


General Gaussian Noise Mechanisms and Their Optimality for Unbiased Mean Estimation

Aleksandar Nikolov  

University of Toronto, Canada

Haohua Tang 

University of Toronto, Canada

Abstract

We investigate unbiased high-dimensional mean estimators in differential privacy. We consider differentially private mechanisms whose expected output equals the mean of the input dataset, for every dataset drawn from a fixed bounded domain K in \mathbb{R}^d . A classical approach to private mean estimation is to compute the true mean and add unbiased, but possibly correlated, Gaussian noise to it. In the first part of this paper, we study the optimal error achievable by a Gaussian noise mechanism for a given domain K , when the error is measured in the ℓ_p norm for some $p \geq 2$. We give algorithms that compute the optimal covariance for the Gaussian noise for a given K under suitable assumptions, and prove a number of nice geometric properties of the optimal error. These results generalize the theory of factorization mechanisms from domains K that are symmetric and finite (or, equivalently, symmetric polytopes) to arbitrary bounded domains.

In the second part of the paper we show that Gaussian noise mechanisms achieve nearly optimal error among all private unbiased mean estimation mechanisms in a very strong sense. In particular, for *every input dataset*, an unbiased mean estimator satisfying concentrated differential privacy introduces approximately at least as much error as the best Gaussian noise mechanism. We extend this result to local differential privacy, and to approximate differential privacy, but for the latter the error lower bound holds either for a dataset or for a neighboring dataset, and this relaxation is necessary.

2012 ACM Subject Classification Theory of computation \rightarrow Theory and algorithms for application domains

Keywords and phrases differential privacy, mean estimation, unbiased estimator, instance optimality

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.85

Related Version *Full Version:* <https://arxiv.org/abs/2301.13850>

Funding This research was supported by an NSERC Discovery Grant.

1 Introduction

Unbiased estimation is a classical topic in statistics, and an elegant theory of the existence and optimality of unbiased estimators, and methods for constructing unbiased estimators and proving lower bounds on their variance have been developed over the last century. We refer the reader to the monographs on this topic by Nikulin and Voinov [43, 44], and other standard statistical texts such as [41]. One highlight of this theory is the existence of uniformly minimum variance unbiased estimators (UMVUE), i.e., estimators whose variance is smaller than any other unbiased estimator for *every* value of the parameter. Such estimators can be derived via the Rao-Blackwell and Lehmann-Scheffé theorems. In addition to their mathematical tractability, unbiased estimators also have the nice property that averaging several estimators decreases the mean squared error. Moreover, in some cases unbiased estimators are also minimax optimal, e.g., the empirical mean in many settings.

In this paper we consider the basic problem of mean estimation under the additional constraint that the privacy of the data must be protected. In particular, we study estimators computed by a randomized algorithm \mathcal{M} (a mechanism), that satisfies (ϵ, δ) -differential



© Aleksandar Nikolov and Haohua Tang;

licensed under Creative Commons License CC-BY 4.0

15th Innovations in Theoretical Computer Science Conference (ITCS 2024).

Editor: Venkatesan Guruswami; Article No. 85; pp. 85:1–85:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

privacy [19]. We focus on the setting where the dataset $X = (x_1, \dots, x_n)$ consists of a sequence of n data points from a domain $K \subseteq \mathbb{R}^d$, and our goal is to estimate the mean $\mu(X) := \frac{1}{n} \sum_{i=1}^n x_i$ using an *unbiased* differentially private mechanism \mathcal{M} . Formally, we use the following definition of unbiased mechanisms.

► **Definition 1.** *A mechanism \mathcal{M} that takes as input datasets of n points in some domain $K \subseteq \mathbb{R}^d$, and outputs a vector in \mathbb{R}^d is unbiased over $K \subseteq \mathbb{R}^d$ if, for every dataset $X \in K^n$, we have $\mathbb{E}[\mathcal{M}(X)] = \mu(X)$, where the expectation is over the randomness of \mathcal{M} .*

We note that this definition captures being *empirically* unbiased, i.e., unbiased with respect to the dataset. This is a desirable property when multiple differentially private analyses are performed on the same dataset, as being unbiased allows decreasing the mean squared error by averaging. Moreover, bias has been raised as a concern when publishing official statistics such as census data. The bias of differentially private mean estimation algorithms in this context was examined, for example, in [48, 47]. It is also possible to define unbiased algorithms in a distributional setting, by assuming that X contains i.i.d. samples from an unknown distribution P , and $\mathbb{E}[\mathcal{M}(X)]$ equals the mean of P , with the expectation taken over the randomness in choosing X , and the randomness used by \mathcal{M} . We give such a definition in the full version of the paper, and extend our results to this setting as well.

The prototypical example of an unbiased differentially private mechanisms is given by oblivious, or noise-adding mechanisms, i.e., mechanisms that output $\mathcal{M}(X) := \mu(X) + Z$ for some mean 0 random variable $Z \in \mathbb{R}^d$ drawn from a fixed distribution that's independent of X . A particularly important oblivious mechanism is the Gaussian noise mechanism [19, 18], for which Z is a mean 0 Gaussian random variable in \mathbb{R}^d with covariance matrix proportional to the identity, and scaled proportionally to the ℓ_2 diameter of the domain K . Another family of oblivious mechanisms is given by the matrix mechanism from [34], and, more generally, factorization mechanisms [40, 22]. Factorization mechanisms are defined for a finite domain $K := \{\pm w_1, \dots, \pm w_N\} \subseteq \mathbb{R}^d$, where a dataset X can be represented by a histogram vector $h \in [0, 1]^N$, defined by letting h_i equal the difference between the fraction of points in X equal to w_i and the fraction of points equal to $-w_i$. Then $\mu(X)$ can be written as Wh , where W is the $d \times N$ matrix with columns w_1, \dots, w_N . A factorization mechanism chooses (usually by solving an optimization problem) matrices L and R for which $W = LR$, and outputs $L(Rh + Z) = \mu(X) + LZ$ for a mean 0 Gaussian noise random variable Z with covariance matrix proportional to the identity, and scaled proportionally to the maximum ℓ_2 norm of a column of the matrix R . Thus, factorization mechanisms can be seen as either a post-processing of the Gaussian noise mechanism, or a method of achieving privacy by adding correlated mean 0 Gaussian noise. Factorization mechanisms have received a lot of attention in differential privacy, both from the viewpoint of theoretical analysis [34, 40, 22, 26, 27, 36], and from a more applied and empirical viewpoint [35, 12].

There are, also, natural and widely used non-oblivious unbiased mechanisms. For example, each iteration of the differentially private stochastic gradient descent algorithm for convex minimization [5] uses an unbiased mechanism to compute a private unbiased estimate of the current gradient. The mechanism first samples a subset of the data points, and then adds Gaussian noise to the mean of the sample. The subsampling effectively adds noise to the mean in a way that's not independent from the dataset X . Another example of a non-oblivious unbiased mechanism is randomized response [45]. In its simplest form, randomized response is defined for the domain $K := \{0, 1\}$, and involves releasing, for each data point x_i in X , x_i with probability $\frac{e^\epsilon}{1+e^\epsilon}$ and $1 - x_i$ with probability $\frac{1}{1+e^\epsilon}$. It is easy to compute an unbiased estimator of $\mu(X)$ from these released points, but the additive error of this estimator is

not independent of X . There are also other mechanisms based on randomized response used for higher-dimensional mean estimation in the local model of differential privacy, see e.g. [16, 8, 22]. They are, likewise, unbiased but not oblivious.

Among the mechanisms mentioned above, in high dimensional settings the mechanisms that add (possibly correlated) Gaussian noise, i.e., the Gaussian noise mechanism and factorization mechanisms based on it, tend to give the lowest error for a given dataset size. Let us call mechanisms that add unbiased but potentially correlated Gaussian noise general Gaussian noise mechanisms. This paper is motivated by the following questions.

1. What is the best error achievable by a general Gaussian noise mechanism for a given domain K and a given measure of error? Relatedly, can this error, and the optimal covariance matrix for the noise be computed efficiently?
2. Are general Gaussian noise mechanisms indeed optimal among all unbiased private mean estimators? Can their error be improved on some “nice” datasets?

To answer the first question, we study the best ℓ_p error achievable by a differentially private Gaussian noise mechanism, i.e., a mechanism $\mathcal{M}_\Sigma(X) := \mu(X) + Z$ for $Z \sim N(0, \Sigma)$, when $X \in K^n$ for a bounded domain K . Here, $N(\mu, \Sigma)$ is the Gaussian distribution with mean μ and covariance matrix Σ . Let us denote the unit Euclidean ball in \mathbb{R}^d by $B_2^d := \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$, where $\|x\|_p := (|x_1|^p + \dots + |x_d|^p)^{1/p}$ is the standard ℓ_p norm on \mathbb{R}^d . For a positive semidefinite matrix $M \in \mathbb{R}^{d \times d}$ and a real number $p \geq 1$, we define $\text{tr}_p(M) := (\sum_i^d M_{ii}^p)^{1/p}$. We define $\text{tr}_\infty(M) := \max_{i=1}^d M_{ii}$. Note that $\text{tr}_p(M)$ is simply the ℓ_p -norm of the diagonal entries of M . Finally, for sets K and L , let us write

$$K \subseteq_{\leftrightarrow} L \iff \exists v \in \mathbb{R}^d, K + v \subseteq L$$

We now define the following key quantity.

► **Definition 2.** For a bounded set $K \subseteq \mathbb{R}^d$ and $p \in [2, \infty]$, we define

$$\Gamma_p(K) := \inf \left\{ \sqrt{\text{tr}_{p/2}(AA^T)} : K \subseteq_{\leftrightarrow} AB_2^d \right\},$$

where the infimum is over $d \times d$ matrices A , and

The definition of Γ_p is motivated by the next theorem. Its proof is given in Section 2.2

► **Theorem 3.** For any $p \in [2, \infty]$, any $\varepsilon > 0$, any $\delta \leq e^{-\varepsilon}$, and any bounded set $K \subseteq \mathbb{R}^d$, there exists a mechanism \mathcal{M} that is unbiased over K , and, for any $X \in K^n$ achieves

$$(\mathbb{E} \|\mathcal{M}(X) - \mu(X)\|_p^2)^{1/2} \lesssim \frac{\sqrt{\min\{p, \log(2d)\} \log(1/\delta)} \Gamma_p(K)}{\varepsilon n},$$

and satisfies (ε, δ) -differential privacy. The mechanism outputs $\mu(X) + Z$, where Z is a mean 0 Gaussian random variable with covariance matrix proportional to AA^T , for a matrix A such that $K \subseteq_{\leftrightarrow} AB_2^d$ and $\sqrt{\text{tr}_{p/2}(AA^T)} \lesssim \Gamma_p(K)$.

