New Lower Bounds for Private Estimation and a Generalized Fingerprinting Lemma

Anonymous Author(s)
Affiliation
Address
email

Abstract

We prove new lower bounds for statistical estimation tasks under the constraint of (ε, δ)-differential privacy. First, we provide tight lower bounds for private covariance estimation of Gaussian distributions. We show that estimating the covariance matrix in Frobenius norm requires $\Omega(d^2)$ samples, and in spectral norm requires $\Omega(d^{3/2})$ samples, both matching upper bounds up to logarithmic factors. We prove these bounds via our main technical contribution, a broad generalization of the fingerprinting method [BUV14] to exponential families. Additionally, using the private Assouad method of Acharya, Sun, and Zhang [ASZ21], we show a tight $\Omega(d/(\alpha^2\varepsilon))$ lower bound for estimating the mean of a distribution with bounded covariance to α-error in ℓ_2-distance. Prior known lower bounds for all these problems were either polynomially weaker or held under the stricter condition of $(\varepsilon, 0)$-differential privacy.

1 Introduction

The last several years have seen a surge of interest in algorithms for statistical estimation under the constraint of *Differential Privacy* (DP) [DMNS06]. We now have a rich algorithmic toolbox for private estimation of mean, covariance, and entire distributions, in a variety of settings.

Despite significant success in designing *algorithms* for estimation tasks, *lower bounds* have been more elusive, and consequently, the existing literature lacks a rigorous understanding of some core problems. Lower bounds against the strong privacy constraint of *pure* DP (i.e., $(\varepsilon, 0)$-DP) are generally not too challenging to prove, most often relying upon the well-known packing technique [HT10]. Such packing lower bounds are optimal in a broad range of settings (see, e.g., [BKSW19]), and thus are frequently effective at providing tight minimax sample complexity lower bounds.

However, pure DP is a very strong privacy constraint: in most practical circumstances, it suffices to only require the weaker privacy notion of *approximate* DP (i.e., (ε, δ)-DP) [DKM$^+$06], where packing lower bounds are no longer applicable. Proving lower bounds under approximate DP is much more challenging, and accordingly, the state of affairs is less satisfying. There exist only a couple of techniques which apply in this setting, including the fingerprinting method [BUV14] and differentially private analogues of Le Cam's method and Assouad's lemma [ASZ18, ASZ21]. These techniques can be quite brittle and thus unable to prove lower bounds for some fairly basic settings. For example, both techniques generally require that the underlying distribution has independent marginals and are thus ineffective at proving tight lower bounds for problems involving correlations such as private covariance estimation. Furthermore, the fingerprinting method often needs the distribution's marginals to have a precise functional form, such as Gaussian or Bernoulli. These restrictions make the approaches somewhat fragile for proving lower bounds which differ too much from existing results. Consequently, we lack approximate DP lower bounds for some fundamental

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

settings which seem to qualitatively differ, including Gaussian covariance estimation and mean estimation of heavy-tailed distributions.

1.1 Our Results

We circumvent these barriers and fill a number of gaps in the literature by providing tight lower bounds for several estimation tasks under the constraint of approximate differential privacy. These include covariance estimation for Gaussians and mean estimation with heavy-tailed data. Our main technical tool for covariance estimation is a novel generalization of the fingerprinting method to exponential families, while we prove our mean estimation lower bound via DP Assouad's lemma.

Our first result is a lower bound for covariance estimation in Frobenius norm.

Theorem 1.1 (Informal)**.** *For any $\alpha \leq \mathcal{O}(1)$, any (ε, δ)-DP mechanism with $\varepsilon \in [0,1]$ and $\delta \leq \mathcal{O}\left(\frac{d^2}{n \ln^2\left(\frac{n}{d^2}\right)}\right)$ that takes n samples from a d-dimensional Gaussian $\mathcal{N}(0, \Sigma)$ with unknown covariance Σ and outputs $M(X)$ satisfying $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \Sigma\|_F^2\right] \leq \alpha^2$ requires $n \geq \Omega\left(\frac{d^2}{\alpha\varepsilon}\right)$.*

This nearly matches the $\widetilde{\mathcal{O}}(d^2)$ upper bound of [KLSU19] up to polylogarithmic factors.[1] Previously, $\Omega(d^2)$ lower bounds for Gaussian covariance estimation were only known under the more restrictive constraint of pure differential privacy [KLSU19, BKSW19].

Our result for Gaussian covariance estimation in Frobenius norm implies the following result for spectral estimation.

Theorem 1.2 (Informal)**.** *For any $\alpha \leq \mathcal{O}\left(\frac{1}{\sqrt{d}}\right)$, any (ε, δ)-DP mechanism with $\varepsilon \in [0,1]$ and $\delta \leq \mathcal{O}\left(\frac{d^2}{n \ln^2\left(\frac{n}{d^2}\right)}\right)$ that takes n samples from a d-dimensional Gaussian $\mathcal{N}(0, \Sigma)$ with unknown covariance Σ and outputs $M(X)$ satisfying $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \Sigma\|_2^2\right] \leq \alpha^2$ requires $n \geq \Omega\left(\frac{d^{1.5}}{\alpha\varepsilon}\right)$.*

Note that the sample complexity of the same problem sans privacy constraints is $\Theta(d)$, and thus the cost of privacy is polynomial in the dimension d. This is in contrast to Gaussian mean estimation in ℓ_2-norm and covariance estimation in Frobenius norm, which maintain their non-private sample complexities of $\mathcal{O}(d)$ and $\mathcal{O}\left(d^2\right)$, respectively.

$\Omega(d^{3/2})$-sample lower bounds for approximate DP covariance estimation in spectral norm were previously known for *worst-case* distributions [DTTZ14]. Our work relaxes this assumption significantly to the Gaussian case while maintaining the same lower bound. A priori, it was not clear that such a result was even true, as there are frequently gaps between the sample complexity of private estimation for worst-case versus well-behaved distributions for even basic settings, including distributions over the hypercube (see, e.g., Remark 6.4 of [BKSW19]).

This lower bound has implications for estimation tasks with scale-dependent error. For example, consider Gaussian mean estimation in Mahalanobis distance. A natural way to approach this problem is to first estimate the covariance matrix spectrally, and then estimate the mean after an appropriate rescaling (see, e.g., [KLSU19]). Our lower bound shows that any such approach must incur the $\Omega(d^{3/2})$ cost of spectral estimation, which is greater than the $\mathcal{O}(d)$ minimax sample complexity of the problem. Indeed, some recent works [BGS+21, LKO21] manage to circumvent this roadblock by adopting more direct (but computationally inefficient) methods.

Finally, we prove new lower bounds for mean estimation of distributions with bounded second moments.

[1] We note that all our lower bounds in this work are stated in terms of mean squared error, while most of the upper bounds guarantee error at most α with probability $1 - \beta$ for some $\beta > 0$. In many natural cases (such as when the estimator's range is naturally bounded), lower bounds against MSE can be converted to constant probability statements via a boosting argument (outlined, e.g., in the proof of Theorem 6.1 of [KLSU19]). The details of such a conversion generally standard when applicable and we do not discuss it further here.

Theorem 1.3 (Informal)**.** *For any $\alpha \leq 1$, any (ε, δ)-DP mechanism with $\delta \leq \varepsilon$ that takes n samples from a distribution over $\mathbb{R}^d$ with second moments bounded by 1 with unknown mean and outputs $M(X)$ satisfying $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \mu\|_2^2\right] \leq \alpha^2$ requires $n \geq \Omega\left(\frac{d}{\alpha^2 \varepsilon}\right)$.*

Similar $\Theta(d/(\alpha^2\varepsilon))$ upper and lower bounds were previously known under the stricter constraint of pure differential privacy [BD14, KSU20, HKM21]. Our result shows that the same bound holds under the weaker constraint of approximate DP, and thus no savings can be obtained by relaxing the privacy notion. In contrast to our other results, we show this lower bound via the DP Assouad method of [ASZ21], thus demonstrating the promise of different approaches for proving lower bounds for differentially private estimation.

1.2 Related Work

The fingerprinting technique for proving lower bounds was introduced by Bun, Ullman, and Vadhan [BUV14], which relied on the existence of *fingerprinting codes* [BS98, Tar08]. Since then, the technique has been significantly simplified and refined in a number of ways [SU17a, DSS+15, BSU17, SU17b, KLSU19, CWZ19], including the removal of fingerprinting codes and distilling the main technical component into the *fingerprinting lemma*. Beyond the aforementioned settings of mean estimation of Gaussians and distributions over the hypercube [BUV14, SU15, DSS+15, KLSU19], fingerprinting lower bounds have also been applied in settings including private empirical risk minimization [BST14] and private spectral estimation [DTTZ14]. Differentially private analogues of Fano, Le Cam, and Assouad were first considered in the local DP setting [DJW13], and more recently under the constraint of central DP [ASZ18, ASZ21].

On the upper bounds side, the most relevant body of work is that which focuses on mean and covariance estimation for distributions satisfying a moment bound, see, e.g., [BD14, KV18, BKSW19, BS19, KLSU19, CWZ19, KSU20, WXDX20, DFM+20, BDKU20, AAK21, BGS+21, KLZ21, HLY21, KMS+21, AL21]. Our lower bounds nearly match these upper bounds for Gaussian covariance estimation [KLSU19] and heavy-tailed mean estimation [KSU20, HKM21]. Previously, these bounds were conjectured to be optimal, albeit on somewhat shaky grounds including evidence from lower bounds proven under pure DP or non-rigorous connections with lower bounds in related settings. Our results rigorously prove optimality and confirm these conjectures. Other works study private statistical estimation in different and more general settings, including mixtures of Gaussians [KSSU19, AAL21], graphical models [ZKKW20], discrete distributions [DHS15], and median estimation [AMB19, TVGZ20]. Some recent directions involve guaranteeing user-level privacy [LSY+20, LSA+21], or a combination of local and central DP for different users [ADK19]. See [KU20] for further coverage of DP statistical estimation.

1.3 Organization

We start by providing our generalized fingerprinting technique in Section 2. Then we show its novel application for proving the lower bound for covariance estimation of Gaussians in Mahalanobis norm, followed by the lower bound for spectral norm estimation in Section 3. Finally, we sketch the proof of the lower bound for mean-estimation of heavy-tailed distributions in Section 4. We describe the notations and preliminaries for privacy and exponential families in Appendix A. After that, we detail the missing proofs from Sections 2, 3, and 4 in Appendices B, C, and D, respectively. Furthermore, we state facts from linear algebra in Appendix E, and facts from probability and statistics in Appendix F. Finally, in Appendix G, we include the proofs of older results obtained in [KLSU19], but using our generalized results in the language of exponential families.

2 Fingerprinting Proofs for Exponential Families

This section is split into three parts. Section 2.1 provides an overview of existing fingerprinting proofs and motivation for our result, and introduces all the necessary notation. Then, in Section 2.2, our new fingerprinting lemma is stated. Finally, in Section 2.3, we give a general recipe of applying it.

2.1 Overview of Fingerprinting Proofs and Intuition for Our Result

At the heart of fingerprinting proofs lies capturing the trade-off between accuracy and privacy. For simplicity, assume we have a univariate data-generating distribution and coming from a family parameterized by its mean μ, which we want to estimate. In contrast to minimax lower bounds in classical statistics (which are commonly obtained by constructing hard testing instances), fingerprinting proofs assume the existence of a subset of the domain of μ (commonly an interval $[-R, R]$) and that a value is drawn $\mu \sim \mathcal{U}([-R, R])$. Then, n independent samples $X_1, \ldots, X_n$ are drawn from $\mathcal{D}_\mu$. Now, given an estimator $f : \mathbb{R}^n \to \mathbb{R}$ for μ, we consider the quantity $Z_i \coloneqq c(R, \mu)\left(f\left(X\right) - \mu\right)\left(X_i - \mu\right), \forall i \in [n]$, which measures the correlation between the output of f and the i-th sample ($c(R, \mu) > 0$ is a scaling factor that exists for purely technical reasons). Defining $Z \coloneqq \sum_{i \in [n]} Z_i$, all proofs first show a lower bound L on $\mathbb{E}[Z + (f\left(X\right) - \mu)^2]$, implying that our estimator cannot at the same time be accurate (with respect to the Mean-Squared-Error (MSE)) and exhibit low correlation with individual samples. This component of the proof is traditionally called the *fingerprinting lemma*.

Privacy comes up in the second step of the proof, where we upper bound the correlation terms $\mathbb{E}\left[Z_i\right]$. We assume that f is a private mechanism with MSE at most $L/2$. Due to privacy, all correlation terms Z_i should be upper-bounded by the same $U = U\left(\varepsilon, \delta\right)$. Applying this for each sample leads to an inequality of the form $nU \geq L/2 \iff n \geq L/(2U)$, so the only thing that remains is to identify the appropriate parameter range for δ to prove the desired bound.

Our work is motivated by a fundamental weakness of existing fingerprinting proofs. To prove lower bounds for high-dimensional distributions via fingerprinting, past works have assumed that the underlying distribution is a product distribution. This made it possible to work component-wise using techniques for single-dimensional distributions. This assumption is somewhat restrictive and makes it difficult to generalize the approach to settings with a richer correlation structure.

Our main technical contribution is providing a broad generalization of the celebrated fingerprinting method to exponential families (see Definition A.4). Distributions belonging in the same exponential family are parameterized by their *natural parameter vector* η, instead of their mean μ, for which it is uncommon to have a general closed-form expression. It is a standard fact that, given n samples $X_1, \ldots, X_n$ from p_η, it suffices to know the values $T\left(X_1\right), \ldots, T\left(X_n\right)$ instead of the samples themselves in order to estimate η (see, e.g., [Jor10]). This justifies the term sufficient statistics used to describe T. We denote the mean and covariance of T by μ_T and Σ_T, respectively.

Based on the above remarks, our fingerprinting lemma provides lower bounds for estimating η instead of μ. We assume the existence of vectors η^1, η^2 with $\eta^1_j \leq \eta^2_j, \forall j \in [k]$ and define the intervals $I_j \coloneqq \left[\eta^1_j, \eta^2_j\right]$ for each coordinate $j \in [k]$. Assuming that the hyperrectangle $\bigotimes_{j \in [k]} I_j$ is a subset of $\mathcal{H}$, we draw a natural parameter vector η uniformly at random from it. Observe that I_j may not be centered around the origin, which poses a minor technical difficulty. To deal with this, we assume that our estimator does not estimate η itself, but rather the deviation of η from the *midpoint* $m \coloneqq (\eta^1 + \eta^2)/2$. The two problems are equivalent because any estimator $f : \mathbb{R}^{n \times d} \to \bigotimes_{j \in [k]} [-(\eta^2_j - \eta^1_j)/2, (\eta^2_j - \eta^1_j)/2]$ for the latter problem can be transformed to an estimator for the former, simply by adding m to it (and vice-versa). Defining R to be the *width vector* $\eta^2 - \eta^1$, the correlation between the estimate for η_j and the sample X_i becomes $Z^j_i \coloneqq [R^2_j/4 - (\eta_j - m_j)^2][f_j\left(X\right) - (\eta_j - m_j)](T_j\left(X_i\right) - \mu_{T,j}), \forall i \in [n], j \in [k]$.

Observe that now, Z^j_i measures the correlation between f and the sufficient statistics of the samples. This is in contrast to older fingerprinting proofs, in which the Z^j_i measure correlation between f and the samples themselves. We consider this to be natural in the context of exponential families. Indeed, we know that, if we want to estimate the natural parameter vector of a distribution that belongs to an exponential family given a dataset $X = (X_1, \ldots, X_n)$, it suffices to know the values $T\left(X_1\right), \ldots, T\left(X_n\right)$ instead of the samples themselves (see [Jor10]). Thus, despite f receiving X, correlation ends up being measured with respect to $T\left(X_i\right), i \in [n]$.

In addition to the above, we introduce three more correlation terms:

$$Z_i \coloneqq \sum_{j\in[k]} Z_i^j = \sum_{j\in[k]} \left[R_j^2/4 - (\eta_j - m_j)^2\right]\left[f_j(X) - (\eta_j - m_j)\right]\left(T_j(X_i) - \mu_{T,j}\right),$$

$$Z^j \coloneqq \sum_{i\in[n]} Z_i^j = \left[R_j^2/4 - (\eta_j - m_j)^2\right]\left[f_j(X) - (\eta_j - m_j)\right]\sum_{i\in[n]}\left(T_j(X_i) - \mu_{T,j}\right),$$

$$Z \coloneqq \sum_{i\in[n]} Z_i = \sum_{j\in[k]} Z^j.$$

We motivate each of the above. Z_i describes the correlation between X_i and our estimates for η. Moreover, Z^j describes the correlation between our estimate for η_j and X. Finally, Z describes the total correlation between the output of f and X. Our fingerprinting lemma for exponential families will involve upper and lower-bounding the quantity $\mathop{\mathbb{E}}_{\eta,X}\left[Z + \|f(X) - (\eta - m)\|_2^2\right]$.

2.2 The Main Lemma and Our Approach to Lower Bounds

We prove our new fingerprinting lemma in two parts. First, we assume we have a distribution p_η that belongs to an exponential family and we want to estimate η_j whereas η_{-j} is fixed (Lemma 2.1). We then use this to get a result when the goal is to estimate all of η (Lemma 2.2).

Lemma 2.1. *Let p_η be a distribution belonging to an exponential family $\mathcal{E}(T,h)$ with natural parameter vector $\eta \in \mathcal{H} \subseteq \mathbb{R}^k$. We assume that η_j is randomly generated by drawing $\eta_j \sim \mathcal{U}(I_j)$ whereas η_{-j} is fixed. Then, for any function $f_j : \mathbb{R}^{n\times d} \to [-R_j/2, R_j/2]$ that takes as input a dataset $X = (X_1,\ldots,X_n) \sim p_\eta^{\bigotimes n}$ we have $\mathop{\mathbb{E}}_{\eta_j,X}\left[Z^j + \left[f_j(X) - (\eta_j - m_j)\right]^2\right] \geq R_j^2/12$.*

Proof. Consider the function $g_j : \mathbb{R}^k \to \mathbb{R}$ where $g_j(\eta) \coloneqq \mathop{\mathbb{E}}_{X}[f_j(X)]$. g_j can equivalently be written as:

$$g_j(\eta) = \int_{\mathbb{R}^d} \cdots \int_{\mathbb{R}^d} \left(\prod_{i\in[n]} p_\eta(x_i)\right) f_j(x_1,\ldots,x_n)\, dx_1 \ldots dx_n.$$

Observe that:

$$\left(\prod_{i\in[n]} p_\eta(x_i)\right) = \left(\prod_{i\in[n]} h(x_i)\right)\exp\left(\left\langle \eta, \sum_{i\in[n]} T(x_i)\right\rangle - nZ(\eta)\right)$$

$$\implies \nabla_\eta\left(\prod_{i\in[n]} p_\eta(x_i)\right) = \left(\prod_{i\in[n]} p_\eta(x_i)\right)\sum_{i\in[n]}\left(T(x_i) - \nabla_\eta Z(\eta)\right)$$

$$= \left(\prod_{i\in[n]} p_\eta(x_i)\right)\sum_{i\in[n]}\left(T(x_i) - \mu_T\right),$$

where the last equality uses the first part of Proposition A.5.

