Local Privacy, Data Processing Inequalities, and Minimax Rates

John C. Duchi[†]   Michael I. Jordan[*]   Martin J. Wainwright[*]
jduchi@stanford.edu   jordan@stat.berkeley.edu   wainwrig@stat.berkeley.edu

Stanford University[†]   University of California, Berkeley[*]
Stanford, CA 94305   Berkeley, CA 94720

Abstract

Working under a model of privacy in which data remains private even from the statistician, we study the tradeoff between privacy guarantees and the utility of the resulting statistical estimators. We prove bounds on information-theoretic quantities, including mutual information and Kullback-Leibler divergence, that depend on the privacy guarantees. When combined with standard minimax techniques, including the Le Cam, Fano, and Assouad methods, these inequalities allow for a precise characterization of statistical rates under local privacy constraints. We provide a treatment of several canonical families of problems: mean estimation, parameter estimation in fixed-design regression, multinomial probability estimation, and nonparametric density estimation. For all of these families, we provide lower and upper bounds that match up to constant factors, and exhibit new (optimal) privacy-preserving mechanisms and computationally efficient estimators that achieve the bounds.

1   Introduction

A major challenge in statistical inference is that of characterizing and balancing statistical utility with the privacy of individuals from whom data is obtained [20, 21, 28]. Such a characterization requires a formal definition of privacy, and *differential privacy* has been put forth as one such formalization [e.g., 24, 10, 25, 34, 35]. In the database and cryptography literatures from which differential privacy arose, early research was mainly algorithmic in focus, and researchers have used differential privacy to evaluate privacy-retaining mechanisms for transporting, indexing, and querying data. More recent work aims to link differential privacy to statistical concerns [22, 51, 33, 48, 16, 46]; in particular, researchers have developed algorithms for private robust statistical estimators, point and histogram estimation, and principal components analysis. Guarantees of optimality in this line of work have often been non-inferential, aiming to approximate a class of statistics under privacy-respecting transformations of the data at hand and not with respect to an underlying population. There has also been recent work within the context of classification problems and the "probably approximately correct" framework of statistical learning theory [e.g. 37, 8] that treats the data as random and aims to recover aspects of the underlying population; we discuss this work in Section 6.

In this paper, we take a fully inferential point of view on privacy, bringing differential privacy into contact with statistical decision theory. Our focus is on the fundamental limits of differentially-private estimation. By treating differential privacy as an abstract constraint on estimators, we obtain independence from specific estimation procedures and privacy-preserving mechanisms. Within this framework, we derive both lower bounds and matching upper bounds on minimax risk. We obtain our lower bounds by integrating differential privacy into the classical
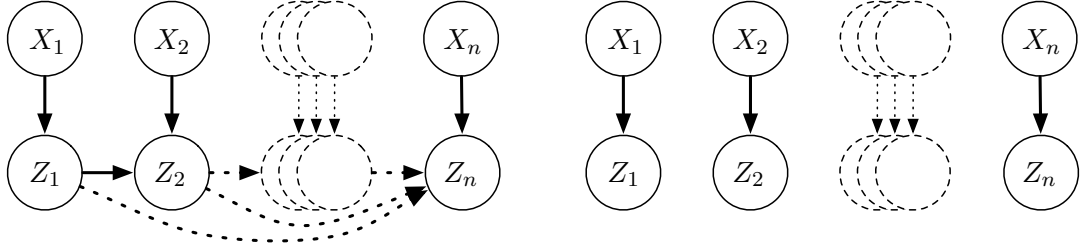
Figure 1. Left: graphical structure of private Z_i and non-private data X_i in interactive case. Right: graphical structure of channel in non-interactive case.

paradigms for bounding minimax risk via the inequalities of Le Cam, Fano, and Assouad, while we obtain matching upper bounds by proposing and analyzing specific private procedures.

We study the setting of *local privacy*, in which providers do not even trust the statistician collecting the data. Although local privacy is a relatively stringent requirement, we view this setting as a natural step in identifying minimax risk bounds under privacy constraints. Indeed, local privacy is one of the oldest forms of privacy: its essential form dates to Warner [50], who proposed it as a remedy for what he termed "evasive answer bias" in survey sampling. We hope that we can leverage deeper understanding of this classical setting to treat other privacy-preserving approaches to data analysis.

More formally, let $X_1, \ldots, X_n \in \mathcal{X}$ be observations drawn according to a distribution P, and let $\theta = \theta(P)$ be a parameter of this unknown distribution. We wish to estimate θ based on access to obscured views $Z_1, \ldots, Z_n \in \mathcal{Z}$ of the original data. The original random variables $\{X_i\}_{i=1}^n$ and the privatized observations $\{Z_i\}_{i=1}^n$ are linked via a family of conditional distributions $Q_i(Z_i \mid X_i = x, Z_{1:i-1} = z_{1:i-1})$. To simplify notation, we typically omit the subscript in Q_i. We refer to Q as a *channel distribution*, as it acts as a conduit from the original to the privatized data, and we assume it is *sequentially interactive*, meaning the channel has the conditional independence structure

$$\{X_i, Z_1, \ldots, Z_{i-1}\} \to Z_i \quad \text{and} \quad Z_i \perp X_j \mid \{X_i, Z_1, \ldots, Z_{i-1}\} \text{ for } j \neq i,$$

illustrated on the left of Figure 1. A special case of such a channel is the *non-interactive* case, in which each Z_i depends only on X_i (Fig. 1, right).

Our work is based on the following definition of privacy. For a given privacy parameter $\alpha \geq 0$, we say that Z_i is an α-*differentially locally private* view of X_i if for all $z_1, \ldots, z_{i-1}$ and $x, x' \in \mathcal{X}$ we have

$$\sup_{S \in \sigma(\mathcal{Z})} \frac{Q_i(Z_i \in S \mid X_i = x, Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1})}{Q_i(Z_i \in S \mid X_i = x', Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1})} \leq \exp(\alpha), \tag{1}$$

where $\sigma(\mathcal{Z})$ denotes an appropriate σ-field on $\mathcal{Z}$. Definition (1) does not constrain Z_i to be a release of data based exclusively on X_i: the channel Q_i may be *interactive* [24], changing based on prior private observations Z_j. We also consider the non-interactive case [50, 27] where Z_i depends only on X_i (see the right side of Figure 1); here the bound (1) reduces to

$$\sup_{S \in \sigma(\mathcal{Z})} \sup_{x, x' \in \mathcal{X}} \frac{Q(Z_i \in S \mid X_i = x)}{Q(Z_i \in S \mid X_i = x')} \leq \exp(\alpha). \tag{2}$$

These definitions capture a type of plausible deniability: no matter what data Z is released, it is nearly equally as likely to have come from one point $x \in \mathcal{X}$ as any other. It is also possible

2

to interpret differential privacy within a hypothesis testing framework, where α controls the error rate in tests for the presence or absence of individual data points in a dataset [51]. Such guarantees against discovery, together with the treatment of issues of side information or adversarial strength that are problematic for other formalisms, have been used to make the case for differential privacy within the computer science literature; see, for example, the papers [27, 24, 6, 30].

Although differential privacy provides an elegant formalism for limiting disclosure and protecting against many forms of privacy breach, it is a stringent measure of privacy, and it is conceivably overly stringent for statistical practice. Indeed, Fienberg et al. [29] criticize the use of differential privacy in releasing contingency tables, arguing that known mechanisms for differentially private data release can give unacceptably poor performance. As a consequence, they advocate—in some cases—recourse to weaker privacy guarantees to maintain the utility and usability of released data. There are results that are more favorable for differential privacy; for example, Smith [48] shows that the non-local form of differential privacy [24] can be satisfied while yielding asymptotically optimal parametric rates of convergence for some point estimators. Resolving such differing perspectives requires investigation into whether particular methods have optimality properties that would allow a general criticism of the framework, and characterizing the trade-offs between privacy and statistical efficiency. Such are the goals of the current paper.

1.1 Our contributions

The main contribution of this work is to provide general techniques for deriving minimax bounds under local privacy constraints and to illustrate these techniques by computing minimax rates for several canonical problems: (a) mean estimation; (b) parameter estimation in fixed design regression; (c) multinomial probability estimation; and (d) density estimation. We now outline our main contributions. (Because a deeper comparison of the current work with prior research requires a formal definition of our minimax framework and presentation of our main results, we defer a full discussion of related work to Section 6. We note here, however, that our minimax rates are for estimation of *population* quantities, in accordance with our connections to statistical decision theory; by way of comparison, most prior work in the privacy literature focuses on accurate approximation of statistics in a conditional analysis in which the data are treated as fixed.

Many methods for obtaining minimax bounds involve information-theoretic quantities relating data-generating distributions [53, 52, 49]. In particular, let P_1 and P_2 denote two distributions on the observations X_i, and for $\nu \in \{1, 2\}$, define the marginal distribution M_ν^n on $\mathcal{Z}^n$ by

$$M_\nu^n(S) := \int Q^n(S \mid x_1, \ldots, x_n) dP_\nu(x_1, \ldots, x_n) \quad \text{for } S \in \sigma(\mathcal{Z}^n). \tag{3}$$

Here $Q^n(\cdot \mid x_1, \ldots, x_n)$ denotes the joint distribution on $\mathcal{Z}^n$ of the private sample $Z_{1:n}$, conditioned on $X_{1:n} = x_{1:n}$. The mutual information of samples drawn according to distributions of the form (3) and the KL divergence between such distributions are key objects in statistical discriminability and minimax rates [36, 9, 53, 52, 49], where they are often applied in one of three lower-bounding techniques: Le Cam's, Fano's, and Assouad's methods.

Keeping in mind the centrality of these information-theoretic quantities, we summarize our main results at a high-level as follows. Theorem 1 bounds the KL divergence between distributions M_1^n and M_2^n, as defined by the marginal (3), by a quantity dependent on the differential privacy parameter α and the total variation distance between P_1 and P_2. The essence of Theorem 1 is that

$$D_{\mathrm{kl}}\left(M_1^n \| M_2^n\right) \lesssim \alpha^2 n \left\| P_1 - P_2 \right\|_{\mathrm{TV}}^2,$$

3

where $\lesssim$ denotes inequality up to numerical constants. When $\alpha^2 < 1$, which is the usual region of interest, this result shows that for statistical procedures whose minimax rate of convergence can be determined by classical information-theoretic methods, the additional requirement of α-local differential privacy causes the *effective sample size* of *any* statistical procedure to be reduced from n to at most $\alpha^2 n$. Section 3.1 contains the formal statement of this theorem, while Section 3.2 provides corollaries showing its application to minimax risk bounds. We follow this in Section 3.3 with applications of these results to estimation of one-dimensional means and fixed-design regression problems, providing corresponding upper bounds on the minimax risk. In addition to our general analysis, we exhibit some striking difficulties of locally private estimation in non-compact spaces: if we wish to estimate the mean of a random variable X satisfying $\mathrm{Var}(X) \leq 1$, the minimax rate of estimation of $\mathbb{E}[X]$ decreases from the parametric $1/n$ rate to $1/\sqrt{n\alpha^2}$.

Theorem 1 is appropriate for many one-dimensional problems, but it does not address difficulties inherent in higher-dimensional problems. With this motivation, our next two main results (Theorems 2 and 3) generalize Theorem 1 and incorporate dimensionality in an essential way: each provides bounds on information-theoretic quantities by dimension-dependent analogues of total variation. More specifically, Theorem 2 provides bounds on mutual information quantities essential in information-theoretic techniques such as Fano's method [53, 52], while Theorem 3 provides analogous bounds on summed pairs of KL-divergences useful in applications of Assouad's method [5, 53, 4].

As a consequence of Theorems 2 and 3, we obtain that for many d-dimensional estimation problems the effective sample size is reduced from n to $n\alpha^2/d$; as our examples illustrate, this dimension-dependent reduction in sample size can have dramatic consequences. We provide the main statement and consequences of Theorem 2 in Section 4, showing its application to obtaining minimax rates for mean estimation in both classical and high-dimensional settings. In Section 5, we present Theorem 3, showing how it provides (sharp) minimax lower bounds for multinomial and probability density estimation. Our results enable us to derive (often new) optimal mechanisms for these problems. One interesting consequence of our results is that Warner's randomized response procedure [50] from the 1960s is an optimal mechanism for multinomial estimation.

Notation: For distributions P and Q defined on a space $\mathcal{X}$, each absolutely continuous with respect to a distribution μ (with corresponding densities p and q), the KL divergence between P and Q is

$$D_{\mathrm{kl}}(P\|Q) := \int_{\mathcal{X}} dP \log \frac{dP}{dQ} = \int_{\mathcal{X}} p \log \frac{p}{q} d\mu.$$

Letting $\sigma(\mathcal{X})$ denote the (an appropriate) σ-field on $\mathcal{X}$, the total variation distance between two distributions P and Q is

$$\|P - Q\|_{\mathrm{TV}} := \sup_{S \in \sigma(\mathcal{X})} |P(S) - Q(S)| = \frac{1}{2} \int_{\mathcal{X}} |p(x) - q(x)| \, d\mu(x).$$

Let P and P_Y denote marginal distributions of random vectors X and Y and $P_Y(\cdot \mid X)$ denote the distribution of Y conditional on X. The mutual information between X and Y is

$$I(X;Y) = \mathbb{E}_P\left[D_{\mathrm{kl}}(P_Y(\cdot \mid X)\|P_Y(\cdot))\right] = \int D_{\mathrm{kl}}(P_Y(\cdot \mid X = x)\|P_Y(\cdot)) \, dP(x).$$

Random variable Y has Laplace(α) distribution if its density is $p_Y(y) = \frac{\alpha}{2}\exp(-\alpha|y|)$. For matrices $A, B \in \mathbb{R}^{d \times d}$, the notation $A \preceq B$ means that $B - A$ is positive semidefinite. For real sequences

$\{a_n\}$ and $\{b_n\}$, we use $a_n \lesssim b_n$ to mean there is a universal constant $C < \infty$ such that $a_n \leq C b_n$ for all n, and $a_n \asymp b_n$ to denote that $a_n \lesssim b_n$ and $b_n \lesssim a_n$.

2 Background and problem formulation

We first establish the minimax framework we use throughout this paper; see references [52, 53, 49] for further background. Let $\mathcal{P}$ denote a class of distributions on the sample space $\mathcal{X}$, and let $\theta(P) \in \Theta$ denote a function defined on $\mathcal{P}$. The space Θ in which the parameter $\theta(P)$ takes values depends on the underlying statistical model (for univariate mean estimation, it is a subset of the real line). Let ρ denote a semi-metric on the space Θ, which we use to measure the error of an estimator for the parameter θ, and let $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ be a non-decreasing function with $\Phi(0) = 0$ (for example, $\Phi(t) = t^2$).

In the classical setting, the statistician is given direct access to i.i.d. observations X_i drawn according to some $P \in \mathcal{P}$. The local privacy setting involves an additional ingredient, namely, a conditional distribution Q that transforms the sample $\{X_i\}_{i=1}^n$ into the private sample $\{Z_i\}_{i=1}^n$ taking values in $\mathcal{Z}$. Based on these Z_i, our goal is to estimate the unknown parameter $\theta(P) \in \Theta$. An estimator $\widehat{\theta}$ is a measurable function $\widehat{\theta} : \mathcal{Z}^n \to \Theta$, and we assess the quality of the estimate $\widehat{\theta}(Z_1, \ldots, Z_n)$ in terms of the risk

$$\mathbb{E}_{P,Q}\left[\Phi\big(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P))\big)\right].$$

For instance, for a univariate mean problem with $\rho(\theta, \theta') = |\theta - \theta'|$ and $\Phi(t) = t^2$, this risk is the mean-squared error. For any fixed conditional distribution Q, the minimax rate is

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, Q) := \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_{P,Q}\left[\Phi\big(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P))\big)\right], \tag{4}$$

where we take the supremum over distributions $P \in \mathcal{P}$, and the infimum is taken over all estimators $\widehat{\theta}$.

For $\alpha > 0$, let $\mathcal{Q}_\alpha$ denote the set of all conditional distributions guaranteeing α-local privacy (1). By minimizing the minimax risk (4) over all $Q \in \mathcal{Q}_\alpha$, we obtain the central object of study for this paper, a functional which characterizes the optimal rate of estimation in terms of the privacy parameter α.

Definition 1. Given a family of distributions $\theta(\mathcal{P})$ and a privacy parameter $\alpha > 0$, the α-*minimax rate* in the metric ρ is

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, \alpha) := \inf_{Q \in \mathcal{Q}_\alpha} \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_{P,Q}\left[\Phi(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P)))\right]. \tag{5}$$

From estimation to testing: A standard first step in proving minimax bounds is to reduce the estimation problem to a testing problem [53, 52, 49]. We use two types of testing problems: one a multiple hypothesis test, the second based on multiple binary hypothesis tests. We begin with the first of the two. Given an index set $\mathcal{V}$ of finite cardinality, consider a family of distributions $\{P_\nu, \nu \in \mathcal{V}\}$ contained within $\mathcal{P}$. This family induces a collection of parameters $\{\theta(P_\nu), \nu \in \mathcal{V}\}$; it is a 2δ-packing in the ρ-semimetric if

$$\rho(\theta(P_\nu), \theta(P_{\nu'})) \geq 2\delta \quad \text{for all } \nu \neq \nu'. \tag{6}$$

We use this family to define the *canonical hypothesis testing problem*:

- first, nature chooses V according to the uniform distribution over $\mathcal{V}$;

- second, conditioned on the choice $V = \nu$, the random sample $X = (X_1, \ldots, X_n)$ is drawn from the n-fold product distribution P_ν^n.

In the classical setting, the statistician directly observes the sample X, while the local privacy constraint means that a new random sample $Z = (Z_1, \ldots, Z_n)$ is generated by sampling Z_i from the distribution $Q(\cdot \mid X_{1:n})$. By construction, conditioned on the choice $V = \nu$, the private sample Z is distributed according to the marginal measure M_ν^n defined in equation (3).

Given the observed vector Z, the goal is to determine the value of the underlying index ν. We refer to any measurable mapping $\psi : \mathcal{Z}^n \to \mathcal{V}$ as a test function. Its associated error probability is $\mathbb{P}(\psi(Z_1, \ldots, Z_n) \neq V)$, where $\mathbb{P}$ denotes the joint distribution over the random index V and Z. The classical reduction from estimation to testing [e.g., 49, Section 2.2] guarantees that the minimax error (4) has lower bound

$$\mathfrak{M}_n(\Theta, \Phi \circ \rho, Q) \geq \Phi(\delta) \inf_\psi \mathbb{P}(\psi(Z_1, \ldots, Z_n) \neq V). \tag{7}$$

The remaining challenge is to lower bound the probability of error in the underlying multi-way hypothesis testing problem. There are a variety of techniques for this, and we focus on bounds on the probability of error (7) due to Le Cam and Fano. The simplest form of Le Cam's inequality [e.g., 53, Lemma 1] is applicable when there are two values ν, ν' in $\mathcal{V}$. In this case,

$$\inf_\psi \mathbb{P}\left(\psi(Z_1, \ldots, Z_n) \neq V\right) = \frac{1}{2} - \frac{1}{2} \left\|M_\nu^n - M_{\nu'}^n\right\|_{\mathrm{TV}}, \tag{8}$$

where the marginal M is defined as in expression (3). More generally, Fano's inequality [52, 32, Lemma 4.2.1] holds when nature chooses uniformly at random from a set $\mathcal{V}$ of cardinality larger than two, and takes the form

$$\inf_\psi \mathbb{P}(\psi(Z_1, \ldots, Z_n) \neq V) \geq \left[1 - \frac{I(Z_1, \ldots, Z_n; V) + \log 2}{\log |\mathcal{V}|}\right]. \tag{9}$$

The second reduction we consider—which transforms estimation problems into multiple binary hypothesis testing problems—uses the structure of the hypercube in an essential way. For some $d \in \mathbb{N}$, we set $\mathcal{V} = \{-1, 1\}^d$. We say that the the family P_ν induces a 2δ-Hamming separation for $\Phi \circ \rho$ if there exists a function $\mathsf{v} : \theta(\mathcal{P}) \to \{-1, 1\}^d$ satisfying

$$\Phi(\rho(\theta, \theta(P_\nu))) \geq 2\delta \sum_{j=1}^d \mathbf{1}\{[\mathsf{v}(\theta)]_j \neq \nu_j\}. \tag{10}$$

Letting $\mathbb{P}_{\pm j}$ denote the joint distribution over the random index V and Z conditional on the jth coordinate $V_j = \pm 1$, we are able to establish the following sharpening of Assouad's lemma [5, 4] (see Appendix F.1 for a proof).

Lemma 1. *Under the conditions of the previous paragraph, we have*

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, Q) \geq \delta \sum_{j=1}^d \inf_\psi \left[\mathbb{P}_{+j}(\psi(Z_{1:n}) \neq +1) + \mathbb{P}_{-j}(\psi(Z_{1:n}) \neq -1)\right].$$

6

With the definition of the marginals $M_{\pm j}^n = 2^{-d+1} \sum_{\nu : \nu_j = \pm 1} M_\nu^n$, expression (8) shows that Lemma 1 is equivalent to the lower bound

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, Q) \geq \delta \sum_{j=1}^d \left[1 - \left\| M_{+j}^n - M_{-j}^n \right\|_{\mathrm{TV}} \right]. \tag{11}$$

As a consequence of the preceding reductions to testing and the error bounds (8), (9), and (11), we obtain bounds on the private minimax rate (5) by controlling variation distances of the form $\| M_1^n - M_2^n \|_{\mathrm{TV}}$ or the mutual information between the random parameter index V and the sequence of random variables $Z_1, \ldots, Z_n$. We devote the following sections to these tasks.

3 Pairwise bounds under privacy: Le Cam and local Fano methods

We begin with results that upper bound the symmetrized Kullback-Leibler divergence under a privacy constraint, developing consequences of this result for both Le Cam's method and a local form of Fano's method. Using these methods, we derive sharp minimax rates under local privacy for estimating 1-dimensional means and for d-dimensional fixed design regression.

3.1 Pairwise upper bounds on Kullback-Leibler divergences

Many statistical problems depend on comparisons between a pair of distributions P_1 and P_2 defined on a common space $\mathcal{X}$. Any conditional distribution Q transforms such a pair of distributions into a new pair (M_1, M_2) via the marginalization (3); that is, $M_j(S) = \int_{\mathcal{X}} Q(S \mid x) dP_j(x)$ for $j = 1, 2$. Our first main result bounds the symmetrized Kullback-Leibler (KL) divergence between these induced marginals as a function of the privacy parameter $\alpha > 0$ associated with the conditional distribution Q and the total variation distance between P_1 and P_2.

Theorem 1. *For any $\alpha \geq 0$, let Q be a conditional distribution that guarantees α-differential privacy. Then for any pair of distributions P_1 and P_2, the induced marginals M_1 and M_2 satisfy the bound*

$$D_{\mathrm{kl}}\left(M_1 \| M_2\right) + D_{\mathrm{kl}}\left(M_2 \| M_1\right) \leq \min\{4, e^{2\alpha}\}(e^\alpha - 1)^2 \left\| P_1 - P_2 \right\|_{\mathrm{TV}}^2. \tag{12}$$

Remarks: Theorem 1 is a type of *strong data processing* inequality [3], providing a quantitative relationship from the divergence $\| P_1 - P_2 \|_{\mathrm{TV}}$ to the KL-divergence $D_{\mathrm{kl}}\left(M_1 \| M_2\right)$ that arises after applying the channel Q. The result of Theorem 1 is similar to a result due to Dwork et al. [25, Lemma III.2], who show that $D_{\mathrm{kl}}\left(Q(\cdot \mid x) \| Q(\cdot \mid x')\right) \leq \alpha(e^\alpha - 1)$ for any $x, x' \in \mathcal{X}$, which implies $D_{\mathrm{kl}}\left(M_1 \| M_2\right) \leq \alpha(e^\alpha - 1)$ by convexity. This upper bound is weaker than Theorem 1 since it lacks the term $\| P_1 - P_2 \|_{\mathrm{TV}}^2$. This total variation term is essential to our minimax lower bounds: more than providing a bound on KL divergence, Theorem 1 shows that differential privacy acts as a contraction on the space of probability measures. This contractivity holds in a strong sense: indeed, the bound (12) shows that even if we start with a pair of distributions P_1 and P_2 whose KL divergence is infinite, the induced marginals M_1 and M_2 always have finite KL divergence.

We provide the proof of Theorem 1 in Section 7. Here we develop a corollary that has useful consequences for minimax theory under local privacy constraints. Suppose that conditionally on $V = \nu$, we draw a sample $X_1, \ldots, X_n$ from the product measure $\prod_{i=1}^n P_{\nu,i}$, and that we draw the

α-locally private sample $Z_1, \ldots, Z_n$ according to the channel $Q(\cdot \mid X_{1:n})$. Conditioned on $V = \nu$, the private sample is distributed according to the measure M_ν^n defined previously (3). Because we allow interactive protocols, the distribution M_ν^n need not be a product distribution in general. Given this setup, we have the following:

Corollary 1. *For any α-locally differentially private (1) conditional distribution Q and any paired sequences of distributions $\{P_{\nu,i}\}$ and $\{P_{\nu',i}\}$,*

$$D_{\mathrm{kl}}\left(M_\nu^n \| M_{\nu'}^n\right) + D_{\mathrm{kl}}\left(M_{\nu'}^n \| M_\nu^n\right) \leq 4(e^\alpha - 1)^2 \sum_{i=1}^n \left\| P_{\nu,i} - P_{\nu',i} \right\|_{\mathrm{TV}}^2. \tag{13}$$

See Section 7.2 for the proof, which requires a few intermediate steps to obtain the additive inequality. Inequality (13) also immediately implies a mutual information bound, which may be useful in applications of Fano's inequality. In particular, if we define the mean distribution $\overline{M}^n = \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} M_\nu^n$, then by the definition of mutual information, we have

$$I(Z_1, \ldots, Z_n; V) = \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} D_{\mathrm{kl}}\left(M_\nu^n \| \overline{M}^n\right) \leq \frac{1}{|\mathcal{V}|^2} \sum_{\nu, \nu'} D_{\mathrm{kl}}\left(M_\nu^n \| M_{\nu'}^n\right)$$

$$\leq 4(e^\alpha - 1)^2 \sum_{i=1}^n \frac{1}{|\mathcal{V}|^2} \sum_{\nu, \nu' \in \mathcal{V}} \left\| P_{\nu,i} - P_{\nu',i} \right\|_{\mathrm{TV}}^2, \tag{14}$$

the first inequality following from the joint convexity of the KL divergence and the final inequality from Corollary 1.

Remarks: Mutual information bounds under local privacy have appeared previously. McGregor et al. [43] study relationships between communication complexity and differential privacy, showing that differentially private schemes allow low communication. They provide a result [43, Prop. 7] guaranteeing $I(X_{1:n}; Z_{1:n}) \leq 3\alpha n$; they strengthen this bound to $I(X_{1:n}; Z_{1:n}) \leq (3/2)\alpha^2 n$ when the X_i are i.i.d. uniform Bernoulli variables. Since the total variation distance is at most 1, our result also implies this scaling (for arbitrary X_i), but it is stronger since it involves the total variation terms $\| P_{\nu,i} - P_{\nu',i} \|_{\mathrm{TV}}$, which are essential in our minimax results. In addition, Corollary 1 allows for *any* (sequentially) interactive channel Q; each Z_i may depend on the private answers $Z_{1:i-1}$ of other data providers.