In addition, it is not hard to also show that this bound on $(\mathbb{E} \|\mathcal{M}(X) - \mu(X)\|_p^2)^{1/2}$ is also tight up to the $\sqrt{\min\{p, \log(2d)\}}$ factor (see the proof of Lemma 23). Thus, $\Gamma_p(K)$ nearly captures the best ℓ_p error we can achieve by a general Gaussian noise mechanism for mean estimation over K .

In the special case when $p \in \{2, \infty\}$ and K is a finite set symmetric around 0, the general Gaussian mechanism above is equivalent to a factorization mechanism, as we show in Section 2.4. Factorization mechanisms, as defined, for example, in [34, 22], can be suboptimal

for non-symmetric K . A trivial example is a singleton $K = \{w\}$ for some $w \neq 0$, for which any factorization mechanism would add non-zero noise, but the optimal mechanism just outputs $\mu(X) := w$ with no noise. This issue is the reason we allow K to be shifted in the definition of $\Gamma_p(K)$, and one can similarly modify factorization mechanisms by allowing a shift of K , i.e., shifting the columns of the matrix W by a fixed vector. It is also not difficult to extend standard factorization mechanism formulations from minimizing ℓ_p error for $p \in \{2, \infty\}$, to minimizing ℓ_p error for any $p \geq 2$. Independently from our work, this was also done in a recent paper by Xiao, He, Zhang, and Kifer, who, in addition, also considered error measures that are convex functions of the per-coordinate variances [46]. These modification allow deriving factorization mechanisms equivalent to the general Gaussian mechanism achieving error $\Gamma_p(K)$ when K is finite, as shown in Theorem 19 below.

Extending factorization mechanisms to infinite K , or K of size exponential in the dimension d , however, is more challenging. A fundamental issue is that the matrices W and R involved in the definition of a factorization mechanism are infinite or exponentially sized in these cases. Heuristic solutions tailored to specific structured K have been proposed, for example in [35, 46]. With the definition of general Gaussian mechanisms, and of $\Gamma_p(K)$, we take a different approach, moving away from factorizations and instead focusing on optimizing the covariance matrix of the noise. This is the role of the matrix A in the definition of $\Gamma_p(K)$: it can be seen as a proxy for the covariance matrix of the noise, which is, per Theorem 3, proportional to AA^T . Equivalently, our definition of $\Gamma_p(K)$ can be thought of as formulating a factorization mechanism only in terms of the left matrix of the factorization, without explicitly writing the right matrix. The benefit in this approach is that the covariance matrix has size $d \times d$, independently of the size of K . Of course, optimizing the covariance matrix may still be computationally expensive, depending on how complicated K is. Nevertheless, we show that finding an optimal A approximately achieving $\Gamma_p(K)$ can be done in polynomial time under natural condition. In particular, we show that $\Gamma_p(K)$ equals the value of the convex optimization problem

$$\begin{aligned} \Gamma_p(K)^2 &= \min \operatorname{tr}_{p/2}(M) \\ &\text{s.t.} \\ &(x+v)^T M^{-1}(x+v) \leq 1 \quad \forall x \in K, \\ &M \succ 0, v \in \mathbb{R}^d. \end{aligned}$$

In the full version of the paper we prove this equivalence, and show that this optimization problem can be approximately solved in polynomial time using the ellipsoid method, assuming the existence of an oracle that approximately solves the quadratic maximization problem $\max_{x \in K} (x+v)^T M^{-1}(x+v)$ for a given $v \in \mathbb{R}^d$ and a given positive definite matrix M . (The notation $M \succ 0$ above means that M is positive definite.) Beyond finite K , such oracles exist for many classes of K , e.g., affine images of ℓ_p and Schatten- p balls when $p \geq 2$, and affine images of other symmetric norms and unitarily invariant matrix norms: see [7] for more information. As one example, we can compute $\Gamma_p(K)$ over zonotopes: sets of the type $K := [u_1, v_1] + \dots + [u_N, v_N]$ where $u_1, \dots, u_N \in \mathbb{R}^d$, and $v_1, \dots, v_N \in \mathbb{R}^d$ are given explicitly, and $[u_i, v_i]$ is the line segment joining u_i and v_i . Such sets are simply affine images of the ℓ_∞ ball $[-1, +1]^N$, and for them the maximization problem $\max_{x \in K} (x+v)^T M^{-1}(x+v)$ can be solved approximately using algorithmic versions of Grothendieck's inequality [24, 1]. These computational results are the first general theoretical guarantees that allow optimizing the Gaussian noise covariance matrix for domains K that are not explicitly presented finite sets or polytopes in vertex representation.

In addition to computational tractability, we also prove a number of properties of $\Gamma_p(K)$ which significantly generalize known facts about factorization norms and factorization mechanisms. Let us highlight some of these properties, whose proofs are given in the full version of the paper:

- $\Gamma_p(K)$ behaves like a norm on convex sets: it is monotone under inclusion, absolutely homogeneous under scaling, i.e., $\Gamma_p(tK) = |t|\Gamma_p(K)$, and satisfies the triangle inequality with respect to Minkowski sum, i.e., $\Gamma_p(K + L) \leq \Gamma_p(K) + \Gamma_p(L)$.
- $\Gamma_p(K)$ admits a nice dual characterization. For example, for $p = 2$ we have

$$\Gamma_2(K) = \sup\{\text{tr}(\text{cov}(P)^{1/2}) : P \in \Delta(K)\}.$$

Above, $\Delta(K)$ is the set of probability measures supported on K , $\text{cov}(P)$ is the covariance matrix of a probability distribution P , and $\text{cov}(P)^{1/2}$ is its positive semidefinite square root of $\text{cov}(P)$. This characterization (and its generalization to $p > 2$, both given in Section 2.5) is particularly useful for proving lower bounds on $\Gamma_p(K)$.

Next we turn to the question of the optimality of general Gaussian noise mechanisms. Our main result is the following theorem.

► **Theorem 4.** *Let $c > 0$ be a small enough absolute constant, let $p \in [2, \infty]$, and let \mathcal{M} be an (ε, δ) -differentially private mechanism that is unbiased over a bounded set $K \subseteq \mathbb{R}^d$. If $\varepsilon \leq c$, and $\delta \leq \min\left\{\frac{c}{n}, \frac{c\varepsilon^2}{d^2}\right\}$, then the following holds. For any dataset $X \in K^n$, there exists a neighboring dataset $X' \in K^n$ (which may equal X) for which*

$$\sqrt{\mathbb{E}\left[\|\mathcal{M}(X') - \mu(X')\|_p^2\right]} \gtrsim \frac{\Gamma_p(K)}{n\varepsilon}.$$

Above, the notation $A \gtrsim B$ for two quantities A and B is used to mean that there exists an absolute constant $c > 0$ such that $A \geq cB$.

Theorem 4, together with Theorem 3, shows that, for any ℓ_p norm for $p \geq 2$, and any unbiased mechanism \mathcal{M} over any domain K , there is a general Gaussian noise mechanism that has ℓ_p error not much larger than that of \mathcal{M} . Moreover, this is true in an *instance optimal* sense: every dataset X has a neighbor X' for which the correlated Gaussian noise mechanism has smaller error (up to small factors). This somewhat complicated version of instance optimality is necessary, since, for any dataset X , there is an unbiased $(0, \delta)$ -differentially private mechanism that has error 0 on X . Roughly speaking, this mechanism outputs $\mu(X)$ on X , and on any other dataset outputs $\mu(X)$ with probability $1 - \delta$ and some other output with probability δ , chosen to make the mechanism unbiased. This illustrates the key difficulty in proving a result such as Theorem 4: one has to rule out the possibility that a mechanism can “cheat” by hard-coding the true answer for one dataset, thus outperforming an “honest” mechanism on that dataset.

We note that Theorem 4 is a simplification of our main result, and what we prove is actually stronger: we show that either the error of \mathcal{M} on X is comparable to that of a general Gaussian noise mechanism, or there is a neighboring dataset X' for which the error grows with $\frac{1}{\sqrt{\delta}}$. Since usually δ is chosen to be very small, this means that, on every input, \mathcal{M} is either dominated by a general Gaussian noise mechanism, or has huge error in the neighborhood of the input.

Theorem 4 strengthens results showing the optimality of Gaussian noise mechanisms among oblivious mechanisms [22] to the more general class of unbiased mechanisms. As mentioned above, some natural unbiased mechanisms fail to be oblivious. We also consider the class of unbiased mechanisms more natural and robust than the class of oblivious mechanisms. It is worth noting further that Gaussian noise mechanisms are known to be

approximately optimal *in the worst case* for sufficiently large datasets among all differentially private mechanisms. This follows, for example, by reductions from general mechanisms to oblivious mechanisms in [6, 22].

Independently from our work, unbiased mechanisms for private mean estimation were also recently investigated by Kamath, Mouzakis, Regehr, Singhal, Steinke, and Ullman [29]. They focus on mechanisms that take as input independent samples from an unknown one-dimensional distribution in some family, e.g., distributions with bounded k -th moments, and study the trade-off between the bias of the mechanism with respect to the distribution's mean, and its variance. Their results are incomparable to ours: on the one hand, they study one-dimensional mean estimation and prove worst-case (i.e., minimax) lower bounds, rather than instance per instance lower bounds; on the other hand, they give a tight trade-off between bias and variance in their setting, whereas we only consider unbiased mechanisms. They also show that for Gaussian mean estimation no purely private (i.e., (ϵ, δ) -differentially private with $\delta = 0$) unbiased mechanism can achieve finite variance.

The kind of instance optimality guarantee shown in Theorem 4 is reminiscent of the theory of uniformly minimum variance unbiased estimators in statistics. Our result is also related to results by Asi and Duchi [3, 2], and by Huang, Liang, and Yi [28] who also study instance optimality for unbiased differentially private algorithms. Asi and Duchi also focus on notions of being unbiased defined with respect to the dataset, but only treat pure differential privacy, i.e., the $\delta = 0$ setting. The results in [3] are tailored to one-dimensional estimation, and the lower bounds proved there are not sufficiently strong to prove the result in Theorem 4. The results in [2] do extend to higher-dimensional problems, but use a non-standard definition of unbiased mechanisms that is incomparable with the more standard definition we use. Their lower bounds also rely on some strong regularity assumptions that we do not make. Finally, the work of Huang, Liang, and Yi [28] is not restricted to unbiased mechanisms, and, similarly to our results, considers mean estimation and optimality for the neighborhood of every dataset (in fact only considering datasets resulting from removing points). In their results, however, optimality is proved only up to a factor of at least the square root of the dimension, making them less interesting in high dimensions.

Let also mention that there are other approaches to instance optimality in differential privacy. One approach is based on local minimax rates, initiated by Ruan and Duchi [17], and Asi and Duchi in [3], and explored further in [15, 37]. The local minimax rates framework is not suitable for proving strong instance optimality for high dimensional problems, as noted in [37]. Another approach, similar to that of [28], based on optimality with respect to subsets of a dataset, was considered in [13]. Their results are also restricted to one-dimensional problems. In general, existing instance optimality results tend to only be meaningful in low dimensional settings, and our work is a rare example of instance optimality up to small factors for a high-dimensional problem.

