbf{1}[\perp \in S] \mathbb{P}[\mathcal{M}_1(Y) = \perp] + \int_{\mathcal{O}_1} \mathbb{P}[\mathcal{M}_2(o_1, Y) \in S] p_Y(o_1) do_1 \\ &\stackrel{(\text{A.2})}{\leq} \mathbf{1}[\perp \in S] \mathbb{P}[\mathcal{M}_1(Y) = \perp] + \int_{\mathcal{O}_1} ((e^{\varepsilon_2} \mathbb{P}[\mathcal{M}_2(o_1, Y') \in S] \wedge 1) + \delta_2) p_Y(o_1) do_1 \\ &\leq \mathbf{1}[\perp \in S] \mathbb{P}[\mathcal{M}_1(Y) = \perp] + \delta_2 \\ &\quad + \int_{\mathcal{O}_1} (e^{\varepsilon_2} \mathbb{P}[\mathcal{M}_2(o_1, Y') \in S] \wedge 1) p_Y(o_1) do_1 \\ &\leq \mathbf{1}[\perp \in S] (e^{\varepsilon_1} \mathbb{P}[\mathcal{M}_1(Y') = \perp] + \mu(\{\perp\})) + \delta_2 \\ &\quad + \int_{\mathcal{O}_1} (e^{\varepsilon_2} \mathbb{P}[\mathcal{M}_2(o_1, Y') \in S] \wedge 1) (e^{\varepsilon_1} p_{Y'}(o_1) do_1 + d\mu(o_1)) \\ &\leq \mathbf{1}[\perp \in S] (e^{\varepsilon_1} \mathbb{P}[\mathcal{M}_1(Y') = \perp]) + \mu(\mathcal{O}_1 \cup \{\perp\}) + \delta_2 \\ &\quad + \int_{\mathcal{O}_1} (e^{\varepsilon_2} \mathbb{P}[\mathcal{M}_2(o_1, Y') \in S] \wedge 1) (e^{\varepsilon_1} p_{Y'}(o_1)) do_1 \\ &\leq \mathbf{1}[\perp \in S] (e^{\varepsilon_1 + \varepsilon_2} \mathbb{P}[\mathcal{M}_1(Y') = \perp]) + \delta_1 + \delta_2 \\ &\quad + \int_{\mathcal{O}_1} e^{\varepsilon_1 + \varepsilon_2} \mathbb{P}[\mathcal{M}_2(o_1, Y') \in S] p_{Y'}(o_1) do_1 \\ &\leq \delta_1 + \delta_2 + e^{\varepsilon_1 + \varepsilon_2} \mathbb{P}[\mathcal{M}(Y') \in S]. \end{aligned}$$

- **$Y$  satisfies  $\Psi_1$  but  $Y'$  does not.** In this case, we have  $\mathbb{P}[\mathcal{M}(Y') = \perp] = 1$ , which implies that

$$\begin{aligned} \mathbb{P}[\mathcal{M}(Y) \in S] - e^{\varepsilon_1 + \varepsilon_2} \mathbb{P}[\mathcal{M}(Y') \in S] &\leq \mathbb{P}[\mathcal{M}(Y) \neq \perp] \\ &= \mathbb{P}[\mathcal{M}(Y) \neq \perp] - e^{\varepsilon_1} \mathbb{P}[\mathcal{M}(Y') \neq \perp] \leq \delta_1, \end{aligned}$$

where the last inequality follows from the fact that  $\mathcal{M}_1$  is  $(\varepsilon_1, \delta_1)$ -DP.

- **$Y'$  satisfies  $\Psi_1$  but  $Y$  does not.** In this case, we have  $\mathbb{P}[\mathcal{M}(Y) = \perp] = 1$ , which implies that

$$\begin{aligned} \mathbb{P}[\mathcal{M}(Y) \in S] - e^{\varepsilon_1 + \varepsilon_2} \mathbb{P}[\mathcal{M}(Y') \in S] &\leq [\mathbb{P}[\mathcal{M}(Y) = \perp] - e^{\varepsilon_1 + \varepsilon_2} \mathbb{P}[\mathcal{M}(Y') = \perp]]_+ \\ &\leq [\mathbb{P}[\mathcal{M}(Y) = \perp] - e^{\varepsilon_1} \mathbb{P}[\mathcal{M}(Y') = \perp]]_+ \\ &\leq \delta_1, \end{aligned}$$

where the last inequality once again follows from the fact that  $\mathcal{M}_1$  is  $(\varepsilon_1, \delta_1)$ -DP.

- **Neither  $Y$  nor  $Y'$  satisfy  $\Psi_1$ .** In this case, both  $\mathcal{M}(Y)$  and  $\mathcal{M}(Y')$  always output  $\perp$ . Therefore, we have  $\mathbb{P}[\mathcal{M}(Y) \in S] = \mathbb{P}[\mathcal{M}(Y') \in S]$ .

Thus, in all cases, we have  $\mathbb{P}[\mathcal{M}(Y) \in S] = e^{\varepsilon_1 + \varepsilon_2} \mathbb{P}[\mathcal{M}(Y') \in S] + \delta_1 + \delta_2$  as desired.  $\blacksquare$

#### A.6.4. APPROXIMATE-DP SELECTION

Finally, we will also use a DP algorithm for the *selection* problem, where the goal is to pick from a (public) set of candidates one which has a high “score”. This problem can be solved using the *exponential mechanism* (McSherry and Talwar, 2007). The version of the algorithm we use deviates slightly from this traditional version in that we also include a check (via the truncated Laplace mechanism of Lemma 35) to make sure that the score is at least a certain threshold  $\kappa$ ; otherwise, the algorithm’s properties are summarized below.

**Theorem 38** *Suppose  $\varepsilon, \delta \in (0, 1]$ . Let  $\mathcal{C}$  be a set of candidates and let  $\text{score} : \mathcal{C} \times \mathcal{Y}$  be a scoring function for candidates as a function of the databases  $Y \in \mathcal{Y}$ , such that its sensitivity (w.r.t.  $Y$ ) is at most  $\Delta$ . There exists an algorithm *Selection* that satisfies the following properties:*

1. *Selection is  $(\varepsilon, \delta)$ -DP.*
2. *If the output of Selection is  $c^* \neq \perp$ , then  $\text{score}(c^*, Y) \geq \kappa$ .*
3. *If there exists  $c \in \mathcal{C}$  such that  $\text{score}(c, Y) \geq \kappa + O\left(\frac{\Delta}{\varepsilon} \cdot \log\left(\frac{|\mathcal{C}|}{\beta\delta}\right)\right)$ , then Selection output  $\perp$  with probability at most  $\beta$ .*

**Proof** Selection works as follows:

1. First, run the  $(\varepsilon/2)$ -DP exponential mechanism (McSherry and Talwar, 2007), i.e. selecting each  $c \in \mathcal{C}$  with probability proportional to  $\exp\left(\frac{\varepsilon}{4\Delta} \cdot \text{score}(c, Y)\right)$ . Let  $c_1$  be the output of this procedure.
2. Sample the noise  $N \sim \text{tLap}\left(-\Delta\left(1 + \frac{2\ln(1/\delta)}{\varepsilon}\right), \frac{2\Delta}{\varepsilon}\right)$  and compute  $\widetilde{\text{score}} = \text{score}(c_1, Y) + N$ . If  $\widetilde{\text{score}} \geq \kappa$ , then output  $c_1$ . Otherwise, output  $\perp$ .

We will now prove each of the claimed properties:

1. The first step satisfies  $(\varepsilon/2)$ -DP via the standard privacy guarantee of the exponential mechanism (McSherry and Talwar, 2007). The second step is  $(\varepsilon/2, \delta)$ -DP due to Lemma 35. Thereby, applying the basic composition theorem implies that Selection is  $(\varepsilon, \delta)$ -DP.
2. Since  $N \leq 0$ , we are guaranteed that if the algorithm outputs  $c^* \in \mathcal{C}$ , we must have  $\text{score}(c, Y) \geq \kappa$  as desired.
3. For any  $c \in \mathcal{C}$ , the standard utility analysis of the exponential mechanism (McSherry and Talwar, 2007) implies that, with probability  $1 - 0.5\beta$ , we have  $\text{score}(c_1, Y) \geq \text{score}(c, Y) - O\left(\frac{\Delta}{\varepsilon} \ln\left(\frac{|C|}{\beta}\right)\right)$ . Moreover, the tail bound of Laplace noise (Lemma 34) implies that with probability  $1 - 0.5\beta$  we have  $N \geq -\Delta\left(1 + \frac{\ln(1/\delta)}{\varepsilon}\right) - O\left(\frac{\Delta}{\varepsilon} \ln(1/\beta)\right) \geq -O\left(\frac{\Delta}{\varepsilon} \ln\left(\frac{1}{\delta\beta}\right)\right)$ . Therefore, if  $\text{score}(c, Y) \geq \kappa + O\left(\frac{\Delta}{\varepsilon} \ln\left(\frac{|C|}{\delta\beta}\right)\right)$ , the probability that the algorithm outputs  $\perp$  is at most  $\beta$ , as desired. ■

## Appendix B. Differentially Private Robust Moment Estimation: Detailed Analysis and Proofs

In this section, we describe a differentially private robust moment estimation algorithm. The following is our main technical result:

**Theorem 39 (Differentially Private Robust Moment Estimation)** *Fix  $C_0 > 0$  and  $k \in \mathbb{N}$ . Then, there exists an  $\eta_0 > 0$  such that for any given outlier rate  $0 < \eta \leq \eta_0$  and  $\varepsilon, \delta > 0$ , there exists a randomized algorithm **Alg** that takes an input of  $n \geq n_0 = \tilde{\Omega}\left(\frac{d^{4k}}{\eta^2} \left(1 + \left(\frac{\ln(1/\delta)}{\varepsilon}\right)^4 + \left(\frac{\ln(1/\delta)}{\varepsilon}\right)^{\frac{2k}{k-1}}\right) \cdot C^{4k} k^{4k+6}\right)$  points  $Y \subseteq \mathbb{Q}^d$  (where  $C = C_0 + \frac{3\ln(3/\delta)}{\varepsilon} + \frac{9}{\varepsilon} + 1$ ), runs in time  $(Bn)^{O(k)}$  (where  $B$  is the bit complexity of the entries of  $Y$ ) and outputs either “reject” or estimates  $\hat{\mu} \in \mathbb{Q}^d$ ,  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$ , and  $\hat{M}^{(t)} \in \mathbb{Q}^{d \times d \times \dots \times d}$  (for all even  $t < 2k$  such that  $t$  divides  $2k$ ) with the following guarantees<sup>2</sup>:*

1. **Privacy:** *Alg* is  $(\varepsilon, \delta)$ -differentially private with respect to the input  $Y$ , viewed as a  $d$ -dimensional database of  $n$  individuals.
2. **Utility:** *Suppose there exists a  $2k$ -certifiably  $C_0$ -subgaussian set  $X \subseteq \mathbb{Q}^d$  of  $n \geq n_0$  points such that  $|Y \cap X| \geq (1 - \eta)n$  with mean  $\mu_*$ , covariance  $\Sigma_* \succeq 2^{-\text{poly}(d)}I$ , and  $t$ -th moments  $M_*^{(t)}$  for  $2 \leq t \leq k$ . Then, with probability at least  $9/10$  over the random choices of the algorithm, **Alg** outputs estimates  $\hat{\mu} \in \mathbb{Q}^d$ ,  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$ , and  $M^{(t)} \in \mathbb{Q}^{d \times d \times \dots \times d}$  (for all even  $t < 2k$  such that  $t$  divides  $2k$ ) satisfying the following guarantees:*

$$\forall u \in \mathbb{R}^d, \langle \hat{\mu} - \mu_*, u \rangle \leq O\left(\sqrt{Ck}\right) \eta^{1-1/2k} \sqrt{u^\top \Sigma_* u},$$

2. The  $\tilde{\Omega}$  notation hides multiplicative logarithmic factors in  $d, C, k, 1/\eta, 1/\varepsilon$ , and  $\ln(1/\delta)$ .

and,

$$\left(1 - O((Ck)^{t/2k})\eta^{1-1/k}\right) \Sigma_* \preceq \hat{\Sigma} \preceq \left(1 + O((Ck)^{t/2k})\eta^{1-1/k}\right) \Sigma_*,$$

and, for every even  $t < 2k$  such that  $t$  divides  $2k$ ,

$$\left(1 - O(Ck)\eta^{1-t/2k}\right) \langle u^{\otimes t}, M_*^{(t)} \rangle \leq \langle u^{\otimes t}, \hat{M}^{(t)} \rangle \leq \left(1 + O(Ck)\eta^{1-t/2k}\right) \langle u^{\otimes t}, M_*^{(t)} \rangle.$$

Moreover, the algorithm succeeds (i.e., does not reject) with probability at least  $9/10$  over the random choices of the algorithm.

Observe that the privacy guarantees of the algorithm are (necessarily) *worst-case*. The utility guarantees, however, hold only under the assumption that  $Y$  is an  $\eta$ -corruption of a good set  $X$ .

The above theorem can also be translated into utility guarantees for points sampled from a given distribution by recalling the well-known fact that points sampled from a certifiably subgaussian distribution are good with high probability:

**Fact 40 (See Section 5 in Kothari and Steurer (2017))** *Suppose  $\mathcal{D}$  is a certifiably  $C$ -subgaussian distribution with mean  $\mu_*$  and covariance  $\Sigma_* \succeq 2^{-\text{poly}(d)}I$  and  $t$ -moment tensors  $M^{(t)}$  for  $t \in \mathbb{N}$ . For any  $k \in \mathbb{N}$ , let  $X = \{x_1, x_2, \dots, x_n\}$  be an i.i.d. sample from  $\mathcal{D}$  of size  $n \geq n_0 = O(d^{2k}/\eta^2)$ . Then, for any  $t \in \mathbb{N}$  such that  $t$  divides  $k$ , with probability at least  $0.99$  over the draw of  $X$ , the following all hold:*

1.  $X$  is  $2k$ -certifiably  $2C$ -subgaussian.
2.  $\left\| \Sigma_*^{-1/2}(\mu(X) - \mu_*) \right\|_2 \leq \eta$ .
3.  $\Sigma(X) \in (1 \pm \eta)\Sigma_*$ .
4.  $\left| \frac{v}{2k} \left\{ \langle v^{\otimes t}, M^{(t)}(X) \rangle \in (1 \pm \eta) \langle v^{\otimes t}, M_*^{(t)} \rangle \right\} \right|$ .