3.2 Consequences for minimax theory under local privacy constraints

We now turn to some consequences of Theorem 1 for minimax theory under local privacy constraints. For ease of presentation, we analyze the case of independent and identically distributed (i.i.d.) samples, meaning that $P_{\nu,i} \equiv P_\nu$ for $i = 1, \ldots, n$. We show that in both Le Cam's inequality and the local version of Fano's method, the constraint of α-local differential privacy reduces the effective sample size (at least) from n to $4\alpha^2 n$.

Consequence for Le Cam's method: The classical non-private version of Le Cam's method bounds the usual minimax risk

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) := \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P\left[\Phi\left(\rho(\widehat{\theta}(X_1, \ldots, X_n), \theta(P)))\right)\right],$$

for estimators $\widehat{\theta} : \mathcal{X}^n \to \Theta$ by a binary hypothesis test. One version of Le Cam's lemma (8) asserts that, for any pair of distributions $\{P_1, P_2\}$ such that $\rho(\theta(P_1), \theta(P_2)) \geq 2\delta$, we have

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \geq \Phi(\delta) \left\{ \frac{1}{2} - \frac{1}{2\sqrt{2}} \sqrt{n D_{\mathrm{kl}}(P_1 \| P_2)} \right\}. \tag{15}$$

Returning to the α-locally private setting, in which the estimator $\widehat{\theta}$ depends only on the private variables $(Z_1, \ldots, Z_n)$, we measure the α-private minimax risk (5). By applying Le Cam's method to the pair (M_1, M_2) along with Corollary 1 in the form of inequality (13), we find:

Corollary 2 (Private form of Le Cam bound)**.** *Given observations from an α-locally differential private channel for some $\alpha \in [0, \frac{22}{35}]$, the α-private minimax risk is lower bounded as*

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, \alpha) \geq \Phi(\delta) \left\{ \frac{1}{2} - \frac{1}{2\sqrt{2}} \sqrt{8 n \alpha^2 \| P_1 - P_2 \|_{\mathrm{TV}}^2} \right\}. \tag{16}$$

Using the fact that $\| P_1 - P_2 \|_{\mathrm{TV}}^2 \leq \frac{1}{2} D_{\mathrm{kl}}(P_1 \| P_2)$, comparison with the original Le Cam bound (15) shows that for $\alpha \in [0, \frac{22}{35}]$, the effect of α-local differential privacy is to reduce the *effective sample size* from n to $4\alpha^2 n$. We illustrate use of this private version of Le Cam's bound in our analysis of the one-dimensional mean problem to follow.

Consequences for local Fano's method: We now turn to consequences for the so-called local form of Fano's method. This method is based on constructing a family of distributions $\{P_\nu, \nu \in \mathcal{V}\}$ that defines a 2δ-packing, meaning $\rho(\theta(P_\nu), \theta(P_{\nu'})) \geq 2\delta$ for all $\nu \neq \nu'$, satisfying

$$D_{\mathrm{kl}}(P_\nu \| P_{\nu'}) \leq \kappa^2 \delta^2 \quad \text{for some fixed } \kappa > 0. \tag{17}$$

We refer to any such construction as a (δ, κ) *local packing*. Recalling Fano's inequality (9), the pairwise upper bounds (17) imply $I(X_1, \ldots, X_n; V) \leq n\kappa^2\delta^2$ by a convexity argument. We thus obtain the local Fano lower bound [36, 9] on the classical minimax risk:

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \geq \Phi(\delta) \left\{ 1 - \frac{n\kappa^2\delta^2 + \log 2}{\log |\mathcal{V}|} \right\}. \tag{18}$$

We now state the extension of this bound to the α-locally private setting.

Corollary 3 (Private form of local Fano inequality)**.** *Consider observations from an α-locally differential private channel for some $\alpha \in [0, \frac{22}{35}]$. Given any (δ, κ) local packing, the α-private minimax risk has lower bound*

$$\mathfrak{M}_n(\Theta, \Phi \circ \rho, \alpha) \geq \Phi(\delta) \left\{ 1 - \frac{4n\alpha^2\kappa^2\delta^2 + \log 2}{\log |\mathcal{V}|} \right\}. \tag{19}$$

Once again, by comparison to the classical version (18), we see that, for all $\alpha \in [0, \frac{22}{35}]$, the price for privacy is a reduction in the effective sample size from n to $4\alpha^2 n$. The proof is again straightfoward using Theorem 1. By Pinsker's inequality, the pairwise bound (17) implies that

$$\| P_\nu - P_{\nu'} \|_{\mathrm{TV}}^2 \leq \frac{1}{2} \kappa^2 \delta^2 \quad \text{for all } \nu \neq \nu'.$$

We find that $I(Z_1, \ldots, Z_n; V) \leq 4n\alpha^2\kappa^2\delta^2$ for all $\alpha \in [0, \frac{22}{35}]$ by combining this inequality with the upper bound (14) from Corollary 1. The claim (19) follows by combining this upper bound with the usual local Fano bound (18).

3.3 Some applications of Theorem 1

In this section, we illustrate the use of the α-private versions of Le Cam's and Fano's inequalities, established in the previous section as Corollaries 2 and 3 of Theorem 1. First, we study the problem of one-dimensional mean estimation. In addition to demonstrating how the minimax rate changes as a function of α, we also reveal some interesting (and perhaps disturbing) effects of enforcing α-local differential privacy: the effective sample size may be even polynomially smaller than $\alpha^2 n$. Our second example studies fixed design linear regression, where we again see the reduction in effective sample size from n to $\alpha^2 n$. We state each of our bounds assuming $\alpha \in [0,1]$; the bounds hold (with different numerical constants) whenever $\alpha \in [0, C]$ for some universal constant C.

3.3.1 One-dimensional mean estimation

For some $k > 1$, consider the family

$$\mathcal{P}_k := \big\{\text{distributions } P \text{ such that } \mathbb{E}_P[X] \in [-1,1] \text{ and } \mathbb{E}_P[|X|^k] \leq 1\big\},$$

and suppose that our goal is to estimate the mean $\theta(P) = \mathbb{E}_P[X]$. The next proposition characterizes the α-private minimax risk in squared ℓ_2-error:

$$\mathfrak{M}_n(\theta(\mathcal{P}_k), (\cdot)^2, \alpha) := \inf_{Q \in \mathcal{Q}_\alpha} \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}_k} \mathbb{E}\left[\big(\widehat{\theta}(Z_1, \ldots, Z_n) - \theta(P)\big)^2\right].$$

Proposition 1. *There exist universal constants $0 < c_\ell \leq c_u < \infty$ such that for all $k > 1$ and $\alpha \in [0,1]$, the minimax error $\mathfrak{M}_n(\theta(\mathcal{P}_k), (\cdot)^2, \alpha)$ is bounded as*

$$c_\ell \min\left\{1, \left(n\alpha^2\right)^{-\frac{k-1}{k}}\right\} \leq \mathfrak{M}_n(\theta(\mathcal{P}_k), (\cdot)^2, \alpha) \leq c_u \min\left\{1, u_k \left(n\alpha^2\right)^{-\frac{k-1}{k}}\right\}, \tag{20}$$

where $u_k = \max\{1, (k-1)^{-2}\}$.

We prove this result using the α-private version (16) of Le Cam's inequality, as stated in Corollary 2. See Section 7.3 for the details.

To understand the bounds (20), it is worthwhile considering some special cases, beginning with the usual setting of random variables with finite variance ($k = 2$). In the non-private setting in which the original sample $(X_1, \ldots, X_n)$ is observed, the sample mean $\widehat{\theta} = \frac{1}{n} \sum_{i=1}^n X_i$ has mean-squared error at most $1/n$. When we require α-local differential privacy, Proposition 1 shows that the minimax rate worsens to $1/\sqrt{n\alpha^2}$. More generally, for any $k > 1$, the minimax rate scales as $\mathfrak{M}_n(\theta(\mathcal{P}_k), (\cdot)^2, \alpha) \asymp (n\alpha^2)^{-\frac{k-1}{k}}$, ignoring k-dependent pre-factors. As $k \uparrow \infty$, the moment condition $\mathbb{E}[|X|^k] \leq 1$ becomes equivalent to the boundedness constraint $|X| \leq 1$ a.s., and we obtain the more standard parametric rate $(n\alpha^2)^{-1}$, where there is no reduction in the exponent.

More generally, the behavior of the α-private minimax rates in (20) helps demarcate situations in which local differential privacy may or may not be acceptable. In particular, for bounded domains—where we may take $k \uparrow \infty$—local differential privacy may be quite reasonable. However, in situations in which the sample takes values in an unbounded space, local differential privacy provides much stricter constraints. Indeed, in Appendix G, we discuss an example that illustrates the pathological consequences of providing (local) differential privacy for non-compact spaces.

3.3.2 Linear regression with fixed design

We turn now to the problem of linear regression. Concretely, for a given design matrix $X \in \mathbb{R}^{n \times d}$, consider the standard linear model

$$Y = X\theta^* + \varepsilon, \tag{21}$$

where $\varepsilon \in \mathbb{R}^n$ is a vector of independent, zero-mean random variables. By rescaling as needed, we may assume that $\theta^* \in \Theta = \mathbb{B}_2(1)$, the Euclidean ball of radius one. Moreover, we assume that a scaling constant $\sigma < \infty$ such that the noise sequence $|\varepsilon_i| \leq \sigma$ for all i. Given the challenges of non-compactness exhibited by the location family estimation problems (cf. Proposition 1), this type of assumption is required for non-trivial results. We also assume that X has rank d; otherwise, the design matrix X has a non-trivial nullspace and θ^* cannot be estimated even when $\sigma = 0$.

With the model (21) in place, let us consider estimation of θ^* in the squared ℓ_2-error, where we provide α-locally differentially private views of the response $Y = \{Y_i\}_{i=1}^n$. By following the outline established in Section 3.2, we provide a sharp characterization of the α-private minimax rate. In stating the result, we let $\rho_j(A)$ denote the jth singular value of a matrix A. (See Section 7.4 for the proof.)

Proposition 2. *In the fixed design regression model where the variables Y_i and are α-locally differentially private for some $\alpha \in [0, 1]$,*

$$\min\left\{1, \frac{\sigma^2 d}{n\alpha^2 \rho_{\max}^2(X/\sqrt{n})}\right\} \lesssim \mathfrak{M}_n\left(\Theta, \|\cdot\|_2^2, \alpha\right) \lesssim \min\left\{1, \frac{\sigma^2 d}{\alpha^2 n \rho_{\min}^2(X/\sqrt{n})}\right\}. \tag{22}$$

To interpret the bounds (22), it is helpful to consider some special cases. First consider the case of an orthonormal design, meaning that $\frac{1}{n}X^\top X = I_{d \times d}$. The bounds (22) imply that $\mathfrak{M}_n(\Theta, \|\cdot\|_2^2, \alpha) \asymp \sigma^2 d/(n\alpha^2)$, so that the α-private minimax rate is fully determined (up to constant pre-factors). Standard minimax rates for linear regression problems scale as $\sigma^2 d/n$; thus, by comparison, we see that requiring differential privacy indeed causes an effective sample size reduction from n to $n\alpha^2$. More generally, up to the difference between the maximum and minimum singular values of the design X, Proposition 2 provides a sharp characterization of the α-private rate for fixed-design linear regression. As the proof makes clear, the upper bounds are attained by adding Laplacian noise to the response variables Y_i and solving the resulting normal equations as in standard linear regression. In this case, the standard Laplacian mechanism [24] is optimal.

4 Mutual information under local privacy: Fano's method

As we have previously noted, Theorem 1 provides indirect upper bounds on the mutual information. However, since the resulting bounds involve pairwise distances only, as in Corollary 1, they must be used with local packings. Exploiting Fano's inequality in its full generality requires a more sophisticated upper bound on the mutual information under local privacy, which is the main topic of this section. We illustrate this more powerful technique by deriving lower bounds for mean estimation problems in both classical as well as high-dimensional settings under the non-interactive privacy model (2).

4.1 Variational bounds on mutual information

We begin by introducing some definitions needed to state the result. Let V be a discrete random variable uniformly distributed over some finite set $\mathcal{V}$. Given a family of distributions $\{P_\nu, \nu \in \mathcal{V}\}$, we define the mixture distribution

$$\overline{P} := \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} P_\nu.$$

A sample $X \sim \overline{P}$ can be obtained by first drawing V from the uniform distribution over $\mathcal{V}$, and then conditionally on $V = \nu$, drawing X from the distribution P_ν. By definition, the mutual information between the random index V and the sample X is

$$I(X; V) = \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} D_{\mathrm{kl}}\left(P_\nu \| \overline{P}\right),$$

a representation that plays an important role in our theory. As in the definition (3), any conditional distribution Q induces the family of marginal distributions $\{M_\nu, \nu \in \mathcal{V}\}$ and the associated mixture $\overline{M} := \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} M_\nu$. Our goal is to upper bound the mutual information $I(Z_1, \dots, Z_n; V)$, where conditioned on $V = \nu$, the random variables Z_i are drawn according to M_ν.

Our upper bound is variational in nature: it involves optimization over a subset of the space $L^\infty(\mathcal{X}) := \left\{ f : \mathcal{X} \to \mathbb{R} \mid \|f\|_\infty < \infty \right\}$ of uniformly bounded functions, equipped with the usual norm $\|f\|_\infty = \sup_{x \in \mathcal{X}} |f(x)|$. We define the 1-ball of the supremum norm

$$\mathbb{B}_\infty(\mathcal{X}) := \left\{ \gamma \in L^\infty(\mathcal{X}) \mid \|\gamma\|_\infty \leq 1 \right\}. \tag{23}$$

We show that this set describes the maximal amount of perturbation allowed in the conditional Q. Since the set $\mathcal{X}$ is generally clear from context, we typically omit this dependence. For each $\nu \in \mathcal{V}$, we define the linear functional $\varphi_\nu : L^\infty(\mathcal{X}) \to \mathbb{R}$ by

$$\varphi_\nu(\gamma) = \int_{\mathcal{X}} \gamma(x)(dP_\nu(x) - d\overline{P}(x)).$$

With these definitions, we have the following result:

Theorem 2. *Let $\{P_\nu\}_{\nu \in \mathcal{V}}$ be an arbitrary collection of probability measures on $\mathcal{X}$, and let $\{M_\nu\}_{\nu \in \mathcal{V}}$ be the set of marginal distributions induced by an α-differentially private distribution Q. Then*

$$\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \left[D_{\mathrm{kl}}\left(M_\nu \| \overline{M}\right) + D_{\mathrm{kl}}\left(\overline{M} \| M_\nu\right) \right] \leq \frac{(e^\alpha - 1)^2}{|\mathcal{V}|} \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{\nu \in \mathcal{V}} (\varphi_\nu(\gamma))^2. \tag{24}$$

It is important to note that, at least up to constant factors, Theorem 2 is never weaker than the results provided by Theorem 1, including the bounds of Corollary 1. By definition of the linear functional φ_ν, we have

$$\sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{\nu \in \mathcal{V}} (\varphi_\nu(\gamma))^2 \overset{(i)}{\leq} \sum_{\nu \in \mathcal{V}} \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} (\varphi_\nu(\gamma))^2 = 4 \sum_{\nu \in \mathcal{V}} \left\| P_\nu - \overline{P} \right\|_{\mathrm{TV}}^2,$$

where inequality (i) follows by interchanging the summation and supremum. Overall, we have

$$I(Z; V) \leq 4(e^{\alpha} - 1)^2 \frac{1}{|\mathcal{V}|^2} \sum_{\nu, \nu' \in \mathcal{V}} \|P_\nu - P_{\nu'}\|_{\mathrm{TV}}^2 .$$

The strength of Theorem 2 arises from the fact that inequality (i)—the interchange of the order of supremum and summation—may be quite loose.

We now present a corollary that extends Theorem 2 to the setting of repeated sampling, providing a tensorization inequality analogous to Corollary 1. Let V be distributed uniformly at random in $\mathcal{V}$, and assume that given $V = \nu$, the observations X_i are sampled independently according to the distribution P_ν for $i = 1, \ldots, n$. For this corollary, we require the non-interactive setting (2) of local privacy, where each private variable Z_i depends only on X_i.

Corollary 4. *Suppose that the distributions $\{Q_i\}_{i=1}^n$ are α-locally differentially private in the non-interactive setting (2). Then*

$$I(Z_1, \ldots, Z_n; V) \leq n(e^{\alpha} - 1)^2 \frac{1}{|\mathcal{V}|} \sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} (\varphi_\nu(\gamma))^2 . \tag{25}$$

We provide the proof of Corollary 4 in Section 8.2. We conjecture that the bound (25) also holds in the fully interactive setting, but given well-known difficulties of characterizing multiple channel capacities with feedback [17, Chapter 15], it may be challenging to verify this conjecture.

Theorem 2 and Corollary 4 relate the amount of mutual information between the random perturbed views Z of the data to geometric or variational properties of the underlying packing $\mathcal{V}$ of the parameter space Θ. In particular, Theorem 2 and Corollary 4 show that if we can find a packing set $\mathcal{V}$ that yields linear functionals φ_ν whose sum has good "spectral" properties—meaning a small operator norm when taking suprema over L^∞-type spaces—we can provide sharper results.

4.2 Applications of Theorem 2 to mean estimation

In this section, we show how Theorem 2, coupled with Corollary 4, leads to sharp characterizations of the α-private minimax rates for classical and high-dimensional mean estimation problems. Our results show that for in d-dimensional mean-estimation problems, the requirement of α-local differential privacy causes a reduction in effective sample size from n to $n\alpha^2/d$. Throughout this section, we assume that the channel Q is *non-interactive*, meaning that the random variable Z_i depends only on X_i, and so that local privacy takes the simpler form (2). We also state each of our results for privacy parameter $\alpha \in [0, 1]$, but note that all of our bounds hold for any constant α, with appropriate changes in the numerical pre-factors.

Before proceeding, we describe two sampling mechanisms for enforcing α-local differential privacy. Our methods for achieving the upper bounds in minimax rates are based on unbiased estimators. Let us assume we wish to construct an α-private unbiased estimate Z for the vector $v \in \mathbb{R}^d$. The following sampling strategies are based on a radius $r > 0$ and a bound $B > 0$ specified for each problem, and they require the Bernoulli random variable

$$T \sim \mathrm{Bernoulli}(\pi_\alpha), \quad \text{where} \quad \pi_\alpha := e^{\alpha}/(e^{\alpha} + 1).$$
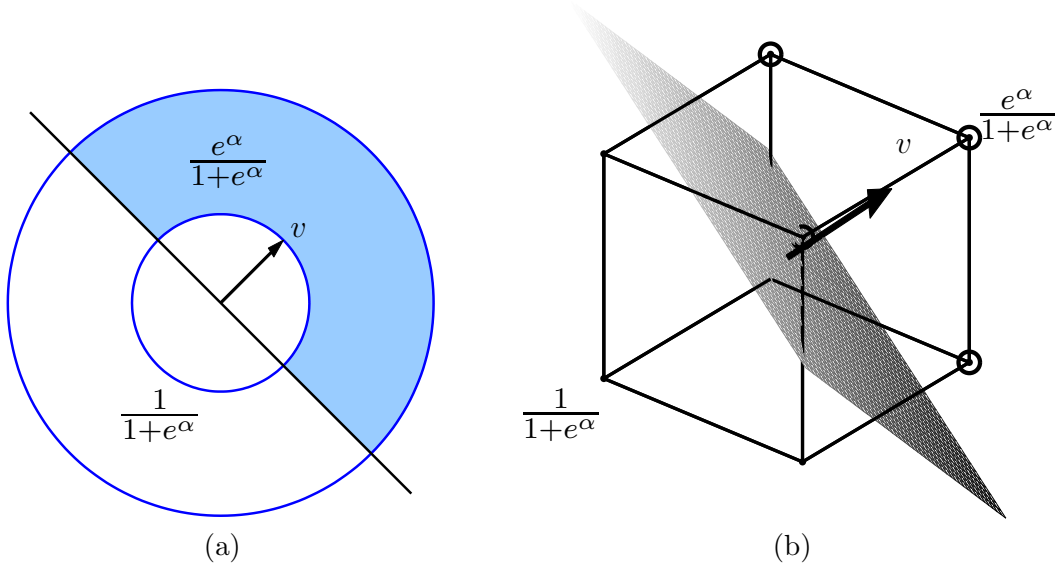
Figure 2. Private sampling strategies. (a) Strategy (26a) for the ℓ_2-ball. Outer boundary of highlighted region sampled uniformly with probability $e^\alpha/(e^\alpha+1)$. (b) Strategy (26b) for the ℓ_∞-ball. Circled point set sampled uniformly with probability $e^\alpha/(e^\alpha+1)$.

Strategy A: Given a vector v with $\|v\|_2 \leq r$, set $\widetilde{v} = rv/\|v\|_2$ with probability $\frac{1}{2} + \|v\|_2/2r$ and $\widetilde{v} = -rv/\|v\|_2$ with probability $\frac{1}{2} - \|v\|_2/2r$. Then sample $T \sim \text{Bernoulli}(\pi_\alpha)$ and set

$$Z \sim \begin{cases} \text{Uniform}(z \in \mathbb{R}^d : \langle z, \widetilde{v} \rangle > 0, \|z\|_2 = B) & \text{if } T = 1 \\ \text{Uniform}(z \in \mathbb{R}^d : \langle z, \widetilde{v} \rangle \leq 0, \|z\|_2 = B) & \text{if } T = 0. \end{cases} \tag{26a}$$

Strategy B: Given a vector v with $\|v\|_\infty \leq r$, construct $\widetilde{v} \in \mathbb{R}^d$ with coordinates $\widetilde{v}_j$ sampled independently from $\{-r, r\}$ with probabilities $1/2 - v_j/(2r)$ and $1/2 + v_j/(2r)$. Then sample $T \sim \text{Bernoulli}(\pi_\alpha)$ and set

$$Z \sim \begin{cases} \text{Uniform}(z \in \{-B, B\}^d : \langle z, \widetilde{v} \rangle > 0) & \text{if } T = 1 \\ \text{Uniform}(z \in \{-B, B\}^d : \langle z, \widetilde{v} \rangle \leq 0) & \text{if } T = 0. \end{cases} \tag{26b}$$

See Figure 2 for visualizations of these sampling strategies. By inspection, each is α-differentially private for any vector satisfying $\|v\|_2 \leq r$ or $\|v\|_\infty \leq r$ for Strategy A or B, respectively. Moreover, each strategy is efficiently implementable: Strategy A by normalizing a sample from the $\mathsf{N}(0, I_{d \times d})$ distribution, and Strategy B by rejection sampling over the scaled hypercube $\{-B, B\}^d$.

Given these sampling strategies, we study the d-dimensional problem of estimating the mean $\theta(P) := \mathbb{E}_P[X]$ of a random vector. We consider a few different metrics for the error of an estimator of the mean to flesh out the testing reduction in Section 2. Due to the difficulties associated with differential privacy on non-compact spaces (recall Section 3.3.1), we focus on distributions with compact support. We defer all proofs to Appendix A; they use a combination of Theorem 2 with Fano's method.

14

4.2.1 Minimax rates

We begin by bounding the minimax rate in the squared ℓ_2-metric. For a parameter $p \in [1, 2]$ and radius $r < \infty$, consider the family

$$\mathcal{P}_{p,r} := \big\{ \text{distributions } P \text{ supported on } \mathbb{B}_p(r) \subset \mathbb{R}^d \big\}. \tag{27}$$

where $\mathbb{B}_p(r) = \{ x \in \mathbb{R}^d \mid \|x\|_p \leq r \}$ is the ℓ_p-ball of radius r.

Proposition 3. *For the mean estimation problem, for all $p \in [1, 2]$ and privacy levels $\alpha \in [0, 1]$,*

$$r^2 \min \left\{ 1, \frac{1}{\sqrt{n\alpha^2}}, \frac{d}{n\alpha^2} \right\} \lesssim \mathfrak{M}_n(\theta(\mathcal{P}_{p,r}), \|\cdot\|_2^2, \alpha) \lesssim r^2 \min \left\{ \frac{d}{n\alpha^2}, 1 \right\}.$$

This bound does not depend on the norm for X so long as $p \in [1, 2]$, which is consistent with the classical mean estimation problem. Proposition 3 demonstrates the substantial difference between d-dimensional mean estimation in private and non-private settings: more precisely, the privacy constraint leads to a multiplicative penalty of d/α^2 in terms of mean-squared error. Indeed, in the non-private setting, the standard mean estimator $\widehat{\theta} = \frac{1}{n} \sum_{i=1}^n X_i$ has mean-squared error at most r^2/n, since $\|X\|_2 \leq \|X\|_p \leq r$ by assumption. Thus, Proposition 3 exhibits an effective sample size reduction of $n \mapsto n\alpha^2/d$.

To show the applicability of the general metric construction in Section 2, we now consider estimation in ℓ_∞-norm; estimation in this metric is natural in scenarios where one wishes only to guarantee that the maximum error of any particular component in the vector θ is small. We focus in this scenario on the family $\mathcal{P}_{\infty,r}$ of distributions P supported on $\mathbb{B}_\infty(r) \subset \mathbb{R}^d$.

Proposition 4. *For the mean estimation problem, for all $\alpha \in [0, 1]$,*

$$\min \left\{ r, \frac{r\sqrt{d \log(2d)}}{\sqrt{n\alpha^2}} \right\} \lesssim \mathfrak{M}_n(\theta(\mathcal{P}_{\infty,r}), \|\cdot\|_\infty, \alpha) \lesssim \min \left\{ r, \frac{r\sqrt{d \log(2d)}}{\sqrt{n\alpha^2}} \right\}.$$

Proposition 4 provides a similar message to Proposition 3 on the loss of statistical efficiency. This is clearest from an example: let X_i be random vectors bounded by one in ℓ_∞-norm. Then classical results on sub-Gaussian random variables [e.g., 12]) immediately imply that the standard non-private mean $\widehat{\theta} = \frac{1}{n} \sum_{i=1}^n X_i$ satisfies $\mathbb{E}[\|\widehat{\theta} - \mathbb{E}[X]\|_\infty] \leq \sqrt{\log(2d)/n}$. Comparing this result to the rate $\sqrt{d \log(2d)/n}$ of Proposition 4, we again see the effective sample size reduction $n \mapsto n\alpha^2/d$.