In addition to Theorem 4, we also show analogous, and, in fact, stronger results for other variants of differential privacy. For concentrated differential privacy [21, 9], our result holds for *every* dataset $X \in K^n$. This is also the case for local differential privacy [23, 30], recalled in the next definition.

► **Definition 5.** A (non-interactive) locally differentially private mechanism \mathcal{M} is defined by a tuple of randomized algorithms $\mathcal{A}, \mathcal{R}_1, \dots, \mathcal{R}_n$, where each \mathcal{R}_i , called a local randomizer, receives a single data point from a domain K , and is $(\epsilon, 0)$ -differentially private with respect to that point, and the algorithm's output on a dataset $X := (x_1, \dots, x_n)$ is defined by

$$\mathcal{M}(X) := \mathcal{A}(\mathcal{R}_1(x_1), \dots, \mathcal{R}_n(x_n)),$$

i.e., the postprocessing of the outputs of the randomizers \mathcal{R}_i by the aggregator \mathcal{A} . Moreover, each algorithm uses independent randomness.

Local differential privacy captures settings in which there is no trusted central authority, and instead every data owner ensures the privacy of their own data. Mean estimation algorithms in local differential privacy are typically based on randomized response, and are not oblivious even when they are unbiased. The recent work [4] studies unbiased mechanisms for mean estimation in this model when $K = B_2^d$, and gives a tight characterization of algorithms that achieve optimal error in ℓ_2 . Here we extend their results to arbitrary bounded domains K , and to error measured in other norms, albeit with a somewhat less tight characterization. Our characterization shows that the instance optimal unbiased locally private mechanism for mean estimation has every local agent add a linear transformation of (non-oblivious) subgaussian noise. The precise results are given in the full version.

As mentioned above, we also prove a result analogous to Theorem 4 for mechanisms that are unbiased in a distributional sense, i.e., where the expectation of the mechanism's output, given an input dataset drawn i.i.d. from some distribution over the domain K , matches the mean of the distribution. In that case we show that, for every distribution P on K , the ℓ_p error of every distributionally unbiased mechanism is at least $\gtrsim \frac{\Gamma_p(K)}{n\varepsilon}$ either on P , or on some distribution Q that is at distance at most $\frac{1}{n}$ from P in total variation distance. Our result thus shows that, in the distributional setting, too, unbiased mechanisms are dominated by general Gaussian noise mechanisms in the neighborhood of every input distribution. The precise results are given in the full version.

Using the properties of the Γ_p function that we establish, we can prove lower bounds on $\Gamma_p(K)$ for domains K that naturally appear in applications. Together with Theorem 4, these lower bounds imply concrete lower bounds on the ℓ_p error of any unbiased mechanism that hold in the neighborhood of every dataset. Analogous results would for distributionally unbiased mechanisms and unbiased locally differentially private mechanisms also follow via the appropriate variant of Theorem 4. First, we state one such lower bound for estimating moment tensors.

► **Theorem 6.** *Let $c > 0$ be a small enough absolute constant, let $p \in [2, \infty]$ and let ℓ be a positive integer. Let \mathcal{M} be an (ε, δ) -differentially private mechanism that takes as input datasets in $(B_2^d)^n$. Suppose that for every dataset $X := (x_1, \dots, x_n)$,*

$$\mathbb{E}[\mathcal{M}(X)] = M_\ell(X) := \frac{1}{n} \sum_{i=1}^n x_i^{\otimes \ell}.$$

If $\varepsilon \leq c$, and $\delta \leq \min\left\{\frac{c}{n}, \frac{c\varepsilon^2}{d^{2\ell}}\right\}$, then, for any dataset $X \in (B_2^d)^n$, there exists a neighboring dataset $X' \in (B_2^d)^n$ (which may equal X) for which

$$\sqrt{\mathbb{E}\left[\|\mathcal{M}(X') - M_\ell(X')\|_p^2\right]} \gtrsim \frac{1}{\varepsilon n} \left(\frac{d}{\ell}\right)^{\ell/p}.$$

This lower bound implies that estimating the ℓ -th moment tensor of a dataset in B_2^d via an (ε, δ) -differentially private unbiased mechanism requires ℓ_2 error at least on the order $\frac{d^{\ell/2}}{\varepsilon n}$ for small enough ε and δ and constant ℓ . This lower bound is nearly matched by the basic Gaussian noise mechanism, which adds independent Gaussian noise to each coordinate of the tensor. At the same time, the projection mechanism [40] allows ℓ_2 error on the scale of $\frac{d^{1/4} \log(1/\delta)^{1/4}}{\sqrt{\varepsilon n}}$ for any constant ℓ , which is much smaller for $n \ll d^{\ell-1/2}$. This follows from an analysis similar to that in [20]: see the recent paper [14] for the argument in the $\ell = 2$ case. Theorem 6 thus illustrates the cost of using private unbiased mechanisms: while they produce answers that are accurate in expectation, they can incur much more error than biased

algorithms, and this is true in the neighborhood of every input. We note that Theorem 6 can easily be extended to unbiased estimates of the covariance, and also to a distributional setting, and show a similar gap between biased and unbiased mechanisms.

Analogous techniques also imply tight upper and lower bounds on estimating ℓ -way marginals on d -dimensional binary data when $\ell = O(1)$. Before we state these bounds, let us recall the general connection between query release and mean estimation. Suppose that $Q = (q_1, \dots, q_k)$ is a sequence of statistical queries on a universe \mathcal{X} , also known as a workload. This means that each q_i is specified by a function $q_i : \mathcal{X} \rightarrow \mathbb{R}$, and, overloading notation, its value on a dataset $X := (x_1, \dots, x_n) \in \mathcal{X}^n$ is defined by $q_i(X) := \frac{1}{n} \sum_{i=1}^n q_i(x_i)$. Overloading notation again, we can define $Q(X) := (q_1(X), \dots, q_k(X))$ to be the sequence of true answers to the queries in Q on the dataset X . When $n = 1$, i.e., X consists of the single data point $x \in \mathcal{X}$, we write $Q(x)$ rather than $Q((x))$. The problem of privately releasing an approximation to $Q(X)$ is equivalent to mean estimation over the set $K_Q := \{Q(x) : x \in \mathcal{X}\}$ in the following sense. Given a dataset $X = (x_1, \dots, x_n) \in \mathcal{X}^n$, we can construct a dataset $f(X) := (Q(x_1), \dots, Q(x_n)) \in K_Q^n$. Clearly, $\mu(f(X)) = Q(X)$, and any differentially private algorithm \mathcal{M} for (unbiased) mean estimation over K_Q gives an (unbiased) differentially private algorithm for releasing $Q(X)$, simply by running $\mathcal{M}(f(X))$. We can also choose an inverse g of f by choosing, for each $y \in K_Q$, some $x \in \mathcal{X}$ such that $Q(x) = y$, and defining $g(Y)$ for $Y = (y_1, \dots, y_n) \in K_Q^n$ as the function that replaces each y_i with the chosen x_i giving $Q(x_i) = y_i$. This shows, in turn, that an (unbiased) differentially private mechanism \mathcal{M} that releases $Q(X)$ gives an (unbiased) private mean estimation algorithm $\mathcal{M}(f(Y))$. These reductions preserve the privacy parameters, the property of being unbiased, and the error.

Let us now specialize this discussion to releasing ℓ -way marginal queries. Let $Q_{d,\ell}^{\text{marg}}$ be the statistical queries over the universe $\mathcal{X} := \{0, 1\}^d$ where each query $q_{s,\beta}$ in $Q_{d,\ell}^{\text{marg}}$ is defined by a sequence of ℓ indices $s := (i_1, \dots, i_\ell) \in [d]^\ell$ and a sequence of ℓ bits $\beta := (\beta_1, \dots, \beta_\ell) \in \{0, 1\}^\ell$, and has value $q_{s,\beta}(x) = \prod_{j=1}^\ell |x_{i_j} - \beta_j|$ on every $x \in \mathcal{X}$. For an easier to read notation, let us write $K_{d,\ell}^{\text{marg}} := K_{Q_{d,\ell}^{\text{marg}}}$. By analyzing $\Gamma_p(K_{d,\ell}^{\text{marg}})$, we derive the following bound on the error necessary to release unbiased estimates of the ℓ -way marginal queries.

► **Theorem 7.** *Let $c > 0$ be a small enough absolute constant, let $p \in [2, \infty]$ and let ℓ be a positive integer. Let \mathcal{M} be an (ε, δ) -differentially private mechanism that takes as input datasets in $(\{0, 1\}^d)^n$. Suppose that for every dataset $X := (x_1, \dots, x_n)$,*

$$\mathbb{E}[\mathcal{M}(X)] = Q_{d,\ell}^{\text{marg}}(X).$$

If $\varepsilon \leq c$, and $\delta \leq \min\left\{\frac{c}{n}, \frac{c\varepsilon^2}{(2d)^{2\ell}}\right\}$, then, for any dataset $X \in (\{0, 1\}^d)^n$, there exists a neighboring dataset $X' \in (\{0, 1\}^d)^n$ (which may equal X) for which

$$\sqrt{\mathbb{E}\left[\left\|\mathcal{M}(X') - Q_{d,\ell}^{\text{marg}}(X')\right\|_p^2\right]} \gtrsim \frac{d^{\frac{\ell}{2} + \frac{\ell}{p}}}{(2\sqrt{2}\ell)^\ell}.$$

Theorem 7 shows a similar gap as Theorem 6 between the optimal error achievable by unbiased and biased mechanisms for releasing marginals. For constant ℓ , the lower bound in Theorem 7 nearly matches the error achievable by adding i.i.d. Gaussian noise. By contrast, the error achieved by the projection mechanism or the private multiplicative weights mechanism, which can be biased, is much smaller for moderate values of n and $\ell > 1$: for example, the projection mechanism achieves error on the order of $\frac{d^{1/4} \log(1/\delta)^{1/4}}{\sqrt{\varepsilon n}}$ for any

constant ℓ [20]. In the case of $\ell = 1$, Theorem 7 gives lower bounds on the error achieved by unbiased mechanisms for one-way marginals that match, up to the dependence on δ , the lower bounds against all (ε, δ) -differentially private algorithms proved via fingerprinting codes [10]. While our lower bounds are for restricted mechanisms, they hold in the neighborhood of every input dataset, rather than for a worst case dataset as the fingerprinting lower bounds.

The proofs of Theorems 6 and 7, and more precise upper and lower bounds on $\Gamma_p(K_{d,\ell}^{\text{marg}})$ are presented in the full version.