Consequently, the gradient of g_j is:

$$\nabla_\eta g_j(\eta) = \int_{\mathbb{R}^d} \cdots \int_{\mathbb{R}^d} \left(\prod_{i\in[n]} p_\eta(x_i)\right) f_j(x_1,\ldots,x_n)\sum_{i\in[n]}\left(T(x_i) - \mu_T\right) dx_1 \ldots dx_n$$

$$= \mathop{\mathbb{E}}_{X}\left[f_j(X)\sum_{i\in[n]}\left(T(X_i) - \mu_T\right)\right]$$

$$= \mathop{\mathbb{E}}_{X}\left[\left[f_j(X) - (\eta_j - m_j)\right]\sum_{i\in[n]}\left(T(X_i) - \mu_T\right)\right],$$

where the last equality relies on the observation that, conditioned on η_j, $\eta_j - m_j$ and $\sum_{i\in[n]} (T(X_i) - \mu_T)$ are independent and $\mathop{\mathbb{E}}_{X}\left[\sum_{i\in[n]} (T(X_i) - \mu_T)\right] = 0$.

We proceed to define the function $G_j : \mathbb{R}^k \to \mathbb{R}$ as:

$$\begin{aligned} G_j(\eta) \coloneqq \mathop{\mathbb{E}}_{X}\left[Z^j\right] &= [R_j^2/4 - (\eta_j - m_j)^2]\mathop{\mathbb{E}}_{X}\left[[f_j(X) - (\eta_j - m_j)]\sum_{i\in[n]} (T_j(X_i) - \mu_{T,j})\right] \\ &= [R_j^2/4 - (\eta_j - m_j)^2]\frac{\partial g_j}{\partial \eta_j}(\eta). \end{aligned}$$

Calculating the expectation of G_j with respect to η_j yields:

$$\begin{aligned} \mathop{\mathbb{E}}_{\eta_j}[G_j(\eta)] &= \mathop{\mathbb{E}}_{\eta_j}\left[[R_j^2/4 - (\eta_j - m_j)^2]\frac{\partial g_j}{\partial \eta_j}(\eta)\right] \\ &= \frac{1}{|I_j|}\int_{\eta_j^1}^{\eta_j^2} [R_j^2/4 - (\eta_j - m_j)^2]\frac{\partial g_j}{\partial \eta_j}(\eta_1, \ldots, t, \ldots, \eta_k)\, dt \\ &= \frac{1}{|I_j|}[R_j^2/4 - (t - m_j)^2] g_j(\eta_1, \ldots, t, \ldots, \eta_k)\Big|_{t=\eta_j^1}^{t=\eta_j^2} \\ &+ 2\frac{1}{|I_j|}\int_{\eta_j^1}^{\eta_j^2} (t - m_j)\, g_j(\eta_1, \ldots, t, \ldots, \eta_k)\, dt \\ &= 2\mathop{\mathbb{E}}_{\eta_j}[(\eta_j - m_j)\, g_j(\eta)]. \end{aligned}$$

where we used integration by parts and the fact that m_j has been defined to be midpoint of I_j.

Based on this, we remark that the inequality we wish to prove can be written in the form:

$$\mathop{\mathbb{E}}_{\eta_j, X}\left[2(\eta_j - m_j) f_j(X) + [f_j(X) - (\eta_j - m_j)]^2\right] \geq R_j^2/12. \tag{1}$$

Appealing to linearity of expectation, we focus on the second term and get:

$$\begin{aligned} \mathop{\mathbb{E}}_{\eta_j, X}\left[[f_j(X) - (\eta_j - m_j)]^2\right] &= \mathop{\mathbb{E}}_{\eta_j, X}\left[f_j^2(X) - 2(\eta_j - m_j) f_j(X) + (\eta_j - m_j)^2\right] \\ &\geq -2\mathop{\mathbb{E}}_{\eta_j, X}[(\eta_j - m_j) f_j(X)] + \mathop{\mathbb{E}}_{\eta_j}\left[(\eta_j - m_j)^2\right] \\ &= -2\mathop{\mathbb{E}}_{\eta_j, X}[(\eta_j - m_j) f_j(X)] + \mathop{\mathrm{Var}}_{\eta_j}(\eta_j) \\ &= -2\mathop{\mathbb{E}}_{\eta_j, X}[(\eta_j - m_j) f_j(X)] + R_j^2/12. \end{aligned}$$

Substituting this to (1) yields the desired result. ■

We now generalize the previous lemma in the setting where all the components of η are drawn uniformly at random from the hyperrectangle $\bigotimes_{j\in[k]} I_j$.

Lemma 2.2. *Let p_η be a distribution belonging to an exponential family $\mathcal{E}(T, h)$ with natural parameter vector $\eta \in \mathcal{H} \subseteq \mathbb{R}^k$. We assume that η is randomly generated by drawing independently $\eta_j \sim \mathcal{U}(I_j), \forall j \in [k]$. Then, for any function $f : \mathbb{R}^{n\times d} \to \bigotimes_{j\in[k]} [-R_j/2, R_j/2]$ that takes as input a dataset $X = (X_1, \ldots, X_n) \sim p_\eta^{\bigotimes n}$ we have $\mathop{\mathbb{E}}_{\eta, X}\left[Z + \|f(X) - (\eta - m)\|_2^2\right] \geq \|R\|_2^2/12$.*

The proof is merely a component-wise application of Lemma 2.2 and is deferred to Appendix B.1.

2.3 From the Lemma to Lower Bounds

We leverage Lemma 2.2 to give a general recipe for proving lower bounds for private estimation. The overall structure of our argument is similar to that of past fingerprinting proofs but uses specific properties of exponential families to deal with the fact that we don't have a product distribution.

We assume the existence of an (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \bigotimes_{j\in[k]} [-R_j/2, R_j/2]$ with $\varepsilon \in [0,1]$ and $\delta > 0$ such that $\mathop{\mathbb{E}}_{\eta,M}\left[\|M(X) - (\eta - m)\|_2^2\right] \le \alpha^2 \le \|R\|_2^2/24$. For the above, Lemma 2.2 with $f \equiv M$ yields:

$$\mathop{\mathbb{E}}_{\eta,X,M}[Z] = \mathop{\mathbb{E}}_{\eta}\left[\mathop{\mathbb{E}}_{X,M}[Z]\right] \ge \|R\|_2^2/12 - \mathop{\mathbb{E}}_{\eta,X,M}\left[\|M(X) - (\eta - m)\|_2^2\right] \ge \|R\|_2^2/24. \quad (2)$$

Thus, our intent is to upper-bound the LHS by a function of n, which will allow us to identify a lower bound to get the desired accuracy. Our strategy consists of working with individual terms Z_i. In particular, we have the upper bound:

Lemma 2.3. *Conditioning on η, for any $\varepsilon \in [0,1], \delta > 0$, and $T > 0$, it holds that:*

$$\mathop{\mathbb{E}}_{X,M}[Z_i] \le 2\delta T + 2\varepsilon\sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^\top \Sigma_T s]} + 2\int_T^\infty \mathop{\mathbb{P}}_{X_i}\left[\left|\sum_{j\in[k]}(T_j(X_i) - \mu_{T,j})\right| > \frac{4t}{\|R\|_\infty^3}\right] dt.$$

The proof involves splitting Z_i in its positive and negative components and applying the definition of DP to each of them. The argument is standard and has appeared in previous works (e.g., see [SU17b]), but we need to apply specific properties of exponential families to deal with the fact that we don't have a product distribution. The proof is lengthy and is deferred to Appendix B.2.

Remark 2.4. The last term in the LHS of Lemma 2.3 involves the tails of the sufficient statistic of the family $\mathcal{E}(T,h)$. It is a known fact that the MGFs of univariate projections of the sufficient statistics of exponential families are well-defined (see the third part of Proposition A.5). This facilitates proving Chernoff-type bounds in the context of this problem. Lemma F.7 offers one recipe for proving such bounds under mild assumptions.

Combining Lemma 2.3 with (2) yields the main result of this section:

Theorem 2.5. *Let p_η be a distribution belonging to an exponential family $\mathcal{E}(T,h)$ with natural parameter vector $\eta \in \mathcal{H} \subseteq \mathbb{R}^k$. Also, let $\eta^1, \eta^2 \in \mathcal{H}$ and let $I_j := \left[\eta_j^1, \eta_j^2\right], \forall j \in [k]$ be a collection of intervals and $R := \eta^2 - \eta^1, m := (\eta^1 + \eta^2)/2$ be the corresponding width and midpoint vectors, respectively. Assume that $\bigotimes_{j\in[k]} I_j \subseteq \mathcal{H}$ and that η is drawn from the distribution $\mathcal{U}\left(\bigotimes_{j\in[k]} I_j\right)$. Moreover, assume that we have a dataset $X \sim p_\eta^{\bigotimes n}$ and an independently drawn point $X_i' \sim p_\eta$ and $X_{\sim i}$ denotes the dataset where X_i has been replaced with X_i'. Finally, let $M : \mathbb{R}^{n\times d} \to \bigotimes_{j\in[k]}[-R_j/2, R_j/2]$ be an (ε,δ)-DP mechanism with $\varepsilon \in [0,1], \delta > 0$ with $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - (\eta - m)\|_2^2\right] \le \alpha^2 \le \|R\|_2^2/24$. Then, for any $T > 0$, it holds that:*

$$n\left(2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{\eta}\left[\sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^\top \Sigma_T s]}\right] + 2\int_T^\infty \mathop{\mathbb{P}}_{X_i}\left[\left|\sum_{j\in[k]}(T_j(X_i) - \mu_{T,j})\right| > \frac{4t}{\|R\|_\infty^3}\right] dt\right) \ge \frac{\|R\|_2^2}{24}, \quad (3)$$

where $s \in \mathbb{R}^k$ with $s_j := [R_j^2/4 - (\eta_j - m_j)^2]\,[M_j(X_{\sim i}) - (\eta_j - m_j)], \forall j \in [k]$.

In order to apply the above theorem, the main idea is to identify values for T and δ such that $\delta T \ge 2\int_T^\infty \mathop{\mathbb{P}}_{X_i}\left[\left|\sum_{j\in[k]}(T_j(X_i) - \mu_{T,j})\right| > 4t/\|R\|_\infty^3\right] dt$, and $3\delta T n \le \|R\|_2^2/48$.

The reasoning behind this is that the first and third terms of (3) both depend on T, with the latter term becoming smaller as T increases. Thus, balancing these two terms involves identifying a value for T that is as small as possible and results in δT dominating the other term. At the same time, we want the sum of those two terms to be smaller than the half of the RHS, so that we get a lower bound that is as tight as possible, leading to the constraint on δ. This results in the inequality $n\mathop{\mathbb{E}}_{\eta}\left[\sqrt{\mathop{\mathbb{E}}_{X,M}[s^\top \Sigma_T s]}\right] \ge \frac{\|R\|_2^2}{96\varepsilon}$. Consequently, obtaining the desired sample complexity lower bounds boils down to upper-bounding the term $\mathbb{E}[s^\top \Sigma_T s]$. In the applications we will see, the upper bounds will be worst-case in terms of η, so the outer expectation $\mathop{\mathbb{E}}_{\eta}[\cdot]$ will be largely ignored.

3 Lower Bounds for Private Gaussian Covariance Estimation

Here, we provide lower bounds for covariance estimation of high-dimensional Gaussians with respect to the Mahalanobis and spectral norms under the constraint of (ε, δ)-DP.

3.1 Estimation with Respect to the Mahalanobis Norm

In this section, we characterize the sample complexity of private Gaussian covariance estimation with respect to the Mahalanobis norm[2] under (ε, δ)-DP. Previous lower bounds either assumed a stricter notion of privacy (e.g., ε-DP - see [KLSU19]) or exhibited significant gaps with known upper bounds (see Section 1.1.1 of [AAK21] for further discussion). We start by stating our main result.

Theorem 3.1. *There exists a distribution $\mathcal{D}$ over $\Sigma \in \mathbb{R}^{d\times d}$ with $\mathbb{I} \preceq \Sigma \preceq 2\mathbb{I}$ such that, given $\Sigma \sim \mathcal{D}$ and $X \sim \mathcal{N}(0, \Sigma)^{\bigotimes n}$, for any $\alpha = \mathcal{O}(1)$ and any (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^{d\times d}$ with $\varepsilon \in [0, 1]$ and $\delta \leq \mathcal{O}\left(\frac{d^2}{n \ln^2\left(\frac{n}{d^2}\right)}\right)$ satisfying $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \Sigma\|_\Sigma^2\right] \leq \alpha^2$, it must hold that $n \geq \Omega\left(\frac{d^2}{\alpha\varepsilon}\right)$.*

Proof Sketch. The outline of the proof involves (1) writing the Gaussian distribution $\mathcal{N}(0, \Sigma)$ as an exponential family, (2) defining a distribution from which η is drawn, (3) showing that estimating the deviation of the natural parameter vector from the midpoint m reduces to covariance estimation (so lower bounds for the former also apply for the latter), (4) upper-bounding the term $\mathbb{E}\left[s^\top \Sigma_T s\right]$, and (5) applying Theorem 2.5 as described at the end of the previous section.

Step (1) is non-trivial because the standard way to write $\mathcal{N}(0, \Sigma)$ as an exponential family involves setting η to be equal to the *canonical flattening* of the *precision matrix* $\left(\Sigma^{-1}\right)^\flat$ (Fact F.3). However, we cannot use this representation, because Σ^{-1} is symmetric, preventing the components of η from being generated independently, as is required by Theorem 2.5. Thus, we consider an alternative parameterization. This involves identifying an upper-triangular matrix U such that $\Sigma^{-1} = U + U^\top$ and setting the parameter vector to be equal to $2U^\flat$. The two parameterizations are equivalent.

Step (2) employs a variant of a construction that has appeared in [KLSU19]. The generative process involves fixing the diagonal of the precision matrix to be constant, while the off-diagonal entries are sampled to be uniform over an interval of width $\Theta(1/d)$ centered at the origin.

Step (3) requires us to describe an algorithm which, given oracle access to (ε, δ)-DP estimate $\widehat{\Sigma}$ of Σ that is α-accurate in Mahalanobis norm, produces an (ε, δ)-DP estimate of the natural parameter vector that is $\mathcal{O}(a)$-accurate in ℓ_2-norm. The idea is to calculate $\widehat{\Sigma}^{-1}$, obtain a projection $\widetilde{\Sigma}^{-1}$ of it onto the support of the output of the algorithm described in step (2), find an upper triangular matrix $\widetilde{U}$ such that $\widetilde{\Sigma}^{-1} = U + U^\top$ and output $2\widetilde{U}^\flat - m$.

Steps (4) and (5) are mostly calculations, so we refer the reader to Appendix C.1 for them. ■

3.2 Estimation with Respect to the Spectral Norm

Here, we prove a lower bound for covariance estimation of high-dimensional Gaussians in spectral norm under the constraint of approximate DP. Employing a reduction-based approach from Mahalanobis estimation to spectral estimation, we directly leverage the results of Section 3.1.

Theorem 3.2. *There exists a distribution $\mathcal{D}$ over $\Sigma \in \mathbb{R}^{d\times d}$ with $\mathbb{I} \preceq \Sigma \preceq 2\mathbb{I}$ such that, given $\Sigma \sim \mathcal{D}$ and $X \sim \mathcal{N}(0, \Sigma)^{\bigotimes n}$, for any $\alpha = \mathcal{O}(1/\sqrt{d})$ and any (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^{d\times d}$ with $\varepsilon \in [0, 1]$ and $\delta \leq \mathcal{O}\left(\frac{d^2}{n \ln^2\left(\frac{n}{d^2}\right)}\right)$ satisfying $\mathop{\mathbb{E}}_{X,M}\left[\left\|\Sigma^{-\frac{1}{2}}\left(M(X) - \Sigma\right)\Sigma^{-\frac{1}{2}}\right\|_2^2\right] \leq \alpha^2$, it must hold that $n \geq \Omega\left(\frac{d^{1.5}}{\alpha\varepsilon}\right)$.*

[2] Recall that the Mahalanobis norm of M with respect to Σ is $\|M\|_\Sigma = \left\|\Sigma^{-1/2} M \Sigma^{-1/2}\right\|_F$. For the lower bound constructions we consider, $\Omega(1)\,\mathbb{I} \preceq \Sigma \preceq \mathcal{O}(1)\,\mathbb{I}$, and in this regime the Mahalanobis and Frobenius distances are equivalent up to a constant factor. Thus, we will often use them interchangeably. The same remark holds for the spectral norms of M and $\Sigma^{-1/2} M \Sigma^{-1/2}$ (which comes up in Theorem 3.2).

The proof is a reduction via the property $\|A\|_F \leq \sqrt{d}\,\|A\|_2$. We defer it to Appendix C.2.

4 Lower Bound for Private Heavy-Tailed Mean Estimation

In this section, we state the lower bound for mean-estimation of distributions with bounded second moment with respect to the ℓ_2 norm. Following that, we sketch the proof of the theorem, but defer the details to Appendix D.

Theorem 4.1. *Let $\mathcal{D}$ be a distribution over $\mathbb{R}^d$ with second moments bounded by 1 and unknown mean μ. Then, for any $\alpha \leq 1$, any (ε, δ)-DP mechanism with $\delta \leq \varepsilon$ that takes $X \sim \mathcal{D}^{\bigotimes n}$ as input and outputs $M(X)$ satisfying $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \mu\|_2^2\right] \leq \alpha^2$ requires $n \geq \Omega\left(\frac{d}{\alpha^2 \varepsilon}\right)$.*

Proof Sketch. We use DP Assouad's Method (Lemma A.3) to prove this theorem. To apply the lemma, we construct a family of distributions indexed by the points of the binary hypercube $\mathcal{E}_d := \{-1, 1\}^d$. Define $p = \frac{2\alpha^2}{d}$ and $t = \frac{\sqrt{d}}{\sqrt{2}\alpha}$. Then for each $v \in \mathcal{E}_d$, we define a distribution $\mathcal{D}_v$ over $\mathbb{R}^d$, such that for $X \sim \mathcal{D}_v$, the following holds.

$$X_i = \begin{cases} v_i t, & \text{with probability } p \\ 0, & \text{with probability } 1-p \end{cases}, \forall i \in [d].$$

It can be shown that the second moment of $\mathcal{D}_v$ is at most $pt^2 = 1$.