Recently, there has been substantial interest in high-dimensional problems, in which the dimension d is larger than the sample size n, but there is a low-dimensional latent structure that makes inference possible. (See the paper by Negahban et al. [44] for a general overview.) Accordingly, let us consider an idealized version of the high-dimensional mean estimation problem, in which we assume that $\theta(P) = \mathbb{E}[X] \in \mathbb{R}^d$ has (at most) one non-zero entry, so $\|\mathbb{E}[X]\|_0 \leq 1$. In the non-private case, estimation of such an s-sparse predictor in the squared ℓ_2-norm is possible at rate $\mathbb{E}[\|\widehat{\theta} - \theta\|_2^2] \leq s \log(d/s)/n$, so that the dimension d can be exponentially larger than the sample size n. With this context, the next result shows that local privacy can have a dramatic impact in the high-dimensional setting. Consider the family

$$\mathcal{P}_{\infty,r}^s := \big\{ \text{distributions } P \text{ supported on } \mathbb{B}_\infty(r) \subset \mathbb{R}^d \text{ with } \|\mathbb{E}_P[X]\|_0 \leq s \big\}.$$

Proposition 5. *For the 1-sparse means problem, for all $\alpha \in [0, 1]$,*

$$\min\left\{r^2, \frac{r^2 d \log(2d)}{n\alpha^2}\right\} \lesssim \mathfrak{M}_n\left(\theta(\mathcal{P}^1_{\infty,r}), \|\cdot\|^2_2, \alpha\right) \lesssim \min\left\{r^2, \frac{r^2 d \log(2d)}{n\alpha^2}\right\}.$$

See Section A.3 for a proof. From Proposition 5, it becomes clear that in locally private but non-interactive (2) settings, high-dimensional estimation is effectively impossible.

4.2.2   Optimal mechanisms: attainability for mean estimation

In this section, we describe how to achieve matching upper bounds in Propositions 3 and 4 using simple and practical algorithms—namely, the "right" type of stochastic perturbation of the observations X_i coupled with a standard mean estimator. We show the optimality of privatizing via the sampling strategies (26a) and (26b); interestingly, we also show that privatizing via Laplace perturbation is strictly sub-optimal. To give a private mechanism, we must specify the conditional distribution Q satisfying α-local differential privacy used to construct Z. In this case, given an observation X_i, we construct Z_i by perturbing X_i in such a way that $\mathbb{E}[Z_i \mid X_i = x] = x$. Each of the strategies (26a) and (26b) also requires a constant B, and we show how to choose B for each strategy to satisfy the unbiasedness condition $\mathbb{E}[Z \mid X = x] = x$.

   We begin with the mean estimation problem for distributions $\mathcal{P}_{p,r}$ in Proposition 3, for which we use the sampling scheme (26a). That is, let $X = x \in \mathbb{R}^d$ satisfy $\|x\|_2 \leq \|x\|_p \leq r$. Then we construct the random vector Z according to strategy (26a), where we set the initial vector $v = x$ in the sampling scheme. To achieve the unbiasedness condition $\mathbb{E}[Z \mid x] = x$, we define the bound

$$B = r\frac{e^\alpha + 1}{e^\alpha - 1}\frac{d\sqrt{\pi}\Gamma(\frac{d-1}{2} + 1)}{\Gamma(\frac{d}{2} + 1)}. \tag{28}$$

(See Appendix F.2 for a proof that $\mathbb{E}[Z \mid x] = x$ with this choice of B). Notably, the choice (28) implies $B \leq cr\sqrt{d}/\alpha$ for a universal constant $c < \infty$, since $d\Gamma(\frac{d-1}{2} + 1)/\Gamma(\frac{d}{2} + 1) \lesssim \sqrt{d}$ and $e^\alpha - 1 = \alpha + \mathcal{O}(\alpha^2)$. As a consequence, generating each Z_i by this perturbation strategy and using the mean estimator $\widehat{\theta} = \frac{1}{n}\sum_{i=1}^n Z_i$, the estimator $\widehat{\theta}$ is unbiased for $\mathbb{E}[X]$ and satisfies

$$\mathbb{E}\left[\|\widehat{\theta} - \mathbb{E}[X]\|^2_2\right] = \frac{1}{n^2}\sum_{i=1}^n \text{Var}(Z_i) \leq \frac{B^2}{n} \leq c\frac{r^2 d}{n\alpha^2}$$

for a universal constant c.

   In Proposition 4, we consider the family $\mathcal{P}_{\infty,r}$ of distributions supported on the ℓ_∞-ball of radius r. In our mechanism for attaining the upper bound, we use the sampling scheme (26b) to generate the private Z_i, so that for an observation $X = x \in \mathbb{R}^d$ with $\|x\|_\infty \leq r$, we resample Z (from the initial vector $v = x$) according to strategy (26b). Again, we would like to guarantee the unbiasedness condition $\mathbb{E}[Z \mid X = x] = x$, for which we use a result of Duchi et al. [19]. That paper shows that taking

$$B = c\frac{r\sqrt{d}}{\alpha} \tag{29}$$

for a (particular) universal constant c, yields the desired unbiasedness [19, Corollary 3]. Since the random variable Z satisfies $Z \in \mathbb{B}_\infty(r)$ with probability 1, each coordinate $[Z]_j$ of Z is sub-Gaussian.

As a consequence, we obtain via standard bounds [12] that

$$\mathbb{E}[\|\widehat{\theta} - \theta\|_\infty^2] \le \frac{B^2 \log(2d)}{n} = c^2 \frac{r^2 d \log(2d)}{n\alpha^2}$$

for a universal constant c, proving the upper bound in Proposition 4.

To conclude this section, we note that the strategy of adding Laplacian noise to the vectors X is sub-optimal. Indeed, consider the the family $\mathcal{P}_{2,1}$ of distributions supported on $\mathbb{B}_2(1) \subset \mathbb{R}^d$ as in Proposition 3. To guarantee α-differential privacy using independent Laplace noise vectors for $x \in \mathbb{B}_2(1)$, we take $Z = x + W$ where $W \in \mathbb{R}^d$ has components W_j that are independent and distributed as $\mathrm{Laplace}(\alpha/\sqrt{d})$. We have the following information-theoretic result: if the Z_i are constructed via the Laplace noise mechanism,

$$\inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\|\widehat{\theta}(Z_1, \ldots, Z_n) - \mathbb{E}_P[X]\|_2^2 \right] \gtrsim \min\left\{ \frac{d^2}{n\alpha^2}, 1 \right\}. \tag{30}$$

See Appendix A.4 for the proof of this claim. The poorer dimension dependence exhibted by the Laplace mechanism (30) in comparison to Proposition 3 demonstrates that sampling mechanisms must be chosen carefully, as in the strategies (26a)–(26b), in order to obtain statistically optimal rates.

5   Bounds on multiple pairwise divergences: Assouad's method

Thus far, we have seen how Le Cam's method and Fano's method, in the form of Theorem 2 and Corollary 4, can give sharp minimax rates for various problems. However, their application appears to be limited to problems whose minimax rates can be controlled via reductions to binary hypothesis tests (Le Cam's method) or for non-interactive channels satisfying the simpler definition (2) of local privacy (Fano's method). In this section, we show that a privatized form of Assouad's method (in the form of Lemma 1) can be used to obtain sharp minimax rates in interactive settings. In particular, it can be applied when the loss is sufficiently "decomposable," so that the coordinate-wise nature of the Assouad construction can be brought to bear. Concretely, we show that an upper bound on a sum of paired KL-divergences, when combined with Assouad's method, provides sharp lower bounds for several problems, including multinomial probability estimation and nonparametric density estimation. Each of these problems can be characterized in terms of an effective dimension d, and our results (paralleling those of Section 4) show that the requirement of α-local differential privacy causes a reduction in effective sample size from n to $n\alpha^2/d$.

5.1   Variational bounds on paired divergences

For a fixed $d \in \mathbb{N}$, we consider collections of distributions indexed using the Boolean hypercube $\mathcal{V} = \{-1, 1\}^d$. For each $i \in [n]$ and $\nu \in \mathcal{V}$, we let the distribution $P_{\nu,i}$ be supported on the fixed set $\mathcal{X}$, and we define the product distribution $P_\nu^n = \prod_{i=1}^n P_{\nu,i}$. Then for $j \in [d]$ we define the paired mixtures

$$P_{+j}^n = \frac{1}{2^{d-1}} \sum_{\nu:\nu_j=1} P_\nu^n, \quad P_{-j}^n = \frac{1}{2^{d-1}} \sum_{\nu:\nu_j=-1} P_\nu^n, \quad P_{\pm j,i} = \frac{1}{2^{d-1}} \sum_{\nu:\nu_j=\pm 1} P_{\nu,i}. \tag{31}$$

(Note that P_{+j}^n is not necessarily a product distribution.) Recalling the marginal channel (3), we may then define the marginal mixtures

$$M_{+j}^n(S) := \frac{1}{2^{d-1}} \sum_{\nu:\nu_j=1} M_\nu^n(S) = \int Q^n(S \mid x_{1:n}) dP_{+j}^n(x_{1:n}) \quad \text{for } j = 1, \dots, d,$$

with the distributions M_{-j}^n defined analogously. For a given pair of distributions (M, M'), we let $D_{\mathrm{kl}}^{\mathrm{sy}}(M \| M') = D_{\mathrm{kl}}(M \| M') + D_{\mathrm{kl}}(M' \| M)$ denote the symmetrized KL-divergence. Recalling the 1-ball of the supremum norm (23), with these definitions we have the following theorem:

Theorem 3. *Under the conditions of the previous paragraph, for any α-locally differentially private* (1) *channel Q, we have*

$$\sum_{j=1}^d D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j}^n \| M_{-j}^n\right) \le 2(e^\alpha - 1)^2 \sum_{i=1}^n \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{j=1}^d \left(\int_{\mathcal{X}} \gamma(x) dP_{+j,i}(x) - dP_{-j,i}(x) \right)^2.$$

Theorem 3 generalizes Theorem 1, which corresponds to the special case $d = 1$, though it also has parallels with Theorem 2, as taking the supremum outside the summation is essential to obtain sharp results. We provide the proof of Theorem 2 in Section 9.

Theorem 3 allows us to prove sharper lower bounds on the minimax risk. A combination of Pinsker's inequality and Cauchy-Schwarz implies

$$\sum_{j=1}^d \left\| M_{+j}^n - M_{-j}^n \right\|_{\mathrm{TV}} \le \frac{1}{2}\sqrt{d} \left(\sum_{j=1}^d D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\mathrm{kl}}\left(M_{-j}^n \| M_{+j}^n\right) \right)^{\frac{1}{2}}.$$

Thus, in combination with the sharper Assouad inequality (11), whenever P_ν induces a 2δ-Hamming separation for $\Phi \circ \rho$ we have

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \ge d\delta \left[1 - \left(\frac{1}{4d} \sum_{j=1}^d D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j}^n \| M_{-j}^n\right) \right)^{\frac{1}{2}} \right]. \tag{32}$$

The combination of inequality (32) with Theorem 3 is the foundation for the remainder of this section.

5.2 Multinomial estimation under local privacy

For our first application of Theorem 3, we return to the original motivation for local privacy [50]: avoiding survey answer bias. Consider the probability simplex

$$\Delta_d := \left\{ \theta \in \mathbb{R}^d \mid \theta \ge 0 \text{ and } \sum_{j=1}^d \theta_j = 1 \right\}.$$

Any vector $\theta \in \Delta_d$ specifies a multinomial random variable taking d states, in particular with probabilities $P_\theta(X = j) = \theta_j$ for $j \in \{1, \dots, d\}$. Given a sample from this distribution, our goal is to estimate the probability vector θ. Warner [50] studied the Bernoulli variant of this problem (corresponding to $d = 2$), proposing a mechanism known as *randomized response*: for a given survey question, respondents answer truthfully with probability $p > 1/2$ and lie with probability $1 - p$. Here we show that an extension of this mechanism is optimal for α-locally differentially private multinomial estimation.

5.2.1 Minimax rates of convergence for multinomial estimation

Our first result provides bounds on the minimax error measured in either the squared ℓ_2-norm or the ℓ_1-norm for (sequentially) interactive channels. The ℓ_1-norm is sometimes more appropriate for probability estimation due to its connections with total variation distance and testing.

Proposition 6. *For the multinomial estimation problem, for any α-locally differentially private channel* (1)*, there exist universal constants $0 < c_\ell \leq c_u < 5$ such that for all $\alpha \in [0, 1]$,*

$$c_\ell \min \left\{ 1, \frac{1}{\sqrt{n\alpha^2}}, \frac{d}{n\alpha^2} \right\} \leq \mathfrak{M}_n \left(\Delta_d, \|\cdot\|_2^2, \alpha \right) \leq c_u \min \left\{ 1, \frac{d}{n\alpha^2} \right\}, \tag{33}$$

and

$$c_\ell \min \left\{ 1, \frac{d}{\sqrt{n\alpha^2}} \right\} \leq \mathfrak{M}_n \left(\Delta_d, \|\cdot\|_1, \alpha \right) \leq c_u \min \left\{ 1, \frac{d}{\sqrt{n\alpha^2}} \right\}. \tag{34}$$

See Appendix B for the proofs of the lower bounds. We provide simple estimation strategies achieving the upper bounds in the next section.

As in the previous section, let us compare the private rates to the classical rate in which there is no privacy. The maximum likelihood estimate sets $\widehat{\theta}_j$ as the proportion of samples taking value j; it has mean-squared error

$$\mathbb{E}\left[\|\widehat{\theta} - \theta\|_2^2 \right] = \sum_{j=1}^{d} \mathbb{E}\left[(\widehat{\theta}_j - \theta_j)^2 \right] = \frac{1}{n} \sum_{j=1}^{d} \theta_j (1 - \theta_j) \leq \frac{1}{n}\left(1 - \frac{1}{d} \right) < \frac{1}{n}.$$

An analogous calculation for the ℓ_1-norm yields

$$\mathbb{E}[\|\widehat{\theta} - \theta\|_1] \leq \sum_{j=1}^{d} \mathbb{E}[|\widehat{\theta}_j - \theta_j|] \leq \sum_{j=1}^{d} \sqrt{\mathrm{Var}(\widehat{\theta}_j)} \leq \frac{1}{\sqrt{n}} \sum_{j=1}^{d} \sqrt{\theta_j(1 - \theta_j)} < \frac{\sqrt{d}}{\sqrt{n}}.$$

Consequently, for estimation in ℓ_1 or ℓ_2-norm, the effect of providing α-differential privacy causes the effective sample size to decrease as $n \mapsto n\alpha^2/d$.

5.2.2 Optimal mechanisms: attainability for multinomial estimation

An interesting consequence of the lower bound (33) is the following: a minor variant of Warner's randomized response strategy is an optimal mechanism. There are also other relatively simple estimation strategies that achieve convergence rate $d/n\alpha^2$; the Laplace perturbation approach [24] is another. Nonetheless, its ease of use, coupled with our optimality results, provide support for randomized response as a desirable probability estimation method.

Let us demonstrate that these strategies attain the optimal rate of convergence. Since there is a bijection between multinomial observations $x \in \{1, \ldots, d\}$ and the d standard basis vectors $e_1, \ldots, e_d \in \mathbb{R}^d$, we abuse notation and represent observations as either when designing estimation strategies. In randomized response, we construct the private vector $Z \in \{0, 1\}^d$ from a multinomial observation $x \in \{e_1, \ldots, e_d\}$ by sampling d coordinates independently via the procedure

$$[Z]_j = \begin{cases} x_j & \text{with probability } \frac{\exp(\alpha/2)}{1 + \exp(\alpha/2)} \\ 1 - x_j & \text{with probability } \frac{1}{1 + \exp(\alpha/2)}. \end{cases} \tag{35}$$

The distribution (35) is α-differentially private: indeed, for $x, x' \in \Delta_d$ and any $z \in \{0,1\}^d$, we have

$$\frac{Q(Z=z \mid x)}{Q(Z=z \mid x')} = \exp\left(\frac{\alpha}{2}\left(\|z-x\|_1 - \|z-x'\|_1\right)\right) \in [\exp(-\alpha), \exp(\alpha)],$$

where the triangle inequality guarantees $\left|\|z-x\|_1 - \|z-x'\|_1\right| \leq 2$. We now compute the expected value and variance of the random variables Z. Using the definition (35), we have

$$\mathbb{E}[Z \mid x] = \frac{e^{\alpha/2}}{1+e^{\alpha/2}}x + \frac{1}{1+e^{\alpha/2}}(\mathbb{1}-x) = \frac{e^{\alpha/2}-1}{e^{\alpha/2}+1}x + \frac{1}{1+e^{\alpha/2}}\mathbb{1}.$$

Since the random variables Z are Bernoulli, we have the variance bound $\mathbb{E}[\|Z\|_2^2] \leq d$. Letting Π_{Δ_d} denote the projection operator onto the simplex, we arrive at the natural estimator

$$\widehat{\theta}_{\text{part}} := \frac{1}{n}\sum_{i=1}^n \left(Z_i - \mathbb{1}/(1+e^{\alpha/2})\right)\frac{e^{\alpha/2}+1}{e^{\alpha/2}-1} \quad \text{and} \quad \widehat{\theta} := \Pi_{\Delta_d}\left(\widehat{\theta}_{\text{part}}\right). \tag{36}$$

The projection of $\widehat{\theta}_{\text{part}}$ onto the probability simplex can be done in time linear in the dimension d of the problem [11], so the estimator (36) is efficiently computable. Since projections onto convex sets are non-expansive, any pair of vectors in the simplex are at most ℓ_2-distance $\sqrt{2}$ apart, and $\mathbb{E}_\theta[\widehat{\theta}_{\text{part}}] = \theta$ by construction, we have

$$\mathbb{E}\left[\|\widehat{\theta}-\theta\|_2^2\right] \leq \min\left\{2, \mathbb{E}\left[\|\widehat{\theta}_{\text{part}}-\theta\|_2^2\right]\right\}$$

$$\leq \min\left\{2, \frac{d}{n}\left(\frac{e^{\alpha/2}+1}{e^{\alpha/2}-1}\right)^2\right\} \lesssim \min\left\{1, \frac{d}{n\alpha^2}\right\}.$$

Similar results hold for the ℓ_1-norm: using the same estimator, since Euclidean projections to the simplex are non-expansive for the ℓ_1 distance,

$$\mathbb{E}\left[\|\widehat{\theta}-\theta\|_1\right] \leq \min\left\{1, \sum_{j=1}^d \mathbb{E}\left[|\widehat{\theta}_{\text{part},j}-\theta_j|\right]\right\} \lesssim \min\left\{1, \frac{d}{\sqrt{n\alpha^2}}\right\}.$$

5.3 Density estimation under local privacy

In this section, we show that the effects of local differential privacy are more severe for nonparametric density estimation: instead of just a multiplicative loss in the effective sample size as in previous sections, imposing local differential privacy leads to a different convergence rate. This result holds even though we solve a problem in which the function estimated and the observations themselves belong to compact spaces.

Definition 2 (Elliptical Sobolev space). For a given orthonormal basis $\{\varphi_j\}$ of $L^2([0,1])$, smoothness parameter $\beta > 1/2$ and radius C, the Sobolev class of order β is given by

$$\mathcal{F}_\beta[C] := \left\{f \in L^2([0,1]) \mid f = \sum_{j=1}^\infty \theta_j\varphi_j \text{ such that } \sum_{j=1}^\infty j^{2\beta}\theta_j^2 \leq C^2\right\}.$$

If we choose the trignometric basis as our orthonormal basis, membership in the class $\mathcal{F}_\beta[C]$ corresponds to smoothness constraints on the derivatives of f. More precisely, for $j \in \mathbb{N}$, consider the orthonormal basis for $L^2([0,1])$ of trigonometric functions:

$$\varphi_0(t) = 1, \quad \varphi_{2j}(t) = \sqrt{2}\cos(2\pi jt), \quad \varphi_{2j+1}(t) = \sqrt{2}\sin(2\pi jt). \tag{37}$$

Let f be a β-times almost everywhere differentiable function for which $|f^{(\beta)}(x)| \leq C$ for almost every $x \in [0,1]$ satisfying $f^{(k)}(0) = f^{(k)}(1)$ for $k \leq \beta - 1$. Then, uniformly over all such f, there is a universal constant $c \leq 2$ such that that $f \in \mathcal{F}_\beta[cC]$ (see, for instance, [49, Lemma A.3]).

Suppose our goal is to estimate a density function $f \in \mathcal{F}_\beta[C]$ and that quality is measured in terms of the squared error (squared $L^2[0,1]$-norm)

$$\|\widehat{f} - f\|_2^2 := \int_0^1 (\widehat{f}(x) - f(x))^2 dx.$$

The well-known [53, 52, 49] (non-private) minimax squared risk scales as

$$\mathfrak{M}_n\left(\mathcal{F}_\beta, \|\cdot\|_2^2, \infty\right) \asymp n^{-\frac{2\beta}{2\beta+1}}. \tag{38}$$

The goal of this section is to understand how this minimax rate changes when we add an α-privacy constraint to the problem. Our main result is to demonstrate that the classical rate (38) is no longer attainable when we require α-local differential privacy.

5.3.1 Lower bounds on density estimation

We begin by giving our main lower bound on the minimax rate of estimation of densities when observations from the density are differentially private. We provide the proof of the following proposition in Section C.1.

Proposition 7. *Consider the class of densities $\mathcal{F}_\beta$ defined using the trigonometric basis (37). There exists a constant $c_\beta > 0$ such that for any α-locally differentially private channel (1) with $\alpha \in [0,1]$, the private minimax risk has lower bound*

$$\mathfrak{M}_n\left(\mathcal{F}_\beta[1], \|\cdot\|_2^2, \alpha\right) \geq c_\beta \left(n\alpha^2\right)^{-\frac{2\beta}{2\beta+2}}. \tag{39}$$

The most important feature of the lower bound (39) is that it involves a *different polynomial exponent* than the classical minimax rate (38). Whereas the exponent in classical case (38) is $2\beta/(2\beta+1)$, it reduces to $2\beta/(2\beta+2)$ in the locally private setting. For example, when we estimate Lipschitz densities ($\beta = 1$), the rate degrades from $n^{-2/3}$ to $n^{-1/2}$.

Interestingly, no estimator based on Laplace (or exponential) perturbation of the observations X_i themselves can attain the rate of convergence (39). This fact follows from results of Carroll and Hall [13] on nonparametric deconvolution. They show that if observations X_i are perturbed by additive noise W, where the characteristic function ϕ_W of the additive noise has tails behaving as $|\phi_W(t)| = \mathcal{O}(|t|^{-a})$ for some $a > 0$, then no estimator can deconvolve $X + W$ and attain a rate of convergence better than $n^{-2\beta/(2\beta+2a+1)}$. Since the characteristic function of the Laplace distribution has tails decaying as t^{-2}, no estimator based on the Laplace mechanism (applied directly to the observations) can attain rate of convergence better than $n^{-2\beta/(2\beta+5)}$. In order to attain the lower bound (39), we must thus study alternative privacy mechanisms.

5.3.2 Achievability by histogram estimators

We now turn to the mean-squared errors achieved by specific practical schemes, beginning with the special case of Lipschitz density functions ($\beta = 1$). In this special case, it suffices to consider a private version of a classical histogram estimate. For a fixed positive integer $k \in \mathbb{N}$, let $\{\mathcal{X}_j\}_{j=1}^{k}$ denote the partition of $\mathcal{X} = [0, 1]$ into the intervals

$$\mathcal{X}_j = [(j-1)/k, j/k) \quad \text{for } j = 1, 2, \ldots, k-1, \text{ and } \mathcal{X}_k = [(k-1)/k, 1].$$

Any histogram estimate of the density based on these k bins can be specified by a vector $\theta \in k\Delta_k$, where we recall $\Delta_k \subset \mathbb{R}_+^k$ is the probability simplex. Letting $\mathbf{1}_E$ denote the characteristic (indicator) function of the set E, any such vector $\theta \in \mathbb{R}^k$ defines a density estimate via the sum

$$f_\theta := \sum_{j=1}^{k} \theta_j \mathbf{1}_{\mathcal{X}_j}.$$

Let us now describe a mechanism that guarantees α-local differential privacy. Given a sample $\{X_1, \ldots, X_n\}$ from the distribution f, consider vectors

$$Z_i := \mathsf{e}_k(X_i) + W_i, \quad \text{for } i = 1, 2, \ldots, n, \tag{40}$$

where $\mathsf{e}_k(X_i) \in \Delta_k$ is a k-vector with j^{th} entry equal to one if $X_i \in \mathcal{X}_j$ and zeroes in all other entries, and W_i is a random vector with i.i.d. Laplace($\alpha/2$) entries. The variables $\{Z_i\}_{i=1}^{n}$ defined in this way are α-locally differentially private for $\{X_i\}_{i=1}^{n}$. Using these private variables, we form the density estimate $\widehat{f} := f_{\widehat{\theta}} = \sum_{j=1}^{k} \widehat{\theta}_j \mathbf{1}_{\mathcal{X}_j}$ based on the vector $\widehat{\theta} := \Pi_k \left(\frac{k}{n} \sum_{i=1}^{n} Z_i \right)$, where Π_k denotes the Euclidean projection operator onto the set $k\Delta_k$. By construction, we have $\widehat{f} \geq 0$ and $\int_0^1 \widehat{f}(x) dx = 1$, so $\widehat{f}$ is a valid density estimate. The following result characterizes its mean-squared estimation error:

Proposition 8. *Consider the estimate $\widehat{f}$ based on $k = (n\alpha^2)^{1/4}$ bins in the histogram. For any 1-Lipschitz density $f : [0, 1] \to \mathbb{R}_+$, the MSE is upper bounded as*

$$\mathbb{E}_f \left[\|\widehat{f} - f\|_2^2 \right] \leq 5(\alpha^2 n)^{-\frac{1}{2}} + \sqrt{\alpha} n^{-3/4}. \tag{41}$$

For any fixed $\alpha > 0$, the first term in the bound (41) dominates, and the $\mathcal{O}((\alpha^2 n)^{-\frac{1}{2}})$ rate matches the minimax lower bound (39) in the case $\beta = 1$. Consequently, the privatized histogram estimator is minimax-optimal for Lipschitz densities, providing a private analog of the classical result that histogram estimators are minimax-optimal for Lipshitz densities. See Section C.2 for a proof of Proposition 8. We remark that a randomized response scheme parallel to that of Section 5.2.2 achieves the same rate of convergence, showing that this classical mechanism is again an optimal scheme.

5.3.3 Achievability by orthogonal projection estimators

For higher degrees of smoothness ($\beta > 1$), standard histogram estimators no longer achieve optimal rates in the classical setting [47]. Accordingly, we now turn to developing estimators based on orthogonal series expansion, and show that even in the setting of local privacy, they can achieve the lower bound (39) for all orders of smoothness $\beta \geq 1$.

Recall the elliptical Sobolev space (Definition 2), in which a function f is represented in terms of its basis expansion $f = \sum_{j=1}^{\infty} \theta_j \varphi_j$. This representation underlies the orthonormal series estimator as follows. Given a sample $X_{1:n}$ drawn i.i.d. according to a density $f \in L^2([0,1])$, compute the empirical basis coefficients

$$\widehat{\theta}_j = \frac{1}{n} \sum_{i=1}^{n} \varphi_j(X_i) \quad \text{for } j \in \{1, \ldots, k\}, \tag{42}$$

where the value $k \in \mathbb{N}$ is chosen either a priori based on known properties of the estimation problem or adaptively, for example, using cross-validation [26, 49]. Using these empirical coefficients, the density estimate is $\widehat{f} = \sum_{j=1}^{k} \widehat{\theta}_j \varphi_j$.