We note that our main theorem for private robust moment estimation, [Theorem 2](#), is an immediate consequence of [Theorem 39](#) and [Fact 40](#).

For the rest of the section, we will work to prove [Theorem 39](#). In [Appendix B.1](#), we will introduce a witness-producing robust moment estimation algorithm that will be used as a subroutine for our main algorithm and present relevant utility guarantees. In [Appendix B.2](#), we will then introduce our main algorithm. After that, we will prove the necessary privacy guarantees in [Appendix B.3](#). Finally, we will put together the pieces to prove our main theorem, [Theorem 39](#), in [Appendix B.4](#).

### B.1. Witness-Producing Version of Robust Moment Estimation Algorithm

As a key building block, we will use the following (non-private) version of the robust moment estimation algorithm of [Kothari and Steurer \(2017\)](#) that uses the same constraint system  $\mathcal{A}$  as in [Kothari and Steurer \(2017\)](#). Our algorithm itself, however, makes one key change (we call our version “witness-producing” for reasons that will soon become clear) to that of [Kothari and Steurer \(2017\)](#) in order to obtain a private robust moment estimation algorithm. Instead of outputting estimates of the moments of the unknown distribution, our algorithm outputs a sequence of non-negative weights  $p_1, p_2, \dots, p_n$  forming a probability distribution on the input set of points  $Y$ . The estimates can then



be obtained by taking moments of the finite set  $Y$  with respect to the probability distribution on  $Y$  defined by the weights  $p_i$ s. This simple change is crucial to our *worst-case* analysis of the resulting algorithm (i.e. even when the distributional assumption that  $Y$  is an  $\eta$ -corruption of some good set  $X$  is not met) and obtaining our privacy guarantees. As we discuss, our blueprint for modifying convex optimization based robust estimation algorithms appears to be broadly applicable beyond the specific setting of robust moment estimation.

The underlying constraint system  $\mathcal{A}$  is shown below, and the witness-producing robust moment estimation algorithm is shown as Algorithm 1.

**Constraint 1 (Constraint System for  $\eta$ -Robust Moment Estimation)**

$\mathcal{A}_{C,k,\eta,n}(\{y_1, y_2, \dots, y_n\})$ :

1.  $w_i^2 = w_i$  for each  $1 \leq i \leq n$ ,
2.  $\sum_{i=1}^n w_i \geq (1 - \eta)n$ ,
3.  $\mu' = \frac{1}{n} \sum_i x'_i$ ,
4.  $w_i(x'_i - y_i) = 0$  for  $1 \leq i \leq n$ ,
5.  $\frac{1}{n} \sum_{i=1}^n \langle x'_i - \mu', v \rangle^k \leq (Ck)^{k/2} \left( \frac{1}{n} \sum_{i=1}^n \langle x'_i - \mu', v \rangle^2 \right)^{k/2}$ .

**Algorithm 1 (Witness-Producing Robust Moment Estimation)**

**Given:** A set of points  $Y = \{y_1, y_2, \dots, y_n\} \subseteq \mathbb{Q}^d$ ,  $\eta > 0$ , a parameter  $k \in \mathbb{N}$ .

**Output:** Either “reject” or non-negative weights  $p_1, p_2, \dots, p_n$  s.t.  $p_i \leq \frac{1}{(1-\eta)n} \forall i$  and  $\sum_i p_i = 1$ .

**Operation:**

1. Find a pseudo-distribution  $\tilde{\zeta}$  of degree  $O(k)$  satisfying the constraint system  $\mathcal{A}_{C,k,\eta,n}(Y)$ . If such a pseudo-distribution does not exist, then return “reject.”
2. Output weights  $p \in [0, 1]^n$  defined by  $p_i = \frac{\tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i]}{\sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i]}$  for each  $i$ .

**Analysis of the witness-producing robust estimation algorithm** Robust estimation algorithms that rely on the use of semidefinite programming are all analyzed under distributional assumptions on the input set of points. Roughly speaking, such algorithms search over set of points that have a large enough intersection with the input corrupted sample and satisfy certain relevant property of the underlying family of distributions. In order to obtain privacy guarantee that holds for worst-case inputs, we need to upgrade the analyses of such algorithms so that they not only provide estimates of the target parameters, but also explicitly produce “witnesses”—these are subsets of the input corrupted sample that define distributions with the estimated parameters and further, satisfy the relevant property of the underlying family of distributions.

In this section, we verify that such a stronger guarantee can be obtained for robust moment estimation algorithm of [Kothari and Steurer \(2017\)](#). Formally, their algorithm succeeds as long as the input is an  $\eta$ -corruption of a certifiably subgaussian set.

The following guarantees for the algorithm above were shown in [Kothari and Steurer \(2017\)](#).

**Fact 41 (Lemmas 4.4, 4.5, and 4.8 in [Kothari and Steurer \(2017\)](#))** *Let  $X \subseteq \mathbb{R}^d$  be a set of size  $n$  that is  $2k$ -certifiably  $C$ -subgaussian with mean  $\mu_*$ , covariance  $\Sigma_*$  and  $t$ -th moment  $M_*^{(t)}$  for  $t$  evenly dividing  $2k$ . Let  $Y$  be an  $\eta$ -corruption of  $X$  for  $\eta$ , a small enough absolute constant. Then, for  $\mu' = \frac{1}{n} \sum_i x'_i$ ,  $\Sigma' = \frac{1}{n} \sum_i (x_i - \mu')(x_i - \mu')^\top$ , and  $M^{(t)'} = \frac{1}{n} \sum_i x_i'^{\otimes t}$ , we have:*

$$\begin{aligned} \mathcal{A} \Big|_{\frac{u}{2k}} \left\{ \langle \mu' - \mu_*, u \rangle^{2k} \leq O(\eta) C^k k^k u^\top \Sigma_* u^k \right\}, \\ \mathcal{A} \Big|_{\frac{u}{2k}} \left\{ \langle \Sigma' - \Sigma_*, u^{\otimes 2} \rangle^k \leq O(\eta) C^k k^k u^\top \Sigma_* u^k \right\}, \\ \mathcal{A} \Big|_{\frac{u}{2k}} \left\{ \langle M^{(t)'} - M_*^{(t)}, u^{\otimes t} \rangle^{2k/t} \leq O(\eta) C^k k^k u^\top \Sigma_* u^k \right\}. \end{aligned}$$

**Lemma 42 (Guarantees for Witness-Producing Robust Moment Estimation Algorithm)**

*Given a subset of  $n$  points  $Y \subseteq \mathbb{Q}^d$  whose entries have bit complexity  $B$ , Algorithm 1 runs in time  $(Bn)^{O(k)}$  and either (a.) outputs “reject,” or (b.) returns a sequence of weights  $0 \leq p_1, p_2, \dots, p_n$  satisfying  $p_1 + p_2 + \dots + p_n = 1$ .*

*Moreover, if  $X \subseteq \mathbb{R}^d$  is  $2k$ -certifiably  $C$ -subgaussian with mean  $\mu_*$ , covariance  $\Sigma_*$  and in general,  $t$ -th moment tensor  $M^{(t)*}$  such that  $|Y \cap X| \geq (1 - \eta)n$ , then [Algorithm 1](#) never rejects, and the corresponding estimates  $\hat{\mu} = \frac{1}{n} \sum_i p_i y_i$  and  $\hat{\Sigma} = \sum_{i=1}^n p_i (y_i - \hat{\mu})(y_i - \hat{\mu})^\top$  satisfy the following guarantees for  $\beta_t = O(C^{t/2} k^{t/2}) \eta^{1-t/2k}$  for  $t \leq k$ :*

1. **Mean Estimation:**

$$\forall u \in \mathbb{R}^d, \langle \hat{\mu} - \mu_*, u \rangle \leq O(\sqrt{Ck}) \eta^{1-1/2k} \sqrt{u^\top \Sigma_* u},$$

2. **Covariance Estimation:**

$$(1 - \beta_2) \Sigma_* \preceq \hat{\Sigma} \preceq (1 + \beta_2) \Sigma_*,$$

3. **Moment Estimation:** For all even  $t < 2k$  such that  $t$  divides  $2k$ ,

$$\forall u \in \mathbb{R}^d, (1 - \beta_t) \langle u^{\otimes t}, M_*^{(t)} \rangle \leq \langle u^{\otimes t}, \hat{M}^{(t)} \rangle \leq (1 + \beta_t) \langle u^{\otimes t}, M_*^{(t)} \rangle$$

4. **Witness:** For  $C' \leq C(1 + O(\eta^{1-1/k}))$ ,

$$\vdash \left\{ \frac{1}{n} \sum_{i=1}^n p_i \langle y_i - \hat{\mu} \rangle^{2k} \leq (C'k)^k \left( \frac{1}{n} \sum_{i=1}^n p_i \langle y_i - \hat{\mu} \rangle^2 \right)^k \right\}$$

The first three properties follow easily from an analysis similar to the one in [Kothari and Steurer \(2017\)](#). We verify the last property below.

**Lemma 43** *Let  $\tilde{\zeta}$  be a pseudo-distribution of degree  $O(k)$  consistent with  $\mathcal{A}$  on input  $Y$  with outlier rate  $\eta \ll 1/k$ . Suppose there exists a  $2k$ -certifiably  $C_1$ -subgaussian distribution  $X \subseteq \mathbb{R}^d$  with mean  $\mu_*$  of size  $n$  such that  $|Y \cap X| \geq (1 - \eta)n$ . Then, for  $\eta \leq \eta_0$  for some absolute constant  $\eta_0$  and for  $\hat{\mu} = \frac{1}{W} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i] y_i$  where  $W = \sum_{i=1}^n \tilde{\mathbb{E}}[w_i]$ , we have:*

$$\Big|_{\frac{u}{2k}} \left\{ \frac{1}{W} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i] \langle y_i - \hat{\mu}, u \rangle^{2k} \leq (C'k)^k \left( \frac{1}{W} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i] \langle y_i - \hat{\mu}, u \rangle^2 \right)^k \right\},$$

for  $C' \leq C(1 + O(\eta^{1-1/2k})k) \leq C + 1$  for small enough  $\eta$ .

**Proof** We have:

$$\left( \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i \langle y_i - \hat{\mu}, u \rangle] \right)^{2k} \leq \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i \langle x'_i - \hat{\mu}, u \rangle^{2k}] \leq \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu' + \mu' - \hat{\mu}, u \rangle^{2k}]$$

The first term on the right-hand side above is at most  $(Ck)^k \tilde{\mathbb{E}}_{\tilde{\zeta}}[(\frac{1}{n} \sum_{i=1}^n \langle x'_i - \mu', u \rangle^2)^k] \leq (C(1 + O(\eta^{1-1/2k})k)^k u^\top \Sigma_* u^k$  using certifiable subgaussianity constraints and Fact 41.

Let us analyze the 2nd term above.

$$\frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu' + \mu' - \hat{\mu}, u \rangle^{2k}] \tag{B.1}$$

$$= \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k}] + 2k \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k-2} \langle \mu' - \hat{\mu}, u \rangle^2] \tag{B.2}$$

$$+ \sum_{j=2}^{2k} \binom{2k}{j} \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k-j} \langle \mu' - \hat{\mu}, u \rangle^j] \tag{B.3}$$

$$\leq \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k}] + 2k \frac{1}{n} \sum_{i=1}^n (\tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k}])^{(2k-2)/2k} (\tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle \mu' - \hat{\mu}, u \rangle^{2k}])^{1/2k} \tag{B.4}$$

$$+ \sum_{j=2}^{2k} \binom{2k}{j} \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k-j} \langle \mu' - \hat{\mu}, u \rangle^j] \tag{B.5}$$

Here, in the 2nd inequality, we used the Hölder's inequality for pseudo-distributions. Let us analyze the 2nd term in the right-hand side above by observing the following that uses the bounds from Fact 41:

$$\tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle \mu' - \hat{\mu}, u \rangle^{2k}] \leq 2^{2k} [\tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle \mu' - \mu_*, u \rangle^{2k}] + 2^{2k} \langle \mu_* - \hat{\mu}, u \rangle^{2k}] \tag{B.6}$$

$$\leq 2^{2k} (Ck)^k \eta^{2k-1} u^\top \Sigma_* u^k + 2^{2k} (Ck)^k \eta^{2k-1} (1 + \beta_2)^k u^\top \Sigma_* u^k \tag{B.7}$$

This allows us to infer that the 2nd term in (B.5) is at most  $\tilde{\mathbb{E}}_{\tilde{\zeta}}[\frac{1}{n} \sum_{i=1}^n \langle x'_i - \mu', u \rangle^{2k}]^{(2k-2)/2k} \cdot (5Ck)^{1/2} \eta^{1-1/2k} \sqrt{u^\top \Sigma_* u} \leq O(k)(Ck)^k \eta^{1-1/2k} u^\top \Sigma_* u^k$  using certifiable subgaussianity constraints and Fact 41.

Let's now analyze the terms corresponding to  $j \geq 2$  in the right-hand side of (B.5). Each of these terms corresponds to a "mixed monomial" in  $\langle x'_i - \mu', u \rangle$  and  $\langle \mu' - \mu, u \rangle$ . Let us first analyze the even individual degree terms.

First observe that by Hölder's inequality for pseudo-distributions again, we have:

$$\frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k-2} \langle \mu' - \hat{\mu}, u \rangle^2] \leq \langle x'_i - \mu', u \rangle^{2k} \binom{2k-2}{k} (\tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle \mu' - \hat{\mu}, u \rangle^{2k}])^{1/k}. \tag{B.8}$$

By an analysis similar to the case of the first term on the right-hand side of (B.5) above, we obtain that the right-hand side is at most:  $O(1)(Ck)^k (\eta^{1-1/2k})^2 u^\top \Sigma_* u^k$ .

Next, let's analyze all terms corresponding to even  $j$ . By Proposition 13, we have:

$$\begin{aligned}
 \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle x'_i - \mu', u \rangle^{2k-2j} \langle \mu' - \hat{\mu}, u \rangle^{2j}] &\leq \frac{1}{n} \sum_{i=1}^n (\tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle \mu' - \hat{\mu}, u \rangle^2 \langle x'_i - \mu', u \rangle^{2k-2j} \langle \mu' - \hat{\mu}, u \rangle^{2j-2}]) \\
 &\leq \frac{2k}{n} \sum_{i=1}^n (\tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle \mu' - \hat{\mu}, u \rangle^2 (\langle x'_i - \mu', u \rangle^{2k-2} + \langle \mu' - \hat{\mu}, u \rangle^{2k-2})])
 \end{aligned}$$

The first term can now be upper bounded by the bound for (B.8) and the 2nd term by an application of Fact 41.