For $u, v \in \mathcal{E}_d$, we define the loss function between $\mathcal{D}_u$ and $\mathcal{D}_v$ to be $\ell(\theta(\mathcal{D}_u), \theta(\mathcal{D}_v)) = \|\theta(\mathcal{D}_u) - \theta(\mathcal{D}_v)\|_2^2$. This implies that (as defined in Lemma A.3) $\tau = 2p^2t^2 = 2p$. On rearranging the terms in Lemma A.3, the above gives us $D \geq \frac{0.04}{\varepsilon+\delta}$.

The final task is to show that there exists a coupling (X, Y) between the mixture distributions p_{+i} and p_{-i} (as defined in Lemma A.3) with the appropriate value of $D = \mathbb{E}\left[d_{\mathrm{Ham}}(X, Y)\right]$. Here, we just construct such a coupling, but verify the bound on D in the appendix. Let (X, Y) be as follows. For every row X_j in X, if $X_j^i \neq 0$, then $Y_j^i = -X_j^i$, otherwise $Y_j^i = X_j^i$. For all other coordinates $k \neq i$, $Y_j^k = X_j^k$. We show that in this case, $D = np$. This gives us $\frac{D}{p} = n \geq \frac{d}{50(\varepsilon+\delta)\alpha^2}$, which proves the theorem. ■

5 Conclusions and Open Problems

In this work, we gave nearly tight lower bounds for private covariance estimation with respect to the Frobenius and spectral norms under the Gaussian distribution and for mean estimation with respect to the ℓ_2-norm under distributions with bounded second moments. This constitutes important progress in the lower bounds literature under approximate differential privacy. On a technical level, we significantly generalized the fingerprinting method. An obvious question is whether the fingerprinting method can be further generalized and what are its limitations. A second obvious question is to generalize Theorem 4.1 for higher-order moments.

Another question involves expanding upon the results of Theorems 3.1 and 3.2 to hold in the regimes $\alpha = \mathcal{O}(\sqrt{d})$ and $\alpha = \mathcal{O}(1)$, respectively. For the Frobenius estimation, the condition $\alpha = \mathcal{O}(1)$ is a consequence of the distribution over covariance matrices that we used to prove the result, whereas for spectral estimation the condition $\alpha = \mathcal{O}(1/\sqrt{d})$ is an artifact of our reduction-based approach. Especially in the latter case, it would be interesting to come up with a "direct" way to obtain the lower bound, instead of reducing from Frobenius estimation.

Finally, an interesting open question may be to obtain lower bounds for the problems considered here in the *high-probability regime*. The guarantees in our lower bounds are expressed in terms of the MSE and in Section 1.1 we sketched how to convert them to constant probability guarantees. It would be interesting come up with techniques to prove tight lower bounds when the probability of success is $1 - \beta$ for $\beta \in (0, 1)$. This was done in the context of pure differential privacy in [HKM21].

References

[AAK21] Ishaq Aden-Ali, Hassan Ashtiani, and Gautam Kamath. On the sample complexity of privately learning unbounded high-dimensional gaussians. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, ALT '21, pages 185–216. JMLR, Inc., 2021.

[AAL21] Ishaq Aden-Ali, Hassan Ashtiani, and Christopher Liaw. Privately learning mixtures of axis-aligned gaussians. In *Advances in Neural Information Processing Systems 34*, NeurIPS '21. Curran Associates, Inc., 2021.

[ADK19] Brendan Avent, Yatharth Dubey, and Aleksandra Korolova. The power of the hybrid model for mean estimation. *Proceedings on Privacy Enhancing Technologies*, 2020(4):48–68, 2019.

[AL21] Hassan Ashtiani and Christopher Liaw. Private and polynomial time algorithms for learning Gaussians and beyond. *arXiv preprint arXiv:2111.11320*, 2021.

[AMB19] Marco Avella-Medina and Victor-Emmanuel Brunel. Differentially private sub-Gaussian location estimators. *arXiv preprint arXiv:1906.11923*, 2019.

[ASZ18] Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private testing of identity and closeness of discrete distributions. In *Advances in Neural Information Processing Systems 31*, NeurIPS '18, pages 6878–6891. Curran Associates, Inc., 2018.

[ASZ21] Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private assouad, fano, and le cam. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, ALT '21, pages 48–78. JMLR, Inc., 2021.

[BD14] Rina Foygel Barber and John C Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.

[BDKU20] Sourav Biswas, Yihe Dong, Gautam Kamath, and Jonathan Ullman. Coinpress: Practical private mean and covariance estimation. In *Advances in Neural Information Processing Systems 33*, NeurIPS '20, pages 14475–14485. Curran Associates, Inc., 2020.

[BGS+21] Gavin Brown, Marco Gaboardi, Adam Smith, Jonathan Ullman, and Lydia Zakynthinou. Covariance-aware private mean estimation without private covariance estimation. *arXiv preprint arXiv:2106.13329*, 2021.

[BKSW19] Mark Bun, Gautam Kamath, Thomas Steinke, and Zhiwei Steven Wu. Private hypothesis selection. In *Advances in Neural Information Processing Systems 32*, NeurIPS '19, pages 156–167. Curran Associates, Inc., 2019.

[BS98] Dan Boneh and James Shaw. Collusion-secure fingerprinting for digital data. *IEEE Transactions on Information Theory*, 44(5):1897–1905, 1998.

[BS19] Mark Bun and Thomas Steinke. Average-case averages: Private algorithms for smooth sensitivity and mean estimation. In *Advances in Neural Information Processing Systems 32*, NeurIPS '19, pages 181–191. Curran Associates, Inc., 2019.

[BST14] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '14, pages 464–473, Washington, DC, USA, 2014. IEEE Computer Society.

[BSU17] Mark Bun, Thomas Steinke, and Jonathan Ullman. Make up your mind: The price of online queries in differential privacy. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 1306–1325, Philadelphia, PA, USA, 2017. SIAM.

[BUV14] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Proceedings of the 46th Annual ACM Symposium on the Theory of Computing*, STOC '14, pages 1–10, New York, NY, USA, 2014. ACM.

[CWZ19] T. Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *arXiv preprint arXiv:1902.04495*, 2019.

[DFM+20] Wenxin Du, Canyon Foot, Monica Moniot, Andrew Bray, and Adam Groce. Differentially private confidence intervals. *arXiv preprint arXiv:2001.02285*, 2020.

[DHS15] Ilias Diakonikolas, Moritz Hardt, and Ludwig Schmidt. Differentially private learning of structured discrete distributions. In *Advances in Neural Information Processing Systems 28*, NIPS '15, pages 2566–2574. Curran Associates, Inc., 2015.

[DJW13] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Local privacy and statistical minimax rates. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '13, pages 429–438, Washington, DC, USA, 2013. IEEE Computer Society.

[DKK+16] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '16, pages 655–664, Washington, DC, USA, 2016. IEEE Computer Society.

[DKM+06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Proceedings of the 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, EUROCRYPT '06, pages 486–503, Berlin, Heidelberg, 2006. Springer.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography*, TCC '06, pages 265–284, Berlin, Heidelberg, 2006. Springer.

[DSS+15] Cynthia Dwork, Adam Smith, Thomas Steinke, Jonathan Ullman, and Salil Vadhan. Robust traceability from trace amounts. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '15, pages 650–669, Washington, DC, USA, 2015. IEEE Computer Society.

[DTTZ14] Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze Gauss: Optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the 46th Annual ACM Symposium on the Theory of Computing*, STOC '14, pages 11–20, New York, NY, USA, 2014. ACM.

[Ger31] S. Gershgorin. Über die abgrenzung der Eigenwerte einer matrix. *Izv. Akad. Nauk. SSSR Ser. Mat.*, 1:749–754, 1931.

[HKM21] Samuel B Hopkins, Gautam Kamath, and Mahbod Majid. Efficient mean estimation with pure differential privacy via a sum-of-squares exponential mechanism. *arXiv preprint arXiv:2111.12981*, 2021.

[HLY21] Ziyue Huang, Yuting Liang, and Ke Yi. Instance-optimal mean estimation under differential privacy. In *Advances in Neural Information Processing Systems 34*, NeurIPS '21. Curran Associates, Inc., 2021.

[HT10] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd Annual ACM Symposium on the Theory of Computing*, STOC '10, pages 705–714, New York, NY, USA, 2010. ACM.

[Jor10] Michael Jordan. Lecture notes for Bayesian modeling and inference, 2010.

[KLSU19] Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. In *Proceedings of the 32nd Annual Conference on Learning Theory*, COLT '19, pages 1853–1902, 2019.

[KLZ21] Gautam Kamath, Xingtu Liu, and Huanyu Zhang. Improved rates for differentially private stochastic convex optimization with heavy-tailed data. *arXiv preprint arXiv:2106.01336*, 2021.

[KMS+21] Gautam Kamath, Argyris Mouzakis, Vikrant Singhal, Thomas Steinke, and Jonathan Ullman. A private and computationally-efficient estimator for unbounded gaussians. *arXiv preprint arXiv:2111.04609*, 2021.

[KSSU19] Gautam Kamath, Or Sheffet, Vikrant Singhal, and Jonathan Ullman. Differentially private algorithms for learning mixtures of separated Gaussians. In *Advances in Neural Information Processing Systems 32*, NeurIPS '19, pages 168–180. Curran Associates, Inc., 2019.

[KSU20] Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In *Proceedings of the 33rd Annual Conference on Learning Theory*, COLT '20, pages 2204–2235, 2020.

[KU20] Gautam Kamath and Jonathan Ullman. A primer on private statistics. *arXiv preprint arXiv:2005.00010*, 2020.

[KV18] Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. In *Proceedings of the 9th Conference on Innovations in Theoretical Computer Science*, ITCS '18, pages 44:1–44:9, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[Li19] Jerry Li. Lecture 18: Private covariance estimation in high dimensions. `https://jerryzli.github.io/robust-ml-fall19/lec18.pdf`, December 2019. Lecture notes.

[LKO21] Xiyang Liu, Weihao Kong, and Sewoong Oh. Differential privacy and robust statistics in high dimensions. *arXiv preprint arXiv:2111.06578*, 2021.

[LM00] Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics*, 28(5):1302–1338, 2000.

[LSA+21] Daniel Levy, Ziteng Sun, Kareem Amin, Satyen Kale, Alex Kulesza, Mehryar Mohri, and Ananda Theertha Suresh. Learning with user-level privacy. In *Advances in Neural Information Processing Systems 34*, NeurIPS '21. Curran Associates, Inc., 2021.

[LSY+20] Yuhan Liu, Ananda Theertha Suresh, Felix Yu, Sanjiv Kumar, and Michael Riley. Learning discrete distributions: User vs item-level privacy. In *Advances in Neural Information Processing Systems 33*, NeurIPS '20. Curran Associates, Inc., 2020.

[SU15] Thomas Steinke and Jonathan Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *Proceedings of the 28th Annual Conference on Learning Theory*, COLT '15, pages 1588–1628, 2015.

[SU17a] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *The Journal of Privacy and Confidentiality*, 7(2):3–22, 2017.

[SU17b] Thomas Steinke and Jonathan Ullman. Tight lower bounds for differentially private selection. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '17, pages 552–563, Washington, DC, USA, 2017. IEEE Computer Society.

[Tar08] Gabor Tardos. Optimal probabilistic fingerprint codes. *Journal of the ACM*, 55(2), 2008.

[TVGZ20] Christos Tzamos, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Ilias Zadik. Optimal private median estimation under minimal distributional assumptions. In *Advances in Neural Information Processing Systems 33*, NeurIPS '20, pages 3301–3311. Curran Associates, Inc., 2020.

[WXDX20] 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 '20, pages 10081–10091. JMLR, Inc., 2020.

[Yu97] Bin Yu. Assouad, fano, and le cam. In *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*, pages 423–435. Springer New York, 1997.

[ZKKW20] 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 '20, pages 11129–11140. JMLR, Inc., 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendices A-G.

3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Preliminaries

Notation. We use $[n]$ to denote the *initial segment of natural numbers* $\{1, 2, \ldots, n\}$ and $\mathbb{R}_+$ to denote the non-negative real numbers. Given any distribution $\mathcal{D}$, $\mathcal{D}^{\bigotimes n}$ denotes the *product measure* where each *marginal distribution* is $\mathcal{D}$. Thus, if we are given n independent samples from $\mathcal{D}$, we write $(X_1, \ldots, X_n) \sim \mathcal{D}^{\bigotimes n}$. An alternative notation that may be used when indexing samples is via superscripts instead of subscripts e.g., X^i. Also, depending on the context, we may use capital Latin characters like X to denote either an individual sample from a distribution or a collection of samples $X = (X_1, \ldots, X_n)$. The *unit ℓ_2-ball* in d-dimensions that is centered at the origin is denoted by $\mathbb{S}^{d-1}$. Finally, for a distribution $\mathcal{D}$ over $\mathbb{R}^d$, we say that it has k-th moments bounded by C if, for any $v \in \mathbb{S}^{d-1}$, $\mathop{\mathbb{E}}_{X\sim\mathcal{D}}\left[\left|\left\langle X - \mathop{\mathbb{E}}_{Y\sim\mathcal{D}}[Y], v\right\rangle\right|^k\right] \leq C$.

Linear Algebra Preliminaries. Given a vector $v \in \mathbb{R}^d$, v_i refers to its i-th component and $v_{-i} \in \mathbb{R}^{d-1}$ describes the vector one gets by removing the i-th component from v. We use $\vec{1}_d$ to refer to the all 1s vector in d dimensions. For a pair of vectors $x, y \in \mathbb{R}^d$, we will write $x \leq y$ if and only if $x_i \leq y_i, \forall i \in [n]$. Given a matrix $M \in \mathbb{R}^{d\times d}$, we denote the element at the i-th row and the j-th column by M_{ij} or $(M)_{ij}$. Also, we denote the *canonical flattening* operation of M into a vector v by $v = M^\flat \in \mathbb{R}^{d^2}$ and the inverse operation by $M = v^\# \in \mathbb{R}^{d\times d}$. We denote its *spectral norm* by $\|M\|_2 \coloneqq \sup_{v\in\mathbb{S}^{d-1}} \|Mv\|_2$. For symmetric matrices M, this reduces to $\|M\|_2 = \sup_{v\in\mathbb{S}^{d-1}} \left|v^\top M v\right|$, which is known to be equal to the largest eigenvalue of M (in absolute value). Also, we denote the *trace* of M by $tr\left(M\right)$ and its *Frobenius norm* by $\|M\|_F \coloneqq \sqrt{tr\left(M^\top M\right)} = \sqrt{\sum_{i=1}^d \sum_{j=1}^d M_{ij}^2}$. For symmetric matrices, this is equal to the square root of the sum of squares of its eigenvalues. Moreover, we consider the *Mahalanobis norm* (with respect to a non-degenerate covariance matrix Σ), which is denoted by $\|M\|_\Sigma \coloneqq \left\|\Sigma^{-\frac{1}{2}} M \Sigma^{-\frac{1}{2}}\right\|_F$. Furthermore, we use the notation $\langle\cdot,\cdot\rangle$ to denote *inner products*, either between vectors $x, y \in \mathbb{R}^d$

(where $\langle x, y\rangle = x^\top y$) or between matrices $A, B \in \mathbb{R}^{d\times d}$ (where $\langle A, B\rangle = tr\left(AB^\top\right)$). Additionally, given a pair of symmetric matrices $A, B \in \mathbb{R}^d$, we write $A \succeq B$ if and only if $x^\top (A - B) x \geq 0, \forall x \in \mathbb{R}^d$. Finally, given matrices $A \in \mathbb{R}^{n\times m}$ and $B \in \mathbb{R}^{k\times \ell}$, the *Kronecker product* of A with B is denoted by $A \otimes B \in \mathbb{R}^{(nk)\times(m\ell)}$. More linear algebra facts appear in Appendix E.

A.1 Privacy Preliminaries

We define differential privacy and state its closure under post-processing property.

Definition A.1 (Differential Privacy (DP) [DMNS06])**.** A randomized algorithm $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies (ε, δ)-differential privacy ((ε, δ)-DP) if for every pair of neighboring datasets $X, X' \in \mathcal{X}^n$ (i.e., datasets that differ in exactly one entry), we have:

$$\mathbb{P}\left[M\left(X\right) \in Y\right] \leq e^{\varepsilon} \cdot \mathbb{P}\left[M\left(X'\right) \in Y\right] + \delta, \forall Y \subseteq \mathcal{Y}.$$

When $\delta = 0$, we say that M satisfies ε-differential privacy or pure differential privacy.

Lemma A.2 (Post Processing [DMNS06])**.** *If $M : \mathcal{X}^n \to \mathcal{Y}$ is (ε, δ)-DP, and $P : \mathcal{Y} \to \mathcal{Z}$ is any randomized function, then the algorithm $P \circ M$ is (ε, δ)-DP.*

Finally, we discuss a recent technique from [ASZ21] for proving lower bounds for estimation problems under approximate DP. It is a private version of the well-known Assouad Lemma [Yu97]. Let $\mathcal{P}$ be a family of distributions over $\mathcal{X}^n$, where $\mathcal{X}$ is the data universe and n is the number of samples. Let $\theta : \mathcal{P} \to \Theta$ be the parameter of the distribution that we want to estimate, and let $\ell : \Theta \to \Theta$ be the pseudo-metric (loss function) for estimating θ. We define the risk of an estimator $\hat{\theta}$ as $\max\limits_{p\in\mathcal{P}} \mathop{\mathbb{E}}\limits_{X\sim p}\left[\ell\left(\hat{\theta}\left(X\right), \theta\left(p\right)\right)\right]$. We use the following notion of minimax risk for (ε, δ)-DP estimators:

$$R\left(\mathcal{P}, \ell, \varepsilon, \delta\right) = \min_{\hat{\theta} \text{ is } (\varepsilon,\delta)\text{-DP}} \max_{p\in\mathcal{P}} \mathop{\mathbb{E}}\limits_{X\sim P}\left[\ell\left(\hat{\theta}\left(X\right), \theta\left(p\right)\right)\right].$$

Let $\mathcal{V} \subseteq \mathcal{P}$ be a subset of distributions indexed by the points in the hypercube $\mathcal{E}_d = \{-1, 1\}^d$. Suppose there exists a $\tau \in \mathbb{R}$, such that for all $u, v \in \mathcal{E}_d$, $\ell\left(\theta\left(p_u\right), \theta\left(p_v\right)\right) \geq 2\tau \cdot \sum\limits_{i=1}^{d} \mathbb{1}\left\{u_i \neq v_i\right\}$. For every coordinate $i \in [d]$, define the following mixture distributions:

$$p_{+i} = \frac{2}{|\mathcal{E}_d|} \sum_{e\in\mathcal{E}_d : e_i = 1} p_e \text{ and } p_{-i} = \frac{2}{|\mathcal{E}_d|} \sum_{e\in\mathcal{E}_d : e_i = -1} p_e.$$

The DP Assouad's method provides a lower bound for the minimax risk by considering the problem of distinguishing between p_{+i} and p_{-i}.