In our local privacy setting, we consider a mechanism that, instead of releasing the vector of coefficients $\big(\varphi_1(X_i), \ldots, \varphi_k(X_i) \big)$ for each data point, employs a random vector $Z_i = (Z_{i,1}, \ldots, Z_{i,k})$ satisfying $\mathbb{E}[Z_{i,j} \mid X_i] = \varphi_j(X_i)$ for each $j \in [k]$. We assume the basis functions are B_0-uniformly bounded, that is, $\sup_j \sup_x |\varphi_j(x)| \leq B_0 < \infty$. This boundedness condition holds for many standard bases, including the trigonometric basis (37) that underlies the classical Sobolev classes and the Walsh basis. We generate the random variables from the vector $v \in \mathbb{R}^k$ defined by $v_j = \varphi_j(X)$ in the hypercube-based sampling scheme (26b), where we assume that the outer bound $B > B_0$. With this sampling strategy, iteration of expectation yields

$$\mathbb{E}[[Z]_j \mid X = x] = c_k \frac{B}{B_0 \sqrt{k}} \left(\frac{e^\alpha}{e^\alpha + 1} - \frac{1}{e^\alpha + 1} \right) \varphi_j(x), \tag{43}$$

where $c_k > 0$ is a constant (which is bounded independently of k). Consequently, it suffices to take $B = \mathcal{O}(B_0 \sqrt{k}/\alpha)$ to guarantee the unbiasedness condition $\mathbb{E}[[Z_i]_j \mid X_i] = \varphi_j(X_i)$.

Overall, the privacy mechanism and estimator perform the following steps:

- given a data point X_i, set the vector $v = [\varphi_j(X_i)]_{j=1}^{k}$;

- sample Z_i according to the strategy (26b), starting from the vector v and using the bound $B = B_0 \sqrt{k}(e^\alpha + 1)/c_k(e^\alpha - 1)$;

- compute the density estimate

$$\widehat{f} := \frac{1}{n} \sum_{i=1}^{n} \sum_{j=1}^{k} Z_{i,j} \varphi_j. \tag{44}$$

The resulting estimate enjoys the following guarantee, which (along with Proposition 8) makes clear that the private minimax lower bound (39) is sharp, providing a variant of the classical rates with a polynomially worse sample complexity. (See Section C.3 for a proof.)

Proposition 9. *Let $\{\varphi_j\}$ be a B_0-uniformly bounded orthonormal basis for $L^2([0,1])$. There exists a constant c (depending only on C and B_0) such that, for any f in the Sobolev space $\mathcal{F}_\beta[C]$, the estimator (44) with $k = (n\alpha^2)^{1/(2\beta+2)}$ has an MSE that is upper bounded as follows:*

$$\mathbb{E}_f \left[\| f - \widehat{f} \|_2^2 \right] \leq c \left(n\alpha^2 \right)^{-\frac{2\beta}{2\beta+2}}. \tag{45}$$

Before concluding our exposition, we make a few remarks on other potential density estimators. Our orthogonal series estimator (44) and sampling scheme (43), while similar in spirit to that proposed by Wasserman and Zhou [51, Sec. 6], is different in that it is locally private and requires a different noise strategy to obtain both α-local privacy and the optimal convergence rate. Lastly, similarly to our remarks on the insufficiency of standard Laplace noise addition for mean estimation, it is worth noting that density estimators that are based on orthogonal series and Laplace perturbation are sub-optimal: they can achieve (at best) rates of $(n\alpha^2)^{-\frac{2\beta}{2\beta+3}}$. This rate is polynomially worse than the sharp result provided by Proposition 9. Again, we see that appropriately chosen noise mechanisms are crucial for obtaining optimal results.

6 Comparison to related work

There has been a substantial amount of work in developing differentially private mechanisms, both in local and non-local settings, and a number of authors have attempted to characterize optimal mechanisms. For example, Kasiviswanathan et al. [37], working within a local differential privacy setting, study Probably-Approximately-Correct (PAC) learning problems and show that the statistical query model [38] and local learning are equivalent up to polynomial changes in the sample size. In our work, we are concerned with a finer-grained assessment of inferential procedures—that of rates of convergence of procedures and their optimality. In the remainder of this section, we discuss further connections of our work to previous research on optimality, global (non-local) differential privacy, as well as error-in-variables models.

6.1 Sample versus population estimation

The standard definition of differential privacy, due to Dwork et al. [24], is somewhat less restrictive than the local privacy formulation considered here. In particular, a conditional distribution Q with output space $\mathcal{Z}$ is α-differentially private if

$$\sup\left\{\frac{Q(S \mid x_{1:n})}{Q(S \mid x'_{1:n})} \mid x_i, x'_i \in \mathcal{X}, S \in \sigma(\mathcal{Z}), d_{\mathrm{ham}}(x_{1:n}, x'_{1:n}) \le 1\right\} \le \exp(\alpha),\tag{46}$$

where d_{ham} denotes the Hamming distance between sets. Several researchers have considered quantities similar to our minimax criteria under local (2) or non-local (46) differential privacy [7, 35, 33, 18]. However, the objective has often been quite different from ours: instead of bounding errors based on population-based quantities, they provide bounds in which the data are assumed to be held fixed. More precisely, let $\theta : \mathcal{X}^n \to \Theta$ denote an estimator, and let $\theta(x_{1:n})$ be a sample quantity based on $x_{1:n}$. Prior work is based on *conditional minimax* risks of the form

$$\mathfrak{M}_n^{\mathsf{cond}}(\theta(\mathcal{X}), \Phi \circ \rho, \alpha) := \inf_Q \sup_{x_{1:n} \in \mathcal{X}^n} \mathbb{E}_Q\left[\Phi\big(\rho\big(\theta(x_{1:n}), \widehat{\theta}\big)\big) \mid X_{1:n} = x_{1:n}\right],\tag{47}$$

where $\widehat{\theta}$ is drawn according to $Q(\cdot \mid x_{1:n})$, the infimum is taken over all α-differentially private channels Q, and the supremum is taken over all possible samples of size n. The only randomness in this conditional minimax risk is provided by the channel; the data are held fixed, so there is no randomness from an underlying population distribution. A partial list of papers that use definitions of this type include Beimel et al. [7, Section 2.4], Hardt and Talwar [35, Definition 2.4], Hall et al. [33, Section 3], and De [18].

The conditional (47) and population minimax risk (5) can differ substantially, and such differences are critical to address within a statistical approach to privacy-constrained inference. The goal of inference is to draw conclusions about the *population-based quantity* $\theta(P)$ based on the sample. Moreover, lower bounds on the conditional minimax risk (47) do not imply bounds on the rate of estimation for the population $\theta(P)$. In fact, the conditional minimax risk (47) involves a supremum over *all possible samples* $x \in \mathcal{X}$, so the opposite is usually true: population risks provide lower bounds on the conditional minimax risk, as we show presently.

An illustrative example is useful to understand the differences. Consider estimation of the mean of a normal distribution with known standard deviation σ^2, in which the mean $\theta = \mathbb{E}[X] \in [-1, 1]$ is assumed to belong to the unit interval. As our Proposition 1 shows, it is possible to estimate the mean of a normally-distributed random variable even under α-local differential privacy (1). In sharp contrast, the following result shows that the conditional minimax risk is infinite for this problem:

Lemma 2. *Consider the normal location family* $\{\mathsf{N}(\theta, \sigma^2) \mid \theta \in [-1, 1]\}$ *under α-differential privacy (46). The conditional minimax risk of the mean statistic is* $\mathfrak{M}_n^{\mathsf{cond}}(\theta(\mathbb{R}), (\cdot)^2, \alpha) = \infty$.

Proof. Assume for sake of contradiction that $\delta > 0$ satisfies

$$Q(|\widehat{\theta} - \theta(x_{1:n})| > \delta \mid x_{1:n}) \leq \frac{1}{2} \quad \text{for all samples } x_{1:n} \in \mathbb{R}^n.$$

Fix $N(\delta) \in \mathbb{N}$ and choose points 2δ-separated points θ_ν, $\nu \in [N(\delta)]$, that is, $|\theta_\nu - \theta_{\nu'}| \geq 2\delta$ for $\nu \neq \nu'$. Then the sets $\{\theta \in \mathbb{R} \mid |\theta - \theta_\nu| \leq \delta\}$ are all disjoint, so for any pair of samples $x_{1:n}$ and $x_{1:n}^\nu$ with $d_{\mathrm{ham}}(x_{1:n}, x_{1:n}^\nu) \leq 1$,

$$Q(\exists \nu \in \mathcal{V} \text{ s.t. } |\widehat{\theta} - \theta_\nu| \leq \delta \mid x_{1:n}) = \sum_{\nu=1}^{N(\delta)} Q(|\widehat{\theta} - \theta_\nu| \leq \delta \mid x_{1:n})$$

$$\geq e^{-\alpha} \sum_{\nu=1}^{N(\delta)} Q(|\widehat{\theta} - \theta_\nu| \leq \delta \mid x_{1:n}^\nu).$$

We may take each sample $x_{1:n}^\nu$ such that $\theta(x_{1:n}^\nu) = \frac{1}{n}\sum_{i=1}^n x_i^\nu = \theta_\nu$ (for example, for each $\nu \in [N(\delta)]$ set $x_1^\nu = n\theta_\nu - \sum_{i=2}^n x_i$) and by assumption,

$$1 \geq Q(\exists \nu \in \mathcal{V} \text{ s.t. } |\widehat{\theta} - \theta_\nu| \leq \delta \mid x_{1:n}) \geq e^{-\alpha} N(\delta) \frac{1}{2}.$$

Taking $N(\delta) > 2e^\alpha$ yields a contradiction. Our argument applies to an arbitrary $\delta > 0$, so the claim follows. $\qquad\square$

There are variations on this result. For instance, even if the output of the mean estimator is restricted to $[-1, 1]$, the conditional minimax risk remains constant. Similar arguments apply to weakenings of differential privacy (e.g., δ-approximate α-differential privacy [23]). Conditional and population risks are very different quantities.

More generally, the population minimax risk usually lower bounds the conditional minimax risk. Suppose we measure minimax risks in some given metric ρ (so the loss $\Phi(t) = t$). Let $\widetilde{\theta}$ be any

estimator based on the original sample $X_{1:n}$, and let $\widehat{\theta}$ be any estimator based on the privatized sample. We then have the following series of inequalities:

$$\mathbb{E}_{Q,P}[\rho(\theta(P), \widehat{\theta})] \leq \mathbb{E}_{Q,P}[\rho(\theta(P), \widetilde{\theta})] + \mathbb{E}_{Q,P}[\rho(\widetilde{\theta}, \widehat{\theta})]$$
$$\leq \mathbb{E}_P[\rho(\theta(P), \widetilde{\theta})] + \sup_{x_{1:n} \in \mathcal{X}^n} \mathbb{E}_{Q,P}[\rho(\widetilde{\theta}(x_{1:n}), \widehat{\theta}) \mid X_{1:n} = x_{1:n}]. \qquad (48)$$

The population minimax risk (5) thus lower bounds the conditional minimax risk (47) via $\mathfrak{M}_n^{\mathrm{cond}}(\widetilde{\theta}(\mathcal{X}), \rho, \alpha) \geq \mathfrak{M}_n(\theta(\mathcal{P}), \rho, \alpha) - \mathbb{E}_P[\rho(\theta(P), \widetilde{\theta})]$. In particular, if there exists an estimator $\widetilde{\theta}$ based on the original (non-private data) such that $\mathbb{E}_P[\rho(\theta(P), \widetilde{\theta})] \leq \frac{1}{2}\mathfrak{M}_n(\theta(\mathcal{P}), \rho, \alpha)$ we are guaranteed that

$$\mathfrak{M}_n^{\mathrm{cond}}(\widetilde{\theta}(\mathcal{X}), \rho, \alpha) \geq \frac{1}{2}\mathfrak{M}_n(\theta(\mathcal{P}), \rho, \alpha),$$

so the conditional minimax risk is lower bounded by a constant multiple of the population minimax risk. This lower bound holds for each of the examples in Sections 3–5; lower bounds on the α-private population minimax risk (5) are stronger than lower bounds on the conditional minimax risk.

To illustrate one application of the lower bound (48), consider the estimation of the sample mean of a data set $x_{1:n} \in \{0, 1\}^n$ under α-local privacy. This problem has been considered before; for instance, Beimel et al. [7] study distributed protocols for this problem. In Theorem 2 of their work, they show that if a protocol has ℓ rounds of communication, the squared error in estimating the sample mean $(1/n)\sum_{i=1}^n x_i$ is $\Omega(1/(n\alpha^2\ell^2))$. The standard mean estimator $\widetilde{\theta}(x_{1:n}) = (1/n)\sum_{i=1}^n x_i$ has error $\mathbb{E}[|\widetilde{\theta}(x_{1:n}) - \theta|] \leq n^{-\frac{1}{2}}$. Consequently, the lower bound (48) with combined with Proposition 1 implies

$$c\frac{1}{\sqrt{n\alpha^2}} - \frac{1}{\sqrt{n}} \leq \mathfrak{M}_n(\theta(\mathcal{P}), |\cdot|, \alpha) - \sup_{\theta \in [-1, 1]} \mathbb{E}[|\widetilde{\theta}(x_{1:n}) - \theta|] \leq \mathfrak{M}_n^{\mathrm{cond}}(\theta(\{-1, 1\}), |\cdot|, \alpha),$$

for some numerical constant $c > 0$. A corollary of our results is such an $\Omega(1/(n\alpha^2))$ lower bound on the conditional minimax risk for mean estimation, allowing for sequential interactivity but not multiple "rounds." An inspection of Beimel et al.'s proof technique [7, Section 4.2] shows that their lower bound also implies a lower bound of $1/n\alpha^2$ for estimation of the population mean $\mathbb{E}[X]$ in one dimension in *non-interactive* (2) settings; it is, however, unclear how to extend their technique to other settings.

6.2 Local versus non-local privacy

It is also worthwhile to make some comparisons to work on non-local forms of differential privacy, mainly to understand the differences between local and global forms of privacy. Chaudhuri and Hsu [15] provide lower bounds for estimation of certain one-dimensional statistics based on a two-point family of problems. Their techniques differ from those of the current paper, and do not appear to provide bounds on the statistic being estimated, but rather one that is near to it. Beimel et al. [8] provide some bounds on sample complexity in the "probably approximate correct" (PAC) framework of learning theory, though extensions to other inferential tasks are unclear. Other work on non-local privacy [e.g., 33, 16, 48] shows that for various types of estimation problems, adding Laplacian noise leads to degraded convergence rates in at most lower-order terms. In contrast, our work shows that the Laplace mechanism may be highly sub-optimal in local privacy.

To understand convergence rates for non-local privacy, let us return to estimation of a multi-nomial distribution in Δ_d, based on observations $X_i \in \{e_j\}_{j=1}^d$. In this case, adding a noise vector $W \in \mathbb{R}^d$ with i.i.d. entries distributed as Laplace(αn) provides differential privacy [23]; the associated mean-squared error is at most

$$\mathbb{E}_\theta\left[\left\|\frac{1}{n}\sum_{i=1}^n X_i + W - \theta\right\|_2^2\right] = \mathbb{E}\left[\left\|\frac{1}{n}\sum_{i=1}^n X_i - \theta\right\|_2^2\right] + \mathbb{E}[\|W\|_2^2] \leq \frac{1}{n} + \frac{d}{n^2\alpha^2}.$$

In particular, in the asymptotic regime $n \gg d$, there is no penalty from providing differential privacy except in higher-order terms. Similar results hold for histogram estimation [33], classification problems [16], and classical point estimation problems [48]; in this sense, local and global forms of differential privacy can be rather different.

6.3 Error-in-variables models

As a final remark on related work, we touch briefly on errors-in-variables models [14, 31], which have been the subject of extensive study. In such problems, one observes a corrupted version Z_i of the true covariate X_i. Privacy analysis is one of the few settings in which it is possible to precisely know the conditional distribution $Q(\cdot \mid X_i)$. However, the mechanisms that are optimal from our analysis—in particular, those in strategies (26a) and (26b)—are more complicated than adding noise directly to the covariates, which leads to complications. Known (statistically) efficient error-in-variables estimation procedures often require either solving certain integrals or estimating equations, or solving non-convex optimization problems [e.g., 39, 41]. Some recent work [40] shows that certain types of non-convex programs arising from errors-in-variables can be solved efficiently. In density estimation (as noted in Section 5.3.1), corrupted observations lead to nonparametric deconvolution problems that appear harder than estimation under privacy constraints. Further investigation of computationally efficient procedures for nonlinear error-in-variables models for privacy-preservation is an interesting direction for future research.

7 Proof of Theorem 1 and related results

We now turn to the proofs of our results, beginning with Theorem 1 and related results. In all cases, we defer the proofs of more technical lemmas to the appendices.

7.1 Proof of Theorem 1

Observe that M_1 and M_2 are absolutely continuous with respect to one another, and there is a measure μ with respect to which they have densities m_1 and m_2, respectively. The channel probabilities $Q(\cdot \mid x)$ and $Q(\cdot \mid x')$ are likewise absolutely continuous, so that we may assume they have densities $q(\cdot \mid x)$ and write $m_i(z) = \int q(z \mid x)dP_i(x)$. In terms of these densities, we have

$$D_{\mathrm{kl}}\left(M_1\|M_2\right) + D_{\mathrm{kl}}\left(M_2\|M_1\right) = \int m_1(z)\log\frac{m_1(z)}{m_2(z)}d\mu(z) + \int m_2(z)\log\frac{m_2(z)}{m_1(z)}d\mu(z)$$

$$= \int \left(m_1(z) - m_2(z)\right)\log\frac{m_1(z)}{m_2(z)}d\mu(z).$$

Consequently, we must bound both the difference $m_1 - m_2$ and the log ratio of the marginal densities. The following two auxiliary lemmas are useful:

Lemma 3. *For any α-locally differentially private conditional, we have*

$$|m_1(z) - m_2(z)| \le c_\alpha \inf_x q(z \mid x) \, (e^\alpha - 1) \, \|P_1 - P_2\|_{\mathrm{TV}}, \tag{49}$$

where $c_\alpha = \min\{2, e^\alpha\}$.

Lemma 4. *Let $a, b \in \mathbb{R}_+$. Then $\left|\log \frac{a}{b}\right| \le \frac{|a-b|}{\min\{a,b\}}$.*

We prove these two results at the end of this section.

With the lemmas in hand, let us now complete the proof of the theorem. From Lemma 4, the log ratio is bounded as

$$\left|\log \frac{m_1(z)}{m_2(z)}\right| \le \frac{|m_1(z) - m_2(z)|}{\min\{m_1(z), m_2(z)\}}.$$

Applying Lemma 3 to the numerator yields

$$\begin{aligned}
\left|\log \frac{m_1(z)}{m_2(z)}\right| &\le \frac{c_\alpha \, (e^\alpha - 1) \, \|P_1 - P_2\|_{\mathrm{TV}} \, \inf_x q(z \mid x)}{\min\{m_1(z), m_2(z)\}} \\
&\le \frac{c_\alpha \, (e^\alpha - 1) \, \|P_1 - P_2\|_{\mathrm{TV}} \, \inf_x q(z \mid x)}{\inf_x q(z \mid x)},
\end{aligned}$$

where the final step uses the inequality $\min\{m_1(z), m_2(z)\} \ge \inf_x q(z \mid x)$. Putting together the pieces leads to the bound

$$\left|\log \frac{m_1(z)}{m_2(z)}\right| \le c_\alpha(e^\alpha - 1) \, \|P_1 - P_2\|_{\mathrm{TV}}.$$

Combining with inequality (49) yields

$$D_{\mathrm{kl}}\left(M_1\|M_2\right) + D_{\mathrm{kl}}\left(M_2\|M_1\right) \le c_\alpha^2 \, (e^\alpha - 1)^2 \, \|P_1 - P_2\|_{\mathrm{TV}}^2 \int \inf_x q(z \mid x) d\mu(z).$$

The final integral is at most one, which completes the proof of the theorem.

It remains to prove Lemmas 3 and 4. We begin with the former. For any $z \in \mathcal{Z}$, we have

$$\begin{aligned}
m_1(z) - m_2(z) &= \int_{\mathcal{X}} q(z \mid x) \left[dP_1(x) - dP_2(x)\right] \\
&= \int_{\mathcal{X}} q(z \mid x) \left[dP_1(x) - dP_2(x)\right]_+ + \int_{\mathcal{X}} q(z \mid x) \left[dP_1(x) - dP_2(x)\right]_- \\
&\le \sup_{x \in \mathcal{X}} q(z \mid x) \int_{\mathcal{X}} \left[dP_1(x) - dP_2(x)\right]_+ + \inf_{x \in \mathcal{X}} q(z \mid x) \int_{\mathcal{X}} \left[dP_1(x) - dP_2(x)\right]_- \\
&= \left(\sup_{x \in \mathcal{X}} q(z \mid x) - \inf_{x \in \mathcal{X}} q(z \mid x)\right) \int_{\mathcal{X}} \left[dP_1(x) - dP_2(x)\right]_+ .
\end{aligned}$$

By definition of the total variation norm, we have $\int \left[dP_1 - dP_2\right]_+ = \|P_1 - P_2\|_{\mathrm{TV}}$, and hence

$$|m_1(z) - m_2(z)| \le \sup_{x,x'} \left|q(z \mid x) - q(z \mid x')\right| \|P_1 - P_2\|_{\mathrm{TV}}. \tag{50}$$

For any $\hat{x} \in \mathcal{X}$, we may add and subtract $q(z \mid \hat{x})$ from the quantity inside the supremum, which implies that

$$
\begin{aligned}
\sup_{x,x'} \left| q(z \mid x) - q(z \mid x') \right| &= \inf_{\hat{x}} \sup_{x,x'} \left| q(z \mid x) - q(z \mid \hat{x}) + q(z \mid \hat{x}) - q(z \mid x') \right| \\
&\leq 2 \inf_{\hat{x}} \sup_{x} \left| q(z \mid x) - q(z \mid \hat{x}) \right| \\
&= 2 \inf_{\hat{x}} q(z \mid \hat{x}) \sup_{x} \left| \frac{q(z \mid x)}{q(z \mid \hat{x})} - 1 \right|.
\end{aligned}
$$

Similarly, we have for any x, x'

$$
\left| q(z \mid x) - q(z \mid x') \right| = q(z \mid x') \left| \frac{q(z \mid x)}{q(z \mid x')} - 1 \right| \leq e^{\alpha} \inf_{\hat{x}} q(z \mid \hat{x}) \left| \frac{q(z \mid x)}{q(z \mid x')} - 1 \right|.
$$

Since for any choice of $x, \hat{x}$, we have $q(z \mid x)/q(z \mid \hat{x}) \in [e^{-\alpha}, e^{\alpha}]$, we find that (since $e^{\alpha} - 1 \geq 1 - e^{-\alpha}$)

$$
\sup_{x,x'} \left| q(z \mid x) - q(z \mid x') \right| \leq \min\{2, e^{\alpha}\} \inf_{x} q(z \mid x) \left(e^{\alpha} - 1 \right).
$$

Combining with the earlier inequality (50) yields the claim (49).

To see Lemma 4, note that for any $x > 0$, the concavity of the logarithm implies that

$$
\log(x) \leq x - 1.
$$

Setting alternatively $x = a/b$ and $x = b/a$, we obtain the inequalities

$$
\log \frac{a}{b} \leq \frac{a}{b} - 1 = \frac{a-b}{b} \quad \text{and} \quad \log \frac{b}{a} \leq \frac{b}{a} - 1 = \frac{b-a}{a}.
$$

Using the first inequality for $a \geq b$ and the second for $a < b$ completes the proof.

7.2 Proof of Corollary 1

Let us recall the definition of the induced marginal distribution (3), given by

$$
M_{\nu}(S) = \int_{\mathcal{X}} Q(S \mid x_{1:n}) dP_{\nu}^{n}(x_{1:n}) \quad \text{for } S \in \sigma(\mathcal{Z}^n).
$$

For each $i = 2, \ldots, n$, we let $M_{\nu,i}(\cdot \mid Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1}) = M_{\nu,i}(\cdot \mid z_{1:i-1})$ denote the (marginal over X_i) distribution of the variable Z_i conditioned on $Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1}$. In addition, use the shorthand notation

$$
D_{\mathrm{kl}} \left(M_{\nu,i} \| M_{\nu',i} \right) := \int_{\mathcal{Z}^{i-1}} D_{\mathrm{kl}} \left(M_{\nu,i}(\cdot \mid z_{1:i-1}) \| M_{\nu',i}(\cdot \mid z_{1:i-1}) \right) dM_{\nu}^{i-1}(z_1, \ldots, z_{i-1})
$$

to denote the integrated KL divergence of the conditional distributions on the Z_i. By the chain-rule for KL divergences [32, Chapter 5.3], we obtain

$$
D_{\mathrm{kl}} \left(M_{\nu}^{n} \| M_{\nu'}^{n} \right) = \sum_{i=1}^{n} D_{\mathrm{kl}} \left(M_{\nu,i} \| M_{\nu',i} \right).
$$

By assumption (1), the distribution $Q_i(\cdot \mid X_i, Z_{1:i-1})$ on Z_i is α-differentially private for the sample X_i. As a consequence, if we let $P_{\nu,i}(\cdot \mid Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1})$ denote the conditional distribution of X_i given the first $i-1$ values $Z_1, \ldots, Z_{i-1}$ and the packing index $V = \nu$, then from the chain rule and Theorem 1 we obtain

$$D_{\mathrm{kl}}\left(M_\nu^n \| M_{\nu'}^n\right) = \sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} D_{\mathrm{kl}}\left(M_{\nu,i}(\cdot \mid z_{1:i-1}) \| M_{\nu',i}(\cdot \mid z_{1:i-1})\right) dM_\nu^{i-1}(z_{1:i-1})$$

$$\leq \sum_{i=1}^n 4(e^\alpha - 1)^2 \int_{\mathcal{Z}^{i-1}} \left\| P_{\nu,i}(\cdot \mid z_{1:i-1}) - P_{\nu',i}(\cdot \mid z_{1:i-1}) \right\|_{\mathrm{TV}}^2 dM_\nu^{i-1}(z_1, \ldots, z_{i-1}).$$

By the construction of our sampling scheme, the random variables X_i are conditionally independent given $V = \nu$; thus the distribution $P_{\nu,i}(\cdot \mid z_{1:i-1}) = P_{\nu,i}$, where $P_{\nu,i}$ denotes the distribution of X_i conditioned on $V = \nu$. Consequently, we have

$$\left\| P_{\nu,i}(\cdot \mid z_{1:i-1}) - P_{\nu',i}(\cdot \mid z_{1:i-1}) \right\|_{\mathrm{TV}} = \left\| P_{\nu,i} - P_{\nu',i} \right\|_{\mathrm{TV}},$$

which gives the claimed result.