The case of odd terms is similar with the first step using Proposition 14.

Altogether, we obtain an upper bound of  $(C(1 + O(\eta^{1-1/2k})k)^k u^\top \Sigma_* u)^k$ .

On the other hand, using the sum-of-squares version of the Cauchy-Schwarz inequality along with the almost triangle inequality and invoking Fact 41 we have:

$$\begin{aligned}
 \mathcal{A} \Big|_{\frac{u}{2k}} \left\{ \left( \frac{1}{n} \sum_{i=1}^n (1 - w_i) \langle x'_i - \hat{\mu}, u \rangle^2 \right)^2 \leq \left( \frac{1}{n} \sum_{i=1}^n (1 - w_i)^2 \right) \frac{1}{n} \sum_{i=1}^n \langle x'_i - \hat{\mu}, u \rangle^4 \right. \\
 \left. \leq 16\eta C^2 \left( \frac{1}{n} \sum_{i=1}^n \langle x'_i - \mu', u \rangle^2 + \frac{1}{n} \sum_{i=1}^n \langle \mu' - \hat{\mu}, u \rangle^2 \right)^2 \leq 20\eta C^2 (1 + \beta_2)^2 u^\top \Sigma_* u^2 \right\}
 \end{aligned}$$

Thus,

$$\begin{aligned}
 \mathcal{A} \Big|_{\frac{u}{2k}} \left\{ \left( \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i \langle y_i - \hat{\mu}, u \rangle^2] \right)^2 \right. \\
 \geq \left( \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[\langle y_i - \hat{\mu}, u \rangle^2] \right)^2 - \left( \frac{1}{n} \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[(1 - w_i) \langle y_i - \hat{\mu}, u \rangle^2] \right)^2 n \\
 \left. \geq (1 - O(Ck)\eta^{1-1/k} - 80\eta C^2) u^\top \Sigma_* u. \right\}
 \end{aligned}$$

The lemma now follows immediately for small enough fixed constant  $\eta$ . ■

## B.2. Private Robust Moment Estimation

We are now ready to present our main algorithm for private robust moment estimation. Our algorithm uses the witness-producing algorithm (Algorithm 1) as a major building block while augmenting it to search for pseudo-distributions that, in addition to satisfying the relevant set of constraints, also minimize an appropriate strongly convex potential function. We define the relevant potential function Pot below in Theorem 44.

**Definition 44 (Potential Function)** *Let  $C > 0$  and  $n, k \in \mathbb{N}$ . For any pseudo-distribution  $\tilde{\zeta}$  of degree 2 consistent with  $\mathcal{A}_{C,k,\eta,n}(Y)$  for outlier rate  $\eta$  and input  $Y \subseteq \mathbb{R}^d$ , let  $\text{Pot}_{\eta, \tilde{\zeta}}^{C,k,n}(Y)$  be defined as  $\left\| \tilde{\mathbb{E}}_{\tilde{\zeta}}[w] \right\|_2^2$ . Furthermore, let  $\text{Pot}_{\eta}^{C,k,n}(Y) = \min_{\tilde{\zeta} \text{ sat } \mathcal{A}_{C,k,\eta,n}(Y)} \text{Pot}_{\eta, \tilde{\zeta}}^{C,k,n}(Y)$  be the*

minimum value of the potential as  $\tilde{\zeta}$  ranges over all pseudo-distributions of degree  $2t$  consistent with  $\mathcal{A}_{C,k,\eta,n}(Y)$ . If no such pseudo-distribution exists, set  $\text{Pot}_\eta(Y) = \infty$ .

When  $C, n, k$  are understood from context, we may suppress these parameters and simply write  $\text{Pot}_\eta$  and  $\text{Pot}_{\eta,\tilde{\zeta}}$ .

Now, we are ready to describe our main private robust moment algorithm, which is listed as Algorithm 2. The algorithm consists of three main steps. In the first step, the randomized DP selection algorithm (Theorem 38) is used to pick an outlier rate (according to a suitable scoring function, as defined below in Theorem 48). The second step invokes the witness-producing algorithm (Algorithm 1) with the outlier rate chosen in step 1, after which one checks that the outputted weights induce a certifiably subgaussian distribution on the input dataset  $Y$ . Finally, in the last step, one takes the estimates of the mean, covariance, and higher moments provided by the resulting weight vector and adds suitable noise to guarantee differential privacy.

### Algorithm 2 (Private Robust Moment Estimation)

**Given:** A set of points  $Y = \{y_1, y_2, \dots, y_n\} \subseteq \mathbb{Q}^d$ , parameters  $C, \eta, \varepsilon, \delta > 0$ ,  $L, k \in \mathbb{N}$ .

**Output:** Estimates  $\hat{\mu}$ ,  $\hat{\Sigma}$ , and  $\hat{M}^{(t)}$  ( $3 \leq t \leq k$ ) for mean, covariance, and  $t$ -moments.

#### Operation:

1. **Stable Outlier Rate Selection:** Use the  $(\varepsilon/3, \delta/3)$ -DP Selection with  $\kappa = L/2$  to sample an integer  $\tau \in [\eta n]$  with the scoring function as defined in Theorem 48. If  $\tau = \perp$ , then reject and halt. Otherwise, let  $\eta' = \tau/n$ .
2. **Witness Checking:** Compute a pseudo-distribution  $\tilde{\zeta}$  of degree  $2k$  satisfying  $\mathcal{A}_{C,k,\eta',n}(Y)$  and minimizing  $\text{Pot}_{\eta',\tilde{\zeta}}(Y)$ . Let  $\gamma \sim \text{tLap}\left(-\left(1 + \frac{3\ln(3/\delta)}{\varepsilon}\right), 3/\varepsilon\right)$  and  $C' = C + \gamma$ . Check that the weight vector  $p = \tilde{\mathbb{E}}_{\tilde{\zeta}}[w]$  induces a  $C'$ -certifiably subgaussian distribution on  $Y$ . If not, reject immediately. Otherwise, let  $\tilde{\mu} = \tilde{\mathbb{E}}_{\tilde{\zeta}}[\mu]$ ,  $\tilde{\Sigma} = \tilde{\mathbb{E}}_{\tilde{\zeta}}[\Sigma]$ , and  $\tilde{M}^{(t)} = \tilde{\mathbb{E}}_{\tilde{\zeta}}[M^{(t)}]$  (for all even  $t < 2k$  such that  $t$  divides  $2k$ ) be the mean, covariance, and  $t^{\text{th}}$  moment estimates, respectively, that are induced by the pseudo-distribution  $\tilde{\zeta}$ .
3. **Noise Addition:** Let  $\gamma_1 = O(C'k)(L/n)^{\frac{1}{2}(1-\frac{1}{2k})}$  and  $\gamma_t = O((C'k)^{t/2})(L/n)^{\frac{1}{2}(1-\frac{t}{2k})}$  for  $t \geq 2$ . Let  $z \sim \mathcal{N}(0, \sigma_1)^d$  and  $Z \sim \mathcal{N}(0, \sigma_2)^{\binom{d+1}{2}}$ , where we interpret  $Z$  as a symmetric  $d \times d$  matrix with i.i.d. entries in the upper triangular portion. Similarly, for  $t \geq 2$ , let  $Z^{(t)} \sim \mathcal{N}(0, \sigma_t)^{\binom{d+(t-1)}{t}}$ , where we interpret  $Z$  as a symmetric  $\underbrace{d \times d \times \dots \times d}_{t \text{ times}}$  tensor with  $\binom{d+(t-1)}{t}$  independent “upper-triangular” entries. Moreover, let

$$\begin{cases} \sigma_j = 6k\varepsilon^{-1}\gamma_j d^{\frac{t-1}{2}} \sqrt{2\ln(7.5k/\delta)}, & \text{for } j = 1, 2 \\ \sigma_j = 6k\varepsilon^{-1}\gamma_j (C'k)^t d^{\frac{t-1}{2}} \sqrt{2\ln(7.5k/\delta)}, & \text{for } j > 2 \end{cases}.$$

Then, output:

- $\hat{\mu} = \tilde{\mu} + \tilde{\Sigma}^{1/2}z.$
- $\hat{\Sigma} = \tilde{\Sigma} + \tilde{\Sigma}^{1/2}Z\tilde{\Sigma}^{1/2}.$
- $\hat{M}^{(t)} = \tilde{M}^{(t)} + ((\tilde{\Sigma} + \tilde{\mu}\tilde{\mu}^T)^{1/2})^{\otimes t}Z^{(t)},$  for all even  $t < 2k$  such that  $t$  divides  $2k.$

### B.3. Privacy Analysis

Our analysis of the privacy of **Alg** is based on a sequence of claims about each of the steps of **Alg** that cumulatively establish the stability of the behavior of **Alg** on adjacent inputs  $Y, Y'$ . We will rely on the following simple but key observation in our analysis. It is easy to verify using the definition of pseudo-distributions.

**Lemma 45 (Adjacent Pseudo-distributions)** *Let  $\tilde{\zeta}$  be a pseudo-distribution of degree  $2k$  that satisfies all the constraints in 1 on input  $Y = \{y_1, y_2, \dots, y_n\}$  with outlier rate  $\eta$ . Let  $Y' \subseteq \mathbb{R}^d$  be adjacent to  $Y$ . Define an adjacent pseudo-distribution  $\tilde{\zeta}'$  (that “zeroes out  $w_i$ ”) by  $\tilde{\mathbb{E}}_{\tilde{\zeta}'}[w_{SP}(X', \dots)] = \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_{SP}(X', \dots)]$  if  $i \notin S$  and  $\tilde{\mathbb{E}}_{\tilde{\zeta}'}[w_{SP}(X', \dots)] = 0$  if  $i \in S$  for every multi-subset  $S$  of  $[n]$  of size at most  $2k$ , every polynomial  $p$  in  $X'$  and other auxiliary indeterminates in  $\mathcal{A}$ . Then,  $\tilde{\zeta}'$  is a pseudo-distribution of degree  $2k$  that satisfies all the constraints in 1 on both inputs  $Y'$  and  $Y$  with outlier parameter  $\eta + 1/n$ .*

This allows us to conclude the following basic calculus of our potential function:

**Lemma 46 (Basic Facts about Pot)** *Suppose that for some  $Y \subseteq \mathbb{R}^d$  of size  $n$ , some  $t \in \mathbb{N}$  and  $\eta' \in [0, \eta_0/4]$ , there is a pseudo-distribution of degree  $2t$  consistent with  $\mathcal{A}$  on input  $Y$ . Then, for every  $\eta \geq \eta'$ , the following holds:*

1. **Monotonicity:**  $\text{Pot}_{\eta+1/n}(Y) \leq \text{Pot}_{\eta}(Y)$ . In particular, **Pot** is monotonically decreasing as its subscript increases.
2. **Lower Bound:**  $\text{Pot}_{\eta}(Y) \geq (1 - \eta)^2 n$ .
3. **Upper Bound:**  $\text{Pot}_{\eta}(Y) \leq (1 - \eta)n$ .

**Proof** The first fact follows immediately from Lemma 45. For the second, observe that any pseudo-distribution  $\tilde{\zeta}$  of degree  $2t$  consistent with  $\mathcal{A}$  on input  $Y$  with outlier rate  $\eta$  must satisfy  $\sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i] \geq (1 - \eta)n$ . Thus, by Cauchy-Schwarz inequality,  $\sum_{i=1}^n \tilde{\mathbb{E}}[w_i]^2 \geq (\sum_{i=1}^n \tilde{\mathbb{E}}[w_i])^2/n = (1 - \eta)^2 n$ . This completes the proof. For the last part, observe that  $\tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i] \leq 1$  for every  $i$ . Thus,  $\sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i]^2 \leq \sum_{i=1}^n \tilde{\mathbb{E}}_{\tilde{\zeta}}[w_i] = (1 - \eta)n$ .  $\blacksquare$

**Analysis of stable outlier rate selection** The goal of the first step of **Alg** is to find an outlier rate  $\eta'$  such that the strongly convex potential function  $\text{Pot}(\tilde{\zeta})$  on the pseudo-distribution we will eventually compute (in Step 3) is close on adjacent input points  $Y, Y'$ . We will later use the strong convexity of the **Pot** and the closeness guarantee on **Pot** on  $Y, Y'$  to infer that the weight vector  $p(Y)$  and  $p(Y')$  output by the algorithm themselves are close.

Our key algorithmic trick to ensure the closeness of the strongly convex potential **Pot** is to find a “stable interval”  $[\eta' - 0.5L/n, \eta' + 0.5L/n]$  of outlier rates  $\eta''$  such that strongly convex potential function at near-optimal solutions must vary slowly as  $\eta''$  varies in the interval. We find such an interval via a variant of the exponential mechanism.

**Definition 47 (Stability)** Fix  $L \in \mathbb{N}$ . Let  $\tau, \gamma \in \{0, \dots, n\}$  such that  $\gamma \leq \tau, n - \tau$ . Suppose for some  $Y \subseteq \mathbb{R}^d$  of size  $n$ , the constraint system  $\mathcal{A}((\tau - \gamma)/n)$  is feasible. We define the stability of the  $2\gamma$  length interval centered at  $\tau$  to be

$$\text{stab}_Y(\tau, \gamma) = \text{Pot}_{(\tau-\gamma)/n}(Y) - \text{Pot}_{(\tau+\gamma)/n}(Y)$$

Observe that if there is a pseudo-distribution consistent with  $\mathcal{A}$  on  $Y$  with outlier rate  $(\tau - \gamma)/n$  then there is a pseudo-distribution consistent with  $\mathcal{A}$  on  $Y$  with any outlier rate  $\geq (\tau - \gamma)/n$ . Thus, stability above is well-defined.