Lemma A.3 (DP Assouad's Method [ASZ21])**.** *For all $i \in [d]$, let $\phi_i : \mathcal{X}^n \to \{-1, 1\}$ be a binary classifier. Then we have the following:*

$$R\left(\mathcal{P}, \ell, \varepsilon, \delta\right) \geq \frac{\tau}{2} \cdot \sum_{i=1}^{d} \min_{\phi_i \text{ is } (\varepsilon,\delta)\text{-DP}} \left(\mathop{\mathbb{P}}\limits_{X\sim p_{+i}}\left[\phi_i\left(X\right) \neq 1\right] + \mathop{\mathbb{P}}\limits_{X\sim p_{-i}}\left[\phi_i\left(X\right) \neq -1\right]\right).$$

Moreover, if $\forall i \in [d]$, there exists a coupling (X, Y) between p_{+i} and p_{-i} with $\mathbb{E}\left[d_{\mathrm{Ham}}\left(X, Y\right)\right] \leq D$, then:

$$R\left(\mathcal{P}, \ell, \varepsilon, \delta\right) \geq \frac{d\tau}{2} \cdot \left(0.9e^{-10\varepsilon D} - 10\delta D\right).$$

A.2 Exponential Families

In this work, we make frequent use of exponential families due to their expressive power. This gives us a unified way of dealing with multiple commonly used distributions (e.g., Gaussians, Bernoullis). We first state the definition of exponential families.

Definition A.4 (Exponential Family)**.** Let $S \subseteq \mathbb{R}^d$ and $h : S \to \mathbb{R}_+, T : S \to \mathbb{R}^k$. Given $\eta \in \mathbb{R}^k$, we say that the density function p_η belongs to the *exponential family* $\mathcal{E}\left(T, h\right)$ if its density can be written in the form:

$$p_\eta\left(x\right) = h\left(x\right) \exp\left(\eta^\top T\left(x\right) - Z\left(\eta\right)\right), \forall x \in S,$$

where:

$$Z(\eta) = \ln\left(\int_S h(x)\exp\left(\eta^\top T(x)\right)\,dx\right).$$

The functions h and T are referred to as the *carrier measure* and the *sufficient statistics* of the family, respectively. Additionally, Z is known as the *log-partition function* and η is the *natural parameter vector*. Finally, we denote the *range of natural parameters* by $\mathcal{H} \subseteq \mathbb{R}^k$, which is the set of values of η for which the log-partition function is well-defined $(Z(\eta) < \infty)$.

We appeal to the following properties of exponential families.

Proposition A.5 (see [Jor10])**.** *Let $\mathcal{E}(T, h)$ be an exponential family parameterized by $\eta \in \mathbb{R}^k$. Then, the following hold:*

- *For all $\eta \in \mathcal{H}$, it holds that:*
$$\mu_T \coloneqq \mathop{\mathbb{E}}_{X\sim p_\eta}[T(X)] = \nabla Z(\eta).$$
- *For all $\eta \in \mathcal{H}$, it holds that:*
$$\Sigma_T \coloneqq \mathop{\mathbb{E}}_{X\sim p_\eta}\left[(T(X)-\mu_T)(T(X)-\mu_T)^\top\right] = \nabla^2 Z(\eta).$$
- *For all $\eta \in \mathcal{H}$ and $s \in \mathbb{R}^k$ such that $\eta + s \in \mathcal{H}$, it holds that:*
$$\mathop{\mathbb{E}}_{X\sim p_\eta}\left[\exp\left(s^\top T(X)\right)\right] = \exp\left(Z(\eta+s) - Z(\eta)\right).$$

B Omitted Proofs from Section 2

B.1 Omitted Proofs from Section 2.2

Proof of Lemma 2.2. We have:

$$\begin{aligned}
\mathop{\mathbb{E}}_{\eta,X}\left[Z + \|f(X) - (\eta - m)\|_2^2\right] &= \mathop{\mathbb{E}}_{\eta,X}\left[\sum_{j\in[k]}\left\{Z^j + [f_j(X) - (\eta_j - m_j)]^2\right\}\right] \\
&= \sum_{j\in[k]}\mathop{\mathbb{E}}_{\eta,X}\left[Z^j + [f_j(X) - (\eta_j - m_j)]^2\right] \\
&= \sum_{j\in[k]}\mathop{\mathbb{E}}_{\eta_{-j}}\left[\mathop{\mathbb{E}}_{\eta_j,X}\left[Z^j + [f_j(X) - (\eta_j - m_j)]^2\right]\right] \\
&\overset{(a)}{\geq} \sum_{j\in[k]}\mathop{\mathbb{E}}_{\eta_{-j}}\left[\frac{R_j^2}{12}\right] \\
&= \frac{\|R\|_2^2}{12},
\end{aligned}$$

where (a) is by a direct application of Lemma 2.1. ■

B.2 Omitted Proofs from Section 2.3

This appendix is dedicated to proving Lemma 2.3 and Theorem 2.5.

To prove the former, in addition to the correlation terms defined in Section 2.1, we introduce one more correlation term. Let X_i' be an independently drawn sample from p_η and let $X_{\sim i}$ be the dataset that is the same as X with X_i replaced by X_i'. We define:

$$s_j \coloneqq \left[\frac{R_j^2}{4} - (\eta_j - m_j)^2\right]\left[M_j(X_{\sim i}) - (\eta_j - m_j)\right], \forall j \in [k], \tag{4}$$

$$\widetilde{Z}_i \coloneqq \sum_{j\in[k]} s_j\left(T_j(X_i) - \mu_{T,j}\right). \tag{5}$$

This measures the effect that replacing an individual sample has on the output of the estimator. Observe that, conditioning on η, $X_{\sim i}$ and X_i are independent. This implies that $\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i\right] = 0$ and $\mathop{\mathrm{Var}}_{X_{\sim i},X_i,M}\left(\widetilde{Z}_i\right) = \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i^2\right]$. Additionally, calculating the variance of $\widetilde{Z}_i$ requires us to reason about the covariance matrix of the sufficient statistics Σ_T, since the calculation involves terms of the form $\left(T_{j_1}\left(X_i\right) - \mu_{T,j_1}\right)\left(T_{j_2}\left(X_i\right) - \mu_{T,j_2}\right), j_1, j_2 \in [k]$.

The previous remarks are summarized in the following lemma.

Lemma B.1. *Conditioning on η, it holds that:*

$$\begin{aligned}
&\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i\right] = 0,\\
&\mathop{\mathrm{Var}}_{X_{\sim i},X_i,M}\left(\widetilde{Z}_i\right) = \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i^2\right] = \mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma_T s\right].
\end{aligned}$$

Proof. We start by computing the mean. We have:

$$\begin{aligned}
\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i\right] &= \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\sum_{j\in[k]} s_j\left(T_j\left(X_i\right) - \mu_{T,j}\right)\right]\\
&= \sum_{j\in[k]} \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[s_j\left(T_j\left(X_i\right) - \mu_{T,j}\right)\right]\\
&\overset{(a)}{=} \sum_{j\in[k]} \mathop{\mathbb{E}}_{X_{\sim i},M}\left[s_j\right] \cancelto{0}{\mathop{\mathbb{E}}_{X_i}\left[T_j\left(X_i\right) - \mu_{T,j}\right]}\\
&= 0,
\end{aligned}$$

where (a) relies on the independence of $s_j = \left[\frac{R_j^2}{4} - \left(\eta_j - m_j\right)^2\right]\left[M_j\left(X_{\sim i}\right) - \left(\eta_j - m_j\right)\right]$ and $T_j\left(X_i\right) - \mu_{T,j}$ conditioned on η.

We now proceed with the variance:

$$\begin{aligned}
\mathop{\mathrm{Var}}_{X_{\sim i},X_i,M}\left(\widetilde{Z}_i\right) &= \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i^2\right]\\
&= \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\sum_{j_1\in[k]}\sum_{j_2\in[k]} s_{j_1}s_{j_2}\left(T_{j_1}\left(X_i\right) - \mu_{T,j_1}\right)\left(T_{j_2}\left(X_i\right) - \mu_{T,j_2}\right)\right]\\
&= \sum_{j_1\in[k]}\sum_{j_2\in[k]} \mathop{\mathbb{E}}_{X_{\sim i},M}\left[s_{j_1}s_{j_2}\right] \mathop{\mathbb{E}}_{X_i}\left[\left(T_{j_1}\left(X_i\right) - \mu_{T,j_1}\right)\left(T_{j_2}\left(X_i\right) - \mu_{T,j_2}\right)\right]\\
&\overset{(a)}{=} \sum_{j_1\in[k]}\sum_{j_2\in[k]} \mathop{\mathbb{E}}_{X_{\sim i},M}\left[s_{j_1}s_{j_2}\right]\Sigma_{T,j_1j_2}\\
&= \mathop{\mathbb{E}}_{X_{\sim i},M}\left[\sum_{j_1\in[k]}\sum_{j_2\in[k]} s_{j_1}\Sigma_{T,j_1j_2}s_{j_2}\right]\\
&= \mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top\Sigma_T s\right].
\end{aligned}$$

where (a) is by the second part of Proposition A.5. ■

Now, we show how to upper-bound $\mathop{\mathbb{E}}_{X,M}\left[Z_i\right]$ using properties of differential privacy. We appeal to the definition of DP to relate $\mathop{\mathbb{E}}_{X,M}\left[Z_i\right]$ with $\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i\right]$ and $\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i^2\right]$.

Proof of Lemma 2.3. Consider the random variables $Z_{i,+} = \max\{Z_i, 0\}$ and $Z_{i,-} = -\min\{Z_i, 0\}$. These random variables are non-negative, so for any $T > 0$, we have:

$$\begin{aligned}
\mathop{\mathbb{E}}_{X,M}[Z_i] &= \mathop{\mathbb{E}}_{X,M}[Z_{i,+}] - \mathop{\mathbb{E}}_{\eta,X,M}[Z_{i,-}] \\
&= \int_0^\infty \mathop{\mathbb{P}}_{X,M}[Z_{i,+} > t]\, dt - \int_0^\infty \mathop{\mathbb{P}}_{X,M}[Z_{i,-} > t]\, dt \\
&= \int_0^\infty \mathop{\mathbb{P}}_{X,M}[Z_i > t]\, dt - \int_0^\infty \mathop{\mathbb{P}}_{X,M}[Z_i < -t]\, dt \\
&= \int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i > t]\, dt - \int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i < -t]\, dt \\
&+ \int_T^\infty \mathop{\mathbb{P}}_{X,M}[Z_i > t]\, dt - \int_T^\infty \mathop{\mathbb{P}}_{X,M}[Z_i < -t]\, dt
\end{aligned} \tag{6}$$

We focus our attention to the first two terms of (6). We consider the sets $S_{+,X_i}, S_{-,X_i} \subseteq \mathbb{R}^k$:

$$\begin{aligned}
S_{+,X_i} &\coloneqq \left\{ Y \in \mathbb{R}^k : \sum_{j\in[k]} \left[\frac{R_j^2}{4} - (\eta_j - m_j)^2\right] [Y_j - (\eta_j - m_j)] \left(T_j(X_i) - \mu_{T,j}\right) \geq t \right\}, \\
S_{-,X_i} &\coloneqq \left\{ Y \in \mathbb{R}^k : \sum_{j\in[k]} \left[\frac{R_j^2}{4} - (\eta_j - m_j)^2\right] [Y_j - (\eta_j - m_j)] \left(T_j(X_i) - \mu_{T,j}\right) \leq -t \right\}.
\end{aligned}$$

By a direct application of the definition of DP on S_{+,X_i} and S_{-,X_i}[3] for the neighboring datasets X and $X_{\sim i}$ we have:

$$\begin{aligned}
&\int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i > t]\, dt - \int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i < -t]\, dt \\
&= \int_0^T \mathop{\mathbb{P}}_{X,M}[M(X) \in S_{+,X_i}]\, dt - \int_0^T \mathop{\mathbb{P}}_{X,M}[M(X) \in S_{-,X_i}]\, dt \\
&\leq \int_0^T \left(e^{\varepsilon} \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}[M(X_{\sim i}) \in S_{+,X_i}] + \delta \right) dt \\
&- \int_0^T e^{-\varepsilon} \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}[M(X_{\sim i}) \in S_{-,X_i}] - \delta \right) dt \\
&\leq \int_0^T \left(e^{\varepsilon} \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] + \delta \right) dt - \int_0^T e^{-\varepsilon} \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right] - \delta \right) dt \\
&= \left(1 + e^{-\varepsilon}\right) \delta T + e^{\varepsilon} \int_0^T \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] dt - e^{-\varepsilon} \int_0^T \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right] dt.
\end{aligned}$$

[3]These are random subsets of $\mathbb{R}^k$, since they depend on the random vector X_i (η is assumed to be fixed). This does not affect the way the definition of privacy is applied.

Under the assumption that $\varepsilon \in [0,1]$, it holds that $e^{\varepsilon} \leq 1 + 2\varepsilon$ and $1 \geq e^{-\varepsilon} \geq 1 - \varepsilon \geq 1 - 2\varepsilon$. Combined with the above, this yields:

$$
\begin{aligned}
&\int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i > t]\ dt - \int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i < -t]\ dt \\
&\leq 2\delta T + (2\varepsilon + 1)\int_0^T \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right]\ dt + (2\varepsilon - 1)\int_0^T \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\ dt \\
&= 2\delta T + 2\varepsilon \int_0^{\infty}\left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] + \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&+ \int_0^T \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&= 2\delta T + 2\varepsilon \int_0^T \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right| > t\right]\ dt + \int_0^T \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&\leq 2\delta T + 2\varepsilon \int_0^{\infty} \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right| > t\right]\ dt + \int_0^T \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&= 2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right|\right] + \int_0^T \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&= 2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right|\right] + \int_0^{\infty} \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&- \int_T^{\infty} \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt.
\end{aligned}
$$

Defining $\widetilde{Z}_{i,+} = \max\left\{\widetilde{Z}_i, 0\right\}$ and $\widetilde{Z}_{i,-} = -\min\left\{\widetilde{Z}_i, 0\right\}$, we get:

$$
\begin{aligned}
&\int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i > t]\ dt - \int_0^T \mathop{\mathbb{P}}_{X,M}[Z_i < -t]\ dt \\
&\leq 2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right|\right] + \left(\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_{i,+}\right] - \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_{i,-}\right]\right) \\
&- \int_T^{\infty} \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&\leq 2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right|\right] + \cancelto{0}{\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i\right]} \\
&- \int_T^{\infty} \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt \\
&= 2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right|\right] - \int_T^{\infty} \left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt, \qquad (7)
\end{aligned}
$$

where we appealed to Lemma B.1 along the way.

The second version of Fact F.4 for $X = \left|\widetilde{Z}_i\right|, Y = 1, p = q = 2$, and Lemma B.1 yield:

$$
\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\left|\widetilde{Z}_i\right|\right] \leq \sqrt{\mathop{\mathbb{E}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i^2\right]} = \sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^{\top}\Sigma_T s]}. \qquad (8)
$$

Substituting (7) and (8) into (6), we get:

$$
\begin{aligned}
\mathop{\mathbb{E}}_{X,M}[Z_i] &\leq 2\delta T + 2\varepsilon\sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^{\top}\Sigma_T s]} + \int_T^{\infty}\left(\mathop{\mathbb{P}}_{X,M}[Z_i > t] - \mathop{\mathbb{P}}_{X,M}[Z_i < -t]\right)\ dt \\
&+ \int_T^{\infty}\left(\mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right] - \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i > t\right]\right)\ dt \\
&\leq 2\delta T + 2\varepsilon\sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^{\top}\Sigma_T s]} + \int_T^{\infty}\left(\mathop{\mathbb{P}}_{X,M}[Z_i > t] + \mathop{\mathbb{P}}_{X_{\sim i},X_i,M}\left[\widetilde{Z}_i < -t\right]\right)\ dt. \qquad (9)
\end{aligned}
$$

We now focus our attention on the last term of (9). We note that:

$$
\begin{aligned}
\{Z_i > t\} &\subseteq \{|Z_i| > t\} \\
&= \left\{ \left| \sum_{j\in[k]} \left[\frac{R_j^2}{4} - (\eta_j - m_j)^2 \right] (M_j(X) - (\eta_j - m_j)) (T_j(X_i) - \mu_{T,j}) \right| > t \right\} \\
&= \{|\langle s, T(X_i) - \mu_T \rangle| > t\} \\
&\overset{(a)}{\subseteq} \{\|s\|_\infty \|T(X_i) - \mu_T\|_1 > t\} \\
&\overset{(b)}{\subseteq} \left\{ \frac{\|R\|_\infty^3}{4} \left| \sum_{j\in[k]} (T_j(X_i) - \mu_{T,j}) \right| > t \right\},
\end{aligned}
$$

where (a) is by the first version of Fact F.4 for $p = \infty$ and $q = 1$ and (b) uses the fact that $|M_j(X) - (\eta_j - m_j)| \leq R_j, \forall j \in [k]$.

The same upper bound holds for the event $\left\{\widetilde{Z}_i < -t\right\}$. Thus, substituting back to (9), we get:

$$
\mathop{\mathbb{E}}_{X,M}[Z_i] \leq 2\delta T + 2\varepsilon \sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^\top \Sigma_T s]} + 2 \int_T^\infty \mathop{\mathbb{P}}_{X_i} \left[\left| \sum_{j\in[k]} (T_j(X_i) - \mu_{T,j}) \right| > \frac{4t}{\|R\|_\infty^3} \right] dt.
$$

This completes our proof. ■

We conclude this appendix with the complete proof of Theorem 2.5.

