arXiv:1305.6000v1 [math.ST] 26 May 2013

Local Privacy and Minimax Bounds: Sharp Rates for Probability Estimation

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

Department of Statistics[*]
Department of Electrical Engineering and Computer Science[†]
University of California, Berkeley
Berkeley, CA, 94720

Abstract

We provide a detailed study of the estimation of probability distributions—discrete and continuous—in a stringent setting in which data is kept private even from the statistician. We give sharp minimax rates of convergence for estimation in these locally private settings, exhibiting fundamental tradeoffs between privacy and convergence rate, as well as providing tools to allow movement along the privacy-statistical efficiency continuum. One of the consequences of our results is that Warner's classical work on randomized response is an optimal way to perform survey sampling while maintaining privacy of the respondents.

1 Introduction

The original motivation for providing privacy in statistical problems, first discussed by Warner [27], was that "for reasons of modesty, fear of being thought bigoted, or merely a reluctance to confide secrets to strangers," respondents to surveys might prefer to be able to answer certain questions non-truthfully, or at least without the interviewer knowing their true response. With this motivation, Warner considered the problem of estimating the fractions of the population belonging to certain strata, which can be viewed as probability estimation within a multinomial model. In this paper, we revisit Warner's probability estimation problem, doing so within a theoretical framework that allows us to characterize optimal estimation under constraints on privacy. We also apply our theoretical tools to a further probability estimation problem—that of nonparametric density estimation.

In the large body of research on privacy and statistical inference [e.g., 27, 18, 12, 13, 19], a major focus has been on the problem of reducing disclosure risk: the probability that a member of a dataset can be identified given released statistics of the dataset. The literature has stopped short, however, of providing a formal treatment of disclosure risk that would permit decision-theoretic tools to be used in characterizing tradeoffs between the utility of achieving privacy and the utility associated with an inferential goal. Recently, a formal treatment of disclosure risk known as "differential privacy" has been proposed and studied in the cryptography, database and theoretical computer science literatures [15, 2, 14]. Differential privacy has strong semantic privacy guarantees make it a good candidate for declaring a statistical procedure or data collection mechanism private, and has been the focus of a growing body of recent work [14, 17, 20, 28, 25, 7, 21, 9, 6, 11].

In this paper, we bring together the formal treatment of disclosure risk provided by differential privacy with the tools of minimax decision theory to provide a theoretical treatment of probability estimation under privacy constraints. Just as in classical minimax theory, we are able to provide

lower bounds on the convergence rates of any estimator, in our case under a restriction to estimators that guarantee privacy. We complement these results with matching upper bounds that are achievable using computationally efficient algorithms. We thus bring classical notions of privacy, as introduced by Warner [27], into contact with differential privacy and statistical decision theory, obtaining quantitative tradeoffs between privacy and statistical efficiency.

1.1 Setting and contributions

Let us develop a bit of basic formalism before describing—at a high level—our main results. We study procedures that receive private views $Z_1, \ldots, Z_n \in \mathcal{Z}$ of an original set of observations, $X_1, \ldots, X_n \in \mathcal{X}$, where $\mathcal{X}$ is the (known) sample space. In our setting, Z_i is drawn conditional on X_i via the *channel distribution* $Q(Z_i \mid X_i = x, Z_j = z_j, j \neq i)$. Note that this channel allows "interactivity" [15], meaning that the distribution of Z_i may depend on X_i as well as the private views Z_j of X_j for $j \neq i$. Allowing interactivity—rather than forcing Z_i to be independent of Z_j—in some cases allows more efficient algorithms, and in our setting means that our lower bounds are stronger.