**Definition 48 (Score Function)** Fix  $n, k \in \mathbb{N}$  and  $C > 0$ . Let  $Y \subseteq \mathbb{R}^d$  be a set of size  $n$ . For a parameter  $L$ , we define the following score function for every integer  $\tau \in [n]$ :

$$\text{score}_{n,C,k}(\tau, Y) = \begin{cases} 0 & \text{if } \text{Alg}(Y, \tau/n) \text{ is infeasible,} \\ \max_{\mathcal{A}_{C,k,(\tau-\gamma)/n,n}(Y) \text{ is feasible}}^{\gamma} \min\{\gamma, 20L - \text{stab}_Y(\tau, \gamma)\} & \text{otherwise.} \end{cases}$$

In the second case, we define  $\gamma_Y^*(\tau) := \arg \max_{\mathcal{A}_{C,k,(\tau-\gamma)/n,n}(Y) \text{ is feasible}}^{\gamma} \min\{\gamma, L - \text{stab}_Y(\tau, \gamma)\}$ .

**Lemma 49** Let  $\tau, \gamma \in [n]$  such that  $\gamma \leq \tau, n - \tau$ . Let  $Y, Y' \subseteq \mathbb{R}^d$  be both of size  $n$  such that they differ on at most one point. Suppose the constraint system  $\mathcal{A}((\tau - \gamma)/n)$  is feasible for both  $Y$  and  $Y'$ . Then, for any  $\tau, \gamma$ ,

$$\text{stab}_{Y'}(\tau, \gamma - 1) \leq \text{stab}_Y(\tau, \gamma)$$

**Proof** Using Lemma 45 and noting that if  $\tilde{\zeta}'$  is adjacent to  $\tilde{\zeta}$  then  $\|\tilde{\mathbb{E}}_{\tilde{\zeta}'}[w]\|_2^2 \leq \|\tilde{\mathbb{E}}_{\tilde{\zeta}}[w]\|_2^2$ , we have:

$$\text{Pot}_{(\tau-\gamma+1)/n}(Y') \leq \text{Pot}_{(\tau-\gamma)/n}(Y),$$

and

$$\text{Pot}_{(\tau+\gamma)/n}(Y) \leq \text{Pot}_{(\tau+\gamma-1)/n}(Y').$$

Combining the two equations yields

$$\begin{aligned} \text{stab}_{Y'}(\tau, \gamma - 1) &= \text{Pot}_{(\tau-\gamma+1)/n}(Y') - \text{Pot}_{(\tau+\gamma-1)/n}(Y') \\ &\leq \text{Pot}_{(\tau-\gamma)/n}(Y) - \text{Pot}_{(\tau+\gamma)/n}(Y) = \text{stab}_Y(\tau, \gamma). \end{aligned}$$

■

**Lemma 50 (Sensitivity of Score Function)** Let  $Y, Y'$  be set of  $n$  points in  $\mathbb{R}^d$  differing at most in one point, and  $\tau \in [n]$ . Then, for every  $\tau > 0$ ,

$$|\text{score}(\tau, Y) - \text{score}(\tau, Y')| \leq 2. \quad (\text{B.9})$$

**Proof** It suffices to prove that  $\text{score}(\tau, Y') \geq \text{score}(\tau, Y) - 2$ . A symmetric argument then proves that  $\text{score}(\tau, Y) \geq \text{score}(\tau, Y') - 2$ , which establishes (B.9).

Consider the following two cases:

- $\text{Alg}(Y, (\tau - 1)/n)$  is infeasible for  $Y$  or  $Y'$ . In this case, we have  $\text{score}(\tau, Y) \leq 2$ , which implies the desired bound.
- $\text{Alg}(Y, (\tau - 1)/n)$  is feasible for both  $Y$  and  $Y'$ .

Let  $\gamma^* := \gamma_Y^*(\tau)$ . From Lemma 49, we know that  $\text{stab}(\tau, \gamma^* - 1, Y') \leq \text{stab}_Y(\tau, \gamma^*) + 2$ . Thus, it follows that

$$\begin{aligned} \text{score}(\tau, Y') &\geq \min\{\gamma^* - 1, 20L - \text{stab}_{Y'}(\tau, \gamma^* - 1)\} \\ &\geq \min\{\gamma^* - 1, 20L - \text{stab}_Y(\tau, \gamma^*)\} \geq \text{score}_Y(\tau) - 1, \end{aligned}$$

as desired. ■

**Lemma 51 (Existence of a Good Stable Interval)** *Suppose  $\mathcal{A}(\eta/2)$  is feasible on  $Y$ . For every  $L \in [0, 0.25\eta n]$ , there is a  $\tau \in [0, \eta n]$  such that  $\text{score}(\tau, Y) \geq L$ .*

**Proof** Consider  $\text{Pot}_{\eta/2}, \text{Pot}_{\eta/2+2L/n}, \dots, \text{Pot}_{\eta/2+2Lr/n}$  where  $r := \lceil 0.25\eta n/L \rceil$ . Observe that  $\text{Pot}_{\eta/2}(Y) - \text{Pot}_\eta(Y) \leq (1 - \eta/2)n - (1 - \eta)^2 n \leq 1.5\eta n$ . Therefore, there must exist  $r^* \in [r]$  such that

$$\text{Pot}_{\eta/2+2L(r^*-1)/n} - \text{Pot}_{\eta/2+2Lr^*/n} \leq \frac{1.5\eta n}{r} \leq 12L.$$

Let  $\tau = \eta/2 + (2Lr^* - 1)/n$  and  $\gamma = L$ . Then, we have  $\text{stab}(\tau, \gamma) \leq 12L$  and, thus,

$$\text{score}(\tau, Y) \geq \max\{\gamma, 20L - 12L\} \geq L. \quad \blacksquare$$

**Lemma 52 (Utility of Score Function)** *Suppose  $\mathcal{A}(\eta/2)$  is feasible on  $Y$ . Let  $\varepsilon, \delta, \beta \in (0, 1]$ . For every  $L \in [0, 0.25\eta n]$ , if  $L \geq O\left(\frac{1}{\varepsilon} \cdot \log\left(\frac{n}{\beta\delta}\right)\right)$ , then with probability  $1 - \beta$ , Theorem 38, invoked with the score function in Theorem 48 and  $\kappa = L/2$ , does not reject, and the output  $\tau$  satisfies  $\text{stab}_Y(\tau, L/2) < 20L$ .*

**Proof** This follows from the guarantee of Selection (Theorem 38), Theorem 51 and the definition of score. ■

**Lemma 53 (Potential Stability Under Good Coupling)** *Let  $\eta, \varepsilon, \delta > 0$  and  $k, L \in \mathbb{N}$  be given input parameters such that  $0.25\eta n \geq L = \Omega\left(\frac{1}{\varepsilon} \cdot \log\left(\frac{n}{\beta\delta}\right)\right)$ . Let  $Y, Y'$  be adjacent subsets of  $\mathbb{Q}^d$ . Suppose Alg does not halt and chooses  $\eta' = \tau/n$  in Step 1 on input  $Y$  and  $Y'$ . Then,*

$$|\text{Pot}_{\eta'}(Y) - \text{Pot}_{\eta'}(Y')| \leq 20L.$$

Consequently, if  $p, p'$  are scalings of  $\tilde{\mathbb{E}}_{\tilde{\zeta}}[w]$  and  $\tilde{\mathbb{E}}_{\tilde{\zeta}'}[w]$  so that  $\|p\|_1 = \|p'\|_1 = 1$ , then,

$$\|p - p'\|_1 \leq 120\sqrt{L/n}.$$



**Proof** It is enough to prove that  $\text{Pot}_{\eta'}(Y) - \text{Pot}_{\eta'}(Y') \leq 20L$  as a symmetric argument proves the other direction and completes the proof.

Let  $\tilde{\zeta}$  be the pseudo-distribution that minimizes  $\left\| \tilde{\mathbb{E}}_{\tilde{\zeta}}[w] \right\|_2^2$  while satisfying  $\mathcal{A}$  on  $Y'$  with outlier rate  $\eta'$  (computed in Step 3 of the algorithm on input  $Y'$ ). Suppose  $Y$  and  $Y'$  differ on  $i$ -th sample point. Let  $\tilde{\zeta}_{adj}$  be the adjacent pseudo-distribution obtained by zeroing out  $w_i$ . Then, from [Theorem 45](#), we know that  $\tilde{\zeta}_{adj}$  is consistent with  $\mathcal{A}$  on input  $Y$  with outlier rate  $\eta' + 1/n$ . Further,  $\left\| \tilde{\mathbb{E}}_{\tilde{\zeta}_{adj}}[w] \right\|_2^2 \leq \left\| \tilde{\mathbb{E}}_{\tilde{\zeta}}[w] \right\|_2^2$ . Thus,  $\text{Pot}_{\eta'+1/n}(Y) \leq \text{Pot}_{\eta'}(Y')$ . Further, [Theorem 52](#) implies that  $|\text{Pot}_{\eta'+1/n}(Y) - \text{Pot}_{\eta'}(Y)| \leq 20L$ . Therefore, we have  $\text{Pot}_{\eta'}(Y) - \text{Pot}_{\eta'}(Y') \leq 20L$  as desired.

Now, by Cauchy-Schwarz inequality, we immediately obtain that:

$$\left\| \tilde{\mathbb{E}}_{\tilde{\zeta}}[w] - \tilde{\mathbb{E}}_{\tilde{\zeta}'}[w] \right\|_1^2 \leq 20nL$$

Thus, from [Lemma 25](#), we have that:

$$\|p - p'\|_1 \leq 120\sqrt{L/n}.$$

■

**Parameter closeness from potential stability** The following lemma observes that if a sequence of weights  $p_i(Y)$  induces a  $2k$ -certifiably  $C'$ -subgaussian distribution on  $Y$  and  $p'_i(Y')$  is a sequence of weights on an adjacent  $Y'$  such that  $p_i(Y)$  is not too far from  $p_i(Y')$ , then,  $p_i(Y')$  must also induce a  $2k$ -certifiably  $C' + 1$ -subgaussian distribution on  $Y'$ .

**Lemma 54** *Let  $0 \leq p_i(Y) \leq \frac{1}{(1-2\eta')}$  be a sequence of non-negative weights adding up to  $n$  that induce a  $2k$ -certifiable  $C'$ -subgaussian distribution on  $Y$ . Let  $p_i(Y')$  be a sequence of non-negative weights adding up to  $n$  on  $Y'$  adjacent to  $Y$  such that  $\|p(Y) - p(Y')\|_1 \leq \beta$  for  $\beta \leq \eta_0$ . Then, for small enough absolute constant  $\eta' > 0$ ,  $p_i(Y')$  induces a  $2k$ -certifiable  $(C' + 1)$ -subgaussian distribution on  $Y$ .*

**Proof** [Proof Sketch] Let's first describe the idea of the proof: the proof of [Lemma 43](#) requires the existence of a certifiably subgaussian distribution that was close (in total variation distance) to the input  $Y$ . Since  $Y$  is adjacent to  $Y'$ , the  $2k$ -certifiably  $C'$ -subgaussian distribution is  $1 - \beta - 2/n$ -close (the additive  $2/n$  comes from “removing” the index of the point where  $Y$  and  $Y'$  differ) in total variation distance to  $Y$ . Thus, the idea is to use the certifiably subgaussian distribution supported on  $Y$  in lieu of  $X$  to repeat the argument. In order to apply [Lemma 43](#), we need a “flat” distribution—but this is easily achieved. Given a distribution with weights (without loss of generality, say, rational numbers  $r_i/s$ ), we can consider a sample expansion to  $ns$  samples that has  $r_i$  copies of sample  $y_i$  for each  $i$  and an analogous transformation to  $Y'$ . And finally, given a pseudo-distribution on  $w_1, w_2, \dots, w_n$  on  $Y \cap Y'$ , we can transform to a pseudo-distribution on  $ns$  variables by each “copying”  $w_i$  for  $i$  such that  $y_i = y'_i r_i$  times. ■

As an immediate corollary of [Lemma 53](#) and [Lemma 54](#), we obtain:

**Corollary 55 (Parameter Closeness from Stability of Potential)** *Let  $\eta, \varepsilon, \delta > 0$  and  $k, L \in \mathbb{N}$  be given input parameters to [Algorithm 2](#) such that  $0.25\eta n \geq L = \Omega\left(\frac{1}{\varepsilon} \cdot \log\left(\frac{n}{\beta\delta}\right)\right)$ . Also, let  $Y, Y'$  be adjacent subsets of  $\mathbb{Q}^d$ . Suppose [Alg](#) does not reject in any of the 3 steps, uses the constant  $C'$  in Step 2 and chooses  $\eta'$  in Step 1 on input  $Y$  and  $Y'$ .*

*Then, for every  $u \in \mathbb{R}^d$  and  $\theta = \sqrt{L/n}$ , we have:*

$$\begin{aligned} \langle \mu_p - \mu_{p'}, u \rangle &\leq O(C'k)\theta^{1-1/2k} \sqrt{u^\top \Sigma_p u}, \\ (1 - O(C'k)\theta^{1-1/k})\Sigma_p &\preceq \Sigma_{p'} \preceq (1 + O(C'k)\theta^{1-1/k})\Sigma_p, \end{aligned}$$

and, for every  $t \leq k$  such that  $t$  divides  $2k$ ,

$$(1 - O(C'^{t/2}k^{t/2}))\theta^{1-t/2k} \langle u^{\otimes t}, M_p^{(t)} \rangle \leq \langle u^{\otimes t}, M_{p'}^{(t)} \rangle \leq (1 + O(C'^{t/2}k^{t/2}))\theta^{1-t/2k} \langle u^{\otimes t}, M_p^{(t)} \rangle,$$

**Proof** Let  $\tilde{\zeta}_{adj}$  be the adjacent pseudo-distribution of degree  $2k$  to  $\tilde{\zeta}$  obtained by zeroing out  $w_i$  where  $i$  is the index of the point that  $Y$  and  $Y'$  differ on. Then, from [Theorem 45](#), we know that  $\tilde{\zeta}_{adj}$  satisfies  $\mathcal{A}$  on both inputs  $Y, Y'$  with outlier rate  $\eta' + 1/n$  and  $|\left| \mathbb{E}_{\tilde{\zeta}_{adj}}[w] - \mathbb{E}_{\tilde{\zeta}}[w] \right|_2^2| \leq 1$ ,  $|\left| \mathbb{E}_{\tilde{\zeta}_{adj}}[w] - \mathbb{E}_{\tilde{\zeta}'}[w] \right|_2^2| \leq 1$ . Let  $p_{adj}$  be the scaling of  $\mathbb{E}_{\tilde{\zeta}_{adj}}[w]$  so that  $\|p_{adj}\|_1 = 1$ . Then, clearly,  $\|p - p_{adj}\|_1 \leq 2/n$  (since  $\eta' \ll 1/2$ ). Further, applying [Lemma 53](#) and triangle inequality, we have that  $\|p_{adj} - p'\|_1 \leq O(\sqrt{L/n})$ . Applying [Fact 27](#) to  $p_{adj}$  and  $p$  on  $Y$  and  $p_{adj}$  and  $p'$  on  $Y'$  and using triangle inequality completes the proof.  $\blacksquare$

**Noise injection in estimate-dependent norms** Our final ingredient for obtaining privacy guarantees for our robust estimation algorithms is a new noise injection mechanism where the distribution of noise depends on the covariance estimated by our algorithm.