7.3 Proof of Proposition 1

The minimax rate characterized by equation (20) involves both a lower and an upper bound, and we divide our proof accordingly. We provide the proof for $\alpha \in (0,1]$, but note that a similar result (modulo different constants) holds for any finite value of α.

Lower bound: We use Le Cam's method to prove the lower bound in equation (20). Fix a given constant $\delta \in (0,1]$, with a precise value to be specified later. For $\nu \in \mathcal{V} \in \{-1,1\}$, define the distribution P_ν with support $\{-\delta^{-1/k}, 0, \delta^{1/k}\}$ by

$$P_\nu(X = \delta^{-1/k}) = \frac{\delta(1+\nu)}{2}, \quad P_\nu(X = 0) = 1 - \delta, \quad \text{and} \quad P_\nu(X = -\delta^{-1/k}) = \frac{\delta(1-\nu)}{2}.$$

By construction, we have $\mathbb{E}[|X|^k] = \delta(\delta^{-1/k})^k = 1$ and $\theta_\nu = \mathbb{E}_\nu[X] = \delta^{\frac{k-1}{k}}\nu$, whence the mean difference is given by $\theta_1 - \theta_{-1} = 2\delta^{\frac{k-1}{k}}$. Applying Le Cam's method (8) and the minimax bound (7) yields

$$\mathfrak{M}_n(\Theta, (\cdot)^2, Q) \geq \left(\delta^{\frac{k-1}{k}}\right)^2 \left(\frac{1}{2} - \frac{1}{2}\left\| M_1^n - M_{-1}^n \right\|_{\mathrm{TV}}\right),$$

where M_ν^n denotes the marginal distribution of the samples $Z_1, \ldots, Z_n$ conditioned on $\theta = \theta_\nu$. Now Pinsker's inequality implies that $\left\| M_1^n - M_{-1}^n \right\|_{\mathrm{TV}}^2 \leq \frac{1}{2}D_{\mathrm{kl}}\left(M_1^n \| M_{-1}^n\right)$, and Corollary 1 yields

$$D_{\mathrm{kl}}\left(M_1^n \| M_{-1}^n\right) \leq 4(e^\alpha - 1)^2 n \left\| P_1 - P_{-1} \right\|_{\mathrm{TV}}^2 = 4(e^\alpha - 1)^2 n \delta^2.$$

Putting together the pieces yields $\left\| M_1^n - M_{-1}^n \right\|_{\mathrm{TV}} \leq (e^\alpha - 1)\delta\sqrt{2n}$. For $\alpha \in (0,1]$, we have $e^\alpha - 1 \leq 2\alpha$, and thus our earlier application of Le Cam's method implies

$$\mathfrak{M}_n(\Theta, (\cdot)^2, \alpha) \geq \left(\delta^{\frac{k-1}{k}}\right)^2 \left(\frac{1}{2} - \alpha\delta\sqrt{2n}\right).$$

Substituting $\delta = \min\{1, 1/\sqrt{32n\alpha^2}\}$ yields the claim (20).

Upper bound: We must demonstrate an α-locally private conditional distribution Q and an estimator that achieves the upper bound in equation (20). We do so via a combination of truncation and addition of Laplacian noise. Define the truncation function $[\cdot]_T : \mathbb{R} \to [-T, T]$ by

$$[x]_T := \max\{-T, \min\{x, T\}\},$$

where the truncation level T is to be chosen. Let W_i be independent Laplace$(\alpha/(2T))$ random variables, and for each index $i = 1, \ldots, n$, define $Z_i := [X_i]_T + W_i$. By construction, the random variable Z_i is α-differentially private for X_i. For the mean estimator $\widehat{\theta} := \frac{1}{n}\sum_{i=1}^{n} Z_i$, we have

$$\mathbb{E}\left[(\widehat{\theta} - \theta)^2\right] = \mathrm{Var}(\widehat{\theta}) + \left(\mathbb{E}[\widehat{\theta}] - \theta\right)^2 = \frac{4T^2}{n\alpha^2} + \frac{1}{n}\mathrm{Var}([X_1]_T) + (\mathbb{E}[Z_1] - \theta)^2. \tag{51}$$

We claim that

$$\mathbb{E}[Z] = \mathbb{E}\left[[X]_T\right] \in \left[\mathbb{E}[X] - \frac{1}{(k-1)T^{k-1}}, \mathbb{E}[X] + \frac{1}{(k-1)T^{k-1}}\right]. \tag{52}$$

Indeed, by the assumption that $\mathbb{E}[|X|^k] \le 1$, we have by a change of variables that

$$\int_T^\infty x\, dP(x) = \int_T^\infty P(X \ge x)\, dx \le \int_T^\infty \frac{1}{x^k}\, dx = \frac{1}{(k-1)T^{k-1}}.$$

Thus

$$\mathbb{E}[[X]_T] \ge \mathbb{E}[\min\{X, T\}] = \mathbb{E}[\min\{X, T\} + [X - T]_+ - [X - T]_+]$$
$$= \mathbb{E}[X] - \int_T^\infty (x - T)\, dP(x) \ge \mathbb{E}[X] - \frac{1}{(k-1)T^{k-1}}.$$

A similar argument yields the upper bound in equation (52).

From the bound (51) and the inequalities that since $[X]_T \in [-T, T]$ and $\alpha^2 \le 1$, we have

$$\mathbb{E}\left[(\widehat{\theta} - \theta)^2\right] \le \frac{5T^2}{n\alpha^2} + \frac{1}{(k-1)^2 T^{2k-2}} \quad \text{valid for any } T > 0.$$

Choosing $T = (5(k-1))^{-\frac{1}{2k}}(n\alpha^2)^{1/(2k)}$ yields

$$\mathbb{E}\left[(\widehat{\theta} - \theta)^2\right] \le \frac{5(5(k-1))^{\frac{-1}{k}}(n\alpha^2)^{\frac{1}{k}}}{n\alpha^2} + \frac{1}{(k-1)^2(5(k-1))^{-1+1/k}(n\alpha^2)^{1-1/k}}$$
$$= 5^{1-1/k}\left(1 + \frac{1}{k-1}\right)\frac{1}{(k-1)^{\frac{1}{k}}(n\alpha^2)^{1-\frac{1}{k}}}.$$

Since $(1+(k-1)^{-1})(k-1)^{-\frac{1}{k}} < (k-1)^{-1}+(k-1)^{-2}$ for $k \in (1,2)$ and is bounded by $1+(k-1)^{-1} \le 2$ for $k \in [2, \infty]$, the upper bound (20) follows.

7.4   Proof of Proposition 2

We now turn to the proof of minimax rates for fixed design linear regression.

Lower bound: We use a slight generalization of the α-private form (19) of the local Fano inequality from Corollary 3. For concreteness, we assume throughout that $\alpha \in [0, \frac{23}{35}]$, but analogous arguments hold for any bounded α with changes only in the constant pre-factors. Consider an instance of the linear regression model (21) in which the noise variables $\{\varepsilon_i\}_{i=1}^n$ are drawn i.i.d. from the uniform distribution on $[-\sigma, +\sigma]$. Our first step is to construct a suitable packing of the unit sphere $\mathbb{S}^{d-1} = \{u \in \mathbb{R}^d : \|u\|_2 = 1\}$ in ℓ_2-norm:

Lemma 5. *There exists a 1-packing $\mathcal{V} = \{\nu^1, \ldots, \nu^N\}$ of the unit sphere $\mathbb{S}^{d-1}$ with $N \geq \exp(d/8)$.*

See Appendix D.1 for the proof of this claim.

For a fixed $\delta \in (0, 1]$ to be chosen shortly, define the family of vectors $\{\theta_\nu, \nu \in \mathcal{V}\}$ with $\theta_\nu := \delta\nu$. Since $\|\nu\|_2 \leq 1$, we have $\|\theta_\nu - \theta_{\nu'}\|_2 \leq 2\delta$. Let $P_{\nu,i}$ denote the distribution of Y_i conditioned on $\theta^* = \theta_\nu$. By the form of the linear regression model (21) and our assumption on the noise variable ε_i, $P_{\nu,i}$ is uniform on the interval $[\langle \theta_\nu, x_i \rangle - \sigma, \langle \theta_\nu, x_i \rangle + \sigma]$. Consequently, for $\nu \neq \nu' \in \mathcal{V}$, we have

$$\|P_{\nu,i} - P_{\nu',i}\|_{\mathrm{TV}} = \frac{1}{2} \int |p_{\nu,i}(y) - p_{\nu',i}(y)| dy$$

$$\leq \frac{1}{2} \left[\frac{1}{2\sigma} |\langle \theta_\nu, x_i \rangle - \langle \theta_{\nu'}, x_i \rangle| + \frac{1}{2\sigma} |\langle \theta_\nu, x_i \rangle - \langle \theta_{\nu'}, x_i \rangle| \right] = \frac{1}{2\sigma} |\langle \theta_\nu - \theta_{\nu'}, x_i \rangle|.$$

Letting V denote a random sample from the uniform distribution on $\mathcal{V}$, Corollary 1 implies that the mutual information is upper bounded as

$$I(Z_1, \ldots, Z_n; V) \leq 4(e^\alpha - 1)^2 \sum_{i=1}^n \frac{1}{|\mathcal{V}|^2} \sum_{\nu,\nu' \in \mathcal{V}} \|P_{\nu,i} - P_{\nu',i}\|_{\mathrm{TV}}^2$$

$$\leq \frac{(e^\alpha - 1)^2}{\sigma^2} \sum_{i=1}^n \frac{1}{|\mathcal{V}|^2} \sum_{\nu,\nu' \in \mathcal{V}} (\langle \theta_\nu - \theta_{\nu'}, x_i \rangle)^2$$

$$= \frac{(e^\alpha - 1)^2}{\sigma^2} \frac{1}{|\mathcal{V}|^2} \sum_{\nu,\nu' \in \mathcal{V}} (\theta_\nu - \theta_{\nu'})^\top X^\top X (\theta_\nu - \theta_{\nu'}).$$

Since $\theta_\nu = \delta\nu$, we have by definition of the maximum singular value that

$$(\theta_\nu - \theta_{\nu'})^\top X^\top X (\theta_\nu - \theta_{\nu'}) \leq \delta^2 \|\nu - \nu'\|_2^2 \rho_{\max}(X^\top X) \leq 4\delta^2 \rho_{\max}^2(X) = 4n\delta^2 \rho_{\max}^2(X/\sqrt{n}).$$

Putting together the pieces, we find that

$$I(Z_1, \ldots, Z_n; V) \leq \frac{4n\delta^2(e^\alpha - 1)^2}{\sigma^2} \rho_{\max}^2(X/\sqrt{n}) \leq \frac{8n\alpha^2\delta^2}{\sigma^2} \rho_{\max}^2(X/\sqrt{n}),$$

where the second inequality is valid for $\alpha \in [0, \frac{23}{35}]$. Consequently, Fano's inequality combined with the packing set $\mathcal{V}$ from Lemma 5 implies that

$$\mathfrak{M}_n \left(\Theta, \|\cdot\|_2^2, \alpha \right) \geq \frac{\delta^2}{4} \left(1 - \frac{8n\delta^2\alpha^2\rho_{\max}^2(X/\sqrt{n})/\sigma^2 + \log 2}{d/8} \right).$$

We split the remainder of the analysis into cases.

Case 1: First suppose that $d \geq 16$. Then setting $\delta^2 = \min\{1, \frac{d\sigma^2}{128n\rho_{\max}^2(X/\sqrt{n})}\}$ implies that

$$\frac{8n\delta^2\alpha^2\rho_{\max}^2(X/\sqrt{n})/\sigma^2 + \log 2}{d/8} \leq 8\left[\frac{\log 2}{d} + \frac{64}{128}\right] < \frac{7}{8}.$$

As a consequence, we have the lower bound

$$\mathfrak{M}_n\left(\Theta, \|\cdot\|_2^2, \alpha\right) \geq \frac{1}{4}\min\left\{1, \frac{d\sigma^2}{128n\rho_{\max}^2(X/\sqrt{n})}\right\} \cdot \frac{1}{8},$$

which yields the claim for $d \geq 16$.

Case 2: Otherwise, we may assume that $d < 16$. In this case, e a lower bound for the case $d = 1$ is sufficient, since apart from constant factors, the same bound holds for all $d < 16$. We use the Le Cam method based on a two point comparison. Indeed, let $\theta_1 = \delta$ and $\theta_2 = -\delta$ so that the total variation distance is at upper bounded $\|P_{1,i} - P_{2,i}\|_{\mathrm{TV}} \leq \frac{\delta}{\sigma}|x_i|$. By Corollary 2, we have

$$\mathfrak{M}_n\left(\Theta, (\cdot)^2, \alpha\right) \geq \delta^2\left(\frac{1}{2} - \delta\frac{(e^\alpha - 1)}{\sigma}\left(\sum_{i=1}^n x_i^2\right)^{\frac{1}{2}}\right).$$

Letting $x = (x_1, \ldots, x_n)$ and setting $\delta^2 = \min\{1, \sigma^2/(16(e^\alpha - 1)^2\|x\|_2^2)\}$ gives the desired result.

Upper bound: We now turn to the upper bound, for which we need to specify a private conditional Q and an estimator $\widehat{\theta}$ that achieves the stated upper bound on the mean-squared error. Let W_i be independent Laplace$(\alpha/(2\sigma))$ random variables. Then the additively perturbed random variable $Z_i = Y_i + W_i$ is α-differentially private for Y_i, since by assumption the response $Y_i \in [\langle\theta, x_i\rangle - \sigma, \langle\theta, x_i\rangle + \sigma]$. We now claim that the standard least-squares estimator of θ^* achieves the stated upper bound. Indeed, the least-squares estimate is given by

$$\widehat{\theta} = (X^\top X)^{-1}X^\top Y = (X^\top X)^{-1}X^\top(X\theta^* + \varepsilon + W).$$

Moreover, from the independence of W and ε, we have

$$\mathbb{E}\left[\|\widehat{\theta} - \theta^*\|_2^2\right] = \mathbb{E}\left[\|(X^\top X)^{-1}X^\top(\varepsilon + W)\|_2^2\right] = \mathbb{E}\left[\|(X^\top X)^{-1}X^\top\varepsilon\|_2^2\right] + \mathbb{E}\left[\|(X^\top X)^{-1}X^\top W)\|_2^2\right].$$

Since $\varepsilon \in [-\sigma, \sigma]^n$, we know that $\mathbb{E}[\varepsilon\varepsilon^\top] \preceq \sigma^2 I_{n\times n}$, and for the given choice of W, we have $\mathbb{E}[WW^\top] = (4\sigma^2/\alpha^2)I_{n\times n}$. Since $\alpha \leq 1$, we thus find

$$\mathbb{E}\left[\|\widehat{\theta} - \theta^*\|_2^2\right] \leq \frac{5\sigma^2}{\alpha^2}\mathrm{tr}\left(X(X^\top X)^{-2}X^\top\right) = \frac{5\sigma^2}{\alpha^2}\mathrm{tr}\left((X^\top X)^{-1}\right).$$

Noting that $\mathrm{tr}((X^\top X)^{-1}) \leq d/\rho_{\min}^2(X) = d/n\rho_{\min}^2(X/\sqrt{n})$ gives the claimed upper bound.

8   Proof of Theorem 2 and related results

In this section, we collect together the proof of Theorem 2 and related corollaries.

8.1   Proof of Theorem 2

Let $\mathcal{Z}$ denote the domain of the random variable Z. We begin by reducing the problem to the case when $\mathcal{Z} = \{1, 2, \ldots, k\}$ for an arbitrary positive integer k. Indeed, in the general setting, we let $\mathcal{K} = \{K_i\}_{i=1}^k$ be any (measurable) finite partition of $\mathcal{Z}$, where for $z \in \mathcal{Z}$ we let $[z]_{\mathcal{K}} = K_i$ for the K_i such that $z \in K_i$. The KL divergence $D_{\mathrm{kl}}\left(M_\nu \| \overline{M}\right)$ can be defined as the supremum of the (discrete) KL divergences between the random variables $[Z]_{\mathcal{K}}$ sampled according to M_ν and $\overline{M}$ over all partitions $\mathcal{K}$ of $\mathcal{Z}$; for instance, see Gray [32, Chapter 5]. Consequently, we can prove the claim for $\mathcal{Z} = \{1, 2, \ldots, k\}$, and then take the supremum over k to recover the general case. Accordingly, we can work with the probability mass functions $m(z \mid \nu) = M_\nu(Z = z)$ and $\overline{m}(z) = \overline{M}(Z = z)$, and we may write

$$D_{\mathrm{kl}}\left(M_\nu \| \overline{M}\right) + D_{\mathrm{kl}}\left(\overline{M} \| M_\nu\right) = \sum_{z=1}^{k} (m(z \mid \nu) - \overline{m}(z)) \log \frac{m(z \mid \nu)}{\overline{m}(z)}. \tag{53}$$

Throughout, we will also use (without loss of generality) the probability mass functions $q(z \mid x) = Q(Z = z \mid X = x)$, where we note that $m(z \mid \nu) = \int q(z \mid x) dP_\nu(x)$.

Now we use Lemma 4 from the proof of Theorem 1 to complete the proof of Theorem 2. Starting with equality (53), we have

$$\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \left[D_{\mathrm{kl}}\left(M_\nu \| \overline{M}\right) + D_{\mathrm{kl}}\left(\overline{M} \| M_\nu\right)\right] \leq \sum_{\nu \in \mathcal{V}} \frac{1}{|\mathcal{V}|} \sum_{z=1}^{k} |m(z \mid \nu) - \overline{m}(z)| \left| \log \frac{m(z \mid \nu)}{\overline{m}(z)} \right|$$

$$\leq \sum_{\nu \in \mathcal{V}} \frac{1}{|\mathcal{V}|} \sum_{z=1}^{k} |m(z \mid \nu) - \overline{m}(z)| \frac{|m(z \mid \nu) - \overline{m}(z)|}{\min\{\overline{m}(z), m(z \mid \nu)\}}.$$

Now, we define the measure m^0 on $\mathcal{Z} = \{1, \ldots, k\}$ by $m^0(z) := \inf_{x \in \mathcal{X}} q(z \mid x)$. It is clear that $\min\{\overline{m}(z), m(z \mid \nu)\} \geq m^0(z)$, whence we find

$$\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \left[D_{\mathrm{kl}}\left(M_\nu \| \overline{M}\right) + D_{\mathrm{kl}}\left(\overline{M} \| M_\nu\right)\right] \leq \sum_{\nu \in \mathcal{V}} \frac{1}{|\mathcal{V}|} \sum_{z=1}^{k} \frac{(m(z \mid \nu) - \overline{m}(z))^2}{m^0(z)}.$$

It remains to bound the final sum. For any constant $c \in \mathbb{R}$, we have

$$m(z \mid \nu) - \overline{m}(z) = \int_{\mathcal{X}} (q(z \mid x) - c) \left(dP_\nu(x) - d\overline{P}(x)\right).$$

We define a set of functions $f : \mathcal{Z} \times \mathcal{X} \to \mathbb{R}$ (depending implicitly on q) by

$$\mathcal{F}_\alpha := \left\{ f \mid f(z, x) \in [1, e^\alpha] m^0(z) \text{ for all } z \in \mathcal{Z} \text{ and } x \in \mathcal{X} \right\}.$$

By the definition of differential privacy, when viewed as a joint mapping from $\mathcal{Z} \times \mathcal{X} \to \mathbb{R}$, the conditional p.m.f. q satisfies $\{(z, x) \mapsto q(z \mid x)\} \in \mathcal{F}_\alpha$. Since constant (with respect to x) shifts do not change the above integral, we can modify the range of functions in $\mathcal{F}_\alpha$ by subtracting $m^0(z)$ from each, yielding the set

$$\mathcal{F}'_\alpha := \left\{ f \mid f(z, x) \in [0, e^\alpha - 1] m^0(z) \text{ for all } z \in \mathcal{Z} \text{ and } x \in \mathcal{X} \right\}.$$

As a consequence, we find that

$$\sum_{\nu \in \mathcal{V}} (m(z \mid \nu) - \overline{m}(z))^2 \le \sup_{f \in \mathcal{F}_\alpha} \left\{ \sum_{\nu \in \mathcal{V}} \left(\int_{\mathcal{X}} f(z, x) \left(dP_\nu(x) - d\overline{P}(x) \right) \right)^2 \right\}$$

$$= \sup_{f \in \mathcal{F}'_\alpha} \left\{ \sum_{\nu \in \mathcal{V}} \left(\int_{\mathcal{X}} \left(f(z, x) - m^0(z) \right) \left(dP_\nu(x) - d\overline{P}(x) \right) \right)^2 \right\}.$$

By inspection, when we divide by $m^0(z)$ and recall the definition of the set $\mathbb{B}_\infty \subset L^\infty(\mathcal{X})$ in the statement of Theorem 2, we obtain

$$\sum_{\nu \in \mathcal{V}} (m(z \mid \nu) - \overline{m}(z))^2 \le \left(m^0(z) \right)^2 (e^\alpha - 1)^2 \sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} \left(\int_{\mathcal{X}} \gamma(x) \left(dP_\nu(x) - d\overline{P}(x) \right) \right)^2.$$

Putting together our bounds, we have

$$\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \left[D_{\mathrm{kl}} \left(M_\nu \| \overline{M} \right) + D_{\mathrm{kl}} \left(\overline{M} \| M_\nu \right) \right]$$

$$\le (e^\alpha - 1)^2 \sum_{z=1}^{k} \frac{1}{|\mathcal{V}|} \frac{\left(m^0(z) \right)^2}{m^0(z)} \sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} \left(\int_{\mathcal{X}} \gamma(x) \left(dP_\nu(x) - d\overline{P}(x) \right) \right)^2$$

$$\le (e^\alpha - 1)^2 \frac{1}{|\mathcal{V}|} \sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} \left(\int_{\mathcal{X}} \gamma(x) \left(dP_\nu(x) - d\overline{P}(x) \right) \right)^2,$$

since $\sum_z m^0(z) \le 1$, which is the statement of the theorem.

8.2   Proof of Corollary 4

In the non-interactive setting (2), the marginal distribution M_ν^n is a product measure and Z_i is conditionally independent of $Z_{1:i-1}$ given V. Thus by the chain rule for mutual information [32, Chapter 5] and the fact (as in the proof of Theorem 2) that we may assume w.l.o.g. that Z has finite range

$$I(Z_1, \ldots, Z_n; V) = \sum_{i=1}^{n} I(Z_i; V \mid Z_{1:i-1}) = \sum_{i=1}^{n} \left[H(Z_i \mid Z_{1:i-1}) - H(Z_i \mid V, Z_{1:i-1}) \right].$$

Since conditioning reduces entropy and $Z_{1:i-1}$ is conditionally independent of Z_i given V, we have $H(Z_i \mid Z_{1:i-1}) \le H(Z_i)$ and $H(Z_i \mid V, Z_{1:i-1}) = H(Z_i \mid V)$. In particular, we have

$$I(Z_1, \ldots, Z_n; V) \le \sum_{i=1}^{n} I(Z_i; V) = \sum_{i=1}^{n} \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} D_{\mathrm{kl}} \left(M_{\nu,i} \| \overline{M}_i \right).$$

Applying Theorem 2 completes the proof.

9   Proof of Theorem 3

The proof of this theorem combines the techniques we used in the proofs of Theorems 1 and 2; the first handles interactivity, while the techniques to derive the variational bounds are reminiscent of those used in Theorem 2. Our first step is to note a consequence of the independence structure in Fig. 1 essential to our tensorization steps. More precisely, we claim that for any set $S \in \sigma(\mathcal{Z})$,

$$M_{\pm j}(Z_i \in S \mid z_{1:i-1}) = \int Q(Z_i \in S \mid Z_{1:i-1} = z_{1:i-1}, X_i = x) dP_{\pm j,i}(x). \tag{54}$$

We postpone the proof of this intermediate claim to the end of this section.

Now consider the summed KL-divergences. Let $M_{\pm j,i}(\cdot \mid z_{1:i-1})$ denote the conditional distribution of Z_i under $P_{\pm j}$, conditional on $Z_{1:i-1} = z_{1:i-1}$. As in the proof of Corollary 1, the chain-rule for KL-divergences [e.g. 32, Chapter 5] implies

$$D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right) = \sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} D_{\mathrm{kl}}\left(M_{+j}(\cdot \mid z_{1:i-1}) \| M_{-j}(\cdot \mid z_{1:i-1})\right) dM_{+j}^{i-1}(z_{1:i-1}).$$

For notational convenience in the remainder of the proof, let us define the symmetrized KL divergence between measures M and M' as $D_{\mathrm{kl}}^{\mathrm{sy}}(M \| M') = D_{\mathrm{kl}}(M \| M') + D_{\mathrm{kl}}(M' \| M)$.

Defining $\overline{P} := 2^{-d} \sum_{\nu \in \mathcal{V}} P_\nu^n$, we have $2\overline{P} = P_{+j} + P_{-j}$ for each j simultaneously, We also introduce $\overline{M}(S) = \int Q(S \mid x_{1:n}) d\overline{M}(x_{1:n})$, and let $\mathbb{E}_{\pm j}$ denote the expectation taken under the marginals $M_{\pm j}$. We then have

$$D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\mathrm{kl}}\left(M_{-j}^n \| M_{+j}^n\right)$$

$$= \sum_{i=1}^n \left(\mathbb{E}_{+j}[D_{\mathrm{kl}}\left(M_{+j,i}(\cdot \mid Z_{1:i-1}) \| M_{-j,i}(\cdot \mid Z_{1:i-1})\right)] + \mathbb{E}_{-j}[D_{\mathrm{kl}}\left(M_{-j,i}(\cdot \mid Z_{1:i-1}) \| M_{+j,i}(\cdot \mid Z_{1:i-1})\right)]\right)$$

$$\leq \sum_{i=1}^n \left(\mathbb{E}_{+j}[D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j,i}(\cdot \mid Z_{1:i-1}) \| M_{-j,i}(\cdot \mid Z_{1:i-1})\right)] + \mathbb{E}_{-j}[D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j,i}(\cdot \mid Z_{1:i-1}) \| M_{-j,i}(\cdot \mid Z_{1:i-1})\right)]\right)$$

$$= 2 \sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j,i}(\cdot \mid z_{1:i-1}) \| M_{-j,i}(\cdot \mid z_{1:i-1})\right) d\overline{M}^{i-1}(z_{1:i-1}),$$

where we have used the definition of $\overline{M}$ and that $2\overline{P} = P_{+j} + P_{-j}$ for all j. Summing over $j \in [d]$ yields

$$\sum_{j=1}^d D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j}^n \| M_{-j}^n\right) \leq 2 \sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} \underbrace{\sum_{j=1}^d D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j,i}(\cdot \mid z_{1:i-1}) \| M_{-j,i}(\cdot \mid z_{1:i-1})\right)}_{=:\mathcal{T}_{j,i}} d\overline{M}^{i-1}(z_{1:i-1}). \tag{55}$$

We bound the underlined expression in inequality (55), whose elements we denote by $\mathcal{T}_{j,i}$.