1.1 Techniques

In terms of techniques for proving Theorem 4 and its extensions, we combine a technique from [22], developed for proving lower bounds on oblivious mechanisms, with classical results from the statistical theory of unbiased estimation. The key insight from [22] is that, in order to prove a theorem like Theorem 4, it is sufficient to show that, for every unit vector $\theta \in \mathbb{R}^d$, the variance of $\theta^T \mathcal{M}(X)$ is bounded below in terms of the width of K in the direction of θ . This reduction is explained in Section 3. While proving a lower bound on the worst-case variance of a one-dimensional private mechanism like $\theta^T \mathcal{M}(X)$ is easy, the challenge is that the lower bound must hold for a fixed X that is not allowed to vary with θ . This is trivial for oblivious mechanisms, but not for unbiased mechanisms. Nevertheless, we show that the classical Hammersley-Chapman-Robins (HCR) bound [25, 11] implies the one-dimensional variance lower bounds we need for pure and concentrated differential privacy. The situation is more subtle for approximate differential privacy, i.e., (ε, δ) -differential privacy for $\delta > 0$. The main technical issue is that applying the HCR bound requires proving an upper bound on the χ^2 divergence between the output distributions $\mathcal{M}(X)$ and $\mathcal{M}(X')$ of a differentially private mechanism \mathcal{M} on two neighboring datasets X and X' . No such finite bound need exist for (ε, δ) -differentially private mechanisms when $\delta > 0$. To get around this issue, we modify one of the output distributions $\mathcal{M}(X)$ and $\mathcal{M}(X')$ so that the χ^2 divergence becomes bounded, and, moreover, the expectations of the two distributions does not change much, unless one of the two distributions already has huge variance. Then we can carry out a win-win analysis: either one of $\mathcal{M}(X)$ or $\mathcal{M}(X')$ has huge variance, or the HCR bound can be applied to them.

1.2 Notation

As already noted, we use $\|x\|_p := (|x_1|^p + \dots + |x_d|^p)^{1/p}$ for the ℓ_p norm of a vector $x \in \mathbb{R}^d$, and $B_p^d := \{x \in \mathbb{R}^d : \|x\|_p \leq 1\}$ for the corresponding unit ball. We write the standard inner product in \mathbb{R}^d as $\langle x, y \rangle := x_1 y_1 + \dots + x_d y_d = x^T y$ for $x, y \in \mathbb{R}^d$. For a $d \times N$ matrix M , we define the $\ell_p \rightarrow \ell_q$ operator norm by $\|M\|_{p \rightarrow q} := \sup_{x \in \mathbb{R}^N : \|x\|_p = 1} \|Mx\|_q$. Note that $\|M\|_{1 \rightarrow 2}$ equals the largest ℓ_2 norm of a column of M , and $\|M\|_{2 \rightarrow \infty}$ equals the largest ℓ_2 norm of a row of M . We also define the Frobenius (or Hilbert-Schmidt) norm $\|M\|_F := \text{tr}(M^T M)^{1/2}$. Note that this is just the ℓ_2 norm of M treated as a vector.

For a $d \times d$ matrix M , we use $M \succeq 0$ to denote that M is positive semidefinite, i.e., M is symmetric and satisfies $x^T M x \geq 0$ for all $x \in \mathbb{R}^d$. If M is also positive definite, i.e., positive semidefinite and non-singular, we write $M \succ 0$. We write $A \succeq B$ and $B \preceq A$ when $A - B \succeq 0$ for two $d \times d$ matrices A and B . We write \sqrt{M} or $M^{1/2}$ for the principle square root of a positive semidefinite matrix M , i.e., \sqrt{M} is a positive semidefinite matrix such that $(\sqrt{M})^2 = M$.

For a probability distribution P , we use $X \sim P$ to denote the fact that the random variable X is distributed according to P . We use $\mathbb{E}_{X \sim P}[f(X)]$ to denote the expectation of the function $f(X)$ when X is a random variable distributed according to P . We use $\text{cov}(P)$ to denote the covariance matrix of a distribution P on \mathbb{R}^d .

2 Gaussian Noise Mechanisms

In this section, we introduce the general Gaussian noise mechanism for estimating means in an arbitrary bounded domain K . This mechanism generalizes known factorization mechanisms, as we discuss later on in the section. The mechanism's error bound generalizes factorization norms: we prove this fact, and some other important properties in this section, as well.

2.1 Preliminaries on Concentrated Differential Privacy

Our mechanism's privacy guarantees are most cleanly stated in the language of concentrated differential privacy. We recall the definition of this variant of differential privacy here, together with some basic properties of it.

Before we define concentrated differential privacy, it is convenient to define a “ratio of probability densities” for general probability distributions. We use the following (standard) definition.

► **Definition 8.** For two probability distribution P and Q over the same ground set, we define $\frac{dP}{dQ}$ as follows. Let $R = \frac{P+Q}{2}$ be a reference distribution, and denote by $\frac{dP}{dR}, \frac{dQ}{dR}$ the Radon-Nykodim derivatives of P and Q with respect to R . Then we take $\frac{dP}{dQ} := \frac{dP/dR}{dQ/dR}$ where the ratio is defined to be ∞ if the denominator is 0 while the numerator is positive.

We also recall the definition of Rényi divergence.

► **Definition 9.** For two probability distributions P and Q over the same ground set, and a real number $\alpha > 1$, the Rényi divergence $D_\alpha(P\|Q)$ of order α is defined by

$$D_\alpha(P\|Q) := \frac{1}{\alpha - 1} \ln \mathbb{E}_{X \sim Q} \left[\left(\frac{dP}{dQ}(X) \right)^\alpha \right].$$

We are now ready to define zero-concentrated differential privacy, following [9].

► **Definition 10.** A mechanism \mathcal{M} satisfies ρ -zero concentrated differential privacy (ρ -zCDP) if, for all neighboring datasets X, X' , we have

$$D_\alpha(\mathcal{M}(X)\|\mathcal{M}(X')) \leq \rho\alpha.$$

Concentrated differential privacy satisfies many nice properties: it has a simple and optimal composition theorem, is invariant under post-processing, and implies some protection to small groups in addition to protecting the privacy of individuals. We refer to [9] for details. The properties we need are stated in the following lemmas, and proofs can be found in [9].

► **Lemma 11.** Suppose that \mathcal{M}_1 is a ρ_1 -zCDP mechanism, and, for every y in the range of \mathcal{M}_1 , $\mathcal{M}_2(y, \cdot)$ is a ρ_2 -zCDP mechanism. Then the composition \mathcal{M} defined on dataset X by $\mathcal{M}(X) := \mathcal{M}_2(\mathcal{M}_1(X), X)$ satisfies $(\rho_1 + \rho_2)$ -zCDP.

In particular, if \mathcal{M} satisfies ρ -zCDP, and \mathcal{A} is a randomized algorithm defined on the range of \mathcal{M} , then the post-processed mechanism defined on dataset X by $\mathcal{A}(\mathcal{M}(X))$ satisfies ρ -zCDP as well.

► **Lemma 12.** If a mechanism \mathcal{M} satisfies ρ -zCDP, then, for any $\delta > 0$, \mathcal{M} also satisfies $(\rho + 2\sqrt{\rho \log(1/\delta)}, \delta)$ -differential privacy.

► **Lemma 13.** Suppose that $f : K^n \rightarrow \mathbb{R}^d$ is a function on size n datasets drawn from the domain K with ℓ_2 sensitivity at most Δ , i.e., for any two neighboring datasets X and X' we have $\|f(X) - f(X')\|_2 \leq \Delta$. Then the mechanism that on input X outputs $\mathcal{M}(X) := f(X) + Z$ for $Z \sim N(0, \sigma^2 I)$ satisfies $\frac{\Delta^2}{2\sigma^2}$ -zCDP.

2.2 A Gaussian Noise Mechanism for General Domains

Recall the notation $\text{tr}_p(M)$, defined in the Introduction: for a positive semidefinite matrix $M \in \mathbb{R}^{d \times d}$ and a real number $p \geq 1$, we have $\text{tr}_p(M) := (\sum_i^d M_{ii}^p)^{1/p}$. Moreover, we have $\text{tr}_\infty(M) := \max_{i=1}^d M_{ii}$. Recall also that $\text{tr}_p(M)$ is simply the ℓ_p -norm of the diagonal entries of M . The next lemma notes a few useful properties of $\text{tr}_p(M)$ which follow from this observation.

► **Lemma 14.** *The function tr_p satisfies the following properties for any $p \geq 1$:*

1. *for any positive semidefinite matrix M , $\text{tr}_p(M) = 0$ implies $M = 0$;*
2. *for any real number $t \geq 0$, and any positive semidefinite matrix M , $\text{tr}_p(tM) = t \text{tr}_p(M)$;*
3. *for any two positive semidefinite matrices M_1 , and M_2 , $\text{tr}_p(M_1 + M_2) \leq \text{tr}_p(M_1) + \text{tr}_p(M_2)$;*
4. *for any two $1 \leq p \leq q \leq \infty$, and any $d \times d$ positive semidefinite matrix M , $\text{tr}_q(M) \leq \text{tr}_p(M) \leq d^{\frac{1}{q} - \frac{1}{p}} \text{tr}_q(M)$.*

Proof. All except the first property are immediate from the observation that when M is a positive semidefinite matrix, $\text{tr}_p(M)$ is the ℓ_p norm of its diagonal entries. This observation also shows that when $\text{tr}_p(M) = 0$ and $M \succeq 0$, the diagonal entries of M are 0. But, since the largest absolute value of any entry of a positive semidefinite matrix is achieved on the diagonal, this also implies that $M = 0$. ◀

Next recall our notation for inclusion of sets up to shifting: for subsets K and L of \mathbb{R}^d , we write

$$K \subseteq_{\leftrightarrow} L \iff \exists v \in \mathbb{R}^d, K + v \subseteq L.$$

Finally, we recall the $\Gamma_p(K)$ function defined in the Introduction as

$$\Gamma_p(K) := \inf \left\{ \sqrt{\text{tr}_{p/2}(AA^T)} : K \subseteq_{\leftrightarrow} AB_2^d \right\},$$

where the infimum is over $d \times d$ matrices A .

The next theorem is the core of the proof of Theorem 3 from the Introduction, and is a slight generalization of Corollary 2.8 from [40]. We defer the proof to the full version.

► **Theorem 15.** *Suppose that $p \in [2, \infty]$, that $K \subseteq \mathbb{R}^d$ is a bounded set, and that for some $d \times d$ matrix A , $K \subseteq_{\leftrightarrow} AB_2^d$. Then the mechanism \mathcal{M} that, on input $X \in K^n$, outputs $\mathcal{M}(X) := \mu(X) + Z$, where $Z \sim N(0, \frac{4}{\varepsilon^2 n^2} AA^T)$, satisfies $\frac{\varepsilon^2}{2}$ -zCDP. In particular, \mathcal{M} is an unbiased $\frac{\varepsilon^2}{2}$ -zCDP mechanism \mathcal{M} that, for any dataset $X \in K^n$ achieves*

$$(\mathbb{E} \|\mathcal{M}(X) - \mu(X)\|_p^2)^{1/2} \lesssim \frac{\sqrt{\min\{p, \log(2d)\}} \text{tr}_{p/2}(AA^T)^{1/2}}{\varepsilon n}.$$

Taking A to achieve $\Gamma_p(K)$ in Theorem 15, and also using Lemma 12, gives Theorem 3. We also have the following corollary for zCDP.