**Lemma 56** *Suppose  $\varepsilon, \delta > 0$ . Let  $A$  be an invertible  $d \times d$  matrix that satisfies  $(1 - \beta)I \preceq AA^T \preceq (1 + \beta)I$ , where  $\beta \leq \frac{\varepsilon}{3d \ln(d/\delta)}$ . Let  $z \in \mathbb{R}^d$  be a vector whose entries are i.i.d. from  $\mathcal{N}(0, 1)$ . Then,*

$$D_{e^\varepsilon}(z, Az) \leq \delta.$$

**Proof** Note that the probability distribution function of  $Az$  at  $u \in \mathbb{R}^d$  is

$$\frac{1}{\det(A)} \frac{1}{(\sqrt{2\pi})^d} e^{-\|A^{-1}u\|_2^2/2}.$$

Moreover,  $\det(A) \leq (1 + \beta)^{d/2}$ , since  $\det(A)^2 = \det(A) \det(A^T) = \det(AA^T) \leq (1 + \beta)^d$ . Thus, the ratio of the probability densities of  $z$  and  $Az$  at  $u$  is

$$\begin{aligned} \det(A) e^{\|A^{-1}u\|_2^2/2 - \|u\|_2^2/2} &\leq (1 + \beta)^{d/2} e^{\|A^{-1}u\|_2^2/2 - \|u\|_2^2/2} \\ &\leq (1 + \beta)^{d/2} e^{\frac{1}{2}u^T((AA^T)^{-1} - I)u} \\ &\leq (1 + \beta)^{d/2} e^{\frac{1}{2}(1-\beta)^{-1}\|u\|_2^2 - \frac{1}{2}\|u\|_2^2} \\ &\leq (1 + \beta)^{d/2} e^{\frac{\beta}{2(1-\beta)}\|u\|_2^2}. \end{aligned}$$

Thus, note that if  $\|u\|_\infty \leq \sqrt{2 \ln(d/\delta)}$ , then  $\|u\|_2 \leq \sqrt{d} \cdot \|u\|_\infty \leq \sqrt{2d \ln(d/\delta)}$ , and so,

$$\det(A) e^{\|A^{-1}u\|_2^2/2 - \|u\|_2^2/2} \leq (1 + \beta)^{d/2} e^{\frac{\beta}{1-\beta} d \ln(d/\delta)} < e^\varepsilon,$$

since  $\beta \leq \frac{\varepsilon}{3d \ln(d/\delta)}$ .

Moreover, by standard tail bounds of the normal distribution, we have that  $\|z\|_\infty > \sqrt{2 \ln(d/\delta)}$  with probability at most  $\delta$ . This proves the claim.  $\blacksquare$

**Lemma 57** *Suppose  $\varepsilon, \delta > 0$ . Let  $A$  be a  $d \times d$  matrix that satisfies  $\|AA^T - I\|_2 \leq \beta$ .*

*Let  $t \in \mathbb{N}$ . Moreover, let  $Z \in \mathbb{R}^{d^t}$  be a random vector indexed by  $[d]^t$ , whose entries  $Z_{i_1, i_2, \dots, i_t}$ , for  $1 \leq i_1 \leq i_2 \leq \dots \leq i_t \leq d$ , are i.i.d. from  $\mathcal{N}(0, 1)$ , and moreover,  $Z_{i_1, i_2, \dots, i_t} = Z_{i_{\pi(1)}, i_{\pi(2)}, \dots, i_{\pi(t)}}$  for any  $i = (i_1, i_2, \dots, i_t)$  and permutation  $\pi$ .*

*If  $\beta \leq \frac{\varepsilon}{8t^2 d^t \ln(d^t/\delta)}$ , then*

$$D_{e^\varepsilon}(Z, A^{\otimes t} Z) \leq \delta.$$

**Proof** Let  $K = \sqrt{2 \ln(d^t/\delta)}$ . By standard tail bounds, note that

$$\mathbb{P}[\|Z\|_\infty > K] \leq \delta. \quad (\text{B.10})$$

Let  $S$  be the subspace of  $\mathbb{R}^{d^t}$  consisting of all symmetric tensors, i.e.,

$$S = \left\{ u \in \mathbb{R}^{d^t} : u_{i_1, i_2, \dots, i_t} = u_{(i_{\pi(1)}, i_{\pi(2)}, \dots, i_{\pi(t)})}, \forall i = (i_1, i_2, \dots, i_t) \in [d]^t, \pi \text{ a permutation on } [t] \right\}.$$

Note that  $S$  is an  $d'$ -dimension invariant subspace of  $A^{\otimes t}$ , where  $d' = \binom{d+t-1}{t} \leq d^t$ . Moreover, let  $R \subseteq [d]^t$  be a representative set of indices of size  $|R| = d'$ , i.e.,  $R$  satisfies the property that for any  $(i_1, i_2, \dots, i_t) \in [d]^t$ , there exists a permutation  $\pi$  on  $[t]$  such that  $(i_{\pi(1)}, i_{\pi(2)}, \dots, i_{\pi(t)}) \in R$ .

Now, let  $M = A^{\otimes t}|_S$  be the restriction of  $A^{\otimes t}$  to the subspace  $S$ . Moreover, let  $Z_R \in \mathbb{R}^{d'}$  denote the projection of  $Z$  to indices in  $R$ .

Note that the probability distribution of  $Z$  can be equivalently viewed as the probability distribution of  $Z_R$ , since  $Z$  is uniquely determined by the projection  $Z_R$ . Let  $p$  be the probability density function of  $Z_R$  over  $\mathbb{R}^{d'}$ . Then, note that the probability distribution of  $MZ_R$  is  $q$ , given by

$$q(v) = \frac{1}{\det(M)} \cdot p(M^{-1}v).$$

for  $v \in \mathbb{R}^{d'}$ . By standard properties, we know that the  $i^{\text{th}}$  singular value of  $M$  is bounded from above by the  $i^{\text{th}}$  singular value of  $A^{\otimes t}$  and bounded from below by the  $(i + d^t - d')$ <sup>th</sup> singular value of  $A^{\otimes t}$ . Moreover, by  $\|AA^T - I\|_2 \leq \beta$ , we know that the singular values of  $A^{\otimes t}$  lie in  $[(1 - \beta)^{t/2}, (1 + \beta)^{t/2}]$ . Hence, the singular values of  $M$  also lie in  $[(1 - \beta)^{t/2}, (1 + \beta)^{t/2}]$ , which, together with  $\beta t \leq \frac{1}{4}$ , implies that

$$\|MM^T - I\|_2 \leq 2\beta t \quad (\text{B.11})$$

and so,

$$\det(M)^2 = \det(M) \det(M^T) = \det(MM^T) \leq (1 + 2\beta t)^{t \cdot d'} \leq (1 + 2\beta t)^{td^t}, \quad (\text{B.12})$$

and so,  $\det(M) \leq (1 + 2\beta t)^{td^t/2}$ .

Let  $u \in \mathbb{R}^{d^t}$ . Note that  $\|u\|_\infty \leq K$  if and only if  $v \in \mathbb{R}^{d^t}$  given by  $v = u|_R$  also satisfies  $\|v\|_\infty \leq K$ . Moreover, note that if  $\|v\|_\infty \leq K$ , then

$$\begin{aligned}
 \frac{p(v)}{q(v)} &\leq \det(M) \cdot \frac{p(v)}{p(M^{-1}v)} \\
 &\leq (1 + 2\beta t)^{td^t/2} \cdot \exp\left(\frac{1}{2} \left(\|M^{-1}v\|_2^2 - \|v\|_2^2\right)\right) \\
 &\leq e^{\beta t^2 d^t} \cdot \exp\left(\frac{1}{2} (v^T((MM^T)^{-1} - I)v)\right) \\
 &\leq e^{\beta t^2 d^t} \cdot \exp\left(\frac{1}{2} \|v\|_2^2 \cdot \|(MM^T)^{-1} - I\|_2\right)
 \end{aligned} \tag{B.13}$$

By (B.11), we have that  $\|(MM^T)^{-1} - I\|_2 \leq \frac{2\beta t}{1-2\beta t} \leq 4\beta t$ , since  $\beta t \leq \frac{1}{4}$ . Therefore, (B.13) is at most

$$e^{\beta t^2 d^t} \cdot \exp(2K^2 d^t \beta t).$$

Thus, if  $\beta \leq \frac{\varepsilon}{2K^2 t d^t} = \frac{\varepsilon}{8t^2 d^t \ln(d^t/\delta)}$ , the above quantity is at most  $e^\varepsilon$ . This, combined with (B.10), proves the desired claim.  $\blacksquare$

**Remark 58** Note that [Theorem 57](#) uses an assumption on the spectral norm of  $\|AA^T - I\|_2$ . However, it is also possible to obtain a version of the lemma under an assumption on the Frobenius norm,  $\|AA^T - I\|_F$ . In particular, if we assume that, instead,  $\|AA^T - I\|_F \leq \beta$ , then [Eq. \(B.12\)](#) instead becomes  $\det(M) \leq \left(1 + \frac{\beta}{\sqrt{d}}\right)^{td^t/2} \leq e^{\beta t d^t - \frac{1}{2}/2}$ : This follows from the fact that (a.) the eigenvalues  $\lambda_1, \lambda_2, \dots, \lambda_d$  of  $AA^T$  satisfy  $\sum_{i=1}^d (\lambda_i - 1)^2 \leq \beta^2$ , (b.) under the aforementioned constraint,  $\lambda_1 \lambda_2 \cdots \lambda_d$  is maximized when  $\lambda_1 = \lambda_2 = \cdots = \lambda_d = 1 + \frac{\beta}{\sqrt{d}}$ , (c.) the eigenvalues of  $(AA^T)^{\otimes t}$  are precisely the  $d^t$   $t$ -fold products of eigenvalues of  $AA^T$ .

**Putting things together** Now, we are ready to prove the main privacy guarantee provided by our robust moment estimation algorithm, [Algorithm 2](#).

**Lemma 59 (Privacy Guarantee)** Suppose  $C, \eta, \varepsilon, \delta > 0$  and  $k \in \mathbb{N}$ . Suppose  $n \geq n_0 = \tilde{\Omega}\left(\left(\frac{Ck^4 d^k}{\varepsilon} (\ln(6kd^k/\delta) + \frac{\varepsilon}{6k})\right)^{\frac{2k}{k-1}}\right)$ . Then, [Alg](#) (given by [Algorithm 2](#)), invoked with  $L = O(\log(n/\delta)/\varepsilon)$ , is  $(\varepsilon, \delta)$ -DP.

**Proof** Let  $\varepsilon' = \varepsilon/3$  and  $\delta' = \delta/3$ . By our adaptive composition theorem under halting ([Theorem 37](#)), it suffices to show that each step of the algorithm is  $(\varepsilon', \delta')$ -DP (given the outputs of the previous steps as parameter<sup>3</sup>). Let  $Y$  and  $Y'$  be any neighboring datasets.

- **Stable Outlier Rate Selection.** Since this step invokes the  $(\varepsilon', \delta')$ -DP Selection algorithm (Selection), it immediately follows from [Theorem 38](#) that this step is  $(\varepsilon', \delta')$ -DP.

3. Note that we may also assume that the algorithm has not halted from the previous steps.

- **Witness Checking.** Let  $C^*(Y)$  denote the smallest  $C^*$  for which  $p_i(Y)$  induces a  $2k$ -certifiable  $C^*$ -subgaussian distribution on  $Y$ . [Theorem 54](#) ensures that  $|C^*(Y) - C^*(Y')| \leq \Delta$  for  $\Delta = 1$ . Therefore, we may apply [Theorem 35](#) with DP parameters  $\varepsilon', \delta'$  to conclude that this step is also  $(\varepsilon', \delta')$ -DP.
- **Noise Addition.** Since the algorithm has not halted in the previous step and the truncated Laplace noise is negative,  $p_i(Y)$  and  $p_i(Y')$  must induce  $2k$ -certifiable  $C'$ -subgaussian distributions on  $Y$  and  $Y'$  respectively. Let  $\tilde{\mu}$  and  $\tilde{\mu}'$  denote the corresponding mean estimates under  $p_i(Y)$  and  $p_i(Y')$ , respectively, and, similarly, let  $\tilde{\Sigma}$  and  $\tilde{\Sigma}'$  denote the corresponding covariance estimates. By [Theorem 55](#), we have that, for all  $u \in \mathbb{R}^d$ ,

$$\langle \tilde{\mu} - \tilde{\mu}', u \rangle \leq \gamma_1 \sqrt{u^\top \tilde{\Sigma} u} \quad (\text{B.14})$$

$$(1 - \gamma_2)\tilde{\Sigma} \preceq \tilde{\Sigma}' \preceq (1 + \gamma_2)\tilde{\Sigma} \quad (\text{B.15})$$

and, for all  $2 \leq t \leq k$ ,

$$(1 - \gamma_t)\langle u^{\otimes t}, \widetilde{M}^{(t)} \rangle \leq \langle u^{\otimes t}, \widetilde{M}^{(t)} \rangle \leq (1 + \gamma_t)\langle u^{\otimes t}, \widetilde{M}^{(t)} \rangle, \quad (\text{B.16})$$

where  $\theta = \sqrt{L/n}$ ,  $\gamma_1 = O(C'k)\theta^{1-1/2k}$ , and  $\gamma_t = O((C'k)^{t/2})\theta^{1-t/2k}$  for  $2 \leq t \leq k$ . Moreover, let  $B = \tilde{\Sigma}^{-1/2}\tilde{\Sigma}'^{1/2}$ .