Without loss of generality (as in the proof of Theorem 2), we may assume $\mathcal{Z}$ is finite, and that $\mathcal{Z} = \{1, 2, \ldots, k\}$ for some positive integer k. Using the probability mass functions $m_{\pm j,i}$ and

omitting the index i when it is clear from context, Lemma 4 implies

$$\mathcal{T}_{j,i} = \sum_{z=1}^{k} \left(m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1})\right) \log \frac{m_{+j}(z \mid z_{1:i-1})}{m_{-j}(z \mid z_{1:i-1})}$$

$$\leq \sum_{z=1}^{k} \left(m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1})\right)^2 \frac{1}{\min\{m_{+j}(z \mid z_{1:i-1}), m_{-j}(z \mid z_{1:i-1})\}}.$$

For each fixed $z_{1:i-1}$, define the infimal measure $m^0(z \mid z_{1:i-1}) := \inf_{x \in \mathcal{X}} q(z \mid X_i = x, z_{1:i-1})$. By construction, we have $\min\{m_{+j}(z \mid z_{1:i-1}), m_{-j}(z \mid z_{1:i-1})\} \geq m^0(z \mid z_{1:i-1})$, and hence

$$\mathcal{T}_{j,i} \leq \sum_{z=1}^{k} \left(m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1})\right)^2 \frac{1}{m^0(z \mid z_{1:i-1})}.$$

Recalling equality (54), we have

$$m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1}) = \int_{\mathcal{X}} q(z \mid x, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x))$$

$$= m^0(z \mid z_{1:i-1}) \int_{\mathcal{X}} \left(\frac{q(z \mid x, z_{1:i-1})}{m^0(z \mid z_{1:i-1})} - 1\right)(dP_{+j,i}(x) - dP_{-j,i}(x)).$$

From this point, the proof is similar to that of Theorem 2. Define the collection of functions

$$\mathcal{F}_\alpha := \{f : \mathcal{X} \times \mathcal{Z}^i \to [0, e^\alpha - 1]\}.$$

Using the definition of differential privacy, we have $\frac{q(z \mid x, z_{1:i-1})}{m^0(z \mid z_{1:i-1})} \in [1, e^\alpha]$, so there exists $f \in \mathcal{F}_\alpha$ such that

$$\sum_{j=1}^{d} \mathcal{T}_{j,i} \leq \sum_{j=1}^{d} \sum_{z=1}^{k} \frac{\left(m^0(z \mid z_{1:i-1})\right)^2}{m^0(z \mid z_{1:i-1})} \left(\int_{\mathcal{X}} f(x, z, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x))\right)^2$$

$$= \sum_{z=1}^{k} m^0(z \mid z_{1:i-1}) \sum_{j=1}^{d} \left(\int_{\mathcal{X}} f(x, z, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x))\right)^2.$$

Taking a supremum over $\mathcal{F}_\alpha$, we find the further upper bound

$$\sum_{j=1}^{d} \mathcal{T}_{j,i} \leq \sum_{z=1}^{k} m^0(z \mid z_{1:i-1}) \sup_{f \in \mathcal{F}_\alpha} \sum_{j=1}^{d} \left(\int_{\mathcal{X}} f(x, z, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x))\right)^2.$$

The inner supremum may be taken independently of z and $z_{1:i-1}$, so we rescale by $(e^\alpha - 1)$ to obtain our penultimate inequality

$$\sum_{j=1}^{d} D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j,i}(\cdot \mid z_{1:i-1}) \| M_{-j,i}(\cdot \mid z_{1:i-1})\right)$$

$$\leq (e^\alpha - 1)^2 \sum_{z=1}^{k} m^0(z \mid z_{1:i-1}) \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{j=1}^{d} \left(\int_{\mathcal{X}} \gamma(x)(dP_{+j,i}(x) - dP_{-j,i}(x))\right)^2.$$

Noting that m^0 sums to a quantity ≤ 1 and substituting the preceding expression in inequality (55) completes the proof.

Finally, we return to prove our intermediate marginalization claim (54). We have that

$$
\begin{aligned}
M_{\pm j}(Z_i \in S \mid z_{1:i-1}) &= \int Q(Z_i \in S \mid z_{1:i-1}, x_{1:n}) dP_{\pm j}(x_{1:n} \mid z_{1:i-1}) \\
&\overset{(i)}{=} \int Q(Z_i \in S \mid z_{1:i-1}, x_i) dP_{\pm j}(x_{1:n} \mid z_{1:i-1}) \\
&\overset{(ii)}{=} \int Q(Z_i \in S \mid Z_{1:i-1} = z_{1:i-1}, X_i = x) dP_{\pm j,i}(x),
\end{aligned}
$$

where equality (i) follows by the assumed conditional independence structure of Q (recall Figure 1) and equality (ii) is a consequence of the independence of X_i and $Z_{1:i-1}$ under $P_{\pm j}$. That is, we have $P_{+j}(X_i \in S \mid Z_{1:i-1} = z_{1:i-1}) = P_{+j,i}(S)$ by the definition of P_ν^n as a product and that $P_{\pm j}$ are a mixture of the products P_ν^n.

10   Conclusions

We have linked minimax analysis from statistical decision theory with differential privacy, bringing some of their respective foundational principles into close contact. Our main technique, in the form of the divergence inequalities in Theorems 1 and 2, and their Corollaries 1–4, shows that applying differentially private sampling schemes essentially acts as a contraction on distributions. These contractive inequalities allow us to give sharp minimax rates for estimation in locally private settings, and we think such results may be more generally applicable. With our examples in Sections 4.2, 5.2, and 5.3, we have developed a framework that shows that roughly, if one can construct a family of distributions $\{P_\nu\}$ on the sample space $\mathcal{X}$ that is not well "correlated" with any member of $f \in L^\infty(\mathcal{X})$ for which $f(x) \in \{-1, 1\}$, then providing privacy is costly: the contraction Theorems 2 and 3 provide is strong.

By providing sharp convergence rates for many standard statistical estimation procedures under local differential privacy, we have developed and explored some tools that may be used to better understand privacy-preserving statistical inference. We have identified a fundamental continuum along which privacy may be traded for utility in the form of accurate statistical estimates, providing a way to adjust statistical procedures to meet the privacy or utility needs of the statistician and the population being sampled.

There are a number of open questions raised by our work. It is natural to wonder whether it is possible to obtain tensorized inequalities of the form of Corollary 4 even for interactive mechanisms. Another important question is whether the results we have provided can be extended to settings in which standard (non-local) differential privacy holds. Such extensions could yield insights into optimal mechanisms for differentially private procedures.

A   Proofs of multi-dimensional mean-estimation results

At a high level, our proofs of these results consist of three steps, the first of which is relatively standard, while the second two exploit specific aspects of the local privacy setting. We outline them here:

(1) The first step is a standard reduction, based on inequalities (7)–(9) in Section 2, from an estimation problem to a multi-way testing problem that involves discriminating between indices ν contained within some subset $\mathcal{V}$ of $\mathbb{R}^d$.

(2) The second step is an appropriate construction of a maximal δ-packing, meaning a set $\mathcal{V} \subset \mathbb{R}^d$ such that each pair is δ-separated and the resulting set is as large as possible. In addition, our arguments require that, for a random variable V uniformly distributed over $\mathcal{V}$, the covariance $\mathrm{Cov}(V)$ has relatively small operator norm.

(3) The final step is to apply Theorem 2 in order to control the mutual information associated with the testing problem. Doing so requires bounding the supremum in Corollary 4 via the the operator norm of $\mathrm{Cov}(V)$.

The estimation to testing reduction of Step 1 was previously described in Section 2. Accordingly, the proofs to follow are devoted to the second and third steps in each case.

A.1   Proof of Proposition 3

We provide a proof of the lower bound, as we provided the argument for the upper bound in Section 4.2.2.

Constructing a good packing: Let k be an arbitrary integer in $\{1, 2, \ldots, d\}$. The following auxiliary result provides a building block for the packing set underlying our proof:

Lemma 6. *For each integer k, there exists a packing $\mathcal{V}_k$ of the k-dimensional hypercube $\{-1, 1\}^k$ with $\|\nu - \nu'\|_1 \geq k/2$ for each $\nu, \nu' \in \mathcal{V}_k$ with $\nu \neq \nu'$ such that $|\mathcal{V}_k| \geq \lceil \exp(k/16) \rceil$, and*

$$\frac{1}{|\mathcal{V}_k|} \sum_{\nu \in \mathcal{V}_k} \nu \nu^\top \preceq 25 I_{k \times k}.$$

See Appendix D.2 for the proof.

For a given $k \leq d$, we extend the set $\mathcal{V}_k \subseteq \mathbb{R}^k$ to a subset of $\mathbb{R}^d$ by setting $\mathcal{V} = \mathcal{V}_k \times \{0\}^{d-k}$. For a parameter $\delta \in (0, 1/2]$ to be chosen, we define a family of probability distributions $\{P_\nu\}_{\nu \in \mathcal{V}}$ constructively. In particular, the random vector $X \sim P_\nu$ (a single observation) is formed by the following procedure:

Choose index $j \in \{1, \ldots, k\}$ uniformly at random and set $X = \begin{cases} re_j & \text{w.p. } \frac{1+\delta\nu_j}{2} \\ -re_j & \text{w.p. } \frac{1-\delta\nu_j}{2}. \end{cases}$ (56)

By construction, these distributions have mean vectors

$$\theta_\nu := \mathbb{E}_{P_\nu}[X] = \frac{\delta r}{k} \nu.$$

Consequently, given the properties of the packing $\mathcal{V}$, we have $X \in \mathbb{B}_1(r)$ with probability 1, and $\|\theta_\nu - \theta_{\nu'}\|_2^2 \geq r^2 \delta^2 / k$. Thus we see that the mean vectors $\{\theta_\nu\}_{\nu \in \mathcal{V}}$ provide us with an $r\delta/\sqrt{k}$-packing of the ball.

Upper bounding the mutual information: Our next step is to bound the mutual information $I(Z_1, \ldots, Z_n; V)$ when the observations X come from the distribution (56) and V is uniform in the set $\mathcal{V}$. We have the following lemma, which applies so long as the channel Q is non-interactive and α-locally private (2). See Appendix E.1 for the proof.

Lemma 7. *Fix $k \in \{1, \ldots, d\}$. Let Z_i be α-locally differentially private for X_i, and let X be sampled according to the distribution (56) conditional on $V = \nu$. Then*

$$I(Z_1, \ldots, Z_n; V) \leq n \frac{25 e^\alpha}{16} \frac{\delta^2}{k} (e^\alpha - e^{-\alpha})^2.$$

Applying testing inequalities: We now show how a combination of the hypercube packing specified by Lemma 6 and the sampling scheme (56) give us our desired lower bound. Fix $k \leq d$ and let $\mathcal{V} = \mathcal{V}_k \times \{0\}^{d-k}$ be the packing of $\{-1,1\}^k \times \{0\}^{d-k}$ defined following Lemma 6. Combining Lemma 7 and the fact that the vectors θ_ν provide an $r\delta/\sqrt{k}$ packing of the ball of cardinality at least $\exp(k/16)$, Fano's inequality implies that for any $k \in \{1, \ldots, d\}$,

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_2^2, \alpha) \geq \frac{r^2 \delta^2}{4k} \left(1 - \frac{25 n e^\alpha \delta^2 (e^\alpha - e^{-\alpha})^2 / (16k) + \log 2}{k/16} \right)$$

Because of the 1-dimensional mean-estimation lower bounds provided in Section 3.3.1, we may assume w.l.o.g. that $k \geq 32$. Setting $\delta_{n,\alpha,k}^2 = \min\{1, k^2 / (50 n e^\alpha (e^\alpha - e^{-\alpha})^2)\}$, we obtain

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_2^2, \alpha) \geq \frac{r^2 \delta_{n,\alpha,k}^2}{4k} \left(1 - \frac{1}{2} - \frac{\log 2}{2} \right) \geq c r^2 \min \left\{ \frac{1}{k}, \frac{k}{n e^\alpha (e^\alpha - e^{-\alpha})^2} \right\}$$

for a universal (numerical) constant c. Since $e^\alpha (e^\alpha - e^{-\alpha})^2 < 16\alpha^2$ for $\alpha \in [0,1]$, we obtain the lower bound

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_2^2, \alpha) \geq c r^2 \max_{k \in [d]} \left\{ \min \left\{ \frac{1}{k}, \frac{k}{n\alpha^2} \right\} \right\}$$

for $\alpha \in [0,1]$ and a universal constant $c > 0$. Setting k in the preceding display to be the integer in $\{1, \ldots, d\}$ nearest $\sqrt{n\alpha^2}$ gives the result of the proposition.

A.2 Proof of Proposition 4

Since the upper bound was established in Section 4.2.2, we focus on the lower bound.

Constructing a good packing: In this case, the packing set is very simple: set $\mathcal{V} = \{\pm e_j\}_{j=1}^d$ so that $|\mathcal{V}| = 2d$. Fix some $\delta \in [0,1]$, and for $\nu \in \mathcal{V}$, define a distribution P_ν supported on $\mathcal{X} = \{-r, r\}^d$ via

$$P_\nu(X = x) = (1 + \delta \nu^\top x/r)/2^d.$$

In words, for $\nu = e_j$, the coordinates of X are independent uniform on $\{-r, r\}$ except for the coordinate j, for which $X_j = r$ with probability $1/2 + \delta \nu_j$ and $X_j = -r$ with probability $1/2 - \delta \nu_j$. With this scheme, we have $\theta(P_\nu) = r\delta\nu$, and since $\|\delta r\nu - \delta r\nu'\|_\infty \geq \delta r$, we have constructed a δr packing in ℓ_∞-norm.

Upper bounding the mutual information: Let V be drawn uniformly from the packing set $\mathcal{V} = \{\pm e_j\}_{j=1}^d$. With the sampling scheme in the previous paragraph, we may provide the following upper bound on the mutual information $I(Z_1, \ldots, Z_n; V)$ for any non-interactive private distribution (2):

Lemma 8. *For any non-interactive α-differentially private distribution Q, we have*

$$I(Z_1, \ldots, Z_n; V) \leq n \frac{e^\alpha}{4d} \left(e^\alpha - e^{-\alpha}\right)^2 \delta^2.$$

See Appendix E.2 for a proof.

Applying testing inequalities: Finally, we turn to application of the testing inequalities. Lemma 8, in conjunction with the standard testing reduction and Fano's inequality (9), implies that

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_\infty, \alpha) \geq \frac{r\delta}{2} \left(1 - \frac{e^\alpha \delta^2 n(e^\alpha - e^{-\alpha})^2/(4d) + \log 2}{\log(2d)}\right).$$

There is no loss of generality in assuming that $d \geq 2$, in which case the choice

$$\delta^2 = \min\left\{1, \frac{d \log(2d)}{e^\alpha (e^\alpha - e^{-\alpha})^2 n}\right\}$$

yields the proposition.

A.3 Proof of Proposition 5

For this proposition, the construction of the packing and lower bound used in the proof of Proposition 4 also apply. Under these packing and sampling procedures, note that the separation of points $\theta(P_\nu) = r\delta\nu$ in ℓ_2-norm is $r\delta$. It thus remains to provide the upper bound. In this case, we use the sampling strategy (26b), as in Proposition 4 and Section 4.2.2, noting that we may take the bound B on $\|Z\|_\infty$ to be $B = c\sqrt{d}r/\alpha$ for a constant c. Let θ^* denote the true mean, assumed to be s-sparse. Now consider estimating θ^* by the ℓ_1-regularized optimization problem

$$\widehat{\theta} := \operatorname*{argmin}_{\theta \in \mathbb{R}^d} \left\{ \frac{1}{2n} \left\|\sum_{i=1}^n (Z_i - \theta)\right\|_2^2 + \lambda \|\theta\|_1 \right\},$$

Defining the error vector $W = \theta^* - \frac{1}{n}\sum_{i=1}^n Z_i$, we claim that

$$\lambda \geq 2\|W\|_\infty \quad \text{implies that} \quad \|\widehat{\theta} - \theta\|_2 \leq 3\lambda\sqrt{s}. \tag{57}$$

This result is a consequence of standard results on sparse estimation (e.g., Negahban et al. [44, Theorem 1 and Corollary 1]).

Now we note if $W_i = \theta^* - Z_i$, then $W = \frac{1}{n} \sum_{i=1}^{n} W_i$, and by construction of the sampling mechanism (26b) we have $\|W_i\|_\infty \leq c\sqrt{d}r/\alpha$ for a constant c. By Hoeffding's inequality and a union bound, we thus have for some (different) universal constant c that

$$\mathbb{P}(\|W\|_\infty \geq t) \leq 2d \exp\left(-c\frac{n\alpha^2 t^2}{r^2 d}\right) \quad \text{for } t \geq 0.$$

By taking $t^2 = r^2 d(\log(2d) + \epsilon^2)/(cn\alpha^2)$, we find that $\|W\|_\infty^2 \leq r^2 d(\log(2d) + \epsilon^2)/(cn\alpha^2)$ with probability at least $1 - \exp(-\epsilon^2)$, which gives the claimed minimax upper bound by appropriate choice of $\lambda = c\sqrt{d \log d/n\alpha^2}$ in inequality (57).

A.4 Proof of inequality (30)

We prove the bound by an argument using the private form of Fano's inequality from Corollary 3. The proof makes use of the classical Varshamov-Gilbert bound (e.g. [53, Lemma 4]):

Lemma 9 (Varshamov-Gilbert). *There is a packing $\mathcal{V}$ of the d-dimensional hypercube $\{-1, 1\}^d$ of size $|\mathcal{V}| \geq \exp(d/8)$ such that*

$$\left\|\nu - \nu'\right\|_1 \geq d/2 \quad \text{for all distinct pairs } \nu, \nu' \in \mathcal{V}.$$

Now, let $\delta \in [0, 1]$ and the distribution P_ν be a point mass at $\delta\nu/\sqrt{d}$. Then $\theta(P_\nu) = \delta\nu/\sqrt{d}$ and $\|\theta(P_\nu) - \theta(P_{\nu'})\|_2^2 \geq \delta^2$. In addition, a calculation implies that if M_1 and M_2 are d-dimensional Laplace(κ) distributions with means θ_1 and θ_2, respectively, then

$$D_{\mathrm{kl}}\left(M_1 \| M_2\right) = \sum_{j=1}^{d} \left(\exp(-\kappa|\theta_{1,j} - \theta_{2,j}|) + \kappa|\theta_{1,j} - \theta_{2,j}| - 1\right) \leq \frac{\kappa^2}{2} \|\theta_1 - \theta_2\|_2^2.$$

As a consequence, we have that under our Laplacian sampling scheme for the Z and with V chosen uniformly from $\mathcal{V}$,

$$I(Z_1, \ldots, Z_n; V) \leq \frac{1}{|\mathcal{V}|^2} n \sum_{\nu, \nu' \in \mathcal{V}} D_{\mathrm{kl}}\left(M_\nu \| M_{\nu'}\right) \leq \frac{n\alpha^2}{2d|\mathcal{V}|^2} \sum_{\nu, \nu' \in \mathcal{V}} \left\|(\delta/\sqrt{d})(\nu - \nu')\right\|_2^2 \leq \frac{2n\alpha^2\delta^2}{d}.$$

Now, applying Fano's inequality (9) in the context of the testing inequality (7), we find that

$$\inf_{\widehat{\theta}} \sup_{\nu \in \mathcal{V}} \mathbb{E}_{P_\nu}\left[\|\widehat{\theta}(Z_1, \ldots, Z_n) - \theta(P_\nu)\|_2^2\right] \geq \frac{\delta^2}{4}\left(1 - \frac{2n\alpha^2\delta^2/d + \log 2}{d/8}\right).$$

We may assume (based on our one-dimensional results in Proposition 1) w.l.o.g. that $d \geq 10$. Taking $\delta^2 = d^2/(48n\alpha^2)$ then implies the result (30).

B   Proofs of multinomial estimation results

In this section, we prove the lower bounds in Proposition 6. Before proving the bounds, however, we outline our technique, which borrows from that in Section A, and which we also use to prove the lower bounds on density estimation. The outline is as follows:

(1) As in step (1) of Section A, our first step is a standard reduction using the sharper version of Assouad's method (Lemma 1) from estimation to a multiple binary hypothesis testing problem. Specifically, we perform a (essentially standard) reduction of the form (10).

(2) Having constructed appropriately separated binary hypothesis tests, we use apply Theorem 3 via inequality (32) to control the testing error in the binary testing problem. Applying the theorem requires bounding certain suprema related to the covariance structure of randomly selected elements of $\mathcal{V} = \{-1, 1\}^d$, as in the arguments in Section A. In this case, though, the symmetry of the binary hypothesis testing problems eliminates the need for carefully constructed packings of step A(2).

With this outline in mind, we turn to the proofs of inequalities (33) and (34). As we proved the upper bounds in Section 5.2.2, this section focuses on the argument for the lower bound. We provide the full proof for the mean-squared Euclidean error, after which we show how the result for the ℓ_1-error follows.

Our first step is to provide a lower bound of the form (10), giving a Hamming separation for the squared error. To that end, fix $\delta \in [0, 1]$, and for simplicity, let us assume that d is even. In this case, we set $\mathcal{V} = \{-1, 1\}^{d/2}$, and for $\nu \in \mathcal{V}$ let P_ν be the multinomial distribution with parameter

$$\theta_\nu := \frac{1}{d}\mathbb{1} + \delta\frac{1}{d}\begin{bmatrix} \nu \\ -\nu \end{bmatrix} \in \Delta_d.$$

For any estimator $\widehat{\theta}$, by defining $\widehat{\nu}_j = \text{sign}(\widehat{\theta}_j - 1/d)$ for $j \in [d/2]$ we have the lower bound

$$\|\widehat{\theta} - \theta_\nu\|_2^2 \geq \frac{\delta^2}{d^2}\sum_{j=1}^{d/2}\mathbf{1}\{\widehat{\nu}_j \neq \nu_j\},$$

so that by the sharper variant (32) of Assouad's Lemma, we obtain

$$\max_{\nu \in \mathcal{V}}\mathbb{E}_{P_\nu}[\|\widehat{\theta} - \theta_\nu\|_2^2] \geq \frac{\delta^2}{4d}\left[1 - \left(\frac{1}{2d}\sum_{j=1}^{d/2}D_{\text{kl}}\left(M_{+j}^n\|M_{-j}^n\right) + D_{\text{kl}}\left(M_{-j}^n\|M_{+j}^n\right)\right)^{\frac{1}{2}}\right]. \tag{58}$$

Now we apply Theorem 3, which requires bounding sums of integrals $\int \gamma(dP_{+j} - dP_{-j})$, where P_{+j} is defined in expression (31). We claim the following inequality:

$$\sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})}\sum_{j=1}^{d/2}\left(\int_{\mathcal{X}}\gamma(x)dP_{+j}(x) - dP_{-j}(x)\right)^2 \leq \frac{8\delta^2}{d}. \tag{59}$$

Indeed, by construction P_{+j} is the multinomial with parameter $(1/d)\mathbb{1} + (\delta/d)[e_j^\top \quad -e_j^\top]^\top \in \Delta_d$ and similarly for P_{-j}, where $e_j \in \{0, 1\}^{d/2}$ denotes the jth standard basis vector. Abusing notation and identifying γ with vectors $\gamma \in [-1, 1]^d$, we have

$$\int_{\mathcal{X}}\gamma(x)dP_{+j}(x) - dP_{-j}(x) = \frac{2\delta}{d}\gamma^\top\begin{bmatrix} e_j \\ -e_j \end{bmatrix},$$

43

whence we find

$$\sum_{j=1}^{d/2} \left(\int_{\mathcal{X}} \gamma(x)dP_{+j}(x) - dP_{-j}(x) \right)^2 = \frac{4\delta^2}{d^2}\gamma^\top \sum_{j=1}^{d/2} \begin{bmatrix} e_j \\ -e_j \end{bmatrix} \begin{bmatrix} e_j \\ -e_j \end{bmatrix}^\top \gamma = \frac{4\delta^2}{d^2}\gamma^\top \begin{bmatrix} I & -I \\ -I & I \end{bmatrix} \gamma \leq \frac{8\delta^2}{d},$$

because the operator norm of the matrix is bounded by 2. This gives the claim (59).

Substituting the bound (59) into the bound (58) via Theorem 3, we obtain

$$\max_{\nu \in \mathcal{V}} \mathbb{E}_{P_\nu}[\|\widehat{\theta} - \theta_\nu\|_2^2] \geq \frac{\delta^2}{4d} \left[1 - \left(4n(e^\alpha - 1)^2 \delta^2/d^2 \right)^{\frac{1}{2}} \right].$$

Choosing $\delta^2 = \min\{1, d^2/(16n(e^\alpha - 1)^2)\}$ gives the lower bound

$$\mathfrak{M}_n(\Delta_d, \|\cdot\|_2^2, \alpha) \geq \min\left\{ \frac{1}{4d}, \frac{d}{64n(e^\alpha - 1)^2} \right\}.$$

To complete the proof, we note that we can prove the preceding upper bound for any even $d_0 \in \{2, \ldots, d\}$; this requires choosing $\nu \in \mathcal{V} = \{-1, 1\}^{d_0/2}$ and constructing the multinomial vectors

$$\theta_\nu = \frac{1}{d_0} \begin{bmatrix} \mathbb{1}_{d_0} \\ 0_{d-d_0} \end{bmatrix} + \frac{\delta}{d_0} \begin{bmatrix} \nu \\ -\nu \\ 0_{d-d_0} \end{bmatrix} \in \Delta_d, \quad \text{where} \quad \mathbb{1}_{d_0} = [1\ 1\ \cdots\ 1]^\top \in \mathbb{R}^{d_0}.$$

Repeating the proof *mutatis mutandis* gives the bound

$$\mathfrak{M}_n(\Delta_d, \|\cdot\|_2^2, \alpha) \geq \max_{d_0 \in \{2,4,\ldots,2\lfloor d/2 \rfloor\}} \min\left\{ \frac{1}{4d_0}, \frac{d_0}{64n(e^\alpha - 1)^2} \right\}.$$

Choosing d_0 to be the even integer closest to $\sqrt{n\alpha^2}$ in $\{1, \ldots, d\}$ and noting that $(e^\alpha - 1)^2 \leq 3\alpha^2$ for $\alpha \in [0, 1]$ gives the claimed result (33).

In the case of measuring error in the ℓ_1-norm, we provide a completely identical proof, except that we have the separation $\|\widehat{\theta} - \theta_\nu\|_1 \geq (\delta/d) \sum_{j=1}^{d/2} \mathbf{1}\{\widehat{\nu}_j \neq \nu_j\}$, and thus inequality (58) holds with the initial multiplier $\delta^2/(4d)$ replaced by $\delta/(4d)$. Parallel reasoning to the ℓ_2^2 case then gives the minimax lower bound

$$\mathfrak{M}_n(\Delta_d, \|\cdot\|_1, \alpha) \geq \frac{\delta}{4d_0} \left[1 - (4n(e^\alpha - 1)^2 \delta^2/d_0^2)^{\frac{1}{2}} \right]$$

for any even $d_0 \in \{2, \ldots, d\}$. Choosing $\delta = \min\{1, d_0^2/(16n(e^\alpha - 1)^2)\}$ gives the claim (34).