Proof of Theorem 2.5. By Lemma 2.2 and (2), we have:

$$
\mathop{\mathbb{E}}_{\eta,X,M}[Z] \geq \frac{\|R\|_2^2}{24}.
$$

Additionally, by Lemma 2.3, we have:

$$
\begin{aligned}
\mathop{\mathbb{E}}_{\eta,X,M}[Z] &= \sum_{i\in[n]} \mathop{\mathbb{E}}_{\eta,X,M}[Z_i] \\
&= \sum_{i\in[n]} \mathop{\mathbb{E}}_{\eta} \left[\mathop{\mathbb{E}}_{X,M}[Z_i] \right] \\
&\leq \sum_{i\in[n]} \mathop{\mathbb{E}}_{\eta} \left[2\delta T + 2\varepsilon \sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^\top \Sigma_T s]} + 2 \int_T^\infty \mathop{\mathbb{P}}_{X_i} \left[\left| \sum_{j\in[k]} (T_j(X_i) - \mu_{T,j}) \right| > \frac{4t}{\|R\|_\infty^3} \right] \right] \\
&\leq n \left(2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{\eta} \left[\sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}[s^\top \Sigma_T s]} \right] + 2 \int_T^\infty \mathop{\mathbb{P}}_{X_i} \left[\left| \sum_{j\in[k]} (T_j(X_i) - \mu_{T,j}) \right| > \frac{4t}{\|R\|_\infty^3} \right] \right).
\end{aligned}
$$

Combining the two, we get the desired result. ■

C Omitted Proofs from Section 3

C.1 Omitted Proofs from Section 3.1

The first step of the proof is to write the Gaussian $\mathcal{N}(0,\Sigma)$ as an exponential family. By Fact F.3, we have that the density can be written in the form:

$$\begin{aligned} p_\eta(x) &= h(x)\exp\left(\eta^\top T(x) - Z(\eta)\right), \\ h(x) &= 1, \\ T(x) &= -\frac{1}{2}\left(xx^\top\right)^\flat, \\ \eta &= \left(\Sigma^{-1}\right)^\flat, \\ Z(\eta) &= \frac{d}{2}\ln(2\pi) - \frac{1}{2}\ln\left(\det\left(\eta^\#\right)\right). \end{aligned}$$

Observe that the natural parameter vector η is the flattening of the *precision matrix* Σ^{-1}. However, we cannot use this representation, because Σ^{-1} is symmetric. Indeed, this implies that the components of η cannot be generated independently, as is required by Theorem 2.5. For that reason, we consider an alternative parameterization.

Lemma C.1. *Let $\mathcal{N}(0,\Sigma)$ be a d-dimensional Gaussian with unknown covariance Σ. Then, it belongs to an exponential family with $h_0 \equiv h, \eta_0 = 2U^\flat, T_0 \equiv T, Z_0(\eta_0) = \frac{d}{2}\ln(2\pi) - \frac{1}{2}\ln\left(\det\left(\frac{\eta_0^\# + \left(\eta_0^\#\right)^\top}{2}\right)\right)$ where $U \in \mathbb{R}^{d\times d}$ is an upper triangular matrix with:*

$$U_{ij} := \begin{cases} \frac{1}{2}\left(\Sigma^{-1}\right)_{ii} & , i = j \\ \left(\Sigma^{-1}\right)_{ij} & , j > i \\ 0 & , i > j \end{cases}.$$

In particular, it holds that $Z(\eta) = Z_0(\eta_0)$ and $\eta^\top T(x) = \eta_0^\top T_0(x)$.

Proof. We observe that $\eta^\# = \Sigma^{-1} = U + U^\top = \frac{\eta_0^\# + \left(\eta_0^\#\right)^\top}{2} \implies Z(\eta) = Z_0(\eta_0)$. Additionally, we have:

$$\begin{aligned} \eta^\top T(x) = \left\langle \left(\Sigma^{-1}\right)^\flat, -\frac{1}{2}\left(xx^\top\right)^\flat \right\rangle = -\frac{1}{2}\left\langle \Sigma^{-1}, xx^\top \right\rangle &= -\frac{1}{2}x^\top\Sigma^{-1}x \\ &= -\frac{1}{2}x^\top\left(U + U^\top\right)x \\ &= -\frac{1}{2}x^\top U x - \frac{1}{2}x^\top U^\top x \\ &\overset{(a)}{=} -x^\top U x \\ &= -\left\langle U, xx^\top \right\rangle \\ &= \left\langle 2U^\flat, -\frac{1}{2}\left(xx^\top\right)^\flat \right\rangle \\ &= \eta_0^\top T_0(x), \end{aligned}$$

where (a) used the fact that $x^\top U^\top x \in \mathbb{R} \implies x^\top U^\top x = \left(x^\top U^\top x\right)^\top = x^\top U x$.

Consequently, we have $p_\eta(x) = h(x)\exp\left(\eta^\top T(x) - Z(\eta)\right) = h_0(x)\exp\left(\eta_0^\top T_0(x) - Z_0(\eta_0)\right)$, verifying the equivalence of the two parameterizations. ■

We now consider the following process to generate η_0. Roughly speaking, modulo the parameterization issues mentioned above, the process fixes the diagonal of the precision matrix to be constant, while the off-diagonal entries are sampled to be uniform over an interval of width $\Theta\left(\frac{1}{d}\right)$ centered at the origin. Similar constructions have been employed previously, e.g., in [KLSU19].

We proceed to identify some properties of the matrices generated by the above process that will help us to prove Theorem 3.1.

Algorithm 1 Covariance Matrix Sampling

Input: $d \in \mathbb{N}$.
Output: A pair (Σ, η_0).

procedure COVSAMP(d)
  for $i \in [d]$ **do**
    for $j \in [d] \setminus [i-1]$ **do**
      if $i = j$ **then**
        Let $\left(\Sigma^{-1}\right)_{ii} = \frac{3}{4}$.
        Let $U_{ii} = \frac{1}{2}\left(\Sigma^{-1}\right)_{ii} = \frac{3}{8}$.
      else
        Draw an independent sample $\left(\Sigma^{-1}\right)_{ij} \sim \mathcal{U}\left(\left[-\frac{1}{4d}, \frac{1}{4d}\right]\right)$.
        Let $\left(\Sigma^{-1}\right)_{ji} = \left(\Sigma^{-1}\right)_{ij}$.
        Let $U_{ij} = \left(\Sigma^{-1}\right)_{ij}$ and $U_{ji} = 0$.
      end if
    end for
  end for
  Let $\eta_0 = 2U^\flat$.
  return (Σ, η_0).
end procedure

Lemma C.2. *Let $(\Sigma, \eta_0) \in \mathbb{R}^{d\times d} \times \mathbb{R}^{d^2}$ be the output of Algorithm 1. We have that:*

1. $\mathbb{I} \preceq \Sigma \preceq 2\mathbb{I}$.
2. *The components of η_0 are supported on intervals of the form:*
$$I_\ell = \begin{cases} \left\{\frac{3}{4}\right\}, & (\ell-1) \bmod d = (\ell-1) \operatorname{div} d \\ \left[-\frac{1}{2d}, \frac{1}{2d}\right], & (\ell-1) \bmod d > (\ell-1) \operatorname{div} d \\ \{0\}, & (\ell-1) \bmod d < (\ell-1) \operatorname{div} d \end{cases}.$$
3. *The components of η_0 are independent and it holds that $\|R\|_\infty = \frac{1}{d}$ and $\|R\|_2^2 = \frac{1}{2}\left(1 - \frac{1}{d}\right)$.*
4. *The vector $m \in \mathbb{R}^{d^2}$ is equal to:*
$$m_\ell = \begin{cases} \frac{3}{4}, & \ell = (k-1)\,d + k, \textit{ for some } k \in [d] \\ 0, & \textit{otherwise} \end{cases}.$$

Proof. We have:

1. By a direct application of Theorem E.5, we get that the eigenvalues of Σ^{-1} all lie in the interval:
$$\left[\frac{3}{4} - \frac{d-1}{4d}, \frac{3}{4} + \frac{d-1}{4d}\right] \subseteq \left[\frac{1}{2}, 1\right].$$
Thus, we have $\frac{1}{2}\mathbb{I} \preceq \Sigma^{-1} \preceq \mathbb{I} \iff \mathbb{I} \preceq \Sigma \preceq 2\mathbb{I}$.
2. Observe that, the first of the three branches in the definition corresponds to the diagonal elements of $\eta_0^\#$, the second corresponds to those above the diagonal, and the third corresponds to those below the diagonal. The result follows directly from this remark.
3. The non-constant components of η_0 are all independent draws from a uniform distribution, which immediately yields the independence property. In addition to that, all $\frac{d(d-1)}{2}$ non-constant components take values in an interval of length $\frac{1}{d}$. This implies that $\|R\|_\infty = \frac{1}{d}$ and $\|R\|_2^2 = \frac{d(d-1)}{2} \cdot \frac{1}{d^2} = \frac{1}{2}\left(1 - \frac{1}{d}\right)$.
4. Observe that the non-diagonal elements of $\eta_0^\# = 2U$ are either constant and equal to 0 or take values in a 0-centered interval. This implies that the only non-zero elements of m are those corresponding to the diagonal elements of $\left(\Sigma^{-1}\right)_{ii}$. This yields the desired result.

■

Remark C.3. Note that Theorem 3.1 holds only for $\alpha = \mathcal{O}(1)$. In contrast, prior fingerprinting lower bounds for mean estimation of product distributions and Gaussians gave non-trivial results in the low-accuracy regime $\alpha = \mathcal{O}\left(\sqrt{d}\right)$ (see, e.g., the results in Section G). In principle, it may be able to achieve similar lower bounds for covariance estimation in Frobenius norm, as the Frobenius diameter of the set of matrices that satisfy $\mathbb{I} \preceq \Sigma \preceq 2\mathbb{I}$ is equal to $\sqrt{d}$. However, by the third part of the above lemma, our construction only has an upper bound on this diameter of $\mathcal{O}(1)$. This makes estimation trivial for any α which is larger, as one could output an arbitrary parameter vector within the set. Thus, to prove a lower bound for all $\alpha = \mathcal{O}\left(\sqrt{d}\right)$, one must consider a different construction, and it is not merely a deficiency of our analysis.

We show that, if one has an estimator for the covariance matrix Σ, then this implies an estimator for the natural parameter vector η_0. Therefore, a lower bound for the latter problem implies a lower bound for the former, allowing us to focus on lower bounds for estimating the natural parameter vector. More precisely, we present an (ε, δ)-DP mechanism $T_M : \mathbb{R}^{n\times d} \to \bigotimes_{j\in[d^2]} I_j$ that satisfies the guarantee $\mathop{\mathbb{E}}_{X,T_M}\left[\|T_M(X) - (\eta_0 - m)\|_2^2\right] \leq \alpha^2$. T_M assumes the existence of an (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^{d\times d}$ such that $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \Sigma\|_\Sigma^2\right] \leq \frac{\alpha^2}{8}$.

Algorithm 2 From Covariance Estimation to Natural Parameter Estimation

Input: $X = (X_1, \ldots, X_n) \sim \mathcal{N}(0, \Sigma)^{\bigotimes n}$ and a mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^{d\times d}$.
Output: $T_M(X) \in \bigotimes_{j\in[d^2]} I_j$.

procedure $T_M(X)$
  Let $\widehat{\Sigma} = M(X)$.
  for $i \in [d]$ **do** ▷ $\widehat{\Sigma}^{-1}$ may be equal to any possible Σ^{-1} generated by Algorithm 1.
    for $j \in [d]$ **do**
      Let $\widetilde{\Sigma}_{ij}^{-1}$ be the projection of $\left(\widehat{\Sigma}^{-1}\right)_{ij}$ onto the support of Σ^{-1}.
    end for
  end for
  Let $\widetilde{U} \in \mathbb{R}^{d\times d}$ be an upper triangular matrix such that $\widetilde{\Sigma}^{-1} = \widetilde{U} + \widetilde{U}^\top$.
  Let $T_M(X) = 2\widetilde{U}^\flat - m$.
  return $T_M(X)$.
end procedure

Lemma C.4. *Let $\Sigma \in \mathbb{R}^{d\times d}$ be a covariance matrix generated by Algorithm 1 for $d \geq 2$ and let $X \sim \mathcal{N}(0, \Sigma)^{\bigotimes n}$ be a dataset. If $M : \mathbb{R}^{n\times d} \to \mathbb{R}^{d\times d}$ is an (ε, δ)-DP mechanism satisfying $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \Sigma\|_\Sigma^2\right] \leq \frac{\alpha^2}{8} \leq \frac{1}{32}$, then $T_M : \mathbb{R}^{n\times d} \to \bigotimes_{j\in[k]} I_j$ is also (ε, δ)-DP and satisfies $\mathop{\mathbb{E}}_{X,T_M}\left[\|T_M(X) - (\eta_0 - m)\|_2^2\right] \leq \alpha^2 \leq \frac{1}{4}$.*

Proof. The privacy guarantee of T_M is an immediate consequence of the guarantee of M and Lemma A.2. Additionally, the only randomness used by T_M is that used by M. Having established that, we focus our attention on the accuracy guarantee. We have:

$$\mathop{\mathbb{E}}_{X,T_M}\left[\|T_M(X) - (\eta_0 - m)\|_2^2\right] = 4 \mathop{\mathbb{E}}_{X,M}\left[\left\|\widetilde{U} - U\right\|_F^2\right].$$

We remark now that:

$$\left\|\widetilde{U} - U\right\|_F^2 = \sum_{i\in[d]}\sum_{j\in[d]}\left(\widetilde{U}_{ij} - U_{ij}\right)^2 = \sum_{j>i}\left(\widetilde{U}_{ij} - U_{ij}\right)^2 = \sum_{j>i}\left(\widetilde{\Sigma}_{ij}^{-1} - \left(\Sigma^{-1}\right)_{ij}\right)^2,$$

where we relied on the fact that the diagonal elements of U can take only a single value.

Consequently, we get:

$$\begin{aligned}
\mathop{\mathbb{E}}_{X,T_M}\left[\left\|T_M\left(X\right)-\left(\eta_0-m\right)\right\|_2^2\right] &= 4\mathop{\mathbb{E}}_{X,M}\left[\sum_{j>i}\left(\widetilde{\Sigma}_{ij}^{-1}-\left(\Sigma^{-1}\right)_{ij}\right)^2\right] \\
&= 2\mathop{\mathbb{E}}_{X,M}\left[\sum_{i\neq j}\left(\widetilde{\Sigma}_{ij}^{-1}-\left(\Sigma^{-1}\right)_{ij}\right)^2\right] \\
&\overset{(a)}{\leq} 2\mathop{\mathbb{E}}_{X,M}\left[\sum_{i=1}^{d}\sum_{j=1}^{d}\left(\left(\widehat{\Sigma}^{-1}\right)_{ij}-\left(\Sigma^{-1}\right)_{ij}\right)^2\right] \\
&= 2\mathop{\mathbb{E}}_{X,M}\left[\left\|\widehat{\Sigma}^{-1}-\Sigma^{-1}\right\|_F^2\right], \qquad (10)
\end{aligned}$$

where (a) relies on the fact the projection we perform in Line 5 ensures that $\left|\widetilde{\Sigma}_{ij}^{-1}-\left(\Sigma^{-1}\right)_{ij}\right| \leq \left|\left(\widehat{\Sigma}^{-1}\right)_{ij}-\left(\Sigma^{-1}\right)_{ij}\right|, \forall i,j\in[d]$.

Now, observe that:

$$\begin{aligned}
\left\|\widehat{\Sigma}^{-1}-\Sigma^{-1}\right\|_F = \left\|\Sigma^{-\frac{1}{2}}\left(\mathbb{I}-\Sigma^{\frac{1}{2}}\widehat{\Sigma}^{-1}\Sigma^{\frac{1}{2}}\right)\Sigma^{-\frac{1}{2}}\right\|_F &\leq \left\|\Sigma^{-\frac{1}{2}}\right\|_2^2\left\|\mathbb{I}-\Sigma^{\frac{1}{2}}\widehat{\Sigma}^{-1}\Sigma^{\frac{1}{2}}\right\|_F \\
&\overset{(a)}{\leq} \left\|\Sigma^{-1}-\widehat{\Sigma}^{-1}\right\|_{\Sigma^{-1}} \\
&= \left\|\Sigma^{-1}-\widehat{\Sigma}^{-1}\right\|_{\Sigma^{-\frac{1}{2}}\widehat{\Sigma}^{-\frac{1}{2}}\widehat{\Sigma}\widehat{\Sigma}^{-\frac{1}{2}}\Sigma^{-\frac{1}{2}}} \\
&\overset{(b)}{=} \left\|\widehat{\Sigma}^{\frac{1}{2}}\Sigma^{\frac{1}{2}}\left(\Sigma^{-1}-\widehat{\Sigma}^{-1}\right)\Sigma^{\frac{1}{2}}\widehat{\Sigma}^{\frac{1}{2}}\right\|_{\widehat{\Sigma}} \\
&= \left\|\widehat{\Sigma}-\widehat{\Sigma}^{\frac{1}{2}}\Sigma^{\frac{1}{2}}\widehat{\Sigma}^{-1}\Sigma^{\frac{1}{2}}\widehat{\Sigma}^{\frac{1}{2}}\right\|_{\widehat{\Sigma}} \\
&\overset{(c)}{=} \left\|\widehat{\Sigma}-\Sigma\right\|_{\widehat{\Sigma}} \\
&\overset{(d)}{\leq} 2\left\|\widehat{\Sigma}-\Sigma\right\|_{\Sigma}, \\
&= 2\left\|T\left(X\right)-\Sigma\right\|_{\Sigma}, \qquad (11)
\end{aligned}$$

where (a) uses the fact that $\Sigma^{-1}\preceq\mathbb{I}$ (which holds by the first part of Lemma C.4), (b) uses Lemma E.2 for $A=\Sigma^{-\frac{1}{2}}\widehat{\Sigma}^{-\frac{1}{2}}$, (c) is by Lemma E.3, and (d) is by Lemma E.4.