We assume each of these private views Z_i is α-differentially private for the original data X_i. To give a precise definition for this type of privacy, known as "local privacy," let $\sigma(\mathcal{Z})$ be the σ-field on $\mathcal{Z}$ over which the channel Q is defined. Then Q provides *α-local differential privacy* if

$$\sup \left\{ \frac{Q(S \mid X_i = x, Z_j = z_j, j \neq i)}{Q(S \mid X_i = x', Z_j = z_j, j \neq i)} \mid S \in \sigma(\mathcal{Z}), z_j \in \mathcal{Z}, \text{ and } x, x' \in \mathcal{X} \right\} \leq \exp(\alpha). \tag{1}$$

In the non-interactive setting (in which we impose the constraint that the providers of the data release a private view independently of the other data providers) the expression (1) simplifies to

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

a formulation of local privacy first proposed by Evfimievski et al. [17]. Although more complex to analyze, the likelihood ratio bound (1) is attractive for many reasons. It means that any individual providing data guarantees his or her own privacy—no further processing or mistakes by a collection agency can compromise one's data—and the individual has plausible deniability about taking a value x, since any outcome z is nearly as likely to have come from some other initial value x'. The likelihood ratio also controls the error rate in tests for the presence of points x in the data [28]. All that is required is that the likelihood ratio (1) be bounded no matter the data provided by other participants.

In the current paper, we study minimax convergence rates when the data provided satisfies the local privacy guarantee (1). Our two main results quantify the penalty that must be paid when local privacy at a level α is provided in multinomial estimation and density estimation problems. At a high level, our first result implies that for estimation of a d-dimensional multinomial probability mass function, the effective sample size of *any* statistical estimation procedure decreases from n to $n\alpha^2/d$ whenever α is a sufficiently small constant. A consequence of our results is that Warner's randomized response procedure [27] enjoys optimal sample complexity; it is interesting to note that even with the recent focus on privacy and statistical inference, the optimal privacy-preserving strategy for problems such as survey collection has been known for almost 50 years.

Our second main result, on density estimation, exhibits an interesting departure from standard minimax estimation results. If the density being estimated has β continuous derivatives, then

classical results on density estimation [e.g., 30, 29, 26]) show that the minimax integrated squared error scales (in the number n of samples) as $n^{-2\beta/(2\beta+1)}$. In the locally private case, we show that—even when densities are bounded and well-behaved—there is a difference in the *polynomial* rate of convergence: we obtain a scaling of $(\alpha^2 n)^{-2\beta/(2\beta+2)}$. We give efficiently implementable algorithms that attain sharp upper bounds as companions to our lower bounds, which in some cases exhibit the necessity of non-trivial sampling strategies to guarantee privacy.

Notation: We summarize here the notation used throughout the paper. Given 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 defined by

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

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

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

For random vectors X and Y, where X is distributed according to the distribution P and $Y \mid X$ is distributed according to $Q(\cdot \mid X)$, let $M(\cdot) = \int Q(\cdot \mid x) dP(x)$ denote the marginal distribution of Y. The mutual information between X and Y is

$$I(X;Y) := \mathbb{E}_P \left[D_{\rm kl}\left(Q(\cdot \mid X) \| M(\cdot)\right)\right] = \int D_{\rm kl}\left(Q(\cdot \mid X = x) \| M(\cdot)\right) dP(x).$$

A random variable Y has Laplace(α) distribution if its density $p_Y(y) = \frac{\alpha}{2} \exp\left(-\alpha |y|\right)$, where $\alpha > 0$. For matrices $A, B \in \mathbb{R}^{d \times d}$, the notation $A \preceq B$ means $B - A$ is positive semidefinite, and $A \prec B$ means $B - A$ is positive definite. We write $a_n \lesssim b_n$ to denote that $a_n = \mathcal{O}(b_n)$ and $a_n \asymp b_n$ to denote that $a_n = \mathcal{O}(b_n)$ and $b_n = \mathcal{O}(a_n)$. For a convex set $C \subset \mathbb{R}^d$, we let Π_C denote the orthogonal projection operator onto C, i.e., $\Pi_C(v) := \operatorname{argmin}_{w \in C} \{\|v - w\|_2\}$.

2 Background and Problem Formulation

In this section, we provide the necessary background on the minimax framework used throughout the paper. Further details on minimax techniques can be found in several standard sources [e.g., 3, 29, 30, 26]. We also reference our companion paper on parametric statistical inference under differential privacy constraints [11]; we make use of two theorems from that earlier paper, but in order to keep the current paper self-contained, we restate them in this section.