C Proofs of density estimation results

In this section, we provide the proofs of the results stated in Section 5.3 on density estimation. We defer the proofs of more technical results to later appendices. Throughout all proofs, we use c to denote a universal constant whose value may change from line to line.
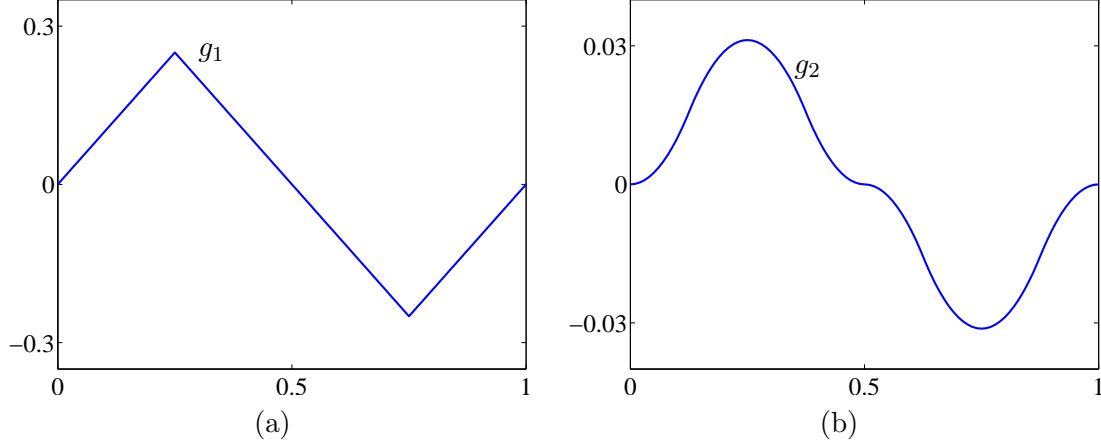
Figure 3. Panel (a): illustration of 1-Lipschitz continuous bump function g_1 used to pack $\mathcal{F}_\beta$ when $\beta = 1$. Panel (b): bump function g_2 with $|g_2''(x)| \leq 1$ used to pack $\mathcal{F}_\beta$ when $\beta = 2$.

C.1   Proof of Proposition 7

As with our proof for multinomial estimation, the argument follows the general outline described at the beginning of Section B. We remark that our proof is based on an explicit construction of densities identified with corners of the hypercube, a more classical approach than the global metric entropy approach of Yang and Barron [52] (cf. [53]). We use the local packing approach since it is better suited to the privacy constraints and information contractions that we have developed. In comparison with our proofs of previous propositions, the construction of a suitable packing of $\mathcal{F}_\beta$ is somewhat more challenging: the identification of densities with finite-dimensional vectors, which we require for our application of Theorem 3, is not immediately obvious. In all cases, we guarantee that our density functions f belong to the trigonometric Sobolev space, so we may work directly with smooth density functions f.

Constructing well-separated densities: We begin by describing a standard framework for defining local packings of density functions. Let $g_\beta : [0,1] \to \mathbb{R}$ be a function satisfying the following properties:

(a) The function g_β is β-times differentiable with

$$0 = g_\beta^{(i)}(0) = g_\beta^{(i)}(1/2) = g_\beta^{(i)}(1) \quad \text{for all } i < \beta.$$

(b) The function g_β is centered with $\int_0^1 g_\beta(x)dx = 0$, and there exist constants $c, c_{1/2} > 0$ such that

$$\int_0^{1/2} g_\beta(x)dx = -\int_{1/2}^1 g_\beta(x)dx = c_{1/2} \quad \text{and} \quad \int_0^1 \left(g_\beta^{(i)}(x)\right)^2 dx \geq c \quad \text{for all } i < \beta.$$

(c) The function g_β is non-negative on $[0, 1/2]$ and non-positive on $[1/2, 1]$, and Lebesgue measure is absolutely continuous with respect to the measures $G_j, j = 1, 2$, given by

$$G_1(A) = \int_{A \cap [0, 1/2]} g_\beta(x)dx \quad \text{and} \quad G_2(A) = -\int_{A \cap [1/2, 1]} g_\beta(x)dx. \tag{60}$$

45

(d) Lastly, for almost every $x \in [0,1]$, we have $|g_\beta^{(\beta)}(x)| \leq 1$ and $|g_\beta(x)| \leq 1$.

As illustrated in Figure 3, the functions g_β are smooth "bump" functions.

Fix a positive integer k (to be specified in the sequel). Our first step is to construct a family of "well-separated" densities for which we can reduce the density estimation problem to one of identifying corners of a hypercube, which allows application of Lemma 1. Specifically, we must exhibit a condition similar to the separation condition (10). For each $j \in \{1,\dots,k$ define the function

$$g_{\beta,j}(x) := \frac{1}{k^\beta} \, g_\beta\left(k\left(x - \frac{j-1}{k}\right)\right) \mathbf{1}\left\{x \in \left[\frac{j-1}{k}, \frac{j}{k}\right]\right\}.$$

Based on this definition, we define the family of densities

$$\left\{f_\nu := 1 + \sum_{j=1}^{k} \nu_j g_{\beta,j} \quad \text{for } \nu \in \mathcal{V}\right\} \subseteq \mathcal{F}_\beta. \tag{61}$$

It is a standard fact [53, 49] that for any $\nu \in \mathcal{V}$, the function f_ν is β-times differentiable, satisfies $|f^{(\beta)}(x)| \leq 1$ for all x. Now, based on some density $f \in \mathcal{F}_\beta$, let us define the sign vector $\mathsf{v}(f) \in \{-1,1\}^k$ to have entries

$$\mathsf{v}_j(f) := \operatorname*{argmin}_{s\in\{-1,1\}} \int_{[\frac{j-1}{k},\frac{j}{k}]} (f(x) - sg_{\beta,j}(x))^2 \, dx.$$

Then by construction of the g_β and v, we have for a numerical constant c (whose value may depend on β) that

$$\|f - f_\nu\|_2^2 \geq c \sum_{j=1}^{k} \mathbf{1}\{\mathsf{v}_j(f) \neq \nu_j\} \int_{[\frac{j-1}{k},\frac{j}{k}]} (g_{\beta,j}(x))^2 dx = \frac{c}{k^{2\beta+1}} \sum_{j=1}^{k} \mathbf{1}\{\mathsf{v}_j(f) \neq \nu_j\}.$$

By inspection, this is the Hamming separation required in inequality (10), whence the sharper version (32) of Assouad's Lemma 1 gives the result

$$\mathfrak{M}_n\left(\mathcal{F}_\beta[1], \|\cdot\|_2^2, \alpha\right) \geq \frac{c}{k^{2\beta}} \left[1 - \left(\frac{1}{4k} \sum_{j=1}^{k} (D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\mathrm{kl}}\left(M_{-j}^n \| M_{+j}^n\right))\right)^{\frac{1}{2}}\right], \tag{62}$$

where we have defined $P_{\pm j}$ to be the probability distribution associated with the averaged densities $f_{\pm j} = 2^{1-k} \sum_{\nu:\nu_j=\pm 1} f_\nu$.

Applying divergence inequalities: Now we must control the summed KL-divergences. To do so, we note that by the construction (61), symmetry implies that

$$f_{+j} = 1 + g_{\beta,j} \quad \text{and} \quad f_{-j} = 1 - g_{\beta,j} \quad \text{for each } j \in [k]. \tag{63}$$

We then obtain the following result, which bounds the averaged KL-divergences.

Lemma 10. *For any α-locally private conditional distribution Q, the summed KL-divergences are bounded as*

$$\sum_{j=1}^{k} \left(D_{\mathrm{kl}} \left(M_{+j}^n \| M_{-j}^n \right) + D_{\mathrm{kl}} \left(M_{+j}^n \| M_{-j}^n \right) \right) \leq 4c_{1/2}^2 \, n \frac{(e^\alpha - 1)^2}{k^{2\beta+1}}.$$

The proof of this lemma is fairly involved, so we defer it to Appendix E.3. We note that, for $\alpha \leq 1$, we have $(e^\alpha - 1)^2 \leq 3\alpha^2$, so we may replace the bound in Lemma 10 with the quantity $cn\alpha^2/k^{2\beta+1}$ for a constant c. We remark that standard divergence bounds using Assouad's lemma [53, 49] provide a bound of roughly $n/k^{2\beta}$; our bound is thus essentially a factor of the "dimension" k tighter.

The remainder of the proof is an application of inequality (62). In particular, by applying Lemma 10, we find that for any α-locally private channel Q, there are constants c_0, c_1 (whose values may depend on β) such that

$$\mathfrak{M}_n \left(\mathcal{F}_\beta, \|\cdot\|_2^2, Q \right) \geq \frac{c_0}{k^{2\beta}} \left[1 - \left(\frac{c_1 n \alpha^2}{k^{2\beta+2}} \right)^{\frac{1}{2}} \right].$$

Choosing $k_{n,\alpha,\beta} = \left(4c_1 n\alpha^2 \right)^{\frac{1}{2\beta+2}}$ ensures that the quantity inside the parentheses is at least $1/2$. Substituting for k in the preceding display proves the proposition.

C.2   Proof of Proposition 8

Note that the operator Π_k performs a Euclidean projection of the vector $(k/n) \sum_{i=1}^n Z_i$ onto the scaled probability simplex, thus projecting $\widehat{f}$ onto the set of probability densities. Given the non-expansivity of Euclidean projection, this operation can only decrease the error $\|\widehat{f} - f\|_2^2$. Consequently, it suffices to bound the error of the unprojected estimator; to reduce notational overhead we retain our previous notation of $\widehat{\theta}$ for the unprojected version. Using this notation, we have

$$\mathbb{E} \left[\|\widehat{f} - f\|_2^2 \right] \leq \sum_{j=1}^{k} \mathbb{E}_f \left[\int_{\frac{j-1}{k}}^{\frac{j}{k}} (f(x) - \widehat{\theta}_j)^2 dx \right].$$

By expanding this expression and noting that the independent noise variables $W_{ij} \sim \mathrm{Laplace}(\alpha/2)$ have zero mean, we obtain

$$\mathbb{E} \left[\|\widehat{f} - f\|_2^2 \right] \leq \sum_{j=1}^{k} \mathbb{E}_f \left[\int_{\frac{j-1}{k}}^{\frac{j}{k}} \left(f(x) - \frac{k}{n} \sum_{i=1}^n [\mathsf{e}_k(X_i)]_j \right)^2 dx \right] + \sum_{j=1}^{k} \int_{\frac{j-1}{k}}^{\frac{j}{k}} \mathbb{E} \left[\left(\frac{k}{n} \sum_{i=1}^n W_{ij} \right)^2 \right]$$

$$= \sum_{j=1}^{k} \int_{\frac{j-1}{k}}^{\frac{j}{k}} \mathbb{E}_f \left[\left(f(x) - \frac{k}{n} \sum_{i=1}^n [\mathsf{e}_k(X_i)]_j \right)^2 \right] dx + k \frac{1}{k} \frac{4k^2}{n\alpha^2}. \tag{64}$$

Next we bound the error term inside the expectation (64). Defining $p_j := \mathbb{P}_f(X \in \mathcal{X}_j) = \int_{\mathcal{X}_j} f(x) dx$, we have

$$k \mathbb{E}_f \left[[\mathsf{e}_k(X)]_j \right] = k p_j = k \int_{\mathcal{X}_j} f(x) dx \in \left[f(x) - \frac{1}{k}, f(x) + \frac{1}{k} \right] \quad \text{for any } x \in \mathcal{X}_j,$$

by the Lipschitz continuity of f. Thus, expanding the bias and variance of the integrated expectation above, we find that

$$\mathbb{E}_f\left[\left(f(x) - \frac{k}{n}\sum_{i=1}^{n}[\mathsf{e}_k(X_i)]_j\right)^2\right] \leq \frac{1}{k^2} + \mathrm{Var}\left(\frac{k}{n}\sum_{i=1}^{n}[\mathsf{e}_k(X_i)]_j\right)$$

$$= \frac{1}{k^2} + \frac{k^2}{n}\mathrm{Var}([\mathsf{e}_k(X)]_j) = \frac{1}{k^2} + \frac{k^2}{n}p_j(1-p_j).$$

Recalling the inequality (64), we obtain

$$\mathbb{E}_f\left[\|\widehat{f}-f\|_2^2\right] \leq \sum_{j=1}^{k}\int_{\frac{j-1}{k}}^{\frac{j}{k}}\left(\frac{1}{k^2} + \frac{k^2}{n}p_j(1-p_j)\right)dx + \frac{4k^2}{n\alpha^2} = \frac{1}{k^2} + \frac{4k^2}{n\alpha^2} + \frac{k}{n}\sum_{j=1}^{k}p_j(1-p_j).$$

Since $\sum_{j=1}^{k}p_j = 1$, we find that

$$\mathbb{E}_f\left[\|\widehat{f}-f\|_2^2\right] \leq \frac{1}{k^2} + \frac{4k^2}{n\alpha^2} + \frac{k}{n},$$

and choosing $k = (n\alpha^2)^{\frac{1}{4}}$ yields the claim.

C.3 Proof of Proposition 9

We begin by fixing $k \in \mathbb{N}$; we will optimize the choice of k shortly. Recall that, since $f \in \mathcal{F}_\beta[C]$, we have $f = \sum_{j=1}^{\infty}\theta_j\varphi_j$ for $\theta_j = \int f\varphi_j$. Thus we may define $\overline{Z}_j = \frac{1}{n}\sum_{i=1}^{n}Z_{i,j}$ for each $j \in \{1,\ldots,k\}$, and we have

$$\|\widehat{f}-f\|_2^2 = \sum_{j=1}^{k}(\theta_j - \overline{Z}_j)^2 + \sum_{j=k+1}^{\infty}\theta_j^2.$$

Since $f \in \mathcal{F}_\beta[C]$, we are guaranteed that $\sum_{j=1}^{\infty}j^{2\beta}\theta_j^2 \leq C^2$, and hence

$$\sum_{j>k}\theta_j^2 = \sum_{j>k}j^{2\beta}\frac{\theta_j^2}{j^{2\beta}} \leq \frac{1}{k^{2\beta}}\sum_{j>k}j^{2\beta}\theta_j^2 \leq \frac{1}{k^{2\beta}}C^2.$$

For the indices $j \leq k$, we note that by assumption, $\mathbb{E}[Z_{i,j}] = \int \varphi_j f = \theta_j$, and since $|Z_{i,j}| \leq B$, we have

$$\mathbb{E}\left[(\theta_j - \overline{Z}_j)^2\right] = \frac{1}{n}\mathrm{Var}(Z_{1,j}) \leq \frac{B^2}{n} = \frac{B_0^2}{c_k}\frac{k}{n}\left(\frac{e^\alpha+1}{e^\alpha-1}\right)^2,$$

where $c_k = \Omega(1)$ is the constant in expression (43). Putting together the pieces, the mean-squared L^2-error is upper bounded as

$$\mathbb{E}_f\left[\|\widehat{f}-f\|_2^2\right] \leq c\left(\frac{k^2}{n\alpha^2} + \frac{1}{k^{2\beta}}\right),$$

where c is a constant depending on B_0, c_k, and C. Choose $k = (n\alpha^2)^{1/(2\beta+2)}$ to complete the proof.

C.4 Insufficiency of Laplace noise for density estimation

Finally, we consider the insufficiency of standard Laplace noise addition for estimation in the setting of this section. Consider the vector $[\varphi_j(X_i)]_{j=1}^k \in [-B_0, B_0]^k$. To make this vector α-differentially private by adding an independent Laplace noise vector $W \in \mathbb{R}^k$, we must take $W_j \sim \text{Laplace}(\alpha/(B_0 k))$. The natural orthogonal series estimator [e.g., 51] is to take $Z_i = [\varphi_j(X_i)]_{j=1}^k + W_i$, where $W_i \in \mathbb{R}^k$ are independent Laplace noise vectors. We then use the density estimator (44), except that we use the Laplacian perturbed Z_i. However, this estimator suffers the following drawback:

Observation 1. *Let $\widehat{f} = \frac{1}{n}\sum_{i=1}^n \sum_{j=1}^k Z_{i,j}\varphi_j$, where the Z_i are the Laplace-perturbed vectors of the previous paragraph. Assume the orthonormal basis $\{\varphi_j\}$ of $L^2([0,1])$ contains the constant function. There is a constant c such that for any $k \in \mathbb{N}$, there is an $f \in \mathcal{F}_\beta[2]$ such that*

$$\mathbb{E}_f\left[\|f - \widehat{f}\|_2^2\right] \geq c(n\alpha^2)^{-\frac{2\beta}{2\beta+3}}.$$

Proof. We begin by noting that for $f = \sum_j \theta_j \varphi_j$, by definition of $\widehat{f} = \sum_j \widehat{\theta}_j \varphi_j$ we have

$$\mathbb{E}\left[\|f - \widehat{f}\|_2^2\right] = \sum_{j=1}^k \mathbb{E}\left[(\theta_j - \widehat{\theta}_j)^2\right] + \sum_{j \geq k+1} \theta_j^2 = \sum_{j=1}^k \frac{B_0^2 k^2}{n\alpha^2} + \sum_{j \geq k+1} \theta_j^2 = \frac{B_0^2 k^3}{n\alpha^2} + \sum_{j \geq k+1} \theta_j^2.$$

Without loss of generality, let us assume $\varphi_1 = 1$ is the constant function. Then $\int \varphi_j = 0$ for all $j > 1$, and by defining the true function $f = \varphi_1 + (k+1)^{-\beta}\varphi_{k+1}$, we have $f \in \mathcal{F}_\beta[2]$ and $\int f = 1$, and moreover,

$$\mathbb{E}\left[\|f - \widehat{f}\|_2^2\right] \geq \frac{B_0^2 k^3}{n\alpha^2} + \frac{1}{(k+1)^{-2\beta}} \geq C_{\beta,B_0}(n\alpha^2)^{-\frac{2\beta}{2\beta+3}},$$

where C_{β,B_0} is a constant depending on β and B_0. This final lower bound comes by minimizing over all k. (If $(k+1)^{-\beta}B_0 > 1$, we can rescale φ_{k+1} by B_0 to achieve the same result and guarantee that $f \geq 0$.) $\square$

This lower bound shows that standard estimators based on adding Laplace noise to appropriate basis expansions of the data fail: there is a degradation in rate from $n^{-\frac{2\beta}{2\beta+2}}$ to $n^{-\frac{2\beta}{2\beta+3}}$. While this is not a formal proof that no approach based on Laplace perturbation can provide optimal convergence rates in our setting, it does suggest that finding such an estimator is non-trivial.

D Packing set constructions

In this appendix, we collect proofs of the constructions of our packing sets.

D.1 Proof of Lemma 5

By the Varshamov-Gilbert bound [e.g., 53, Lemma 4], there is a packing $\mathcal{H}_d$ of the d-dimensional hypercube $\{-1, 1\}^d$ of size $|\mathcal{H}_d| \geq \exp(d/8)$ satisfying $\|u - v\|_1 \geq d/2$ for all $u, v \in \mathcal{H}_d$ with $u \neq v$. For each $u \in \mathcal{H}_d$, set $\nu_u = u/\sqrt{d}$, so that $\|\nu_u\|_2 = 1$ and $\|\nu_u - \nu_v\|_2^2 \geq d/d = 1$ for $u \neq v \in \mathcal{H}_d$. Setting $\mathcal{V} = \{\nu_u \mid u \in \mathcal{H}_d\}$ gives the desired result.

D.2 Proof of Lemma 6

We use the probabilistic method [2], showing that for random draws from the Boolean hypercube, a collection of vectors as claimed in the lemma exists with positive probability. Consider a set of N vectors $\nu^i \in \{-1, 1\}^k$ sampled uniformly at random from the Boolean hypercube, and for a fixed $t > 0$, define the two "bad" events

$$
\mathcal{B}_1 := \big\{ \exists\, i \neq j \mid \big\| \nu^i - \nu^j \big\|_1 < k/2 \big\}, \quad \text{and} \quad \mathcal{B}_2(t) := \bigg\{ \frac{1}{N} \sum_{i=1}^N \nu^i (\nu^i)^\top \npreceq (t+1) I_{k \times k} \bigg\}.
$$

We begin by analyzing $\mathcal{B}_1$. Letting $\{W_\ell\}_{\ell=1}^k$ denote a sequence of i.i.d. Bernoulli $\{0, 1\}$ variables, for any $i \neq j$, the event $\{\| \nu^i - \nu^j \|_1 < k/2\}$ is equivalent to the event $\{\sum_{\ell=1}^k W_\ell < k/4\}$. Consequently, by combining the union bound with the the Hoeffding bound, we find

$$
\mathbb{P}(\mathcal{B}_1) \leq \binom{N}{2} \mathbb{P}\big(\| \nu_i - \nu_j \|_1 < k/2 \big) \leq \binom{N}{2} \exp(-k/8). \tag{65}
$$

Turning to the event $\mathcal{B}_2(t)$, we have $\frac{1}{N} \sum_{i=1}^N \nu_i (\nu^i)^\top \npreceq (t+1) I_{k \times k}$ if and only if the maximum eigenvalue $\lambda_{\max}(\frac{1}{N} \sum_{i=1}^N \nu^i (\nu^i)^\top - I_{k \times k})$ is larger than t. Using sharp versions of the Ahlswede-Winter inequalities [1] (see Corollary 4.2 in the paper [42]), we obtain

$$
\mathbb{P}(\mathcal{B}_2(t)) \leq k \exp\left(-\frac{Nt^2}{k^2} \right). \tag{66}
$$

Finally, combining the union bound with inequalities (65) and (66), we find that

$$
\mathbb{P}(\mathcal{B}_1 \cup \mathcal{B}_2(t)) \leq \frac{N(N-1)}{2} \exp(-k/8) + d \exp\left(-\frac{Nt^2}{k^2} \right).
$$

By inspection, if we choose $t = 24$ and $N = \lceil \exp(k/16) \rceil$, the above bound is strictly less than 1, so a packing satisfying the constraints must exist.

E Information bounds

In this appendix, we collect the proofs of lemmas providing mutual information and KL-divergence bounds.

E.1 Proof of Lemma 7

Our strategy is to apply Theorem 2 to bound the mutual information. Without loss of generality, we may assume that $r = 1$ so the set $\mathcal{X} = \{\pm e_j\}_{j=1}^k$, where $e_j \in \mathbb{R}^d$. Thus, under the notation of Theorem 2, we may identify vectors $\gamma \in L^\infty(\mathcal{X})$ by vectors $\gamma \in \mathbb{R}^{2k}$. If we define $\overline{\nu} = \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \nu$ to be the mean element of the packing set, the linear functional φ_ν defined in Theorem 2 is

$$
\varphi_\nu(\gamma) = \frac{1}{2k} \bigg[\sum_{j=1}^k \gamma(e_j) \frac{1 + \nu_j \delta}{2} + \sum_{j=1}^k \gamma(-e_j) \frac{1 - \nu_j \delta}{2} \bigg] - \frac{1}{2k} \bigg[\sum_{j=1}^k \gamma(e_j) \frac{1 + \overline{\nu}_j \delta}{2} + \sum_{j=1}^k \gamma(-e_j) \frac{1 - \overline{\nu}_j \delta}{2} \bigg]
$$

$$
= \frac{1}{2k} \sum_{j=1}^k \bigg[\frac{\delta}{2} \gamma(e_j)(\nu_j - \overline{\nu}_j) - \frac{\delta}{2} \gamma(-e_j)(\nu_j - \overline{\nu}_j) \bigg] = \frac{\delta}{4k} \gamma^\top \begin{bmatrix} I_{k \times k} & 0_{k \times d-k} \\ -I_{k \times k} & 0_{k \times d-k} \end{bmatrix} (\nu - \overline{\nu}).
$$

Define the matrix

$$A := \begin{bmatrix} I_{k \times k} & 0_{k \times d-k} \\ -I_{k \times k} & 0_{k \times d-k} \end{bmatrix} \in \{-1, 0, 1\}^{2k \times d}.$$

Then we have that

$$\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \varphi_\nu(\gamma)^2 = \frac{\delta^2}{(4k)^2} \gamma^\top A \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} (\nu - \overline{\nu})(\nu - \overline{\nu})^\top A^\top \gamma$$

$$= \frac{\delta^2}{(4k)^2} \gamma^\top A \left(\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \nu\nu^\top - \overline{\nu}\overline{\nu}^\top \right) A^\top \gamma \le \frac{\delta^2}{(4k)^2} \gamma^\top A \left(\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \nu\nu^\top \right) A^\top \gamma$$

$$\le \frac{25}{16} \frac{\delta^2}{k^2} \gamma^\top A I_{2d \times 2d} A^\top \gamma = \left(\frac{5\delta}{4k} \right)^2 \gamma^\top \begin{bmatrix} I_{k \times k} & -I_{k \times k} \\ -I_{k \times k} & I_{k \times k} \end{bmatrix} \gamma. \tag{67}$$

Here the final inequality used our assumption on the sum of outer products in $\mathcal{V}$.

We complete our proof using the bound (67). The operator norm of the matrix specified in (67) is 2. As a consequence, since we have the containment

$$\mathbb{B}_\infty = \left\{ \gamma \in \mathbb{R}^{2k} : \|\gamma\|_\infty \le 1 \right\} \subset \left\{ \gamma \in \mathbb{R}^{2k} : \|\gamma\|_2^2 \le 2k \right\}$$

we have the inequality

$$\sup_{\gamma \in \mathbb{B}_\infty} \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \varphi_\nu(\gamma)^2 \le \frac{25\delta^2}{16k^2} \cdot 2 \cdot 2k = \frac{25}{4} \frac{\delta^2}{k}$$

Applying Theorem 2 completes the proof.