Substituting based on (11) into (10), we get:

$$\mathop{\mathbb{E}}_{X,T_M}\left[\left\|T_M\left(X\right)-\left(\eta_0-m\right)\right\|_2^2\right] \leq 8\mathop{\mathbb{E}}_{X,M}\left[\left\|M\left(X\right)-\Sigma\right\|_{\Sigma}^2\right],$$

which yields the desired result. ■

An immediate consequence of the above is the following:

Corollary C.5. *Let $\Sigma\in\mathbb{R}^{d\times d}$ be a covariance matrix generated by Algorithm 1 for $d\geq 2$ and let $X\sim\mathcal{N}\left(0,\Sigma\right)^{\bigotimes n}$. If any (ε,δ)-DP mechanism $T:\mathbb{R}^{n\times d}\to\bigotimes_{j\in[k]}I_j$ satisfying $\mathop{\mathbb{E}}_{X,T}\left[\left\|T\left(X\right)-\left(\eta_0-m\right)\right\|_2^2\right]\leq 8\alpha^2\leq\frac{1}{4}$ requires at least $n\geq n_{\eta_0}$ samples, the same sample complexity lower bound holds for any (ε,δ)-DP mechanism $M:\mathbb{R}^{n\times d}\to\mathbb{R}^{d\times d}$ that satisfies $\mathop{\mathbb{E}}_{X,M}\left[\left\|M\left(X\right)-\Sigma\right\|_{\Sigma}^2\right]\leq\alpha^2$.*

Having established the reduction from estimating η_0 with respect to the ℓ_2-norm to estimating Σ with respect to the Mahalanobis norm, it remains now to apply Theorem 2.5 to the former problem. For that reason, we work towards bounding the various quantities involved in it.

First, in accordance with Theorem 2.5, we assume the existence of an (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \bigotimes_{j\in[d^2]} I_j$ such that:

$$\mathop{\mathbb{E}}_{X,M}\left[\left\|M\left(X\right)-\left(\eta_0-m\right)\right\|_2^2\right] \leq 8\alpha^2 \leq \frac{1}{96} \leq \frac{\left\|R\right\|_2^2}{24}, \tag{12}$$

where we appealed to the third part of Lemma C.2 and the assumption that $d \geq 2$.

The next step to establishing our lower bound is reasoning about the quantity $\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma_{T_0} s\right]$. To do that, we identify an expression for Σ_{T_0} for $\mathcal{N}\left(0,\Sigma\right)$ and use a result from [DKK+16]. By Lemma C.1, we have $T_0 \equiv T \implies \Sigma_{T_0} = \Sigma_T$. This yields:

$$\begin{aligned}
\Sigma_T &= \mathop{\mathbb{E}}_{Y\sim\mathcal{N}(0,\Sigma)}\left[T\left(Y\right)T\left(Y\right)^\top\right] - \mu_T\mu_T^\top \\
&= \frac{1}{4}\mathop{\mathbb{E}}_{Y\sim\mathcal{N}(0,\Sigma)}\left[\left(YY^\top\right)^\flat\left(\left(YY^\top\right)^\flat\right)^\top\right] - \frac{1}{4}\Sigma^\flat\left(\Sigma^\flat\right)^\top \\
&= \frac{1}{4}\mathop{\mathbb{E}}_{Y\sim\mathcal{N}(0,\Sigma)}\left[\left(Y\otimes Y\right)\left(Y\otimes Y\right)^\top\right] - \frac{1}{4}\Sigma^\flat\left(\Sigma^\flat\right)^\top
\end{aligned} \tag{13}$$

Observe now that, based on the definition we gave for s in Section 2 and the characteristics of the particular exponential family that we are considering here, the matrix $s^\# \in \mathbb{R}^{d\times d}$ is symmetric. We use this observation to connect our problem with the following result:

Proposition C.6. *[Theorem 4.2 from [DKK+16]] Let $\mathcal{S}_{\mathrm{sym}} = \left\{M^\flat \in \mathbb{R}^{d^2} : M = M^\top\right\}$ and $X \sim \mathcal{N}\left(0,\Sigma\right)$. Let M be the $d^2 \times d^2$ matrix given by $M = \mathbb{E}\left[\left(X\otimes X\right)\left(X\otimes X\right)^\top\right]$. Then, as an operator on $\mathcal{S}_{\mathrm{sym}}$, we have:*

$$M = 2\Sigma^{\otimes 2} + \left(\Sigma^\flat\right)\left(\Sigma^\flat\right)^\top.$$

Using the above, we prove the following lemma:

Lemma C.7. *Let $\Sigma \in \mathbb{R}^{d\times d}$ be a covariance matrix generated by Algorithm 1 for $d \geq 2$ and let $X \sim \mathcal{N}\left(0,\Sigma\right)^{\bigotimes n}$ be a dataset. Also, let $M : \mathbb{R}^{n\times d} \to \bigotimes_{j\in[d^2]} I_j$ be an (ε,δ)-DP mechanism such that $\mathop{\mathbb{E}}_{X,M}\left[\left\|M\left(X\right)-\left(\eta_0-m\right)\right\|_2^2\right] \leq 8\alpha^2$. Then, it holds that:*

$$\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top\Sigma_T s\right] \leq \frac{\alpha^2}{d^4}.$$

Proof. By Proposition C.6 and (13), we get that:

$$\begin{aligned}
\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma_T s\right] &= \frac{1}{4}\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \mathop{\mathbb{E}}_{Y\sim\mathcal{N}(0,\Sigma)}\left[\left(Y\otimes Y\right)\left(Y\otimes Y\right)^\top\right] s\right] - \frac{1}{4}\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma^\flat \left(\Sigma^\flat\right)^\top s\right] \\
&= \frac{1}{4}\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \left(2\Sigma^{\otimes 2} + \left(\Sigma^\flat\right)\left(\Sigma^\flat\right)^\top\right) s\right] - \frac{1}{4}\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma^\flat \left(\Sigma^\flat\right)^\top s\right] \\
&= \frac{1}{2}\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma^{\otimes 2} s\right] \\
&= \frac{1}{2}\mathop{\mathbb{E}}_{X_{\sim i},M}\left[\left\|s\right\|_2^2 \frac{s^\top}{\left\|s\right\|_2}\Sigma^{\otimes 2}\frac{s^\top}{\left\|s\right\|_2}\right] \\
&\overset{(a)}{\leq} \frac{1}{2}\left\|\Sigma^{\otimes 2}\right\|_2 \mathop{\mathbb{E}}_{X_{\sim i},M}\left[\left\|s\right\|_2^2\right] \\
&\overset{(b)}{=} \frac{1}{2}\left\|\Sigma\right\|_2^2 \mathop{\mathbb{E}}_{X_{\sim i},M}\left[\left\|s\right\|_2^2\right] \\
&\overset{(c)}{\leq} 2\mathop{\mathbb{E}}_{X_{\sim i},M}\left[\sum_{j\in[d^2]}\left(\frac{R_j^2}{4} - \left(\eta_j - m_j\right)^2\right)^2 \left(M_j\left(X_{\sim i}\right) - \left(\eta_j - m_j\right)\right)^2\right] \\
&\leq \frac{\left\|R\right\|_\infty^4}{32} 2\mathop{\mathbb{E}}_{X_{\sim i},M}\left[\left\|M\left(X_{\sim i}\right) - \left(\eta - m\right)\right\|_2^2\right] \\
&\overset{(d)}{\leq} \frac{\left\|R\right\|_\infty^4}{8}\cdot 8\alpha^2 \\
&= \left\|R\right\|_\infty^4 \alpha^2 \\
&= \frac{\alpha^2}{d^4},
\end{aligned}$$

where (a) is by the definition of the spectral norm, (b) is by Fact E.1, (c) is by the assumption $\Sigma \preceq 2\mathbb{I}$, (d) is because $\mathop{\mathbb{E}}_{X,M}\left[\left\|M\left(X\right) - \left(\eta_0 - m\right)\right\|_2^2\right] \leq 8\alpha^2$ and the last equality is by the third part of Lemma C.2. ■

Now, it remains to reason about the tail probability for $Y \sim \mathcal{N}\left(0,\Sigma\right)$:

$$\mathbb{P}\left[\left|\sum_{j\in[d^2]}\left(T_{0,j}\left(Y\right) - \mu_{T,0,j}\right)\right| > \frac{4t}{\left\|R\right\|_\infty^3}\right] = \mathbb{P}\left[\left|\sum_{j\in[k]}\left(T_j\left(Y\right) - \mu_{T,j}\right)\right| > 4d^3 t\right] \tag{14}$$

as well as the integral:

$$\int_T^\infty \mathop{\mathbb{P}}_{Y\sim\mathcal{N}(0,\Sigma)}\left[\left|\sum_{j\in[d^2]}\left(T_j\left(Y\right) - \mu_{T,j}\right)\right| > 4d^3 t\right] dt. \tag{15}$$

We do so in the following lemma:

Lemma C.8. *Let $\Sigma \in \mathbb{R}^{d\times d}$ be a covariance matrix with $\mathbb{I} \preceq \Sigma \preceq 2\mathbb{I}$. We have that:*

$$\begin{aligned}
&\int_T^\infty \mathop{\mathbb{P}}_{Y\sim\mathcal{N}(0,\Sigma)}\left[\left|\sum_{j\in[d^2]}\left(T_j\left(Y\right) - \mu_{T,j}\right)\right| > 4d^3 t\right] dt \\
&\leq \frac{1}{4d^2}\left[\left(1 + \sqrt{\frac{\pi}{2}}\right) e^{-\frac{\left(\sqrt{8d^2T+1}-1\right)^2}{2}} + \sqrt{\pi} e^{-4d^4T^2}\right].
\end{aligned}$$

Proof. We start by remarking that:

$$\sum_{j\in[k]} T_j\left(Y\right) = -\frac{1}{2}\sum_{i=1}^d\sum_{j=1}^d Y_i Y_j = -\frac{1}{2}\left(\sum_{j=1}^d Y_j\right)^2 = -\frac{1}{2}\left(\vec{1}_d^\top Y\right)^2,$$

and:

$$\sum_{j\in[k]} \mu_{T,j} = -\frac{1}{2}\vec{1}_d^\top \Sigma \vec{1}_d.$$

Using standard facts about multivariate Gaussians, we get that:

$$Y \sim \mathcal{N}(0,\Sigma) \implies \vec{1}_d^\top Y \sim \mathcal{N}\left(0, \vec{1}_d^\top \Sigma \vec{1}_d\right) \iff \frac{\vec{1}_d^\top Y}{\sqrt{\vec{1}_d^\top \Sigma \vec{1}_d}} \sim \mathcal{N}(0,1)$$
$$\implies \frac{\left(\vec{1}_d^\top Y\right)^2}{\vec{1}_d^\top \Sigma \vec{1}_d} \sim \chi^2(1).$$

Based on the above, (14) becomes:

$$\mathop{\mathbb{P}}_{Y\sim\mathcal{N}(0,\Sigma)}\left[\left|\frac{\left(\vec{1}_d^\top Y\right)^2}{\vec{1}_d^\top \Sigma \vec{1}_d} - 1\right| > \frac{8d^3 t}{\vec{1}_d^\top \Sigma \vec{1}_d}\right] \leq \mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[|Z-1| > 4d^2 t\right],$$

where the last inequality uses the fact that $\Sigma \preceq 2\mathbb{I} \implies \vec{1}_d^\top \Sigma \vec{1}_d \leq 2\left\|\vec{1}_d\right\|_2^2 = 2d$.

Thus, upper-bounding this probability boils down to using bounds on the tail probabilities of χ^2 distributions. For that reason, we appeal to Fact F.6 for $k=1$, which gives us that:

$$\mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[Z-1 \geq 2\sqrt{\ln\left(\frac{1}{\beta}\right)} + 2\ln\left(\frac{1}{\beta}\right)\right] \leq \beta,$$

and:

$$\mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[1-Z \geq 2\sqrt{\ln\left(\frac{1}{\beta}\right)}\right] \leq \beta.$$

In our case, the bounds become:

$$\mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[Z-1 \geq t\right] \leq e^{-\frac{\left(\sqrt{2t+1}-1\right)^2}{2}}, t \geq 0,$$

and:

$$\mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[1-Z \geq t\right] \leq e^{-\frac{t^2}{4}}, t \geq 0.$$

We have:

$$\begin{aligned}
\int_T^\infty \mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[|Z-1| > 4d^2 t\right]\, dt &= \int_T^\infty \left(\mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[Z-1 > 4d^2 t\right] + \mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[1-Z > 4d^2 t\right]\right) dt \\
&\overset{(a)}{=} \frac{1}{4d^2}\int_{4d^2T}^\infty \left(\mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[Z-1 > t'\right] + \mathop{\mathbb{P}}_{Z\sim\chi^2(1)}\left[1-Z > t'\right]\right) dt' \\
&= \frac{1}{4d^2}\int_{4d^2T}^\infty \left(e^{-\frac{\left(\sqrt{2t'+1}-1\right)^2}{2}} + e^{-\frac{t'^2}{4}}\right) dt', \qquad (16)
\end{aligned}$$

where in (a) we set $t' = 4d^2 t$.

For the first term of (16), we get:

$$\begin{aligned}
\frac{1}{4d^2}\int_{4d^2T}^\infty e^{-\frac{\left(\sqrt{2t'+1}-1\right)^2}{2}}\, dt' &\overset{(a)}{=} \frac{1}{4d^2}\int_{\sqrt{8d^2T+1}-1}^\infty (u+1)\, e^{-\frac{u^2}{2}}\, du \\
&= \frac{1}{4d^2}\left(\sqrt{\frac{\pi}{2}}\mathrm{erfc}\left(\frac{\sqrt{8d^2T+1}-1}{\sqrt{2}}\right) + e^{-\frac{\left(\sqrt{8d^2T+1}-1\right)^2}{2}}\right) \\
&\leq \frac{1}{4d^2}\left(1+\sqrt{\frac{\pi}{2}}\right) e^{-\frac{\left(\sqrt{8d^2T+1}-1\right)^2}{2}}, \qquad (17)
\end{aligned}$$

where in (a) we set $u = \sqrt{2t'+1} - 1$ and in the rest we used Fact F.5.

For the second term of (16), we have the upper bound:

$$\frac{1}{4d^2}\int_{4d^2T}^{\infty} e^{-\frac{t'^2}{4}}\,dt' = \frac{\sqrt{\pi}}{4d^2}\mathrm{erfc}\left(2d^2T\right) \leq \frac{\sqrt{\pi}}{4d^2}e^{-4d^4T^2}, \tag{18}$$

where we again appealed to Fact F.5.

Substituting (17) and (18) into (16), we get the desired result. ■

We are now ready to prove the main result of this section.

Proof of Theorem 3.1. We assume that the process generating Σ is that of Algorithm 1. By the first part of Lemma C.2, the condition $\mathbb{I} \preceq \Sigma \preceq 2\mathbb{I}$ is satisfied. We assume that there exists an (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^{d^2}$ that satisfies $\underset{X,M}{\mathbb{E}}\left[\|M(X) - (\eta_0 - m)\|_2^2\right] \leq 8\alpha^2 \leq \frac{1}{96}$. Combining the results of Lemmas C.7 and C.8 with Theorem 2.5, we get:

$$n\left\{2\delta T + 2\frac{\alpha\varepsilon}{d^2} + \frac{1}{2d^2}\left[\left(1+\sqrt{\frac{\pi}{2}}\right)e^{-\frac{\left(\sqrt{8d^2T+1}-1\right)^2}{2}} + \sqrt{\pi}e^{-4d^4T^2}\right]\right\} \geq \frac{1}{48}\left(1-\frac{1}{d}\right) \geq \frac{1}{96}. \tag{19}$$

We want to identify a range of values for T and δ such that:

$$\delta T \geq \frac{1}{2d^2}\left(1+\sqrt{\frac{\pi}{2}}\right)e^{-\frac{\left(\sqrt{8d^2T+1}-1\right)^2}{2}} + \frac{\sqrt{\pi}}{2d^2}e^{-4d^4T^2}. \tag{20}$$

We set $T = \max\left\{\frac{1}{8d^2}\left[\left(2\ln\left(\frac{1}{\delta}\right)+1\right)^2 - 1\right], \frac{1}{2d^2}\sqrt{\ln\left(\frac{1}{\delta}\right)}\right\}$. Thus, for (20) to hold, it suffices to satisfy the stronger inequality:

$$\begin{aligned}
&\delta\max\left\{\frac{1}{8d^2}\left[\left(2\ln\left(\frac{1}{\delta}\right)+1\right)^2-1\right], \frac{1}{2d^2}\sqrt{\ln\left(\frac{1}{\delta}\right)}\right\} \geq \frac{1}{2d^2}\left(1+\sqrt{\frac{\pi}{2}}\right)\delta + \frac{\sqrt{\pi}}{2d^2}\delta \\
\iff\ &\max\left\{\left(2\ln\left(\frac{1}{\delta}\right)+1\right)^2-1, 4\sqrt{\ln\left(\frac{1}{\delta}\right)}\right\} \geq 4\left(1+\sqrt{\frac{\pi}{2}}\right)+4\sqrt{\pi} \\
\iff\ &\sqrt{\ln\left(\frac{1}{\delta}\right)}\max\left\{\sqrt{\ln\left(\frac{1}{\delta}\right)}\left(\ln\left(\frac{1}{\delta}\right)+1\right), 1\right\} \geq 1+\frac{2+\sqrt{2}}{2}\sqrt{\pi}.
\end{aligned}$$

The above holds for $\delta \leq 0.1$ and we have $T = \frac{1}{8d^2}\left[\left(2\ln\left(\frac{1}{\delta}\right)+1\right)^2-1\right]$. Thus, (19) becomes:

$$n\left\{\frac{3\delta}{8d^2}\left[\left(2\ln\left(\frac{1}{\delta}\right)+1\right)^2-1\right] + 2\frac{\alpha\varepsilon}{d^2}\right\} \geq \frac{1}{96}. \tag{21}$$

It suffices to identify a range for the values of δ such that:

$$\begin{aligned}
&\frac{3n\delta}{8d^2}\left[\left(2\ln\left(\frac{1}{\delta}\right)+1\right)^2-1\right] \leq \frac{1}{192} \\
\iff\ &\delta\ln\left(\frac{1}{\delta}\right)\left(\ln\left(\frac{1}{\delta}\right)+1\right) \leq \frac{d^2}{288n}.
\end{aligned}$$

Since we assumed that $\delta \leq 0.1$, we have $1+\log\left(\frac{1}{\delta}\right) \leq 2\log\left(\frac{1}{\delta}\right)$. Therefore, it suffices to have:

$$\delta\ln^2\left(\frac{1}{\delta}\right) \leq \frac{d^2}{576n}. \tag{22}$$

Let $\phi := \frac{d^2}{576n}$. We argue that the above inequality holds for $\delta \leq \frac{\phi}{16\ln^2\left(\frac{1}{\phi}\right)}$. We remark that it may not necessarily be the case that $\frac{\phi}{16\ln^2\left(\frac{1}{\phi}\right)} \leq 0.1$. This is true when $\phi \leq 0.55$. We will proceed assuming

that this holds and verify the assumption at the end. Now, we observe that the function $\delta \ln^2\left(\frac{1}{\delta}\right)$ is strictly increasing for $\delta \leq e^{-2} \approx 0.13$. Thus, it suffices to show the result for $\delta = \frac{\phi}{16\ln^2\left(\frac{1}{\phi}\right)} \leq 0.1$. Indeed, we have:

$$\begin{aligned}
\delta \ln^2\left(\frac{1}{\delta}\right) &= \frac{\phi}{16}\left(\frac{\ln\left(\frac{1}{\phi}\right) + \ln\left(16\ln^2\left(\frac{1}{\phi}\right)\right)}{\ln\left(\frac{1}{\phi}\right)}\right)^2 \\
&= \frac{\phi}{16}\left(1 + \frac{\ln\left(16\ln^2\left(\frac{1}{\phi}\right)\right)}{\ln\left(\frac{1}{\phi}\right)}\right)^2 \\
&= \frac{\phi}{16}\left(1 + \frac{\ln\left(16x^2\right)}{x}\right)^2,
\end{aligned}$$

where we set $x := \ln\left(\frac{1}{\phi}\right)$.