► **Corollary 16.** *For any $p \in [2, \infty]$, any $\varepsilon > 0$, and any bounded set $K \subseteq \mathbb{R}^d$, there exists a mechanism \mathcal{M} that is unbiased over K , for any $X \in K^n$ achieves*

$$(\mathbb{E} \|\mathcal{M}(X) - \mu(X)\|_p^2)^{1/2} \lesssim \frac{\sqrt{\min\{p, \log(2d)\}} \Gamma_p(K)}{\varepsilon n},$$

and satisfies $\frac{\varepsilon^2}{2}$ -zCDP.

In the full version, we also give a variant of this result for local differential privacy.

2.3 Basic Properties of $\Gamma_p(K)$

In this subsection we give some important properties of the Γ_p function. Proofs and further properties are deferred to the full version of the paper. The properties are given in the following theorem.

► **Theorem 17.** *The function Γ_p satisfies the following properties:*

1. (invariance with respect to convex hulls) for any bounded set $K \subseteq \mathbb{R}^d$, $\Gamma_p(K) = \Gamma_p(\overline{\text{conv}} K)$;
2. (monotonicity) whenever $K \subseteq_{\leftrightarrow} \overline{\text{conv}} L$, we have $\Gamma_p(K) \leq \Gamma_p(L)$;
3. (homogeneity) for any bounded set $K \subseteq \mathbb{R}^d$, and any $t \in \mathbb{R}$, $\Gamma_p(tK) = |t|\Gamma_p(K)$;
4. (triangle inequality) for any bounded sets $K, L \subseteq \mathbb{R}^d$, and their Minkowski sum $K + L$, $\Gamma_p(K + L) \leq \Gamma_p(K) + \Gamma_p(L)$.

2.4 Connection to Factorization

In this subsection we show that, in the case when $K = \{\pm w_1, \dots, \pm w_N\} \subseteq \mathbb{R}^d$ is a finite symmetric set, the quantities $\Gamma_2(K)$ and $\Gamma_\infty(K)$ can be equivalently formulated in terms of the factorization norms $\gamma_F(W)$ and $\gamma_2(W)$ of the matrix $W := (w_i)_{i=1}^N$. These norms have been studied in prior work on factorization mechanisms in differential privacy. The γ_2 norm is classical in functional analysis: see, e.g., the book by Tomczak-Jaegermann [42]. It was first applied to differential privacy implicitly in [39, 40], and more explicitly in [38]. The γ_F norm is implicit in the work on the matrix mechanism [33], and the notation we use is from [22], albeit with different normalization. We define natural analogs of these quantities that correspond to Γ_p for any $p \in [2, \infty]$. Our general formulation of Gaussian noise mechanisms thus generalizes factorization mechanisms to more general domains and more general measures of error.

First we recall the definitions of the γ_F and γ_2 factorization norms, and introduce a definition of a family of factorization norms parameterized by $p \in [2, \infty]$ that we later show correspond to Γ_p .

► **Definition 18.** *The γ_2 and the γ_F factorization norms of a $d \times N$ real matrix W are defined¹ as*

$$\gamma_2(W) := \inf\{\|A\|_{2 \rightarrow \infty} \|C\|_{1 \rightarrow 2} : AC = W\}$$

$$\gamma_F(W) := \inf\{\|A\|_F \|C\|_{1 \rightarrow 2} : AC = W\}.$$

More generally, we define, for $p \in [2, \infty]$,

$$\gamma_{(p)}(W) := \inf\left\{\sqrt{\text{tr}_{p/2}(AA^T)} \|C\|_{1 \rightarrow 2} : AC = W\right\},$$

where $\gamma_{(2)}(W) = \gamma_F(W)$ and $\gamma_{(\infty)}(W) = \gamma_2(W)$.

We have the following connection between Γ_p and these factorization norms. The proof of the theorem is deferred to the full version of the paper.

► **Theorem 19.** *For any $p \in [2, \infty]$, and any $d \times N$ real matrix W with columns w_1, \dots, w_N , for the set $K^{\text{sym}} := \{\pm w_1, \dots, \pm w_N\}$ we have*

$$\Gamma_p(K^{\text{sym}}) = \gamma_{(p)}(W).$$

Moreover, for the set $K := \{w_1, \dots, w_N\}$ we have

$$\Gamma_p(K) = \inf_{v \in \mathbb{R}^d} \gamma_{(p)}(W + v\mathbf{1}^T),$$

where $\mathbf{1}$ is the N -dimensional all-ones vector.

¹ In [22] the γ_F norm is normalized differently.

Notice that $(W_1 + W_2)B_1^N \subseteq W_1B_1^N + W_2B_2^N$. Then, a triangle inequality for $\gamma_{(p)}$, as well as homogeneity, follow from Theorem 17. The fact that $\gamma_{(p)}(W) = 0$ only if $W = 0$ follows from the observation that $\text{tr}_{p/2}(AA^T) = 0$ implies the diagonal of AA^T is 0, which implies that $A = 0$. This verifies that $\gamma_{(p)}$ is, indeed, a norm on matrices.

2.5 Duality

Our goal in this section is to derive a dual characterization of $\Gamma_p(K)$ as a maximization problem over probability distributions on K . We first carry this out for $p = 2$, and then reduce the general case to $p = 2$. This dual characterization is useful in the proofs of some of our later results.

Let us introduce some notation before we state our main duality result.

► **Definition 20.** For a compact set $K \subseteq \mathbb{R}^d$, we define $\Delta(K)$ to be the set of Borel regular measures supported on K .

► **Definition 21.** For a probability measure P over \mathbb{R}^d , we use the notation $\text{cov}(P)$ for the covariance matrix of P , i.e.,

$$\text{cov}(P) := \mathbb{E}_{Y \sim P}[(Y - \mathbb{E}[Y])(Y - \mathbb{E}[Y])^T].$$

The following theorem is our dual characterization of $\Gamma_p(K)$ as a problem of maximizing the covariance of probability distributions over K . It generalizes the known dual characterizations of the γ_2 and γ_F factorization norms [32, 27].

► **Theorem 22.** Let $K \subseteq \mathbb{R}^d$ be a bounded set, and let $p \in (2, \infty]$. Then, for $q := \frac{p}{p-2}$ we have the identity

$$\Gamma_p(K) = \sup\{\text{tr}((D \text{cov}(P)D)^{1/2}) : D \text{ diagonal}, D \succeq 0, \text{tr}_q(D^2) = 1, P \in \Delta(K)\}. \quad (1)$$

Moreover, if K is symmetric around 0 (i.e., $K = -K$), then

$$\Gamma_p(K) = \sup\{\text{tr}((D \mathbb{E}_{X \sim P}[XX^T]D)^{1/2}) : D \text{ diagonal}, D \succeq 0, \text{tr}_q(D^2) = 1, P \in \Delta(K)\}. \quad (2)$$

Our proof of Theorem 22 uses Sion's minimax theorem and the compactness of $\Delta(K)$ in the weak* topology. We defer the proof to the full version of the paper.

3 High-dimensional Lower Bound from One-dimensional Marginals

In this section we give a framework for deriving lower bounds on the ℓ_p error of an unbiased mean estimation mechanism from lower bounds on the variance of its one-dimensional marginals. This framework is essentially the same as the one proposed in [22] for oblivious mechanisms, with some small improvements. In particular, here we generalize the framework in [22] to not necessarily symmetric domains, and to error measured in the ℓ_p norm for $p \in [2, \infty]$. We also give slightly different, easier proofs of some of the main claims.

In the following, we use the notation $\text{cov}(\mathcal{M}(X))$ for the covariance matrix of the output distribution of the mechanism \mathcal{M} on input dataset X . The next lemma gives a lower bound on the ℓ_p error in terms of a function of the covariance matrix. We defer the (easy) proof to the full version of the paper.

► **Lemma 23.** For any $p \in [2, \infty]$, and any unbiased mechanism \mathcal{M} over $K \subseteq \mathbb{R}^d$, and any input dataset $X \in K^n$, we have

$$\mathbb{E} \left[\|\mathcal{M}(X) - \mu(X)\|_p^2 \right]^{1/2} \geq \sqrt{\text{tr}_{p/2}(\text{cov}(\mathcal{M}(X)))}$$

85:14 Private Unbiased Mean Estimation

The next lemma is key to our framework. Before we state it, we introduce notation for the support function and the width of a set in a given direction.

► **Definition 24.** We define the support function $h_K : \mathbb{R}^d \rightarrow \mathbb{R}$ of a set $K \subseteq \mathbb{R}^d$ by $h_K(\theta) := \sup_{x \in K} \langle x, \theta \rangle$. We define the width function $w_K : \mathbb{R}^d \rightarrow \mathbb{R}$ of a set $K \subseteq \mathbb{R}^d$ by

$$w_K(\theta) := \sup_{x \in K} \langle x, \theta \rangle - \inf_{x \in K} \langle x, \theta \rangle = h_K(\theta) + h_K(-\theta).$$

► **Lemma 25.** Let $c > 0$, let $K \subseteq \mathbb{R}^d$, let $X \in K^n$ be a dataset, and let \mathcal{M} be an unbiased mechanism over K . If, for all $\theta \in \mathbb{R}^d$, \mathcal{M} satisfies that

$$\sqrt{\text{Var}[\theta^T \mathcal{M}(X)]} \geq c w_K(\theta),$$

then $K \subseteq \frac{1}{c} \sqrt{\text{cov}(\mathcal{M}(X))} B_2^d$.

Proof. Let us denote $E := \sqrt{\text{cov}(\mathcal{M}(X))} B_2^d$. Then, for each $\theta \in \mathbb{R}^d$, we have

$$h_E(\theta) = \left\| \sqrt{\text{cov}(\mathcal{M}(X))} \theta \right\|_2 = \sqrt{\theta^T \text{cov}(\mathcal{M}(X)) \theta} = \sqrt{\text{Var}[\theta^T \mathcal{M}(X)]}.$$

Together with the assumption of the lemma, this means that $h_E(\theta) \geq c w_K(\theta)$ for all $\theta \in \mathbb{R}^d$. Let then $v \in K$ be arbitrary, and note that

$$h_{K-v}(\theta) = h_K(\theta) - \langle v, \theta \rangle = \max_{x \in K} \langle x, \theta \rangle - \langle v, \theta \rangle \geq 0,$$

for all $\theta \in \mathbb{R}^d$. Therefore,

$$h_{K-v}(\theta) \leq h_{K-v}(\theta) + h_{K-v}(-\theta) = w_{K-v}(\theta) = w_K(\theta).$$

Then, for all $\theta \in \mathbb{R}^d$, $h_{K-v}(\theta) \leq \frac{1}{c} h_E(\theta) = h_{(1/c)E}(\theta)$ which is equivalent to $K - v \subseteq \frac{1}{c} E$. ◀

Combining the two lemmas, we have the following lemma that allows us to reduce proving our lower bounds to proving one-dimensional lower bounds on variance.