E.2 Proof of Lemma 8

It is no loss of generality to assume the radius $r = 1$. We use the notation of Theorem 2, recalling the linear functionals $\varphi_\nu : L^\infty(\mathcal{X}) \to \mathbb{R}$. Because the set $\mathcal{X} = \{-1, 1\}^d$, we can identify vectors $\gamma \in L^\infty(\mathcal{X})$ with vectors $\gamma \in \mathbb{R}^{2^d}$. Moreover, we have (by construction) that

$$\varphi_\nu(\gamma) = \sum_{x \in \{-1,1\}^d} \gamma(x) p_\nu(x) - \sum_{x \in \{-1,1\}^d} \gamma(x) \overline{p}(x)$$

$$= \frac{1}{2^d} \sum_{x \in \mathcal{X}} \gamma(x)(1 + \delta\nu^\top x - 1) = \frac{\delta}{2^d} \sum_{x \in \mathcal{X}} \gamma(x)\nu^\top x.$$

For each $\nu \in \mathcal{V}$, we may construct a vector $u_\nu \in \{-1, 1\}^{2^d}$, indexed by $x \in \{-1, 1\}^d$, with

$$u_\nu(x) = \nu^\top x = \begin{cases} 1 & \text{if } \nu = \pm e_j \text{ and } \operatorname{sign}(\nu_j) = \operatorname{sign}(x_j) \\ -1 & \text{if } \nu = \pm e_j \text{ and } \operatorname{sign}(\nu_j) \ne \operatorname{sign}(x_j). \end{cases}$$

For $\nu = e_j$, we see that $u_{e_1}, \ldots, u_{e_d}$ are the first d columns of the standard Hadamard transform matrix (and u_{-e_j} are their negatives). Then we have that $\sum_{x \in \mathcal{X}} \gamma(x)\nu^\top x = \gamma^\top u_\nu$, and

$$\varphi_\nu(\gamma) = \gamma^\top u_\nu u_\nu^\top \gamma.$$

Note also that $u_\nu u_\nu^\top = u_{-\nu} u_{-\nu}^\top$, and as a consequence we have

$$\sum_{\nu \in \mathcal{V}} \varphi_\nu(\gamma)^2 = \frac{\delta^2}{4^d} \gamma^\top \sum_{\nu \in \mathcal{V}} u_\nu u_\nu^\top \gamma = \frac{2\delta^2}{4^d} \gamma^\top \sum_{j=1}^d u_{e_j} u_{e_j}^\top \gamma. \tag{68}$$

But now, studying the quadratic form (68), we note that the vectors u_{e_j} are orthogonal. As a consequence, the vectors (up to scaling) u_{e_j} are the only eigenvectors corresponding to positive eigenvalues of the positive semidefinite matrix $\sum_{j=1}^d u_{e_j} u_{e_j}^\top$. Thus, since the set

$$\mathbb{B}_\infty = \left\{ \gamma \in \mathbb{R}^{2^d} : \|\gamma\|_\infty \le 1 \right\} \subset \left\{ \gamma \in \mathbb{R}^{2^d} : \|\gamma\|_2^2 \le 2^d \right\},$$

we have via an eigenvalue calculation that

$$\sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} \varphi_\nu(\gamma)^2 \le \frac{2\delta^2}{4^d} \sup_{\gamma : \|\gamma\|_2^2 \le 2^d} \gamma^\top \sum_{j=1}^d u_{e_j} u_{e_j}^\top \gamma$$

$$= \frac{2\delta^2}{4^d} \|u_{e_1}\|_2^4 = 2\delta^2$$

since $\|u_{e_j}\|_2^2 = 2^d$ for each j. Applying Theorem 2 and Corollary 4 completes the proof.

E.3 Proof of Lemma 10

This result relies on Theorem 3, along with a careful argument to understand the extreme points of $\gamma \in L^\infty([0,1])$ that we use when applying the result. First, we take the packing $\mathcal{V} = \{-1, 1\}^\beta$ and densities f_ν for $\nu \in \mathcal{V}$ as in the construction (61). Overall, our first step is to show for the purposes of applying Theorem 3, it is no loss of generality to identify $\gamma \in L^\infty([0,1])$ with vectors $\gamma \in \mathbb{R}^{2k}$, where γ is constant on intervals of the form $[i/2k, (i+1)/2k]$. With this identification complete, we can then provide a bound on the correlation of any $\gamma \in \mathbb{B}_\infty$ with the densities $f_{\pm j}$ defined in (63), which completes the proof.

With this outline in mind, let the sets D_i, $i \in \{1, 2, \ldots, 2k\}$, be defined as $D_i = [(i-1)/2k, i/2k)$ except that $D_{2k} = [(2k-1)/2k, 1]$, so the collection $\{D_i\}_{i=1}^{2k}$ forms a partition of the unit interval $[0,1]$. By construction of the densities f_ν, the sign of $f_\nu - 1$ remains constant on each D_i. Let us define (for shorthand) the linear functionals $\varphi_j : L^\infty([0,1]) \to \mathbb{R}$ for each $j \in \{1, \ldots, k\}$ via

$$\varphi_j(\gamma) := \int \gamma(dP_{+j} - dP_{-j}) = \sum_{i=1}^{2k} \int_{D_i} \gamma(x)(f_{+j}(x) - f_{-j}(x))dx = 2 \int_{D_{2j-1} \cup D_{2j}} \gamma(x) g_{\beta, j}(x)dx,$$

where we recall the definitions (63) of the mixture densities $f_{\pm j} = 1 \pm g_{\beta, j}$. Since the set $\mathbb{B}_\infty$ from Theorem 3 is compact, convex, and Hausdorff, the Krein-Milman theorem [45, Proposition 1.2] guarantees that it is equal to the convex hull of its extreme points; moreover, since the functionals $\gamma \mapsto \varphi_j^2(\gamma)$ are convex, the supremum in Theorem 3 must be attained at the extreme points of $\mathbb{B}_\infty([0,1])$. As a consequence, when applying the divergence bound

$$\sum_{j=1}^k \left(D_{\mathrm{kl}} \left(M_{+j}^n \| M_{-j}^n \right) + D_{\mathrm{kl}} \left(M_{-j}^n \| M_{+j}^n \right) \right) \le 2n(e^\alpha - 1)^2 \sup_{\gamma \in \mathbb{B}_\infty} \sum_{j=1}^k \varphi_j^2(\gamma), \tag{69}$$

we can restrict our attention to $\gamma \in \mathbb{B}_\infty$ for which $\gamma(x) \in \{-1, 1\}$.

Now we argue that it is no loss of generality to assume that γ, when restricted to D_i, is a constant (apart from a measure zero set). Fix $i \in [2k]$, and assume for the sake of contradiction that there exist sets $B_i, C_i \subset D_i$ such that $\gamma(B_i) = \{1\}$ and $\gamma(C_i) = \{-1\}$, while $\mu(B_i) > 0$ and $\mu(C_i) > 0$ where μ denotes Lebesgue measure.[1] We will construct vectors γ_1 and $\gamma_2 \in \mathbb{B}_\infty$ and a value $\lambda \in (0, 1)$ such that

$$\int_{D_i} \gamma(x) g_{\beta,j}(x) dx = \lambda \int_{D_i} \gamma_1(x) g_{\beta,j}(x) dx + (1 - \lambda) \int_{D_i} \gamma_2(x) g_{\beta,j}(x) dx$$

simultaneously for all $j \in [k]$, while on $D_i^c = [0, 1] \setminus D_i$, we will have the equivalence

$$\gamma_1|_{D_i^c} \equiv \gamma_2|_{D_i^c} \equiv \gamma|_{D_i^c}.$$

Indeed, set $\gamma_1(D_i) = \{1\}$ and $\gamma_2(D_i) = \{-1\}$, otherwise setting $\gamma_1(x) = \gamma_2(x) = \gamma(x)$ for $x \notin D_i$. For the unique index $j \in [k]$ such that $[(j - 1)/k, j/k] \supset D_i$, we define

$$\lambda := \frac{\int_{B_i} g_{\beta,j}(x) dx}{\int_{D_i} g_{\beta,j}(x) dx} \quad \text{so} \quad 1 - \lambda = \frac{\int_{C_i} g_{\beta,j}(x) dx}{\int_{D_i} g_{\beta,j}(x) dx}.$$

By the construction of the function g_β, the functions $g_{\beta,j}$ do not change signs on D_i, and the absolute continuity conditions on g_β specified in equation (60) guarantee $1 > \lambda > 0$, since $\mu(B_i) > 0$ and $\mu(C_i) > 0$. We thus find that for any $j \in [k]$,

$$\int_{D_i} \gamma(x) g_{\beta,j}(x) dx = \int_{B_i} \gamma_1(x) g_{\beta,j}(x) dx + \int_{C_i} \gamma_2(x) g_{\beta,j}(x) dx$$

$$= \int_{B_i} g_{\beta,j}(x) dx - \int_{C_i} g_{\beta,j}(x) dx = \lambda \int_{D_i} g_{\beta,j}(x) dx - (1 - \lambda) \int_{D_i} g_{\beta,j}(x) dx$$

$$= \lambda \int \gamma_1(x) g_{\beta,j}(x) dx + (1 - \lambda) \int \gamma_2(x) g_{\beta,j}(x) dx.$$

(Notably, for j such that $g_{\beta,j}$ is identically 0 on D_i, this equality is trivial.) By linearity and the strong convexity of the function $x \mapsto x^2$, then, we find that for sets $E_j := D_{2j-1} \cup D_{2j}$,

$$\sum_{j=1}^k \varphi_j^2(\gamma) = \sum_{j=1}^k \left(\int_{E_j} \gamma(x) g_{\beta,j}(x) dx \right)^2$$

$$< \lambda \sum_{j=1}^k \left(\int_{E_j} \gamma_1(x) g_{\beta,j}(x) dx \right)^2 + (1 - \lambda) \sum_{\nu \in \mathcal{V}} \left(\int_{E_j} \gamma_2(x) g_{\beta,j}(x) dx \right)^2.$$

Thus one of the densities γ_1 or γ_2 must have a larger objective value than γ. This is our desired contradiction, which shows that (up to measure zero sets) any γ attaining the supremum in the information bound (69) must be constant on each of the D_i.

Having shown that γ is constant on each of the intervals D_i, we conclude that the supremum (69) can be reduced to a finite-dimensional problem over the subset

$$\mathcal{B}_{1,2k} := \left\{ u \in \mathbb{R}^{2k} \mid \|u\|_\infty \leq 1 \right\}$$

---

[1] For a function f and set A, the notation $f(A)$ denotes the image $f(A) = \{f(x) \mid x \in A\}$.

of $\mathbb{R}^{2k}$. In terms of this subset, the supremum (69) can be rewritten as the the upper bound

$$\sup_{\gamma \in \mathbb{B}_\infty} \sum_{j=1}^{k} \varphi_j(\gamma)^2 \leq \sup_{\gamma \in \mathcal{B}_{1,2k}} \sum_{j=1}^{k} \left(\gamma_{2j-1} \int_{D_{2j-1}} g_{\beta,j}(x) dx + \gamma_{2j} \int_{D_{2j}} g_{\beta,j}(x) dx \right)^2$$

By construction of the function g_β, we have the equality

$$\int_{D_{2j-1}} g_{\beta,j}(x) dx = -\int_{D_{2j}} g_{\beta,j}(x) dx = \int_0^{\frac{1}{2k}} g_{\beta,1}(x) dx = \int_0^{\frac{1}{2k}} \frac{1}{k^\beta} g_\beta(kx) dx = \frac{c_{1/2}}{k^{\beta+1}}.$$

This implies that

$$\frac{1}{2e^\alpha(e^\alpha-1)^2 n} \sum_{j=1}^{k} \left(D_{kl} \left(M_{+j}^n \| M_{-j}^n \right) + D_{kl} \left(M_{+j}^n \| M_{-j}^n \right) \right) \leq \sup_{\gamma \in \mathbb{B}_\infty} \sum_{j=1}^{k} \varphi_j(\gamma)^2$$

$$\leq \sup_{\gamma \in \mathcal{B}_{1,2k}} \sum_{j=1}^{k} \left(\frac{c_{1/2}}{k^{\beta+1}} \gamma^\top (e_{2j-1} - e_{2j}) \right)^2 = \frac{c_{1/2}^2}{k^{2\beta+2}} \sup_{\gamma \in \mathcal{B}_{1,2k}} \gamma^\top \sum_{j=1}^{k} (e_{2j-1} - e_{2j})(e_{2j-1} - e_{2j})^\top \gamma, \quad (70)$$

where $e_j \in \mathbb{R}^{2k}$ denotes the jth standard basis vector. Rewriting this using the Kronecker product $\otimes$, we have

$$\sum_{j=1}^{k} (e_{2j-1} - e_{2j})(e_{2j-1} - e_{2j})^\top = I_{k \times k} \otimes \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \preceq 2I_{2k \times 2k}.$$

Combining this bound with our inequality (70), we obtain

$$\sum_{j=1}^{k} \left(D_{kl} \left(M_{+j}^n \| M_{-j}^n \right) + D_{kl} \left(M_{+j}^n \| M_{-j}^n \right) \right) \leq 4n(e^\alpha-1)^2 \frac{c_{1/2}^2}{k^{2\beta+2}} \sup_{\gamma \in \mathcal{B}_{1,2k}} \|\gamma\|_2^2 = 4c_{1/2}^2 \frac{n(e^\alpha-1)^2}{k^{2\beta+1}}.$$

F Technical arguments

In this appendix, we collect proofs of technical lemmas and results needed for completeness.

F.1 Proof of Lemma 1

Fix an (arbitrary) estimator $\widehat{\theta}$. By assumption (10), we have

$$\Phi(\rho(\theta, \theta(P_\nu))) \geq 2\delta \sum_{j=1}^{d} \mathbf{1}\{[\mathsf{v}(\theta)]_j \neq \nu_j\}.$$

Taking expectations, we see that

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\Phi(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P))) \right] \geq \max_{\nu \in \mathcal{V}} \mathbb{E}_{P_\nu} \left[\Phi(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta_\nu)) \right]$$

$$\geq \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \mathbb{E}_{P_\nu} \left[\Phi(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta_\nu)) \right]$$

$$\geq \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} 2\delta \sum_{j=1}^{d} \mathbb{E}_{P_\nu} \left[\mathbf{1}\left\{ [\psi(\widehat{\theta})]_j \neq \nu_j \right\} \right]$$

as the average is smaller than the maximum of a set and using the separation assumption (10). Recalling the definition (31) of the mixtures $P_{\pm j}$, we swap the summation orders to see that

$$\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}} P_\nu\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right) = \frac{1}{|\mathcal{V}|}\sum_{\nu:\nu_j=1} P_\nu\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right) + \frac{1}{|\mathcal{V}|}\sum_{\nu:\nu_j=-1} P_\nu\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right)$$

$$= \frac{1}{2}P_{+j}\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right) + \frac{1}{2}P_{-j}\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right).$$

This gives the statement claimed in the lemma, while taking an infimum over all testing procedures $\psi : \mathcal{Z}^n \to \{-1,+1\}$ gives the claim (11).

F.2 Proof of unbiasedness for sampling strategy (26a)

We compute the expectation of a random variable Z sampled according to the strategy (26a), i.e. we compute $\mathbb{E}[Z \mid v]$ for a vector $v \in \mathbb{R}^d$. By scaling, it is no loss of generality to assume that $\|v\|_2 = 1$, and using the rotational symmetry of the ℓ_2-ball, we see it is no loss of generality to assume that $v = e_1$, the first standard basis vector.

Let the function s_d denote the surface area of the sphere in $\mathbb{R}^d$, so that

$$s_d(r) = \frac{d\pi^{d/2}}{\Gamma(d/2+1)}r^{d-1}$$

is the surface area of the sphere of radius r. (We use s_d as a shorthand for $s_d(1)$ when convenient.) Then for a random variable W sampled uniformly from the half of the ℓ_2-ball with first coordinate $W_1 \geq 0$, symmetry implies that by integrating over the radii of the ball,

$$\mathbb{E}[W] = e_1 \frac{2}{s_d}\int_0^1 s_{d-1}\left(\sqrt{1-r^2}\right)r\,dr.$$

Making the change of variables to spherical coordinates (we use ϕ as the angle), we have

$$\frac{2}{s_d}\int_0^1 s_{d-1}\left(\sqrt{1-r^2}\right)r\,dr = \frac{2}{s_d}\int_0^{\pi/2} s_{d-1}(\cos\phi)\sin\phi\,d\phi = \frac{2s_{d-1}}{s_d}\int_0^{\pi/2}\cos^{d-2}(\phi)\sin(\phi)\,d\phi.$$

Noting that $\frac{d}{d\phi}\cos^{d-1}(\phi) = -(d-1)\cos^{d-2}(\phi)\sin(\phi)$, we obtain

$$\frac{2s_{d-1}}{s_d}\int_0^{\pi/2}\cos^{d-2}(\phi)\sin(\phi)\,d\phi = -\left.\frac{\cos^{d-1}(\phi)}{d-1}\right|_0^{\pi/2} = \frac{1}{d-1},$$

or that

$$\mathbb{E}[W] = e_1\frac{(d-1)\pi^{\frac{d-1}{2}}\Gamma(\frac{d}{2}+1)}{d\pi^{\frac{d}{2}}\Gamma(\frac{d-1}{2}+1)}\frac{1}{d-1} = e_1\underbrace{\frac{\Gamma(\frac{d}{2}+1)}{\sqrt{\pi}d\Gamma(\frac{d-1}{2}+1)}}_{=:c_d}, \tag{71}$$

where we define the constant c_d to be the final ratio.

Allowing again $\|v\|_2 \leq r$, with the expression (71), we see that for our sampling strategy for Z, we have

$$\mathbb{E}[Z \mid v] = v\frac{B}{r}c_d\left(\frac{e^\alpha}{e^\alpha+1} - \frac{1}{e^\alpha+1}\right) = \frac{B}{r}c_d\frac{e^\alpha-1}{e^\alpha+1}.$$

Consequently, the choice

$$B = \frac{e^\alpha + 1}{e^\alpha - 1}\frac{r}{c_d} = \frac{e^\alpha + 1}{e^\alpha - 1}\frac{r\sqrt{\pi}d\Gamma(\frac{d-1}{2} + 1)}{\Gamma(\frac{d}{2} + 1)}$$

yields $\mathbb{E}[Z \mid v] = v$. Moreover, we have

$$\|Z\|_2 = B \le r\frac{e^\alpha + 1}{e^\alpha - 1}\frac{3\sqrt{\pi}\sqrt{d}}{2}$$

by Stirling's approximation to the Γ-function. By noting that $(e^\alpha + 1)/(e^\alpha - 1) \le 3/\alpha$ for $\alpha \le 1$, we see that $\|Z\|_2 \le 8r\sqrt{d}/\alpha$.

G   Effects of differential privacy in non-compact spaces

In this appendix, we present a somewhat pathological example that demonstrates the effects of differential privacy in non-compact spaces. Let us assume only that $\theta \in \mathbb{R}$ and $\alpha < \infty$, and we denote $\mathcal{P}_\theta$ to be the collection of probability measures with variance 1 having θ as a mean. In contrast to the non-private case, where the risk of the sample mean scales as $1/n$, we obtain

$$\mathfrak{M}_n(\mathbb{R}, (\cdot)^2, \alpha) = \infty \tag{72}$$

for all $n \in \mathbb{N}$. To see this, consider the Fano inequality version (9). Fix $\delta > 0$ and choose $\{\theta_1 = 0, \theta_2 = 2\delta, \dots, \theta_N = 2N\delta\}$ where $N = N(\delta, n) = \max\{\lceil \exp(64(e^\alpha - 1)^2 n)\rceil, 2^4\}$. Then by applying Corollary 1, we have for $\mathcal{V} = [N]$ that

$$\mathfrak{M}_n(\mathbb{R}, (\cdot)^2, \alpha) \ge \delta^2 \left(1 - \frac{4(e^\alpha - 1)^2 n \sum_{\nu,\nu'\in\mathcal{V}} \|P_\nu - P_{\nu'}\|_{\text{TV}}^2 / |\mathcal{V}|^2 + \log 2}{\log N(\delta, n)}\right).$$

We have $\|P_\nu - P_{\nu'}\|_{\text{TV}} \le 1$ for any two distributions P_ν and $P_{\nu'}$, which implies

$$\mathfrak{M}_n(\mathbb{R}, (\cdot)^2, \alpha) \ge \delta^2 \left(1 - \frac{16(e^\alpha - 1)^2 n + \log 2}{\log N(\delta, n)}\right) \ge \delta^2 \left(1 - \frac{1}{2}\right) = \frac{1}{2}\delta^2.$$

Since $\delta > 0$ was arbitrary, this proves the infinite minimax risk bound (72). The construction to achieve (72) is somewhat contrived, but it suggests that care is needed when designing differentially private inference procedures, and shows that even in cases when it is possible to attain a parametric rate of convergence, there may be no (locally) differentially private inference procedure.

References

[1] R. Ahlswede and A. Winter. Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48(3):569–579, March 2002.

[2] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley-Interscience, second edition, 2000.

[3] V. Anantharam, A. Gohari, S. Kamath, and C. Nair. On maximal correlation, hypercontractivity, and the data processing inequality studied by Erkip and Cover. *arXiv:1304.6133 [cs.IT]*, 2013. URL http://arxiv.org/abs/1304.6133.

[4] E. Arias-Castro, E. Candés, and M. Davenport. On the fundamental limits of adaptive sensing. *IEEE Transactions on Information Theory*, 59(1):472–481, 2013.

[5] P. Assouad. Deux remarques sur l'estimation. *C. R. Academy Scientifique Paris Séries I Mathematics*, 296(23):1021–1024, 1983.

[6] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar. Privacy, accuracy, and consistency too: A holistic solution to contingency table release. In *Proceedings of the 26th ACM Symposium on Principles of Database Systems*, 2007.

[7] A. Beimel, K. Nissim, and E. Omri. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology*, volume 5157 of *Lecture Notes in Computer Science*, pages 451–468. Springer, 2008.

[8] A. Beimel, S. P. Kasiviswanathan, and K. Nissim. Bounds on the sample complexity for private learning and private data release. In *Proceedings of the 7th Theory of Cryptography Conference*, pages 437–454, 2010.

[9] L. Birgé. Approximation dans les espaces métriques et théorie de l'estimation. *Zeitschrift für Wahrscheinlichkeitstheorie und verwebte Gebiet*, 65:181–238, 1983.

[10] A. Blum, K. Ligett, and A. Roth. A learning theory approach to non-interactive database privacy. In *Proceedings of the Fourtieth Annual ACM Symposium on the Theory of Computing*, 2008.

[11] P. Brucker. An $O(n)$ algorithm for quadratic knapsack problems. *Operations Research Letters*, 3(3):163–166, 1984.

[12] V. Buldygin and Y. Kozachenko. *Metric Characterization of Random Variables and Random Processes*, volume 188 of *Translations of Mathematical Monographs*. American Mathematical Society, 2000.

[13] R. Carroll and P. Hall. Optimal rates of convergence for deconvolving a density. *Journal of the American Statistical Association*, 83(404):1184–1186, 1988.

[14] R. Carroll, D. Ruppert, L. Stefanski, and C. Crainiceanu. *Measurement Error in Nonlinear Models: A Modern Perspective*. Chapman and Hall, second edition, 2006.

[15] K. Chaudhuri and D. Hsu. Convergence rates for differentially private statistical estimation. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.

[16] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.

[17] T. M. Cover and J. A. Thomas. *Elements of Information Theory, Second Edition*. Wiley, 2006.

[18] A. De. Lower bounds in differential privacy. In *Proceedings of the Ninth Theory of Cryptography Conference*, 2012. URL http://arxiv.org/abs/1107.2183.

[19] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy aware learning. *arXiv:1210.2085 [stat.ML]*, 2012. URL http://arxiv.org/abs/1210.2085.

[20] G. T. Duncan and D. Lambert. Disclosure-limited data dissemination. *Journal of the American Statistical Association*, 81(393):10–18, 1986.

[21] G. T. Duncan and D. Lambert. The risk of disclosure for microdata. *Journal of Business and Economic Statistics*, 7(2):207–217, 1989.

[22] C. Dwork and J. Lei. Differential privacy and robust statistics. In *Proceedings of the Fourty-*

First Annual ACM Symposium on the Theory of Computing, 2009.

[23] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology (EUROCRYPT 2006)*, 2006.

[24] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference*, pages 265–284, 2006.

[25] C. Dwork, G. N. Rothblum, and S. P. Vadhan. Boosting and differential privacy. In *51st Annual Symposium on Foundations of Computer Science*, pages 51–60, 2010.

[26] S. Efromovich. *Nonparametric Curve Estimation: Methods, Theory, and Applications.* Springer-Verlag, 1999.

[27] A. V. Evfimievski, J. Gehrke, and R. Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the Twenty-Second Symposium on Principles of Database Systems*, pages 211–222, 2003.

[28] S. E. Fienberg, U. E. Makov, and R. J. Steele. Disclosure limitation using perturbation and related methods for categorical data. *Journal of Official Statistics*, 14(4):485–502, 1998.

[29] S. E. Fienberg, A. Rinaldo, and X. Yang. Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *The International Conference on Privacy in Statistical Databases*, 2010.

[30] S. R. Ganta, S. Kasiviswanathan, and A. Smith. Composition attacks and auxiliary information in data privacy. In *Proceedings of the 14th ACM SIGKDD Conference on Knowledge and Data Discovery (KDD)*, 2008.

[31] L. J. Gleser. Estimation in a multivariate "errors in variables" regression model: large sample results. *Annals of Statistics*, 9(1):24–44, 1981.

[32] R. M. Gray. *Entropy and Information Theory.* Springer, 1990.

[33] R. Hall, A. Rinaldo, and L. Wasserman. Random differential privacy. *arXiv:1112.2680 [stat.ME]*, 2011. URL http://arxiv.org/abs/1112.2680.

[34] M. Hardt and G. N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *51st Annual Symposium on Foundations of Computer Science*, 2010.

[35] M. Hardt and K. Talwar. On the geometry of differential privacy. In *Proceedings of the Fourty-Second Annual ACM Symposium on the Theory of Computing*, pages 705–714, 2010. URL http://arxiv.org/abs/0907.3754.

[36] R. Z. Has'minskii. A lower bound on the risks of nonparametric estimates of densities in the uniform metric. *Theory of Probability and Applications*, 23:794–798, 1978.

[37] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

[38] M. Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the Association for Computing Machinery*, 45(6):983–1006, 1998.

[39] H. Ling and R. Li. Variable selection for partially linear models with measurement errors. *Journal of the American Statistical Association*, 104(485):234–248, 2009.

[40] P.-L. Loh and M. J. Wainwright. High-dimensional regression with noisy and missing data: provable guarantees with nonconvexity. *Annals of Statistics*, 40(3):1637–1664, 2012.

[41] Y. Ma and R. Li. Variable selection in measurement error models. *Bernoulli*, 16(1):274–300, 2010.

[42] L. W. Mackey, M. I. Jordan, R. Y. Chen, B. Farrell, and J. A. Tropp. Matrix concentration inequalities via the method of exchangeable pairs. *arXiv:1201.6002 [math.PR]*, 2012. URL `http://arxiv.org/abs/1201.6002`.

[43] A. McGregor, I. Mironov, T. Pitassi, O. Reingold, K. Talwar, and S. Vadhan. The limits of two-party differential privacy. In *51st Annual Symposium on Foundations of Computer Science*, 2010.

[44] S. Negahban, P. Ravikumar, M. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27 (4):538–557, 2012.

[45] R. R. Phelps. *Lectures on Choquet's Theorem, Second Edition*. Springer, 2001.

[46] B. I. P. Rubinstein, P. L. Bartlett, L. Huang, and N. Taft. Learning in a large function space: privacy-preserving mechanisms for SVM learning. *Journal of Privacy and Confidentiality*, 4 (1):65–100, 2012.

[47] D. Scott. On optimal and data-based histograms. *Biometrika*, 66(3):605–610, 1979.

[48] A. Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Fourty-Third Annual ACM Symposium on the Theory of Computing*, 2011.

[49] A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009.

[50] S. Warner. Randomized response: a survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

[51] L. Wasserman and S. Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.

[52] Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, 27(5):1564–1599, 1999.

[53] B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer-Verlag, 1997.