2.1 Minimax framework

Let $\mathcal{P}$ denote a class of distributions on the sample space $\mathcal{X}$, and let $\theta : \mathcal{P} \to \Theta$ denote a function defined on $\mathcal{P}$. The range Θ depends on the underlying statistical model; for example, for density estimation, Θ may consist of the set of probability densities defined on $[0, 1]$. We let ρ denote the semi-metric on the space Θ that we use to measure the error of an estimator for θ, and $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ be a non-decreasing function with $\Phi(0) = 0$ (for example, $\Phi(t) = t^2$).

Recalling that $\mathcal{Z}$ is the domain of the private variables Z_i, let $\widehat{\theta} : \mathcal{Z}^n \to \Theta$ denote an arbitrary estimator for θ. Let $\mathcal{Q}_\alpha$ denote the set of conditional (or channel) distributions guaranteeing α-local privacy (1); then for any $Q \in \mathcal{Q}_\alpha$ we can define the minimax rate

$$\mathfrak{M}_n\left(\theta(\mathcal{P}), \Phi \circ \rho, Q\right) := \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_{P,Q}\left[\Phi\left(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P))\right)\right] \tag{3a}$$

associated with estimating θ based on the private samples $(Z_1, \ldots, Z_n)$. In the definition (3a), the expectation is taken both with respect to the distribution P on the variables $X_1, \ldots, X_n$ and the α-private channel Q. By taking the infimum over all possible channels $Q \in \mathcal{Q}_\alpha$, we obtain the central object of interest for this paper, the *α-private minimax rate* for the family $\theta(\mathcal{P})$, defined as

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, \alpha) := \inf_{\widehat{\theta}, Q \in \mathcal{Q}_\alpha} \sup_{P \in \mathcal{P}} \mathbb{E}_{P,Q}\left[\Phi\left(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P))\right)\right]. \tag{3b}$$

A standard route for lower bounding the minimax risk (3a) is by reducing the estimation problem to the testing problem of identifying a point $\theta \in \Theta$ from a finite collection of well-separated points [30, 29]. Given an index set $\mathcal{V}$ of finite cardinality, the indexed family of distributions $\{P_\nu, \nu \in \mathcal{V}\} \subset \mathcal{P}$ is said to be a 2δ-packing of Θ if $\rho(\theta(P_\nu), \theta(P_{\nu'})) \geq 2\delta$ for all $\nu \neq \nu'$ in $\mathcal{V}$. The setup is that of a standard hypothesis testing problem: nature chooses $V \in \mathcal{V}$ uniformly at random, then data $(X_1, \ldots, X_n)$ are drawn from the n-fold conditional product distribution P_ν^n, conditioning on $V = \nu$. The problem is to identify the member ν of the packing set $\mathcal{V}$.

In this work we have the additional complication that all the statistician observes are the private samples $Z_1, \ldots, Z_n$. To that end, if we let $Q^n(\cdot \mid x_{1:n})$ denote the conditional distribution of $Z_1, \ldots, Z_n$ given that $X_1 = x_1, \ldots, X_n = x_n$, we define the marginal channel M_ν^n via the expression

$$M_\nu^n(A) := \int Q^n(A \mid x_1, \ldots, x_n) dP_\nu(x_1, \ldots, x_n) \quad \text{for } A \in \sigma(\mathcal{Z}^n). \tag{4}$$

Letting $\psi : \mathcal{Z}^n \to \mathcal{V}$ denote an arbitrary testing procedure—a measureable mapping $\mathcal{Z}^n \to \mathcal{V}$—we have the following minimax risk bound, whose two parts are known as Le Cam's two-point method and Fano's inequality. In the lemma, we let $\mathbb{P}$ denote the joint distribution of the random variable V and the samples Z_i.