Note that in order to show that the noise addition step is  $(\varepsilon', \delta')$ -DP, it suffices to show that

$$D_{e^{\varepsilon''}}(\tilde{\mu} + \tilde{\Sigma}^{1/2}z, \tilde{\mu}' + \tilde{\Sigma}'^{1/2}z) \leq \delta'' \quad (\text{B.17})$$

$$D_{e^{\varepsilon''}}(\tilde{\Sigma} + \tilde{\Sigma}^{1/2}Z\tilde{\Sigma}^{1/2}, \tilde{\Sigma}' + \tilde{\Sigma}'^{1/2}Z\tilde{\Sigma}'^{1/2}) \leq \delta'' \quad (\text{B.18})$$

$$\forall 2 < t \leq k, \quad D_{e^{\varepsilon''}}(\widetilde{M}^{(t)} + (\tilde{\Sigma}^{1/2})^{\otimes t}Z^{(t)}, \widetilde{M}^{(t)} + (\tilde{\Sigma}'^{1/2})^{\otimes t}Z^{(t)}) \leq \delta'' \quad (\text{B.19})$$

for  $\varepsilon'' = \varepsilon'/k$  and  $\delta'' = \delta'/k$ , since [Fact 30](#) and standard DP composition ([Dwork et al., 2006a](#)) then imply that the entire noise addition step is  $(\varepsilon', \delta')$ -DP. We now establish each of the above inequalities.

**Noise addition for mean:** We first show [\(B.17\)](#). Note that

$$\begin{aligned} D_{e^{\varepsilon''}}(\tilde{\mu} + \tilde{\Sigma}^{1/2}z, \tilde{\mu}' + \tilde{\Sigma}'^{1/2}z) &= D_{e^{\varepsilon''}}(\tilde{\Sigma}^{1/2}z, (\tilde{\mu}' - \tilde{\mu}) + \tilde{\Sigma}^{1/2}z) \\ &= D_{e^{\varepsilon''}}(z, \tilde{\Sigma}^{-1/2}(\tilde{\mu}' - \tilde{\mu}) + Bz) \\ &= D_{e^{\varepsilon''/2}}(z, z + \tilde{\Sigma}^{-1/2}(\tilde{\mu}' - \tilde{\mu})) \\ &\quad + e^{\varepsilon''/2}D_{e^{\varepsilon''/2}}(z + \tilde{\Sigma}^{-1/2}(\tilde{\mu}' - \tilde{\mu}), \tilde{\Sigma}^{-1/2}(\tilde{\mu}' - \tilde{\mu}) + Bz) \end{aligned} \quad (\text{B.20})$$

$$= D_{e^{\varepsilon''/2}}(z, z + \tilde{\Sigma}^{-1/2}(\tilde{\mu}' - \tilde{\mu})) + D_{e^{\varepsilon''/2}}(z, Bz), \quad (\text{B.21})$$

where [\(B.20\)](#) follows from [Theorem 31](#). For the first term on the right-hand side of [\(B.21\)](#), we note that  $\left\| \tilde{\Sigma}^{-1/2}(\tilde{\mu}' - \tilde{\mu}) \right\|_2 \leq \gamma_1$  (which follows from plugging in  $u = \tilde{\Sigma}^{-1}(\tilde{\mu} - \tilde{\mu}')$  into [\(B.14\)](#)). Thus, by the standard hockey-stick divergence calculation for the Gaussian mechanism ([Dwork and Roth, 2014, Appendix A](#)), we have that

$$D_{e^{\varepsilon''/2}}(z, z + \tilde{\Sigma}^{-1/2}(\tilde{\mu}' - \tilde{\mu})) < \delta''/2, \quad (\text{B.22})$$

provided that

$$\sigma_1 \geq \frac{2\gamma_1 \sqrt{2 \ln(2.5/\delta'')}}{\varepsilon''},$$

For the second term in (B.21), note that (B.15) implies that  $(1 - \gamma_2)I \preceq BB^T \preceq (1 + \gamma_2)I$ . Moreover,  $\gamma_2 \leq \frac{\varepsilon''}{3d \ln(2d/\delta'')}$ , by the condition  $n \geq n_0$ . Thus, by Theorem 56,

$$D_{e^{\varepsilon''/2}}(z, Bz) \leq \delta''/2. \quad (\text{B.23})$$

Therefore, (B.22), (B.23), and (B.21) imply (B.17), as desired.

**Noise addition for covariance:** Next, we establish (B.18). Observe that

$$\begin{aligned} & D_{e^{\varepsilon''}}(\tilde{\Sigma} + \tilde{\Sigma}^{1/2} Z \tilde{\Sigma}^{1/2}, \tilde{\Sigma}' + \tilde{\Sigma}'^{1/2} Z \tilde{\Sigma}'^{1/2}) \\ &= D_{e^{\varepsilon''}}(I + Z, BB^T + BZB^T) \\ &= D_{e^{\varepsilon''}}(Z, (BB^T - I) + BZB^T) \\ &\leq D_{e^{\varepsilon''/2}}(Z, Z + (BB^T - I)) \\ &\quad + e^{\varepsilon''/2} D_{e^{\varepsilon''/2}}(Z + (BB^T - I), (BB^T - I) + BZB^T) \end{aligned} \quad (\text{B.24})$$

$$\leq D_{e^{\varepsilon''/2}}(Z, Z + (BB^T - I)) + e^{\varepsilon''/2} D_{e^{\varepsilon''/2}}(Z, BZB^T), \quad (\text{B.25})$$

where (B.24) follows from Theorem 31. To bound the right-hand side of (B.25), note that the first term is precisely the hockey-stick divergence computation corresponding to the Gaussian mechanism (restricted to the upper triangular portion). Moreover, by (B.15),

$$\|BB^T - I\|_F \leq \sqrt{d} \cdot \|BB^T - I\|_2 \leq \gamma_2 \sqrt{d}. \quad (\text{B.26})$$

Therefore ((Dwork and Roth, 2014, Appendix A)), as long as

$$\sigma_2 \geq \frac{2\gamma_2 \sqrt{2d \ln(2.5/\delta'')}}{\varepsilon''},$$

we have that

$$D_{e^{\varepsilon''/2}}(Z, Z + (BB^T - I)) \leq \delta''/2. \quad (\text{B.27})$$

For the second term in (B.25), note that  $\sigma_2^{-1}Z$  has entries distributed in  $\mathcal{N}(0, 1)$ . Moreover, let  $Z' = \text{vec}(Z)$  be the  $d^2$ -dimensional vector given by the flattening of  $Z$  (see Theorem 19). By Fact 21, we know that  $BZB^T = B^{\otimes 2}Z'$ . Thus, by Theorem 57 applied with  $t = 2$ ,

$$\begin{aligned} D_{e^{\varepsilon''/2}}(Z, BZB^T) &= D_{e^{\varepsilon''/2}}(Z', B^{\otimes 2}Z') \\ &= D_{e^{\varepsilon''/2}}(\sigma_2^{-1}Z', B^{\otimes 2}(\sigma_2^{-1}Z')) \\ &\leq \delta''/2e^{\varepsilon''/2}, \end{aligned} \quad (\text{B.28})$$

as long as

$$\gamma_2 < \frac{\varepsilon''}{32d^2 \ln(2d^2 e^{\varepsilon''/2}/\delta'')},$$

which is true, since  $n \geq n_0$  by the conditions of the theorem. Thus, (B.27) and (B.28) imply that (B.25) is at most  $\delta''/2 + e^{\varepsilon''/2}(\delta''/2e^{\varepsilon''/2}) = \delta''/2$ , which establishes (B.18).

**Noise addition for higher-order moments:** Let  $2 < t \leq k$ . We write  $R = \widetilde{\Sigma} + \widetilde{\mu}\widetilde{\mu}^T$  and  $R' = \widetilde{\Sigma}' + \widetilde{\mu}'\widetilde{\mu}'^T$  for simplicity.

Observe that the injective/spectral norm  $\|\cdot\|_\sigma$  of  $(R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)})$  can be bounded as

$$\begin{aligned}
 \left\| (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) \right\|_\sigma &= \sup_{\substack{v \in \mathbb{R}^d \\ \|v\|_2=1}} \left| (v^{\otimes t})^T (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) \right| \\
 &\leq \sup_{\substack{v \in \mathbb{R}^d \\ \|v\|_2=1}} \left| \langle (R^{-1/2}v)^{\otimes t}, \widetilde{M}'^{(t)} - \widetilde{M}^{(t)} \rangle \right| \\
 &\leq \gamma_t \cdot \sup_{\substack{v \in \mathbb{R}^d \\ \|v\|_2=1}} \left| \langle (R^{-1/2}v)^{\otimes t}, \widetilde{M}^{(t)} \rangle \right| \\
 &\leq \gamma_t \cdot (C'k)^t \cdot \sup_{\substack{v \in \mathbb{R}^d \\ \|v\|_2=1}} \left| \left( (R^{-1/2}v)^T R (R^{-1/2}v) \right)^{t/2} \right| \quad (\text{B.29}) \\
 &= \gamma_t \cdot (C'k)^t \cdot \sup_{\substack{v \in \mathbb{R}^d \\ \|v\|_2=1}} \|v\|_2^t \\
 &= \gamma_t \cdot (C'k)^t,
 \end{aligned}$$

where (B.29) follows from the  $C'$ -subgaussianity property of the distribution induced by the weight vector at the end of Step 2. Therefore, the Frobenius norm (or Hilbert-Schmidt norm) can be bounded as (see Corollary 4.10 of Wang et al. (2017))

$$\left\| (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) \right\|_F \leq d^{\frac{t-1}{2}} \cdot \left\| (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) \right\|_\sigma \leq \gamma_t \cdot (C'k)^t \cdot d^{\frac{t-1}{2}}. \quad (\text{B.30})$$

Moreover, letting  $W = R^{-1/2}R^{1/2}$ , we have

$$\begin{aligned}
 &D_{e^{\varepsilon''}} (\widetilde{M}^{(t)} + (R^{1/2})^{\otimes t} Z^{(t)}, \widetilde{M}'^{(t)} + (R^{1/2})^{\otimes t} Z^{(t)}) \\
 &= D_{e^{\varepsilon''}} \left( (R^{1/2})^{\otimes t} Z^{(t)}, \widetilde{M}'^{(t)} - \widetilde{M}^{(t)} + (R^{1/2})^{\otimes t} Z^{(t)} \right) \\
 &= D_{e^{\varepsilon''}} \left( Z^{(t)}, (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) + W^{\otimes t} Z^{(t)} \right) \\
 &\leq D_{e^{\varepsilon''/2}} \left( Z^{(t)}, Z^{(t)} + (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) \right) \\
 &\quad + e^{\varepsilon''/2} D_{e^{\varepsilon''/2}} \left( Z^{(t)} + (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}), \right. \\
 &\quad \left. (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) + W^{\otimes t} Z^{(t)} \right) \quad (\text{B.31}) \\
 &\leq D_{e^{\varepsilon''/2}} \left( Z^{(t)}, Z^{(t)} + (R^{-1/2})^{\otimes t} (\widetilde{M}'^{(t)} - \widetilde{M}^{(t)}) \right) \\
 &\quad + e^{\varepsilon''/2} D_{e^{\varepsilon''/2}} (Z^{(t)}, W^{\otimes t} Z^{(t)}), \quad (\text{B.32})
 \end{aligned}$$

where again we have used Theorem 31 in (B.31). In order to bound the right-hand side of (B.32), note that the first term is again the hockey-stick divergence computation corresponding to the

Gaussian mechanism (restricted according to symmetry conditions). Recalling (B.30), we see that ((Dwork and Roth, 2014, Appendix A)) as long as

$$\sigma_t \geq \frac{2\gamma_t(C'k)^t d^{\frac{t-1}{2}} \sqrt{2 \ln(2.5/\delta'')}}{\varepsilon''},$$

we have that

$$D_{e^{\varepsilon''/2}} \left( Z^{(t)}, Z^{(t)} + (R^{-1/2})^{\otimes t} \left( \widetilde{M}^{(t)} - \widetilde{M}^{(t)} \right) \right) \leq \delta''/2. \quad (\text{B.33})$$

For the second term in (B.32), note that  $\sigma_t^{-1} Z^{(t)}$  has entries distributed in  $\mathcal{N}(0, 1)$ . Moreover, note that  $\|WW^T - I\|_F \leq \|BB^T - I\|_F \leq \gamma_2 \sqrt{d}$  by (B.26) and the fact that (B.15) implies

$$(1 - \gamma_2)R \preceq R' \preceq (1 + \gamma_2)R.$$

Thus, by Theorem 57, we have that

$$\begin{aligned} D_{e^{\varepsilon''/2}} \left( Z^{(t)}, W^{\otimes t} Z^{(t)} \right) &= D_{e^{\varepsilon''/2}} \left( \sigma_t^{-1} Z^{(t)}, W^{\otimes t} (\sigma_t^{-1} Z^{(t)}) \right) \\ &\leq \delta''/2 e^{\varepsilon''/2}, \end{aligned} \quad (\text{B.34})$$

as long as

$$\gamma_2 < \frac{\varepsilon''}{16t^2 d^t \ln(2d^t e^{\varepsilon''/2}/\delta'')},$$

which is true since  $n \geq n_0$ , by the hypothesis of the lemma. Thus, (B.33) and (B.34) imply that (B.32) is at most  $\delta''/2 + e^{\varepsilon''/2}(\delta''/2e^{\varepsilon''/2}) = \delta''$ , which establishes (B.19), as desired.  $\blacksquare$

#### B.4. Proof of Theorem 39

We are now ready to prove our main theorem, Theorem 39.

**Proof** [Proof of Theorem 39] Choose  $\beta = 1/30$ . Choose  $L = \Omega\left(\frac{1}{\varepsilon} \cdot \log\left(\frac{n}{\beta\delta}\right)\right)$  (according to the condition in Theorem 52). Moreover, let  $C = C_0 + \frac{3 \ln(3/\delta)}{\varepsilon} + \frac{9}{\varepsilon} + 1$ . Then, we claim that setting Alg to be Algorithm 2 with parameters  $C, \eta, \varepsilon, \delta, L, k$  satisfies the desired conditions, as long as  $\eta \leq \eta_0$ , where we set  $\eta_0$  later.

Note that the desired privacy guarantees follow immediately from Theorem 59.

It remains to prove the utility guarantees. Suppose that there indeed exists a good set  $X \subseteq \mathbb{Q}^d$  with mean  $\mu_*$ , covariance  $\Sigma_*$ , and  $t$ -th moments  $M_*^{(t)}$  for  $3 \leq t \leq k$ , such that  $|Y \cap X| \geq (1 - \eta)n$ .