It suffices to have $\left(1 + \frac{\ln\left(16x^2\right)}{x}\right)^2 \leq 16 \iff \frac{\ln\left(16x^2\right)}{x} \leq 3$. This holds for any $x > 0$, which is satisfied by $0 < \phi \leq 0.55$. Consequently, we have for δ the range $\delta \leq \frac{\phi}{16\ln^2\left(\frac{1}{\phi}\right)} = \frac{d^2}{9216n\ln^2\left(\frac{576n}{d^2}\right)}$.

Thus, (21) becomes:

$$2n\frac{\alpha\varepsilon}{d^2} \geq \frac{1}{192} \iff n \geq \frac{d^2}{384\alpha\varepsilon}.$$

The above also satisfies the condition on ϕ, since we have assumed that $\varepsilon \leq 1$ and $\alpha \leq \frac{1}{16\sqrt{3}}$.

Appealing to Corollary C.5 completes the proof. ■

C.2 Omitted Proofs from Section 3.2

Proof of Theorem 3.2. First, assume that Σ is generated by Algorithm 1 (as was assumed in all of Section 3.1). Additionally, assume that there exists an (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^{d\times d}$ with $\varepsilon \in [0,1]$ and $\delta \leq \mathcal{O}\left(\frac{d^2}{n\ln^2\left(\frac{n}{d^2}\right)}\right)$ that satisfies $\mathop{\mathbb{E}}_{X,M}\left[\left\|\Sigma^{-\frac{1}{2}}\left(M\left(X\right) - \Sigma\right)\Sigma^{-\frac{1}{2}}\right\|_2^2\right] \leq \alpha^2$ and uses $n = o\left(\frac{d^{1.5}}{\alpha\varepsilon}\right)$ samples. Using the standard matrix-norm inequality $\|A\|_F \leq \sqrt{d}\,\|A\|_2, \forall A \in \mathbb{R}^{k\times\ell}$, we get that:

$$\begin{aligned}
\mathop{\mathbb{E}}_{X,M}\left[\|M\left(X\right) - \Sigma\|_\Sigma^2\right] &= \mathop{\mathbb{E}}_{X,M}\left[\left\|\Sigma^{-\frac{1}{2}}\left(M\left(X\right) - \Sigma\right)\Sigma^{-\frac{1}{2}}\right\|_F^2\right] \\
&\leq d\mathop{\mathbb{E}}_{X,M}\left[\left\|\Sigma^{-\frac{1}{2}}\left(M\left(X\right) - \Sigma\right)\Sigma^{-\frac{1}{2}}\right\|_2^2\right] \\
&\leq d\alpha^2.
\end{aligned}$$

By Theorem 3.1 for $\alpha \to \sqrt{d}\alpha$, we get that any mechanism with the characteristics of M that estimates a Gaussian covariance matrix in Mahalanobis error $\sqrt{d}\alpha$ requires $n \geq \Omega\left(\frac{d^{1.5}}{\alpha\varepsilon}\right)$ samples. This leads to a contradiction, implying the desired sample complexity lower bound. ■

D Omitted Proof from Section 4

Proof of Theorem 4.1. We define the following family of distributions, such that the second moment of each distribution is upper-bounded by 1. The loss function between the true distribution and the estimate is the squared ℓ_2-distance between their means. For each $v \in \mathcal{E}_d = \{-1,1\}^d$, we define the distribution $\mathcal{D}_v$ over $\mathbb{R}^d$ as follows. Let $t > 0$ and $0 < p < 1$. For any $X \sim \mathcal{D}_v$, we have:

$$X_i = \begin{cases} v_i t, & \text{with probability } p \\ 0, & \text{with probability } 1-p \end{cases}, \forall i \in [d].$$

Then $\mathop{\mathbb{E}}_{X\sim\mathcal{D}_v}[X] = ptv$. We show that the second moment of this distribution is upper-bounded by pt^2, that is, $\mathop{\mathbb{E}}_{X\sim\mathcal{D}_v}\left[\langle X - ptv, u\rangle^2\right] \leq pt^2, \forall u \in \mathbb{S}^{d-1}$. Note that, for any $i \in [d]$, we have $\mathbb{E}\left[(X_i - ptv_i)^2\right] = p(1-p)t^2$. We omit the subscript of the expectation when the context is clear. We have:

$$\begin{aligned}
\mathbb{E}\left[\langle X - ptv, u\rangle^2\right] &= \mathbb{E}\left[\left(\sum_{i=1}^{d}(X_i - ptv_i)u_i\right)^2\right] \\
&= \sum_{i=1}^{d}\mathbb{E}\left[(X_i - ptv_i)^2\right]u_i^2 + \sum_{i\neq j}\mathbb{E}\left[X_i - ptv_i\right]\mathbb{E}\left[X_j - ptv_j\right]u_iu_j \\
&= \sum_{i=1}^{d} p(1-p)t^2u_i^2 \\
&= p(1-p)t^2 \\
&\leq pt^2.
\end{aligned}$$

We want this to be upper-bounded by 1.

$$pt^2 = 1. \tag{23}$$

Next, for any $u, v \in \mathcal{E}_d$, we bound the loss function between $\mathcal{D}_{v_1}$ and $\mathcal{D}_{v_2}$. As mentioned before, we define $\ell(\theta(\mathcal{D}_u), \theta(\mathcal{D}_v)) = \|\theta(\mathcal{D}_u) - \theta(\mathcal{D}_v)\|_2^2$. We have the following:

$$\begin{aligned}
\|\theta(\mathcal{D}_u) - \theta(\mathcal{D}_v)\|_2^2 &= \sum_{i=1}^{d} p^2t^2(v_i - u_i)^2 \\
&= p^2t^2\sum_{i:u_i\neq v_i}(v_i - u_i)^2 \\
&= 4p^2t^2\sum_{i=1}^{d}\mathbb{1}\{v_i \neq u_i\}.
\end{aligned}$$

Using the notation in Lemma A.3, we have $\tau = 2p^2t^2$. Using the same lemma, we have:

$$\begin{aligned}
R(\mathcal{P}, \ell, \varepsilon, \delta) &\geq \frac{d\tau}{2}\cdot\left(0.9e^{-10\varepsilon D} - 10\delta D\right) \\
&\geq \frac{d\tau}{2}\cdot\left[0.9(1 - 10\varepsilon D) - 10\delta D\right] \\
&= \frac{d\tau}{2}\cdot\left[0.9 - D(9\varepsilon + 10\delta)\right] \\
&\geq \frac{d\tau}{2}\cdot\left[0.9 - 10D(\varepsilon + \delta)\right].
\end{aligned}$$

Setting $R(\mathcal{P}, \ell, \varepsilon, \delta) \leq \alpha^2$ and rearranging the above, we get:

$$D \geq \frac{1}{10(\varepsilon + \delta)}\cdot\left(0.9 - \frac{2\alpha^2}{d\tau}\right). \tag{24}$$

Now, from Equation (23), $\tau = 2p^2t^2 = 2p$. Using this and Inequality (24), and setting $p = \frac{2\alpha^2}{d}$, we have the following.

$$D \geq \frac{0.04}{\varepsilon + \delta}. \tag{25}$$

For $i \in [d]$, we define the mixture distributions p_{+i} and p_{-i} as in Lemma A.3. We define a coupling (X, Y) between p_{+i} and p_{-i} as follows. For every row X_j in X, if $X_j^i \neq 0$, then $Y_j^i = -X_j^i$,

otherwise $Y_j^i = X_j^i$. For all other coordinates $k \neq i$, $Y_j^k = X_j^k$. Notice that, for all $k \neq i$, the distributions for coordinate k for both p_{+i} and p_{-i} are identical because all coordinates are independent. Therefore, by setting $Y_j^k = X_j^k$, these coordinates of X_j and Y_j have identical distributions for all $j \in [n]$. The coordinate i of X_j is t with probability p, and 0 with probability $1-p$. Hence, the coordinate i of Y_j is $-t$ with probability p, and 0 with probability $1-p$. This is true for all $j \in [n]$. Since all n rows in both X (hence, in Y) are chosen independently at random, the distribution of X is p_{+i} and the distribution of Y is p_{-i}. Therefore, (X, Y) is a valid coupling of p_{+i} and p_{-i}. Next, we have to determine the value of $D = \mathbb{E}\left[d_{\mathrm{Ham}}\left(X, Y\right)\right]$. Note that for each $j \in [n]$, X_j can only differ from Y_j when $X_j^i \neq 0$. This can only happen with probability p. Therefore, by linearity of expectations, $D = np$.

Finally, using the fact that $p = \frac{2\alpha^2}{d}$ and $t = \frac{\sqrt{d}}{\sqrt{2}\alpha}$, and substituting this in Inequality (25), we have:

$$np \geq \frac{0.04}{\varepsilon + \delta} \implies n \geq \frac{d}{50\left(\varepsilon + \delta\right)\alpha^2}.$$

When $\varepsilon \geq \delta$, this gives us the desired lower bound. ■

E Facts from Linear Algebra

In this appendix we collect some results from linear algebra which are referenced in this work.

Fact E.1. *Let $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{k \times \ell}$. Then, it holds that $\|A \otimes B\|_2 = \|A\|_2 \|B\|_2$.*

Lemma E.2. *Let $A, \Sigma \in \mathbb{R}^{d \times d}$ be non-singular matrices and let $X \in \mathbb{R}^{d \times d}$. Then, it holds that $\|X\|_{A\Sigma A^\top} = \left\|A^{-1} X \left(A^\top\right)^{-1}\right\|_\Sigma$.*

For a proof of this statement for symmetric A (the proof does not change), we refer the reader to Section 2.2 of [Li19].

Lemma E.3. *Let $A, B \in \mathbb{R}^{d \times d}$ be symmetric positive definite matrices. Then, it holds that:*

$$\left\|A - A^{\frac{1}{2}} B^{\frac{1}{2}} A^{-1} B^{\frac{1}{2}} A^{\frac{1}{2}}\right\|_A = \|A - B\|_A .$$

Proof. We have that:

$$\begin{aligned}
\left\|A - A^{\frac{1}{2}} B^{\frac{1}{2}} A^{-1} B^{\frac{1}{2}} A^{\frac{1}{2}}\right\|_A^2 &= \left\|\mathbb{I} - B^{\frac{1}{2}} A^{-1} B^{\frac{1}{2}}\right\|_F^2 \\
&= \operatorname{tr}\left(\left(\mathbb{I} - B^{\frac{1}{2}} A^{-1} B^{\frac{1}{2}}\right)^2\right) \\
&= \operatorname{tr}\left(\mathbb{I} - 2B^{\frac{1}{2}} A^{-1} B^{\frac{1}{2}} + B^{\frac{1}{2}} A^{-1} B A^{-1} B^{\frac{1}{2}}\right) \\
&\overset{(a)}{=} \operatorname{tr}\left(\mathbb{I} - 2A^{-\frac{1}{2}} B A^{-\frac{1}{2}} + \left(A^{-\frac{1}{2}} B A^{-\frac{1}{2}}\right)^2\right) \\
&= \operatorname{tr}\left(\left(\mathbb{I} - A^{-\frac{1}{2}} B A^{-\frac{1}{2}}\right)^2\right) \\
&= \left\|\mathbb{I} - A^{-\frac{1}{2}} B A^{-\frac{1}{2}}\right\|_F^2 \\
&= \|A - B\|_A^2 ,
\end{aligned}$$

where in (a) we use the cyclic property of the trace in individual terms. ■

Lemma E.4. *Let $\Sigma_1, \Sigma_0 \in \mathbb{R}^{d \times d}$ be a pair of positive definite matrices. If $\|\Sigma_1 - \Sigma_0\|_{\Sigma_0} < \frac{1}{2}$, it holds that $\|\Sigma_1 - \Sigma_0\|_{\Sigma_1} \leq 2\|\Sigma_1 - \Sigma_0\|_{\Sigma_0}$.*

Proof. By the definition of the Mahalanobis norm, we have:

$$\begin{aligned}\|\Sigma_1-\Sigma_0\|_{\Sigma_0}^2 &= \left\|\mathbb{I}-\Sigma_0^{-\frac{1}{2}}\Sigma_1\Sigma_0^{-\frac{1}{2}}\right\|_F \\ &= tr\left(\left(\mathbb{I}-\Sigma_0^{-\frac{1}{2}}\Sigma_1\Sigma_0^{-\frac{1}{2}}\right)^2\right) \\ &= tr\left(\mathbb{I}-2\Sigma_0^{-\frac{1}{2}}\Sigma_1\Sigma_0^{-\frac{1}{2}}+\Sigma_0^{-\frac{1}{2}}\Sigma_1\Sigma_0^{-1}\Sigma_1\Sigma_0^{-\frac{1}{2}}\right) \\ &= d-2tr\left(\Sigma_0^{-1}\Sigma_1\right)+tr\left(\left(\Sigma_0^{-1}\Sigma_1\right)^2\right). \end{aligned} \tag{26}$$

Similarly, we have:

$$\|\Sigma_1-\Sigma_0\|_{\Sigma_1}^2 = d-2tr\left(\Sigma_1^{-1}\Sigma_0\right)+tr\left(\left(\Sigma_1^{-1}\Sigma_0\right)^2\right). \tag{27}$$

The eigenvalues of the matrix $\Sigma_0^{-1}\Sigma_1$ are positive. Indeed, the previous matrix can be written as $\Sigma_0^{-\frac{1}{2}}\left(\Sigma_0^{-\frac{1}{2}}\Sigma_1\Sigma_0^{-\frac{1}{2}}\right)\Sigma_0^{\frac{1}{2}}$, implying that it similar to the symmetric matrix $\Sigma_0^{-\frac{1}{2}}\Sigma_1\Sigma_0^{-\frac{1}{2}}$. This last matrix is positive definite, since $u^\top\Sigma_0^{-\frac{1}{2}}\Sigma_1\Sigma_0^{-\frac{1}{2}}u=\left(\Sigma_0^{-\frac{1}{2}}u\right)^\top\Sigma_1\left(\Sigma_0^{-\frac{1}{2}}u\right)$ and $\Sigma_1\succ 0$. Thus, if $\lambda_1,\ldots,\lambda_d>0$ are the eigenvalues of $\Sigma_0^{-1}\Sigma_1$, (26) and (27) can be written as:

$$\begin{aligned}\|\Sigma_1-\Sigma_0\|_{\Sigma_0}^2 &= d-2\cdot\sum_{i=1}^d\lambda_i+\sum_{i=1}^d\lambda_i^2=\sum_{i=1}^d\left(\lambda_i-1\right)^2, \\ \|\Sigma_1-\Sigma_0\|_{\Sigma}^2 &= d-2\cdot\sum_{i=1}^d\frac{1}{\lambda_i}+\sum_{i=1}^d\frac{1}{\lambda_i^2}=\sum_{i=1}^d\left(\frac{1}{\lambda_i}-1\right)^2.\end{aligned}$$

Therefore, proving the desired upper bound reduces to the following maximization problem:

$$\begin{aligned}\max_{\lambda_i}\quad & \sum_{i=1}^d\left(\frac{1}{\lambda_i}-1\right)^2 \\ \text{s.t.}\quad & \sum_{i=1}^d\left(\lambda_i-1\right)^2\le\|\Sigma_1-\Sigma_0\|_{\Sigma_0}^2<\frac{1}{4},\end{aligned}$$

where we have omitted the constraint $\lambda_i>0$, since we necessarily have $\lambda_i\ge 1-\|\Sigma_1-\Sigma_0\|_{\Sigma_0}>\frac{1}{2},\forall i\in[d]$. The above problem is solvable exactly via the KKT conditions, yielding:

$$\|\Sigma_1-\Sigma_0\|_{\Sigma_1}\le\frac{\|\Sigma_1-\Sigma_0\|_{\Sigma_0}\sqrt{d}}{\sqrt{d}-\|\Sigma_1-\Sigma_0\|_{\Sigma_0}}\le 2\|\Sigma_1-\Sigma_0\|_{\Sigma_0},$$

for $\|\Sigma_1-\Sigma_0\|_{\Sigma_0}<\frac{1}{2}$ and $d\ge 1$. ■

Theorem E.5 (Gershgorin Circle Theorem (see [Ger31]))**.** *Let $A\in\mathbb{R}^{d\times d}$. For any eigenvalue λ of A, there exists an $i\in[d]$, such that:*

$$|\lambda-a_{ii}|\le\sum_{j\ne i}|a_{ij}|.$$

F Facts from Probability & Statistics

In this appendix we collect some results from probability and statistics which are frequently referenced in this work.

Fact F.1. *The Bernoulli distribution* $\mathrm{Be}(p)$ *is an exponential family with support* $S=\{0,1\}$ *and* $h(x)=1,\eta=\ln\left(\frac{1-p}{p}\right),T(x)=x,Z(\eta)=\ln(1+e^\eta)$.