► **Lemma 26.** Let $c > 0$, let $K \subseteq \mathbb{R}^d$, let $X \in K^n$ be a dataset, and let \mathcal{M} be an unbiased mechanism over K . If, for all $\theta \in \mathbb{R}^d$, \mathcal{M} satisfies that

$$\sqrt{\text{Var}[\theta^T \mathcal{M}(X)]} \geq c w_K(\theta),$$

then we have that

$$\mathbb{E} [\|\mathcal{M}(X) - \mu(X)\|_p^2]^{1/2} \geq c \Gamma_p(K).$$

Proof. By Lemma 25, and the definition of $\Gamma_p(\cdot)$, we have

$$\Gamma_p(K) \leq \frac{1}{c} \sqrt{\text{tr}_{p/2}(\text{cov}(\mathcal{M}(X)))}.$$

On the other hand, by Lemma 23,

$$\sqrt{\text{tr}_{p/2}(\text{cov}(\mathcal{M}(X)))} \leq \mathbb{E} \left[\|\mathcal{M}(X) - \mu(X)\|_p^2 \right]^{1/2}.$$

Combining the two inequalities and multiplying through by c gives the lemma. ◀

4 Lower Bound for Pure and Concentrated Differential Privacy

We first show that for any dataset X we can find a neighboring dataset X' so $\mu(X)$ and $\mu(X')$ are far in a given direction. The following simple geometric lemma is helpful for that goal.

► **Lemma 27.** *For any $\beta > 0$, any bounded $K \subseteq \mathbb{R}^d$, any $x \in K$, and any $\theta \in \mathbb{R}^d$, there exists a point $x' \in K$ such that $|\langle \theta, x - x' \rangle| \geq \frac{(1-\beta)w_K(\theta)}{2}$.*

Proof. Let x_+ be such that $\langle \theta, x_+ \rangle \geq h_K(\theta) - \frac{\beta w_K(\theta)}{2}$ and let x_- be such that $\langle -\theta, x_+ \rangle \geq h_K(-\theta) - \frac{\beta w_K(\theta)}{2}$. Thus,

$$|\langle \theta, x - x_+ \rangle| + |\langle \theta, x - x_- \rangle| \geq \langle \theta, x_+ - x_- \rangle \geq (1 - \beta)w_K(\theta).$$

Then, it must be true that either $|\langle \theta, x - x_+ \rangle| \geq \frac{(1-\beta)w_K(\theta)}{2}$, or $|\langle \theta, x - x_- \rangle| \geq \frac{(1-\beta)w_K(\theta)}{2}$, and we can choose $x' \in \{x_+, x_-\}$ accordingly. ◀

Next we use Lemma 27 to construct a neighboring dataset X' for any dataset X so that the means of X and X' are far in the direction of θ .

► **Lemma 28.** *For any $\beta > 0$, any bounded $K \subseteq \mathbb{R}^d$, any $\theta \in \mathbb{R}^d$, and any dataset $X \in K^n$, there exists a neighboring dataset X' such that*

$$|\langle \theta, \mu(X) - \mu(X') \rangle| \geq \frac{(1 - \beta)w_K(\theta)}{2n}.$$

Proof. For any given X , we change only the first data point to construct X' . Let $x_1 \in K$ be the first data point of X , and, take x'_1 to be the point x' guaranteed by Lemma 27 used with $x := x_1$. Then we set $X' := (x'_1, x_2, \dots, x_n)$. We have

$$|\langle \theta, \mu(X) - \mu(X') \rangle| = |(\mu(\theta^T X) - \mu(\theta^T X'))| = \left| \frac{1}{n} \langle \theta, x_1 - x'_1 \rangle \right| \geq \frac{(1 - \beta)w_K(\theta)}{2n},$$

where we use the notation

$$\theta^T X := (\theta^T x_1, \dots, \theta^T x_n), \quad \theta^T X' := (\theta^T x'_1, \dots, \theta^T x_n).$$

This completes the proof. ◀

Recall that, for two probability distributions P and Q , defined on the same ground set, the χ^2 -divergence between them is defined by

$$\chi^2(P||Q) := \mathbb{E}_{X \sim Q} \left[\left(\frac{dP}{dQ}(X) - 1 \right)^2 \right].$$

The following lemma, bounding the χ^2 -divergence between the output distributions of an ε -differentially private mechanism run on two neighboring datasets, is likely well-known. We omit the proof from this version of the paper.

► **Lemma 29.** *Suppose that \mathcal{M} is an ε -differentially private mechanism, and X, X' are two neighboring datasets. Let P be the probability distribution of $\mathcal{M}(X)$, and Q the probability distribution of $\mathcal{M}(X')$. Then $\chi^2(P||Q) \leq e^{-\varepsilon}(e^\varepsilon - 1)^2$.*

For our one dimensional lower bounds, we use the classical Hammersley-Chapman-Robins bound [25, 11], stated in the next lemma.

► **Lemma 30.** For any two probability distributions P and Q over the reals, and for random variables X, Y distributed, respectively, according to P and Q , we have

$$\sqrt{\text{Var}(Y)} \geq \frac{|\mathbb{E}[X] - \mathbb{E}[Y]|}{\sqrt{\chi^2(P\|Q)}}.$$

A (not tight) lower bound for pure differential privacy follows immediately from Lemmas 28, 29, 30, and Lemma 26. Instead, we state a lower bound for zCDP, which is nearly tight for small ε .

► **Theorem 31.** For any $\frac{\varepsilon^2}{2}$ -zCDP mechanism \mathcal{M} that is unbiased over $K \subseteq \mathbb{R}^d$, and any dataset $X \in K^n$, its error is bounded below as

$$\sqrt{\mathbb{E} \left[\|\mathcal{M}(X) - \mu(X)\|_p^2 \right]} \gtrsim \frac{\Gamma_p(K)}{n\sqrt{e^{\varepsilon^2} - 1}}.$$

Proof. Let $c < 2$, and let X' be any dataset that is neighboring with X and satisfies $|\theta^T(\mu(X) - f(X'))| \geq \frac{c \cdot w_K(\theta)}{n}$. Such a dataset exists by Lemma 28. Let Q be the probability distribution of $\theta^T \mathcal{M}(X)$ and P the probability distribution $\theta^T \mathcal{M}(X')$. From the definition of zCDP, the 2-Renyi divergence of P and Q is bounded as $D_2(P\|Q) \leq \varepsilon^2$. We then use the relationship between the 2-Renyi divergence and the χ^2 -divergence to write $\chi^2(P\|Q) = 2^{D_2(P\|Q)} - 1 \leq e^{\varepsilon^2} - 1$. The theorem now follows from Lemmas 26 and 30. ◀

5 Lower Bound for Approximate Differential Privacy

Our lower bounds for approximate differential privacy do not follow directly from the Hammersley-Chapman-Robins bound, because the probability distributions of $\mathcal{M}(X)$ and $\mathcal{M}(X')$, for two neighboring datasets X and X' , and an (ε, δ) -differentially private mechanism \mathcal{M} , may not have the same support. For this reason, the χ^2 -divergence between the distributions can be infinite. This leads to some complications, both for the one-dimensional, and for the higher-dimensional lower bounds, presented in this section.

5.1 One-dimensional Lower Bound

Let us introduce notation for the subset of the ground set where the ratio of densities is small. In our context, this will be the subset of possible outputs of a mechanism for which the mechanism satisfies pure differential privacy for a pair of neighboring inputs.

► **Definition 32.** For two probability distribution P and Q over the same ground set Ω , we define

$$S_{P,Q,\varepsilon} := \left\{ \omega \in \Omega : e^{-\varepsilon} \leq \frac{dP}{dQ}(\omega) \leq e^\varepsilon \right\}.$$

We restate Case 2 of Lemma 3.3 from [31] here. The lemma captures the fact that an approximately differentially private mechanism is “purely differentially private with high probability”.

► **Lemma 33.** Let \mathcal{M} be an (ε, δ) -differentially private mechanism, let X, X' be neighboring datasets, and define P to be the probability distribution of $\mathcal{M}(X)$, and Q to be the probability distribution of $\mathcal{M}(X')$. Then,

$$\max \{ \Pr \{ \mathcal{M}(X) \notin S_{P,Q,2\varepsilon} \}, \Pr \{ \mathcal{M}(X') \notin S_{P,Q,2\varepsilon} \} \} \leq \delta' := \frac{2\delta}{1 - e^{-\varepsilon}}.$$

The main challenge in applying our techniques to approximate differential privacy is that the χ^2 -divergence between the output distributions of an (ε, δ) -differentially private mechanism run on two neighboring datasets is not necessarily bounded. To get around this issue, we modify one of the two distributions slightly, so that we can fall back on bounds on the χ^2 -divergence for pure differential privacy. Here we will use Lemma 33 crucially.

► **Lemma 34.** *Let \mathcal{M} be an (ε, δ) -differentially private mechanism with range \mathbb{R}^d , let X be an arbitrary dataset, and let X' be any dataset that is neighboring with X . Let Q be the probability distribution of $\theta^T \mathcal{M}(X)$, and P the probability distribution of $\theta^T \mathcal{M}(X')$. Define $S := S_{P,Q,2\varepsilon}$. There exists a probability distribution \widehat{P} s.t.*

$$\left| \log \frac{d\widehat{P}}{dQ}(y) \right| \leq 2\varepsilon - \log(1 - \delta') \quad \text{for any } y \text{ in the range of } \mathcal{M} \quad (3)$$

$$|\mathbb{E}_{Y \sim \widehat{P}}[Y] - \mathbb{E}_{Y \sim Q}[Y]| = |C_\delta \mathbb{E}_{Y \sim P}[Y \mathbf{1}\{Y \in S\}] - \mathbb{E}_{Y \sim Q}[Y \mathbf{1}\{Y \in S\}]| \quad (4)$$

where δ' is as in Lemma 33, and $C_\delta := \frac{Q(S)}{P(S)}$.

While we omit the proof of Lemma 34, we note here that \widehat{P} is defined, for any (measurable) subset T of the range of \mathcal{M} , by $\widehat{P}(T) := C_\delta P(T \cap S) + Q(T \setminus S)$.

Combining Lemma 29 and Lemma 34, we get the following bound on the χ^2 -divergence.