By Theorem 38, we have that Step 1 (stable outlier rate selection) rejects and halts with probability at most  $\beta = 1/30$ , and the resulting output  $\tau$  satisfies  $\text{score}(\tau, Y) \geq L/2$ . In particular, the latter condition implies that there exists some  $\gamma \geq L/2$  for which  $\mathcal{A}\left(\frac{\tau - \gamma}{n}\right)$  is feasible. By monotonicity,  $\mathcal{A}(\eta')$  is also feasible, where we let  $\eta' = \tau/n$ .

Hence, the invocation of Algorithm 1 in Step 2 does not yield “reject.” Moreover, note that by Theorem 34, we have that  $C' = C + \gamma \geq C_0$  with probability at least  $29/30$ . In this case, the computed weight vector  $p$  induces a  $C'$ -certifiably subgaussian distribution on  $Y$ . Hence, the probability of rejection in Step 2 is at most  $1/30$ .



Let  $\tilde{\mu} = \tilde{\mathbb{E}}_{\tilde{z}}[\mu]$ ,  $\tilde{\Sigma} = \tilde{\mathbb{E}}_{\tilde{z}}[\Sigma]$ , and  $\tilde{M}^{(t)} = \tilde{\mathbb{E}}_{\tilde{z}}[\Sigma]$  (for  $2 \leq t \leq k$ ) be the estimates of the mean, covariance, and  $t$ -th moments, respectively, that are outputted by the [Algorithm 1](#) subroutine in [Step 2](#) of [Algorithm 2](#). Then, by [Theorem 42](#), we have

$$\begin{aligned} \forall u \in \mathbb{R}^d, \langle \tilde{\mu} - \mu_*, u \rangle &\leq O\left(\sqrt{Ck}\right) \eta^{1-1/2k} \sqrt{u^\top \Sigma_* u} \\ (1 - \beta_2) \Sigma_* &\preceq \tilde{\Sigma} \preceq (1 + \beta_2) \Sigma_* \end{aligned}$$

and, for all even  $2 \leq t \leq k$  such that  $t$  divides  $2k$ ,

$$\forall u \in \mathbb{R}^d, (1 - \beta_t) \langle u^{\otimes t}, M_*^{(t)} \rangle \leq \langle u^{\otimes t}, \tilde{M}^{(t)} \rangle \leq (1 + \beta_t) \langle u^{\otimes t}, M_*^{(t)} \rangle,$$

where  $\beta_t = \beta_t(\eta) = O((Ck)^{t/2}) \eta^{1-t/2k}$ . We now set  $\eta_0$  such that  $\beta_t(\eta_0) \leq \frac{1}{2}$  for all aforementioned  $t$ . Note that this guarantees that  $\beta_t = \beta_t(\eta) \leq \frac{1}{2}$ , since we are assuming  $\eta \leq \eta_0$ .

Now, consider the noise addition step, i.e., [Step 3](#) of [Algorithm 2](#). Note that by the Cauchy-Schwarz Inequality, for any  $u \in \mathbb{R}^d$ , we have

$$\begin{aligned} \langle \tilde{\Sigma}^{1/2} z, u \rangle &= \langle \Sigma_*^{-1/2} \tilde{\Sigma}^{1/2} z, \Sigma_*^{1/2} u \rangle \\ &\leq \left\| \Sigma_*^{-1/2} \tilde{\Sigma}^{1/2} z \right\|_2 \cdot \left\| \Sigma_*^{1/2} u \right\|_2 \\ &= (z^\top \tilde{\Sigma}^{1/2} \Sigma_*^{-1} \tilde{\Sigma}^{1/2} z) \cdot \sqrt{u^\top \Sigma_* u} \\ &\leq \|z\|_2^2 \cdot \left(1 + \left\| \tilde{\Sigma}^{1/2} \Sigma_*^{-1} \tilde{\Sigma}^{1/2} - I \right\|_2\right) \cdot \sqrt{u^\top \Sigma_* u} \\ &\leq \|z\|_2^2 \cdot (1 + 2\beta_2) \cdot \sqrt{u^\top \Sigma_* u} \\ &\leq 2 \|z\|_2^2 \cdot \sqrt{u^\top \Sigma_* u}, \end{aligned}$$

since  $\beta_2 \leq \frac{1}{2}$  by our choice of  $\eta_0$ . Now, note that with probability at least  $1 - \frac{1}{30k}$ , we have that  $\|z\|_2 \leq O\left(\sigma_1 \sqrt{d \ln(kd)}\right)$ , in which case it follows that

$$\|z\|_2^2 = O(d) \cdot \sigma_1^2 \ln(kd) = O(\sqrt{Ck}) \eta^{1-1/2k},$$

by our choice of  $n \geq n_0$ . Thus, the mean estimate  $\hat{\mu}$  outputted by the [Step 3](#) satisfies

$$\begin{aligned} \langle \hat{\mu} - \mu_*, u \rangle &= \langle \hat{\mu} - \tilde{\mu}, u \rangle + \langle \tilde{\mu} - \mu_*, u \rangle \\ &= \langle \tilde{\Sigma}^{1/2} z, u \rangle + \langle \tilde{\mu} - \mu_*, u \rangle \\ &= O\left(\sqrt{Ck}\right) \eta^{1-1/2k} \sqrt{u^\top \Sigma_* u}. \end{aligned} \tag{B.35}$$

Next, we consider the utility guarantee for the covariance. Note that  $\|Z\|_2 \leq \nu_2 = O\left(\sigma_2 \sqrt{d \ln(kd^2)}\right)$  with probability at least  $1 - \frac{1}{30k}$  (this follows from standard spectral properties of Wigner matrices; see, for instance, [Tao \(2012\)](#)), in which case, it follows that  $-\nu_2 \tilde{\Sigma} \preceq \tilde{\Sigma}^{1/2} Z \tilde{\Sigma}^{1/2} \preceq \nu_2 \tilde{\Sigma}$ . Moreover, by our choice of  $n \geq n_0$  as well as  $\eta_0$ , we have that  $\nu_2 \leq \beta_2 \leq \frac{1}{2}$ . Thus, it follows that

$$\hat{\Sigma} \preceq (1 + \beta_2) \tilde{\Sigma}$$

$$\begin{aligned}
 &\preceq (1 + \beta_2)^2 \Sigma_* \\
 &\preceq \left(1 + \frac{5}{2} \beta_2\right) \Sigma_* \\
 &\preceq \left(1 + O(Ck) \cdot \eta^{1-1/k}\right) \Sigma_*,
 \end{aligned}$$

and by a similar argument, we also have  $\hat{\Sigma} \succeq (1 - O(Ck) \cdot \eta^{1-1/k}) \Sigma_*$ , thus implying that

$$\left(1 - O(Ck) \cdot \eta^{1-1/k}\right) \Sigma_* \preceq \hat{\Sigma} \preceq \left(1 + O(Ck) \cdot \eta^{1-1/k}\right) \Sigma_*. \quad (\text{B.36})$$

Finally, we consider the utility guarantee for moment estimation. Suppose  $2 \leq t \leq k$  and  $t$  is an even number dividing  $2k$ . Let  $A = \tilde{\Sigma} + \tilde{\mu}\tilde{\mu}^T$  and  $A_* = \Sigma_* + \mu_*\mu_*^T$ . Note that for any  $2 < t \leq k$ , we have  $\|Z^{(t)}\|_F = O\left(\sigma_t d^{t/2} \sqrt{\ln(kd^t)}\right)$  with probability at least  $1 - \frac{1}{30k}$ . In this case, note that for any  $u \in \mathbb{R}^d$ , we have the following (recall that  $\|\cdot\|_\sigma$  indicates the injective norm of a tensor):

$$\begin{aligned}
 \langle u^{\otimes t}, (A^{1/2})^{\otimes t} Z^{(t)} \rangle &= \langle (A^{1/2}u)^{\otimes t}, Z^{(t)} \rangle \\
 &\leq \|Z^{(t)}\|_\sigma \cdot \|A^{1/2}u\|_2^t \\
 &\leq \|Z^{(t)}\|_F \cdot \|A^{1/2}u\|_2^t \\
 &= O(\sigma_t d^{t/2} \sqrt{\ln(kd^t)}) \cdot \|A^{1/2}u\|_2^t \\
 &= O(\sigma_t d^{t/2} \sqrt{\ln(kd^t)}) \cdot (u^T A u)^{t/2} \\
 &= O(\sigma_t (d^{t/2} \sqrt{\ln(kd^t)}) \cdot ((1 + \beta_2)u^T A_* u)^{t/2} \\
 &= O(\sigma_t (d(1 + \beta_2))^{t/2} \sqrt{\ln(kd^t)}) \cdot \|A_*^{1/2}u\|_2^t \\
 &= O(\sigma_t (de^{\beta_2})^{t/2} \sqrt{\ln(kd^t)}) \cdot \langle u^{\otimes t}, M_*^{(t)} \rangle \quad (\text{B.37}) \\
 &= O((Ck)^{t/2} \eta^{1-t/2k}) \cdot \langle u^{\otimes t}, M_*^{(t)} \rangle, \quad (\text{B.38})
 \end{aligned}$$

where (B.37) follows from Jensen's Inequality, and (B.38) follows from our choice of  $n \geq n_0$ . Thus, the moment estimate  $\hat{M}^{(t)} = \tilde{M}^{(t)} + (A^{1/2})^{\otimes t} Z^{(t)}$  outputted by our algorithm satisfies

$$\begin{aligned}
 \langle u^{\otimes t}, \hat{M}^{(t)} \rangle &\leq \langle u^{\otimes t}, \tilde{M}^{(t)} \rangle + \langle u^{\otimes t}, (A^{1/2})^{\otimes t} Z^{(t)} \rangle \\
 &\leq (1 + \beta_t) \langle u^{\otimes t}, M_*^{(t)} \rangle + O((Ck)^{t/2} \eta^{1-t/2k}) \cdot \langle u^{\otimes t}, M_*^{(t)} \rangle \\
 &= \left(1 + O((Ck)^{t/2} \eta^{1-t/2k})\right) \langle u^{\otimes t}, M_*^{(t)} \rangle.
 \end{aligned}$$

In a similar fashion, we also get that  $\langle u^{\otimes t}, \hat{M}^{(t)} \rangle \geq (1 - O((Ck)^{t/2} \eta^{1-t/2k})) \langle u^{\otimes t}, M_*^{(t)} \rangle$ , thus implying that

$$\left(1 - O((Ck)^{t/2} \eta^{1-t/2k})\right) \langle u^{\otimes t}, M_*^{(t)} \rangle \leq \langle u^{\otimes t}, \hat{M}^{(t)} \rangle \leq \left(1 + O((Ck)^{t/2} \eta^{1-t/2k})\right) \langle u^{\otimes t}, M_*^{(t)} \rangle. \quad (\text{B.39})$$

Hence, (B.35), (B.36), and (B.39) imply the desired utility guarantees.

Moreover, recall that the rejection probabilities at Steps 1 and 2 are each at most  $\frac{1}{30}$ , and it is not possible to reject in Step 3. Moreover, the  $\leq k$  utility guarantees each fail with probability at most  $\frac{1}{30k}$ . Thus, by a union bound, it follows that the algorithm does not reject and, moreover, outputs estimates satisfying the desired utility guarantees with probability at least  $1 - \frac{1}{30} - \frac{1}{30} - k \cdot \frac{1}{30k} = \frac{9}{10}$ .

Finally, note that the running time of  $(Bn)^{O(k)}$  follows from the time complexity guarantee in Theorem 42, as the invocation of Algorithm 1 in Step 2 is the bottleneck. Steps 1 and 3 are easily seen to run in  $(Bn)^{O(k)}$  time. This completes the proof.  $\blacksquare$

### Appendix C. Robust Mean and Covariance Estimation for Certifiably Hypercontractive Distributions

In this section, we observe that we can upgrade our guarantees from the previous section for robust estimation of moments of distributions that have certifiably hypercontractive degree 2 polynomials.

**Definition 60** A distribution  $cD$  on  $\mathbb{R}^d$  with mean  $\mu_*$  and covariance  $\Sigma_*$  is said to have  $2h$ -certifiably  $C$ -hypercontractive degree 2 polynomials if for a  $d \times d$  matrix-valued indeterminate  $Q$  and  $\bar{x} = x - \mu_*$ ,

$$\left| \frac{Q}{2h} \left\{ \mathbb{E}_{x \sim D} (\bar{x}^\top Q \bar{x}) - \mathbb{E}_{x \sim D} \bar{x}^\top Q \bar{x} \right\}^{2h} \leq (Ch)^{2h} \left\| \Sigma_*^{1/2} Q \Sigma_*^{1/2} \right\|_F^{2h} \right\}.$$

The Gaussian distribution (Kauers et al., 2014), uniform distribution on the hypercube and more generally other product domains and their affine transforms are known to satisfy  $2t$ -certifiably  $C$ -hypercontractivity with an absolute constant  $C$  for every  $t$ .

In order to derive this conclusion, we note the following analog of the witness-producing algorithm and its guarantees:

**Witness-producing version of the robust moment estimation algorithm** We will use the following (non-private) guarantees for the robust moment estimation algorithm in the previous section that hold for a strengthening of the constraint system  $\mathcal{A}$  with certifiable hypercontractivity constraints. Using the analysis of Diakonikolas et al. (2016), the following guarantees were recently shown in Kothari et al. (2021) for the case when the unknown distribution is Gaussian.

For any  $d \times d$  matrix-valued indeterminate  $Q$ , let  $\bar{x}'_i{}^\top Q \bar{x}'_i = x'^\top Q x' - \frac{1}{n} \sum_{i=1}^n x'_i{}^\top Q x'$ .

#### $\mathcal{A}$ : Constraint System for $\eta$ -Robust Moment Estimation

1.  $w_i^2 = w_i$  for each  $1 \leq i \leq n$ ,
2.  $\Pi^2 = \frac{1}{n} \sum_{i=1}^n (x'_i - \mu')(x'_i - \mu')^\top$ ,
3.  $\sum_{i=1}^n w_i \geq (1 - \eta)n$ ,
4.  $\mu' = \frac{1}{n} \sum_i x'_i$ ,
5.  $w_i(x'_i - y_i) = 0$  for  $1 \leq i \leq n$ ,

$$6. \frac{1}{n} \sum_{i=1}^n \bar{x}'_i{}^\top Q \bar{x}'_i{}^2 \leq C \|\Pi Q \Pi\|_F^2.$$

The following guarantees for the algorithm above were shown in [Bakshi and Kothari \(2020a\)](#).