Fact F.2. *The d-dimensional Gaussian distribution with known covariance matrix and unknown mean $\mathcal{N}(\mu, \mathbb{I})$ is an exponential family with support $S = \mathbb{R}^d$ and $h(x) = \frac{1}{(2\pi)^{\frac{d}{2}}} e^{-\frac{\|x\|_2^2}{2}}, \eta = \mu, T(x) = x, Z(\eta) = \frac{\|\mu\|_2^2}{2}$.*

Fact F.3. *The d-dimensional Gaussian distribution with known mean and unknown covariance matrix $\mathcal{N}(0, \Sigma)$ is an exponential family with support $S = \mathbb{R}^d$ and $h(x) = 1, \eta = \left(\Sigma^{-1}\right)^\flat, T(x) = -\frac{1}{2}\left(xx^\top\right)^\flat, Z(\eta) = \frac{d}{2}\ln(2\pi) - \frac{1}{2}\ln\left(\det\left(\eta^\#\right)\right)$.*

Fact F.4 (Hölder's)**.** *We consider the following variants of Hölder's inequality:*

1. *Let $x, y \in \mathbb{R}^d$ and $p, q \geq 1$ such that $\frac{1}{p} + \frac{1}{q} = 1$. Then:*
$$|\langle x, y\rangle| \leq \|x\|_p \|y\|_q$$
2. *Let $X_1, \ldots, X_n$ be random variables over $\mathbb{R}$, and let $p_1, \ldots, p_n \geq 1$ with $\sum_{i=1}^n \frac{1}{p_i} = 1$. Then:*
$$\mathbb{E}\left[\left|\prod_{i=1}^n X_i\right|\right] \leq \prod_{i=1}^n \mathbb{E}\left[|X_i|^{p_i}\right]^{\frac{1}{p_i}},$$
assuming all the quantities involved are well-defined.

Fact F.5. *We denote the* complementary error function *by:*
$$\operatorname{erfc}(x) \coloneqq \frac{2}{\sqrt{\pi}} \int_z^\infty e^{-t^2}\, dt, x \geq 0.$$
It holds that $\operatorname{erfc}(x) \leq e^{-x^2}, \forall x \geq 0$.

Fact F.6 (Lemma 1 of [LM00])**.** *If $X \sim \chi^2(k)$, then:*
$$\mathbb{P}\left[X - k \geq 2\sqrt{k\log\left(\frac{1}{\beta}\right)} + 2\log\left(\frac{1}{\beta}\right)\right] \leq \beta,$$
and:
$$\mathbb{P}\left[k - X \geq 2\sqrt{k\log\left(\frac{1}{\beta}\right)}\right] \leq \beta.$$
Thus, if $Y \sim \mathcal{N}(0, \mathbb{I})$, then $\mathbb{P}\left[\|Y\|_2^2 \geq d + 2\sqrt{d\log\left(\frac{1}{\beta}\right)} + 2\log\left(\frac{1}{\beta}\right)\right] \leq \beta$.

Lemma F.7. *Let p_η be a density over $S \subseteq \mathbb{R}^d$ belonging to an exponential family $\mathcal{E}(T, h)$ that is parameterized by $\eta \in \mathcal{H} = \mathbb{R}^k$. We have:*
$$\mathop{\mathbb{P}}_{X \sim p_\eta}\left[\sum_{j \in [k]} \left(T_j(X) - \mu_{T,j}\right) > t\right] \leq \exp\left[Z\left(\eta + \lambda_1^\star \vec{1}_k\right) - Z(\eta) - \lambda_1^\star\left(t + \vec{1}_k^\top \mu_T\right)\right],$$
and:
$$\mathop{\mathbb{P}}_{X \sim p_\eta}\left[\sum_{j \in [k]} \left(T_j(X) - \mu_{T,j}\right) < -t\right] \leq \exp\left[Z\left(\eta + \lambda_2^\star \vec{1}_k\right) - Z(\eta) - \lambda_2^\star\left(t + \vec{1}_k^\top \mu_T\right)\right],$$
where $\lambda_1^\star > 0$ and $\lambda_2^\star < 0$ are the unique solutions to the equations:
$$\vec{1}_k^\top\left(\nabla_\eta Z\left(\eta + \lambda_1^\star \vec{1}_k\right) - \nabla_\eta Z(\eta)\right) = t,$$
and:
$$\vec{1}_k^\top\left(\nabla_\eta Z\left(\eta + \lambda_2^\star \vec{1}_k\right) - \nabla_\eta Z(\eta)\right) = t,$$
respectively (assuming that the solutions exist).

Proof. We will prove only the first of the two statements, since the other proof is similar by symmetry. Observe that:

$$\sum_{j\in[k]} \left(T_j\left(X\right) - \mu_{T,j}\right) = \vec{1}_k^\top \left(T\left(X\right) - \mu_T\right).$$

For $\lambda > 0$, the third part of Proposition A.5 yields:

$$\begin{aligned}\mathbb{P}\left[\vec{1}_k^\top T\left(X\right) > t + \vec{1}_k^\top \mu_T\right] &\leq \frac{\mathbb{E}\left[\exp\left[\left(\lambda\vec{1}_k\right)^\top T\left(X\right)\right]\right]}{\exp\left[\lambda\left(t + \vec{1}_k^\top \mu_T\right)\right]} \\ &= \exp\left[Z\left(\eta + \lambda\vec{1}_k\right) - Z\left(\eta\right) - \lambda\left(t + \vec{1}_k^\top \mu_T\right)\right].\end{aligned}$$

Let $g\left(\lambda\right) \coloneqq Z\left(\eta + \lambda\vec{1}_k\right) - Z\left(\eta\right) - \lambda\left(t + \vec{1}_k^\top \mu_T\right), \forall\lambda > 0$. Getting a tight upper bound for the above boils down to minimizing this function.

Taking the derivative with respect to λ, we get:

$$g'\left(\lambda\right) = \vec{1}_k^\top \nabla_\eta Z\left(\eta + \lambda\vec{1}_k\right) - \vec{1}_k^\top \mu_T - t = \vec{1}_k^\top \left(\nabla_\eta Z\left(\eta + \lambda\vec{1}_k\right) - \nabla_\eta Z\left(\eta\right)\right) - t,$$

where the last equality used the first part of Proposition A.5.

The second derivative of g is:

$$g''\left(\lambda\right) = \vec{1}_k^\top \nabla_\eta^2 Z\left(\eta + \lambda\vec{1}_k\right)\vec{1}_k.$$

By the second part of Proposition A.5, we get that $\nabla_\eta^2 Z\left(\eta + \lambda\vec{1}_k\right)$ is equal to the covariance matrix of a distribution in the family that has parameter vector $\eta + \lambda\vec{1}_k$ and is hence a PSD matrix. This implies that $g''\left(\lambda\right) \geq 0, \forall\lambda > 0$, so g is convex. Thus, if the equation $g'\left(\lambda\right) = 0$ has a solution $\lambda_1^\star > 0$, it is the unique minimizer of g. ■

G Basic Applications: Recovering Existing Lower Bounds

We show how to use Theorem 2.5 to recover existing lower bounds from [KLSU19]: (1) mean estimation of binary product distributions, and (2) mean estimation of high-dimensional Gaussians. Note that both classes of distributions are exponential families. The error metric for these lower bounds would be the mean-squared-error (MSE), as opposed to high-probability-error, as in [KLSU19], but the bounds could be converted to high-probability bounds at the cost of a $\log\left(d\right)$ factor in the sample complexity.

G.1 Mean Estimation of Binary Product Distributions

We start by stating the theorem for mean estimation of binary product distributions over $\{0,1\}^d$.

Theorem G.1 (Product Distribution Mean Estimation)**.** *There exists a distribution $\mathcal{D}$ over vectors $p \in \left[\frac{1}{3}, \frac{2}{3}\right]^d$ such that, given $p \sim \mathcal{D}$ and n independent samples $X = \left(X_1, \ldots, X_n\right)$ from a binary product distribution P over $\{0,1\}^d$ with mean p, for any $\alpha = \mathcal{O}\left(\sqrt{d}\right)$ and any (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^d$ with $\varepsilon \in [0,1]$ and $\delta = \mathcal{O}\left(\frac{1}{n}\right)$ that satisfies $\mathop{\mathbb{E}}_{X,M}\left[\left\|M\left(X\right) - p\right\|_2^2\right] \leq \alpha^2$, it holds that $n = \Omega\left(\frac{d}{\alpha\varepsilon}\right)$.*

Fact F.1 establishes that Bernoulli distributions are an exponential family. Using that, we have that the mass function of binary product distributions can be written in the form $p_\eta =$

$h(x)\exp\left(\eta^\top T(x) - Z(\eta)\right), \forall x \in \{0,1\}^d$ with:

$$\begin{aligned} h(x) &= 1, \\ T(x) &= x, \\ \eta &= \left(\ln\left(\frac{1-p_1}{p_1}\right), \ldots, \ln\left(\frac{1-p_d}{p_d}\right)\right)^\top, \\ Z(\eta) &= \prod_{j\in[d]} \ln\left(1+e^{\eta_j}\right). \end{aligned}$$

Our process to generate η involves drawing independently $\eta_j \sim \mathcal{U}([-\ln(2), \ln(2)])$. We have that $\eta_j = \ln\left(\frac{1}{p_j} - 1\right) \iff p_j = \frac{1}{1+e^{\eta_j}}$, yielding $\eta_j \in [-\ln(2), \ln(2)] \iff p_j \in \left[\frac{1}{3}, \frac{2}{3}\right]$. Thus, we have $I_j = [-\ln(2), \ln(2)], R_j = 2\ln(2), \|R\|_2^2 = 4\ln^2(2)\, d, \|R\|_\infty = 2\ln(2)$, and $m = 0$.

We now show how to reduce estimating η with an ℓ_2-guarantee under (ε, δ)-DP to estimating p with an ℓ_2-guarantee under the same constraint.

Lemma G.2. *Let $p \in \left[\frac{1}{3}, \frac{2}{3}\right]^d$ be a randomly generated vector and let $X \sim P^{\bigotimes n}$ be a dataset drawn from a binary product distribution over $\{0,1\}^d$ with mean p. If $M : \mathbb{R}^{n\times d} \to \left[\frac{1}{3}, \frac{2}{3}\right]^d$ is an (ε, δ)-DP mechanism satisfying $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - p\|_2^2\right] \le \alpha^2 \le \frac{d}{9}$, then there exists a $T_M : \mathbb{R}^{n\times d} \to [-\ln(2), \ln(2)]^d$ that is also (ε, δ)-DP and satisfies $\mathop{\mathbb{E}}_{X,T_M}\left[\|T_M(X) - \eta\|_2^2\right] \le \frac{81}{4}\alpha^2$.*

Proof. We define T_M to be $T_{M,j}(X) = \ln\left(\frac{1-M_j(X)}{M_j(X)}\right), \forall j \in [d]$. T_M satisfies (ε, δ)-DP by Lemma A.2 and the only randomness employed by T_M is that of M. Consider now the function $g : (0,1) \to \mathbb{R}$ with $g(x) = \ln\left(\frac{1}{x} - 1\right)$ and $g'(x) = -\frac{1}{x(1-x)}$. We have $\eta_j = g(p_j)$ and $T_{M,j}(X) = g(T_j(X))$. Observe now that, by the Mean Value Theorem, we have for some ξ_j between η_j and $M_j(X)$ (implying that $\xi_j \in \left[\frac{1}{3}, \frac{2}{3}\right]$):

$$|g(M_j(X)) - g(p_j)| = |g'(\xi_j)|\,|M_j(X) - p_j| \le \frac{9}{2}|M_j(X) - p_j|.$$

Applying this coordinate-wise, we get $\mathop{\mathbb{E}}_{X,T_M}\left[\|T_M(X) - \eta\|_2^2\right] \le \frac{81}{4}\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - p\|_2^2\right]$, yielding the desired result. ■

We now present a restatement of the previous lemma from the point of view of lower bounds.

Corollary G.3. *Let $p \in \left[\frac{1}{3}, \frac{2}{3}\right]^d$ be a randomly generated vector and let $X \sim P^{\bigotimes n}$ be a dataset drawn from a binary product distribution over $\{0,1\}^d$ with mean p. If any (ε, δ)-DP mechanism $T : \{0,1\}^{n\times d} \to [-\ln(2), \ln(2)]^d$ satisfying $\mathop{\mathbb{E}}_{X,T}\left[\|T(X) - \eta\|_2^2\right] \le \frac{81}{4}\alpha^2 \le \frac{9}{4}d$ requires at least $n \ge n_\eta$ samples, the same sample complexity lower bound holds for any (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \left[\frac{1}{3}, \frac{2}{3}\right]^d$ that satisfies $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - p\|_2^2\right] \le \alpha^2$.*

We are now ready to prove Theorem G.1.

Proof of Theorem G.1. We assume that we have an (ε, δ)-DP mechanism $M : \{0,1\}^{n\times d} \to [-\ln(2), \ln(2)]^d$ with:

$$\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \eta\|_2^2\right] \le \frac{81}{4}\alpha^2 \le \frac{\ln^2(2)\, d}{6} \le \frac{\|R\|_2^2}{24}.$$

By Theorem 2.5 for product distributions with $\eta_j \in [-\ln(2), \ln(2)], \forall j \in [d]$, we have:

$$n\left(2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{\eta}\left[\sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma_T s\right]}\right] + 2\int_T^\infty \mathop{\mathbb{P}}_{X_i}\left[\left|\sum_{j\in[d]}\left(T_j\left(X_i\right) - \mu_{T,j}\right)\right| > \frac{t}{2\ln^3(2)}\right] dt\right)$$
$$\geq \frac{\ln^2(2)\, d}{6}. \tag{28}$$

Observe that, since we have a product distribution and $T(x) = x$, it holds that $\mu_T = p$ and:

$$\Sigma_T = \Sigma = \operatorname{diag}\left\{p_1\left(1-p_1\right), \ldots, p_d\left(1-p_d\right)\right\} \preceq \frac{1}{4}\mathbb{I},$$

for $p_i \in \left[\frac{1}{3}, \frac{2}{3}\right], \forall i \in [d]$.

Thus, we have:

$$\mathop{\mathbb{E}}_{X_{\sim i},M}\left[s^\top \Sigma_T s\right] = \frac{1}{4}\mathop{\mathbb{E}}_{X_{\sim i},M}\left[\|s\|_2^2\right] \leq \frac{\|R\|_\infty^4}{64}\alpha^2 = \frac{\ln^4(2)\,\alpha^2}{4}. \tag{29}$$

Furthermore, we remark that:

$$\left|\sum_{j\in[d]}\left(T_j\left(X_i\right) - \mu_{T,j}\right)\right| = \left|\sum_{j\in[d]}\left(X_i^j - p_j\right)\right| \leq \frac{2d}{3}.$$

Thus, setting $T = \frac{4\ln^3(2)d}{3}$ yields:

$$\int_T^\infty \mathop{\mathbb{P}}_{X_i}\left[\left|\sum_{j\in[d]}\left(T_j\left(X_i\right) - \mu_{T,j}\right)\right| > \frac{t}{2\ln^3(2)}\right] dt = 0. \tag{30}$$

Substituting (29) and (30) into (28), we get:

$$n\left(\frac{8\ln^3(2)\, d\delta}{3} + \frac{\ln^2(2)\,\alpha\varepsilon}{2}\right) \geq \frac{\ln^2(2)\, d}{6}.$$

Setting $\delta \leq \frac{1}{32\ln(2)n}$ results in $n \geq \frac{d}{12\alpha\varepsilon}$ and appealing to Corollary G.3 completes the proof. ■

G.2 Mean Estimation of High-Dimensional Gaussians

Our techniques also apply to recovering lower bounds for mean estimation of a Gaussian with known covariance. By Fact F.2, we have that this class of distributions is an exponential family with $\eta = \mu$ and $T(x) = x$ (implying that $\mu_T = \mu, \Sigma_T = \Sigma$). Therefore, we do not need to resort to a reduction-based approach (as we did with product distributions).

Theorem G.4 (Gaussian Mean Estimation)**.** *Given $\mu \sim \mathcal{U}\left([-1,1]^d\right)$ and $X \sim \mathcal{N}(\mu, \mathbb{I})^{\otimes n}$, for any $\alpha = \mathcal{O}\left(\sqrt{d}\right)$ and any (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to \mathbb{R}^d$ with $\varepsilon \in [0,1]$ and $\delta = \mathcal{O}\left(\frac{\sqrt{d}}{n\sqrt{\ln\left(\frac{n}{\sqrt{d}}\right)}}\right)$ that satisfies $\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \mu\|_2^2\right] \leq \alpha^2$, it holds that $n = \Omega\left(\frac{d^2}{\alpha\varepsilon}\right)$.*

Proof. We assume that the existence of an (ε, δ)-DP mechanism $M : \mathbb{R}^{n\times d} \to [-1,1]^d$ with:

$$\mathop{\mathbb{E}}_{X,M}\left[\|M(X) - \mu\|_2^2\right] \leq \alpha^2 \leq \frac{d}{6} \leq \frac{\|R\|_2^2}{24}.$$

By Theorem 2.5 for Gaussians $\mathcal{N}(\mu, \mathbb{I})$ with $\mu^j \in [-1,1], \forall j \in [d]$, we have:

$$n\left(2\delta T + 2\varepsilon \mathop{\mathbb{E}}_{\mu}\left[\sqrt{\mathop{\mathbb{E}}_{X_{\sim i},M}\left[\|s\|_2^2\right]}\right] + 2\int_T^\infty \mathop{\mathbb{P}}_{X_i}\left[\left|\sum_{j\in[d]}\left(X_i^j - \mu^j\right)\right| > \frac{t}{2}\right] dt\right) \geq \frac{d}{6}. \tag{31}$$

Thus, we have:

$$\mathop{\mathbb{E}}_{X_{\sim i},M}\left[\|s\|_2^2\right] \leq \frac{\|R\|_\infty^4}{16}\alpha^2 = \alpha^2. \tag{32}$$

We remark now that, $X_i^j - \mu^j \sim \mathcal{N}(0,1) \implies Y = \sum_{j\in[d]}\left(X_i^j - \mu^j\right) \sim \mathcal{N}(0,d)$ due to the independence of $X_i^j, j \in [d]$. Thus, we have for the last term of the LHS of (31):

$$\begin{aligned}\int_T^\infty \mathop{\mathbb{P}}_{Y\sim\mathcal{N}(0,d)}\left[|Y| > \frac{t}{2}\right] dt \leq \int_T^\infty e^{-\frac{t^2}{8d}}\, dt = \sqrt{8d}\int_{\frac{T}{\sqrt{8d}}}^\infty e^{-u^2}\, du = \sqrt{2\pi d}\cdot \operatorname{erfc}\left(\frac{T}{\sqrt{8d}}\right) \\ = \sqrt{2\pi d}e^{-\frac{T^2}{8d}},\end{aligned} \tag{33}$$

Substituting (32) and (33) into (31), we get:

$$n\left(2\delta T + 2\alpha\varepsilon + 2\sqrt{2\pi d}e^{-\frac{T^2}{8d}}\right) \geq \frac{d}{6}.$$

It remains to set the values of T and δ appropriately so that $\delta T \geq 2\sqrt{2\pi d}e^{-\frac{T^2}{8d}}$ and $3n\delta T \leq \frac{d}{12}$. We do not perform these calculations here, since they are similar to those in the proof of Theorem 6.5 from [KLSU19]. ■