► **Lemma 35.** *Let \widehat{P} and Q be as in Lemma 34. Then we have*

$$\chi^2(\widehat{P} \| Q) \leq \widehat{\varepsilon}^2 := e^{-(2\varepsilon - \log(1 - \delta'))} (e^{2\varepsilon - \log(1 - \delta')} - 1)^2.$$

The next lower bound on the variance of an approximately differentially private mechanism in a given direction is crucial to our argument. The proof is deferred to the full version. The main idea is to show that, unless the variance of $\theta^T \mathcal{M}(X')$ is very large, its distribution has similar mean as \widehat{P} above, and we can use Lemma 35 and the HCR bound.

► **Lemma 36.** *Let $c > 0$ be a small enough absolute constant, and let \mathcal{M} be (ε, δ) -differentially private mechanism that is unbiased over $K \subseteq \mathbb{R}^d$. Define $\widehat{\varepsilon}$ and δ' as in Lemmas 35 and 33. If $\delta' \leq \frac{c}{n}$, then for every $\theta \in \mathbb{R}^d$, and for every dataset $X \in K^n$, either*

$$\sqrt{\text{Var}[\theta^T \mathcal{M}(X)]} \gtrsim \frac{w_K(\theta)}{n\widehat{\varepsilon}},$$

or there exist some X' neighboring with X such that

$$\sqrt{\text{Var}[\theta^T \mathcal{M}(X')]} \gtrsim \frac{w_K(\theta)}{n\sqrt{\delta'}}.$$

5.2 High-dimensional Lower Bound for ℓ_2

We first state a lower bound on ℓ_2 error that is nearly tight when δ is small with respect to the minimum width of the domain K . Then we will show that we can always ensure the minimum width is not too small.

The next lemma follows immediately from Lemmas 26 and 36, and an application of the Cauchy-Schwarz inequality.

► **Lemma 37.** *Let $c > 0$ be a small enough absolute constant, and let \mathcal{M} be (ε, δ) -differentially private mechanism that is unbiased over $K \subseteq \mathbb{R}^d$. Define $\widehat{\varepsilon}$ and δ' as in Lemmas 35 and 33. If $\delta' \leq \frac{c}{n}$, and K satisfies $\min_{\theta \in \mathbb{R}^d: \|\theta\|_2=1} w_K(\theta) \geq w_0$, then for every dataset $X \in K^n$, either*

$$\sqrt{\mathbb{E} \left[\|\mathcal{M}(X) - \mu(X)\|_2^2 \right]} \gtrsim \frac{\Gamma_2(K)}{n\widehat{\varepsilon}},$$

85:18 Private Unbiased Mean Estimation

or there exists a neighboring dataset X' to X s.t.

$$\sqrt{\mathbb{E} \left[\|\mathcal{M}(X') - \mu(X')\|_2^2 \right]} \gtrsim \frac{w_0}{n\sqrt{\delta'}}.$$

Proof. By Lemma 36, one of the following two cases will hold:

Case 1. For all $\theta \in \mathbb{R}^d$, $\|\theta\|_2 = 1$, $\sqrt{\text{Var}[\theta^T \mathcal{M}(X)]} \gtrsim \frac{w_K(\theta)}{n\varepsilon}$. Then, by Lemma 26,

$$\mathbb{E} [\|\mathcal{M}(X) - \mu(X)\|_2] \gtrsim \frac{\Gamma(K)}{n\widehat{\varepsilon}}.$$

Case 2. There exists a $\theta^* \in \mathbb{R}^d$, $\|\theta^*\|_2 = 1$ such that there exists some X' neighboring to X for which $\sqrt{\text{Var}[\theta^{*T} \mathcal{M}(X')]} \gtrsim \frac{w_K(\theta^*)}{n\sqrt{\delta'}}$. Then, by the Cauchy-Schwarz inequality, we have

$$\begin{aligned} \sqrt{\mathbb{E} \left[\|\mathcal{M}(X') - \mu(X')\|_2^2 \right]} &\geq \sqrt{\mathbb{E} \left[\langle \theta^*, \mathcal{M}(X') - \mu(X') \rangle^2 \right]} \\ &= \sqrt{\text{Var}[\theta^{*T} \mathcal{M}(X')]} \\ &\gtrsim \frac{w_K(\theta^*)}{n\sqrt{\delta'}} \geq \frac{w_0}{n\sqrt{\delta'}}. \end{aligned}$$

This completes the proof. \blacktriangleleft

In general, the minimum width of K can be 0 even if $\Gamma_2(K)$ is large. The next lemma shows that, nevertheless, any K has a projection $P(K)$ for which the minimum width (within the image of P) and $\Gamma_2(P(K))$ are within a factor linear in the dimension, and $\Gamma_2(P(K))$ is comparable to $\Gamma_2(K)$.

► **Lemma 38.** For any $K \subseteq \mathbb{R}^d$, there exists an orthogonal projection $P : \mathbb{R}^d \rightarrow \mathbb{R}^d$, such that $P(K)$ satisfies both of the following two conditions:

$$\begin{aligned} \Gamma_2(P(K)) &\geq \frac{\Gamma_2(K)}{2} \\ \min_{\theta \in \text{Im}(P): \|\theta\|_2=1} w_{P(K)}(\theta) &\geq \frac{\Gamma_2(K)}{2d} \end{aligned}$$

Above, $\text{Im}(P)$ is the image of P .

Proof. Let $K_0 := K$, and consider the following procedure. Set, initially, $i := 1$. While there is a direction θ_i , $\|\theta_i\|_2 = 1$ which is orthogonal to $\theta_1, \dots, \theta_{i-1}$ (if $i > 1$), and is such that $w_{K_{i-1}}(\theta_i) < \frac{\Gamma_2(K)}{2d}$, we set K_i to be the projection of K_{i-1} orthogonal to θ_i , and set $i := i + 1$. Continue until no such direction can be found, or until we have made d projections, after which K_d is a point. Suppose that this procedure terminates after $k \leq d$ projections.

Let $P(K) := K_k$ be the new set at the end of the procedure. P is the orthogonal projection onto the subspace orthogonal to the span of $\theta_1, \dots, \theta_k$. From the construction it is clear that for any θ in the range of P such that $\|\theta\|_2 = 1$, $w_{P(K)}(\theta) \geq \frac{\Gamma_2(K)}{2d}$, or the procedure would not have terminated. It remains to analyze $\Gamma_2(P(K))$, and here we use the triangle inequality for Γ_2 . Define the segment $L_i := \left[-\frac{w_{K_{i-1}}(\theta_i)}{2}\theta_i, \frac{w_{K_{i-1}}(\theta_i)}{2}\theta_i \right]$. Then $K_{i-1} \subseteq_{\leftrightarrow} K_i + L_i$, and, by induction, we have $K \subseteq_{\leftrightarrow} K_k + \sum_{i=1}^k L_i$. By the choice of θ_i ,

$$\Gamma_2(L_i) \leq w_{K_{i-1}}(\theta_i) \leq \frac{\Gamma_2(K)}{2d},$$

and, using the monotonicity and triangle inequality properties from Theorem 17, we have

$$\Gamma_2(K) \leq \Gamma_2(K_k) + \sum_{i=1}^k \Gamma_2(L_i) \leq \Gamma_2(P(K)) + \frac{\Gamma_2(K)}{2}.$$

Rearranging gives us that $\Gamma_2(P(K)) \geq \frac{\Gamma_2(K)}{2}$. \blacktriangleleft

Combining these two lemmas above we can prove a lower bound on the ℓ_2 error for general K that is tight as long as δ is sufficiently small with respect to d .

► **Theorem 39.** *Let $c > 0$ be a small enough absolute constant, and let \mathcal{M} be an (ε, δ) -differentially private mechanism that is unbiased over $K \subseteq \mathbb{R}^d$. Define $\hat{\varepsilon}$ and δ' as in Lemmas 35 and 33. If $\delta' \leq \frac{c}{n}$, then for every dataset $X \in K^n$, either*

$$\sqrt{\mathbb{E} \left[\|\mathcal{M}(X) - \mu(X)\|_2^2 \right]} \gtrsim \frac{\Gamma_2(K)}{n\hat{\varepsilon}}$$

or there exists a dataset X' neighboring with X s.t.

$$\sqrt{\mathbb{E} \left[\|\mathcal{M}(X') - \mu(X')\|_2^2 \right]} \gtrsim \frac{\Gamma_2(K)}{\sqrt{\delta'nd}}$$

Proof. The key observation is that we can define an (ε, δ) -differentially private mechanism that's unbiased over $P(K)$ using \mathcal{M} , so that the ℓ_2 error does not increase. To do so, we can fix, for any $x \in P(K)$ a preimage $f(x)$ so that $P(f(x)) = x$. Then we apply f pointwise to any dataset $\tilde{X} := (\tilde{x}_1, \dots, \tilde{x}_n) \in (P(K))^n$ to get $f(\tilde{X}) := (f(\tilde{x}_1), \dots, f(\tilde{x}_n))$ in K^n so that $P(f(\tilde{X})) := (P(f(\tilde{x}_1)), \dots, P(f(\tilde{x}_n))) = \tilde{X}$. Moreover, for a fixed dataset $X \in K^n$, we can make sure that $f(P(X)) = X$. Then, given \mathcal{M} , we define $\mathcal{M}'(\tilde{X}) := P(\mathcal{M}(f(\tilde{X})))$. Since f maps neighboring datasets to neighboring datasets, and \mathcal{M}' is a postprocessing of $\mathcal{M}(f(\tilde{X}))$, \mathcal{M}' is (ε, δ) -differentially private.

Because orthogonal projection does not increase the ℓ_2 norm, and since we ensured $f(P(X)) = X$, we have

$$\begin{aligned} \sqrt{\mathbb{E} \left[\|\mathcal{M}(X) - \mu(X)\|_2^2 \right]} &\geq \\ &\sqrt{\mathbb{E} \left[\|P(\mathcal{M}(X)) - P(\mu(X))\|_2^2 \right]} = \sqrt{\mathbb{E} \left[\|\mathcal{M}'(P(X)) - \mu(P(X))\|_2^2 \right]}. \end{aligned} \quad (5)$$

An analogous analysis works for a dataset X' that is neighboring to X .

It is clear that $\text{Im}(P)$ is isometric with \mathbb{R}^{d-k} , both endowed with the ℓ_2 metric. The theorem then follows from Lemmas 37 and 38. \blacktriangleleft

5.3 High-dimensional Lower Bound for ℓ_p , $p > 2$

To prove a lower bound for the non-Euclidean case ℓ_p , $p > 2$, we reduce to the ℓ_2 case. We do so via Theorem 22. The details are deferred to the full version.