**Fact 61 ((Bakshi and Kothari, 2020a))** *Let  $X \subseteq \mathbb{R}^d$  be an i.i.d. sample of size  $n \geq n_0 = \tilde{O}(d^2/\eta)$  from  $\mathcal{N}(\mu_*, \Sigma_*)$ . Let  $Y$  be an  $\eta$ -corruption of  $X$ . Then, for  $\mu' = \frac{1}{n} \sum_i x'_i$ ,  $\Sigma' = \frac{1}{n} \sum_i (x_i - \mu')(x_i - \mu')^\top$ , we have:*

$$\begin{aligned} \mathcal{A} \Big|_{O(k)}^u & \left\{ \langle \mu' - \mu_*, u \rangle \leq O(\eta^{1-1/2k}) u^\top \Sigma_* u^2 \right\}, \\ \mathcal{A} \Big|_{O(k)}^u & \left\{ \langle u, (\Sigma' - \Sigma_*) u \rangle \leq O(\eta^{1-1/k}) u^\top \Sigma_* u \right\}, \\ \mathcal{A} \Big|_{O(k)} & \left\{ \left\| \Sigma_*^{-1/2} \Sigma' \Sigma_*^{-1/2} - I \right\|_F^2 \leq O(\eta^{1-1/k}) \right\}. \end{aligned}$$

The first two guarantees of the lemma below were shown in [Bakshi and Kothari \(2020a\)](#). The third guarantee follows from an argument similar to that of [Lemma 43](#). Notice that the key difference in the guarantees below (compared to the ones in [Lemma 42](#)) is the bound on the Frobenius (instead of the weaker spectral) distance between the estimated covariance and true unknown covariance.

**Lemma 62 (Guarantees for Witness-Producing Robust Moment Estimation Algorithm)**

*Given a subset of  $n$  points  $Y \subseteq \mathbb{Q}^d$  whose entries have bit complexity  $B$ , [Algorithm 1](#) runs in time  $(Bn)^{O(1)}$  and either (a.) outputs “reject,” or (b.) returns a sequence of weights  $0 \leq p_1, p_2, \dots, p_n$  satisfying  $p_1 + p_2 + \dots + p_n = 1$ .*

*Moreover, if there exists a set  $X \subseteq \mathbb{R}^d$  of points with 4-certifiably  $C$ -hypercontractive degree 2 polynomials with mean  $\mu_*$ , covariance  $\Sigma_*$ , then [Algorithm 1](#) does not reject, and the corresponding estimates  $\hat{\mu} = \frac{1}{n} \sum_i p_i y_i$  and  $\hat{\Sigma} = \sum_{i=1}^n p_i (y_i - \hat{\mu})(y_i - \hat{\mu})^\top$  satisfy the following guarantees:*

**1. Mean Estimation:**

$$\forall u \in \mathbb{R}^d, \langle \hat{\mu} - \mu_*, u \rangle \leq O(\sqrt{C}) \eta^{3/4} \sqrt{u^\top \Sigma_* u},$$

**2. Covariance Estimation:**

$$\left\| \Sigma_*^{-1/2} \hat{\Sigma} \Sigma_*^{-1/2} - I \right\|_F \leq O(C \eta^{1/2}),$$

**3. Witness:** For  $C' \leq C(1 + O(\eta^{1/2}))$ ,

$$\left| \frac{Q}{n} \left\{ \frac{1}{n} \sum_{i=1}^n p_i \left( \langle y_i - \hat{\mu}, Q(y_i - \hat{\mu}) \rangle - \frac{1}{n} \sum_{i=1}^n p_i \langle y_i - \hat{\mu}, Q(y_i - \hat{\mu}) \rangle \right) \right\}^2 \leq C' \left\| \hat{\Sigma}^{1/2} Q \hat{\Sigma}^{1/2} \right\|_F^2 \right\}$$

We can now use the above witness-producing algorithm to obtain a stronger Frobenius norm estimation guarantee with  $(\varepsilon, \delta)$ -privacy for Gaussian distributions. Notice that the only change from the previous section is in the choice of the constraint system  $\mathcal{A}$  and the corresponding change in the witness checking step.

The parameter closeness from potential stability is also upgraded from [Theorem 55](#):

---

**Algorithm 1:** Private Robust Moment Estimation

---

Given: A set of points  $Y = \{y_1, y_2, \dots, y_n\} \subseteq \mathbb{Q}^d$ , parameters  $\eta, \varepsilon, \delta > 0, L \in \mathbb{N}$ .

Output: Estimates  $\hat{\mu}$  and  $\hat{\Sigma}$  for mean and covariance.

Operation:

- (a) **Stable Outlier Rate Selection:** Use the  $(\varepsilon/3, \delta/3)$ -DP Selection with  $\kappa = L/2$  to sample an integer  $\tau \in [\eta n]$  with the scoring function as defined in [Theorem 48](#). If  $\tau = \perp$ , then reject and halt. Otherwise, let  $\eta' = \tau/n$ .
  - (b) **Witness Checking:** Compute a pseudo-distribution  $\tilde{\zeta}$  of degree  $O(1)$  satisfying  $\mathcal{A}$  on input  $Y$  with outlier rate  $\eta'$  and minimizing  $\text{Pot}_{\eta', \tilde{\zeta}}(Y)$ . Let  $\gamma \sim \text{tLap}\left(-\left(1 + \frac{3 \ln(3/\delta)}{\varepsilon}\right), 3/\varepsilon\right)$ . Check that the weight vector  $p = \tilde{\mathbb{E}}_{\tilde{\zeta}}[w]$  induces a distribution on  $Y$  that has  $(C + \gamma)$ -certifiably hypercontractive polynomials. If not, reject immediately. Otherwise, let  $\tilde{\mu} = \tilde{\mathbb{E}}_{\tilde{\zeta}}[\mu]$  and  $\tilde{\Sigma} = \tilde{\mathbb{E}}_{\tilde{\zeta}}[\Sigma]$ .
  - (c) **Noise Addition:** Let  $\gamma_1 = O(C')(L/n)^{\frac{1}{4}}$  and  $\gamma_2 = O(C')(L/n)^{\frac{1}{4}}$ . Let  $z \sim \mathcal{N}(0, \sigma_1)^d$  and  $Z \sim \mathcal{N}(0, \sigma_2)^{\binom{d+1}{2}}$ , where we interpret  $Z$  has a symmetric  $d \times d$  matrix with independent lower-triangular entries, and  $\sigma_j = 12\varepsilon^{-1}\gamma_j\sqrt{2\ln(15/\delta)}$  for  $1 \leq j \leq 2$ . Then, output:
    - $\hat{\mu} = \tilde{\mu} + \tilde{\Sigma}^{1/2}z$ .
    - $\hat{\Sigma} = \tilde{\Sigma} + \tilde{\Sigma}^{1/2}Z\tilde{\Sigma}^{1/2}$ .
-

**Lemma 63 (Parameter Closeness from Stability of Potential)** *Let  $\eta, \varepsilon, \delta > 0$  and  $L \in \mathbb{N}$  be given input parameters to [Algorithm 1](#) such that  $0.25\eta n \geq L = \Omega\left(\frac{1}{\varepsilon} \cdot \log\left(\frac{n}{\beta\delta}\right)\right)$ . Also, let  $Y, Y'$  be adjacent subsets of  $\mathbb{Q}^d$ . Suppose [Alg](#) does not reject in any of the 3 steps, uses the constant  $C'$  in Step 2 and chooses  $\eta'$  in Step 1 on input  $Y$  and  $Y'$ .*

*Then, for every  $u \in \mathbb{R}^d$  and  $\theta = \sqrt{L/n}$ , we have:*

$$\langle \mu_p - \mu_{p'}, u \rangle \leq O(C')\theta^{3/4}\sqrt{u^\top \Sigma_p u}$$

and

$$\left\| \Sigma_p^{-1/2} \Sigma_{p'} \Sigma_p^{-1/2} - I \right\|_F \leq O(C')\theta^{1/2}.$$

The following theorem summarizes our privacy and utility guarantees for the algorithm above. We specialize to the “base case assumption” of 4-certifiably  $C$ -hypercontractivity of degree 2 polynomials in order to derive explicit bounds here. Our analysis of the algorithm above follows *mutatis mutandis* with the key upgrade being the stronger Frobenius norm guarantees in [Lemma 54](#) that hold under certifiably hypercontractivity constraints in our constraint system  $\mathcal{A}$  (this requires us to use a version of [Theorem 57](#) that makes use of a bound on  $\|AA^T - I\|_F$  instead of  $\|AA^T - I\|_2$ ; see the remark at the end of [Theorem 57](#)). As before, the  $\tilde{\Omega}$  notation hides logarithmic multiplicative factors in  $d, C, 1/\eta, 1/\varepsilon$ , and  $\ln(1/\delta)$ .

**Theorem 64 (Private Robust Mean and Covariance Estimation for Certifiably Hypercontractive Distributions)** *Fix  $C_0 > 0$ . Then, there exists an  $\eta_0 > 0$  such that for any given outlier rate  $0 < \eta \leq \eta_0$  and  $\varepsilon, \delta > 0$ , there exists a randomized algorithm [Alg](#) that takes an input of  $n \geq n_0 = \tilde{\Omega}\left(\frac{d^8}{\eta^2} \left(1 + \frac{\ln(1/\delta)}{\varepsilon}\right)^4 \cdot C^4\right)$  points  $Y = \{y_1, y_2, \dots, y_n\} \subseteq \mathbb{Q}^d$  (where  $C = C_0 + \frac{3\ln(3/\delta)}{\varepsilon} + \frac{9}{\varepsilon} + 1$ ), runs in time  $(Bn)^{O(1)}$  (where  $B$  is the bit complexity of the entries of  $Y$ ) and outputs either “reject” or estimates  $\hat{\mu} \in \mathbb{Q}^d$  and  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$  with the following guarantees:*

1. **Privacy:** *[Alg](#) is  $(\varepsilon, \delta)$ -differentially private with respect to the input  $Y$ , viewed as a  $d$ -dimensional database of  $n$  individuals.*
2. **Utility:** *Suppose there exists a 4-certifiably  $C_0$ -subgaussian set  $X = \{x_1, x_2, \dots, x_n\} \subseteq \mathbb{Q}^d$  such that  $|Y \cap X| \geq (1 - \eta_0)n$  with mean  $\mu_*$  and covariance  $\Sigma_* \succeq 2^{-\text{poly}(d)}I$ . Then, with probability at least 9/10 over the random choices of the algorithm, [Alg](#) outputs estimates  $\hat{\mu} \in \mathbb{Q}^d$  and  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$  satisfying the following guarantees:*

$$\forall u \in \mathbb{R}^d, \langle \hat{\mu} - \mu_*, u \rangle \leq O(\sqrt{C}\eta^{3/4})\sqrt{u^\top \Sigma_* u},$$

and,

$$\left\| \Sigma_*^{-1/2} \hat{\Sigma} \Sigma_*^{-1/2} - I \right\|_F \leq O(C\sqrt{\eta}).$$

*Moreover, the algorithm succeeds (i.e., does not reject) with probability at least 9/10 over the random choices of the algorithm.*

When specialized to Gaussian distributions, the Frobenius guarantee above is suboptimal—the robust estimation algorithms of [Diakonikolas et al. \(2016\)](#) allow estimating the mean and covariance of the unknown Gaussian distribution to an error  $\tilde{O}(\eta)$ . We can in fact recover the stronger guarantees by relying on the analysis in [Kothari et al. \(2021\)](#)[Theorem 1 and 2] of the same constraint system above for the case of Gaussian distributions (in the “utility case”). This yields the following corollary:

**Theorem 3 (Mean and Covariance Estimation for Gaussian Distributions)** *Fix  $\varepsilon, \delta > 0$ . Then, there exists an absolute constant  $\eta_0 > 0$  such that for any given outlier rate  $0 < \eta \leq \eta_0$ , there exists a randomized algorithm **Alg** that takes an input of  $n \geq n_0 = \tilde{\Omega}\left(\frac{d^8}{\eta^4}\left(1 + \frac{\ln(1/\delta)}{\varepsilon}\right)^4\right)$  points  $Y \subseteq \mathbb{Q}^d$ , runs in time  $(Bn)^{O(1)}$  (where  $B$  is the bit complexity of the entries of  $Y$ ) and outputs either “reject” or estimates  $\hat{\mu} \in \mathbb{Q}^d$  and  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$  with the following guarantees:*

1. **Privacy:** *Alg* is  $(\varepsilon, \delta)$ -differentially private with respect to the input  $Y$ , viewed as a  $d$ -dimensional database of  $n$  individuals.
2. **Utility:** *Let  $X = \{x_1, x_2, \dots, x_n\}$  be an i.i.d. sample of size  $n \geq n_0$  from a Gaussian distribution with mean  $\mu_*$  and covariance  $\Sigma_* \succeq 2^{-\text{poly}(d)}I$  such that  $Y$  is an  $\eta$ -corruption of  $X$ . Then, with probability at least  $9/10$  over the random choices of the algorithm, **Alg** outputs estimates  $\hat{\mu} \in \mathbb{Q}^d$  and  $\hat{\Sigma} \in \mathbb{Q}^{d \times d}$  satisfying the following guarantees:*

$$\forall u \in \mathbb{R}^d, \langle \hat{\mu} - \mu_*, u \rangle \leq \tilde{O}\left(\eta \cdot \frac{\log(1/\delta)}{\varepsilon}\right) \sqrt{u^\top \Sigma_* u},$$

and,

$$\left\| \Sigma_*^{-1/2} \hat{\Sigma} \Sigma_*^{-1/2} - I \right\|_F \leq \tilde{O}\left(\eta \cdot \sqrt{\frac{\log(1/\delta)}{\varepsilon}}\right),$$

where the  $\tilde{O}$  hides multiplicative logarithmic factors in  $1/\eta$ . In particular,  $d_{\text{TV}}(\mathcal{N}(\hat{\mu}, \hat{\Sigma}), \mathcal{N}(\mu_*, \Sigma_*)) < \tilde{O}(\eta \log(1/\delta)/\varepsilon)$ .