► **Theorem 40.** *Let $c > 0$ be a small enough absolute constant, let $p \in [2, \infty]$, and let \mathcal{M} be an (ε, δ) -differentially private mechanism that is unbiased over $K \subseteq \mathbb{R}^d$. Define $\hat{\varepsilon}$ and δ' as in Lemmas 35 and 33. If $\delta' \leq \frac{c}{n}$, then for every dataset $X \in K^n$, either*

$$\sqrt{\mathbb{E} \left[\|\mathcal{M}(X) - \mu(X)\|_p^2 \right]} \gtrsim \frac{\Gamma_p(K)}{n\hat{\varepsilon}}$$

or there exists a dataset X' neighboring with X s.t.

$$\sqrt{\mathbb{E} \left[\|\mathcal{M}(X') - \mu(X')\|_p^2 \right]} \gtrsim \frac{\Gamma_p(K)}{\sqrt{\delta'nd}}$$

Theorem 4 in the Introduction follows from Theorems 39 and 40.

References

- 1 Noga Alon and Assaf Naor. Approximating the cut-norm via Grothendieck’s inequality. In *ACM Symposium on Theory of Computing*, pages 72–80, 2004. doi:10.1145/1007352.1007371.
- 2 Hilal Asi and John C. Duchi. Instance-optimality in differential privacy via approximate inverse sensitivity mechanisms. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/a267f936e54d7c10a2bb70dbe6ad7a89-Abstract.html>.
- 3 Hilal Asi and John C. Duchi. Near instance-optimality in differential privacy. *CoRR*, abs/2005.10630, 2020. arXiv:2005.10630.
- 4 Hilal Asi, Vitaly Feldman, and Kunal Talwar. Optimal algorithms for mean estimation under local differential privacy. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 1046–1056. PMLR, 2022. URL: <https://proceedings.mlr.press/v162/asi22b.html>.
- 5 Raef Bassily, Adam D. Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 464–473. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.56.
- 6 Aditya Bhaskara, Daniel Dadush, Ravishankar Krishnaswamy, and Kunal Talwar. Unconditional differentially private mechanisms for linear queries. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 1269–1284. ACM, 2012. doi:10.1145/2213977.2214089.
- 7 Vijay Bhattiprolu, Euiwoong Lee, and Assaf Naor. A framework for quadratic form maximization over convex sets through nonconvex relaxations. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC ’21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 870–881. ACM, 2021. doi:10.1145/3406325.3451128.
- 8 Jarosław Błasiok, Mark Bun, Aleksandar Nikolov, and Thomas Steinke. Towards instance-optimal private query release. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms. SODA 2019*, pages 2480–2497. SIAM, Philadelphia, PA, 2019. doi:10.1137/1.9781611975482.152.
- 9 Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pages 635–658, 2016. doi:10.1007/978-3-662-53641-4_24.
- 10 Mark Bun, Jonathan Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 1–10. ACM, 2014. doi:10.1145/2591796.2591877.
- 11 Douglas G. Chapman and Herbert Robbins. Minimum variance estimation without regularity assumptions. *Ann. Math. Statistics*, 22:581–586, 1951. doi:10.1214/aoms/1177729548.
- 12 Christopher A. Choquette-Choo, H. Brendan McMahan, Keith Rush, and Abhradeep Thakurta. Multi-epoch matrix factorization mechanisms for private machine learning. *CoRR*, abs/2211.06530, 2022. doi:10.48550/arXiv.2211.06530.

- 13 Travis Dick, Alex Kulesza, Ziteng Sun, and Ananda Theertha Suresh. Subset-based instance optimality in private estimation. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 7992–8014. PMLR, 2023. URL: <https://proceedings.mlr.press/v202/dick23a.html>.
- 14 Wei Dong, Yuting Liang, and Ke Yi. Differentially private covariance revisited. *CoRR*, abs/2205.14324, 2022. doi:10.48550/arXiv.2205.14324.
- 15 Wei Dong and Ke Yi. A nearly instance-optimal differentially private mechanism for conjunctive queries. In *PODS '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, pages 213–225. ACM, 2022. doi:10.1145/3517804.3524143.
- 16 John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Minimax optimal procedures for locally private estimation. *J. Amer. Statist. Assoc.*, 113(521):182–201, 2018. doi:10.1080/01621459.2017.1389735.
- 17 John C. Duchi and Feng Ruan. The right complexity measure in locally private estimation: It is not the fisher information. *CoRR*, abs/1806.05756, 2018. arXiv:1806.05756.
- 18 Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, pages 486–503, 2006.
- 19 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Conference on Theory of Cryptography, TCC'06*, pages 265–284, Berlin, Heidelberg, 2006. Springer-Verlag. doi:10.1007/11681878_14.
- 20 Cynthia Dwork, Aleksandar Nikolov, and Kunal Talwar. Efficient algorithms for privately releasing marginals via convex relaxations. *Discrete Comput. Geom.*, 53(3):650–673, 2015. doi:10.1007/s00454-015-9678-x.
- 21 Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016. URL: <http://arxiv.org/abs/1603.01887>, arXiv:1603.01887.
- 22 Alexander Edmonds, Aleksandar Nikolov, and Jonathan Ullman. The power of factorization mechanisms in local and central differential privacy. In *STOC'20—Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 425–438. ACM, 2020. doi:10.1145/3357713.3384297.
- 23 Alexandre V. Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In *PODS*, pages 211–222. ACM, 2003.
- 24 Alexandre Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques. *Bol. Soc. Mat. Sao Paulo*, 8(1-79):88, 1953.
- 25 J. M. Hammersley. On estimating restricted parameters. *J. Roy. Statist. Soc. Ser. B*, 12:192–229; discussion, 230–240, 1950. URL: [http://links.jstor.org/sici?sici=0035-9246\(1950\)12:2<192:0ERP>2.0.CO;2-M&origin=MSN](http://links.jstor.org/sici?sici=0035-9246(1950)12:2<192:0ERP>2.0.CO;2-M&origin=MSN).
- 26 Monika Henzinger and Jalaj Upadhyay. Constant matters: Fine-grained complexity of differentially private continual observation using completely bounded norms. *CoRR*, abs/2202.11205, 2022. arXiv:2202.11205.
- 27 Monika Henzinger, Jalaj Upadhyay, and Sarvagya Upadhyay. Almost tight error bounds on differentially private continual counting. *CoRR*, abs/2211.05006, 2022. doi:10.48550/arXiv.2211.05006.
- 28 Ziyue Huang, Yuting Liang, and Ke Yi. Instance-optimal mean estimation under differential privacy. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 25993–26004, 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/da54dd5a0398011cdfa50d559c2c0ef8-Abstract.html>.

- 29 Gautam Kamath, Argyris Mouzakis, Matthew Regehr, Vikrant Singhal, Thomas Steinke, and Jonathan R. Ullman. A bias-variance-privacy trilemma for statistical estimation. *CoRR*, abs/2301.13334, 2023. doi:10.48550/arXiv.2301.13334.
- 30 Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? In *FOCS*, pages 531–540. IEEE, October 25–28 2008.
- 31 Shiva Prasad Kasiviswanathan and Adam D. Smith. On the ‘semantics’ of differential privacy: A bayesian formulation. *J. Priv. Confidentiality*, 6(1), 2014. doi:10.29012/jpc.v6i1.634.
- 32 Troy Lee, Adi Shraibman, and Robert Spalek. A direct product theorem for discrepancy. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity, CCC 2008, 23-26 June 2008, College Park, Maryland, USA*, pages 71–80. IEEE Computer Society, 2008. doi:10.1109/CCC.2008.25.
- 33 Chao Li, Michael Hay, Vibhor Rastogi, Gerome Miklau, and Andrew McGregor. Optimizing linear counting queries under differential privacy. In *Proceedings of the 29th ACM Symposium on Principles of Database Systems, PODS’10*, pages 123–134. ACM, 2010.
- 34 Chao Li, Gerome Miklau, Michael Hay, Andrew McGregor, and Vibhor Rastogi. The matrix mechanism: optimizing linear counting queries under differential privacy. *VLDB J.*, 24(6):757–781, 2015. doi:10.1007/s00778-015-0398-x.
- 35 Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. Optimizing error of high-dimensional statistical queries under differential privacy. *Proc. VLDB Endow.*, 11(10):1206–1219, 2018. doi:10.14778/3231751.3231769.
- 36 Brendan McMahan, Keith Rush, and Abhradeep Guha Thakurta. Private online prefix sums via optimal matrix factorizations. *CoRR*, abs/2202.08312, 2022. arXiv:2202.08312.
- 37 Audra McMillan, Adam D. Smith, and Jonathan R. Ullman. Instance-optimal differentially private estimation. *CoRR*, abs/2210.15819, 2022. doi:10.48550/arXiv.2210.15819.
- 38 Aleksandar Nikolov. *New Computational Aspects of Discrepancy Theory*. PhD thesis, Rutgers, The State University of New Jersey, 2014. doi:doi:10.7282/T3RN3749.
- 39 Aleksandar Nikolov, Kunal Talwar, and Li Zhang. The geometry of differential privacy: the sparse and approximate cases. In *STOC’13—Proceedings of the 2013 ACM Symposium on Theory of Computing*, pages 351–360. ACM, New York, 2013. doi:10.1145/2488608.2488652.
- 40 Aleksandar Nikolov, Kunal Talwar, and Li Zhang. The geometry of differential privacy: The small database and approximate cases. *SIAM J. Comput.*, 45(2):575–616, 2016. doi:10.1137/130938943.
- 41 Kandethody M. Ramachandran and Chris P. Tsokos. *Mathematical statistics with applications*. Elsevier/Academic Press, Amsterdam, 2009.
- 42 N. Tomczak-Jaegermann. *Banach-Mazur Distances and Finite-Dimensional Operator Ideals*. Pitman Monographs and Surveys in Pure and Applied Mathematics 38. J. Wiley, New York, 1989.
- 43 V. G. Voinov and M. S. Nikulin. *Unbiased estimators and their applications. Vol. 1*, volume 263 of *Mathematics and its Applications*. Kluwer Academic Publishers, Dordrecht, 1993. Univariate case, Translated from the 1989 Russian original by L. E. Strautman and revised by the authors. doi:10.1007/978-94-011-1970-2.
- 44 V. G. Voinov and M. S. Nikulin. *Unbiased estimators and their applications. Vol. 2*, volume 362 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1996. Multivariate case.
- 45 Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- 46 Yingtai Xiao, Guanlin He, Danfeng Zhang, and Daniel Kifer. An optimal and scalable matrix mechanism for noisy marginals under convex loss functions. *CoRR*, abs/2305.08175, 2023. doi:10.48550/arXiv.2305.08175.
- 47 Keyu Zhu, Ferdinando Fioretto, Pascal Van Hentenryck, Saswat Das, and Christine Task. Privacy and bias analysis of disclosure avoidance systems. *CoRR*, abs/2301.12204, 2023. doi:10.48550/arXiv.2301.12204.

- 48 Keyu Zhu, Pascal Van Hentenryck, and Ferdinando Fioretto. Bias and variance of post-processing in differential privacy. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual Event, February 2-9, 2021*, pages 11177–11184. AAAI Press, 2021. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/17333>, doi:10.1609/AAAI.V35I12.17333.