Lemma 1 (Minimax risk bound)**.** *For the previously described estimation and testing problems, we have the lower bound*

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, Q) \geq \Phi(\delta) \inf_{\psi} \mathbb{P}(\psi(Z_1, \ldots, Z_n) \neq V), \tag{5}$$

where the infimum is taken over all testing procedures. For a binary test specified by $\mathcal{V} = \{\nu, \nu'\}$,

$$\inf_{\psi} \mathbb{P}\left(\psi(Z_1, \ldots, Z_n) \neq V\right) = \frac{1}{2} - \frac{1}{2} \left\| M_\nu^n - M_{\nu'}^n \right\|_{\mathrm{TV}}, \tag{6a}$$

and more generally,

$$\inf_{\psi} \mathbb{P}(\psi(Z_1, \ldots, Z_n) \neq V) \geq \left[1 - \frac{I(Z_1, \ldots, Z_n; V) + \log 2}{\log |\mathcal{V}|}\right]. \tag{6b}$$

For Le Cam's inequality (6a), see, e.g., Lemma 1 of Yu [30] or Theorem 2.2 of Tsybakov [26]; for Fano's inequality (6b), see Eq. (1) of Yang and Barron [29] or Chapter 2 of Cover and Thomas [8].

2.2 Information bounds

The main step in proving minimax lower bounds is to control the divergences involved in the lower bounds (6a) and (6b). In our companion paper [11], we present two results, which we now review, in which bounds on $\left\|M_\nu^n - M_{\nu'}^n\right\|_{\mathrm{TV}}$ and $I(Z_1, \ldots, Z_n; V)$ are obtained as a function of the amount of privacy provided and the distances between the underlying distributions P_ν. The first result [11, Theorem 1 and Corollary 1] gives control over pairwise KL-divergences between the marginals (4), which lends itself to application of Le Cam's method (6a) and simple applications of Fano's inequality (6b). The second result [11, Theorem 2 and Corollary 4] provides a variational upper bound on the mutual information $I(Z_1, \ldots, Z_n; V)$—variational in the sense that it requires optimization over the set of functions

$$\mathcal{G}_\alpha := \left\{\gamma \in L^\infty(\mathcal{X}) \mid \sup_{x\in\mathcal{X}} |\gamma(x)| \le \frac{1}{2}(e^\alpha - e^{-\alpha})\right\}.$$

Here $L^\infty(\mathcal{X}) := \{f : \mathcal{X} \to \mathbb{R} \mid \sup_{x\in\mathcal{X}} |f(x)| < \infty\}$ denotes the space of bounded functions on $\mathcal{X}$.

Our bounds apply to any channel distribution Q that is α-locally private (1). For each $i \in \{1, \ldots, n\}$, let $P_{\nu,i}$ be the distribution of X_i conditional on the random packing element $V = \nu$, and let M_ν^n be the marginal distribution (4) induced by passing X_i through Q. Define the mixture distribution $\overline{P}_i = \frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}} P_{\nu,i}$, and the linear functionals $\varphi_{\nu,i} : L^\infty(\mathcal{X}) \to \mathbb{R}$ by

$$\varphi_{\nu,i}(\gamma) := \int_{\mathcal{X}} \gamma(x)\left(dP_{\nu,i}(x) - d\overline{P}_i(x)\right).$$

With this notation we can state the following proposition, which summarizes the results that we will need from Duchi et al. [11]:

Proposition 1 (Information bounds). *(a) For all $\alpha \ge 0$,*

$$D_{\mathrm{kl}}\left(M_\nu^n \| M_{\nu'}^n\right) + D_{\mathrm{kl}}\left(M_{\nu'}^n \| M_\nu^n\right) \le 4(e^\alpha - 1)^2 \sum_{i=1}^n \left\|P_{\nu,i} - P_{\nu',i}\right\|_{\mathrm{TV}}^2. \tag{7}$$

(b) For all $\alpha \in [0, \log(\frac{1}{2} + \frac{1}{2}\sqrt{3})]$,

$$I(Z_1, \ldots, Z_n; V) \le C_\alpha \sum_{i=1}^n \frac{1}{|\mathcal{V}|} \sup_{\gamma\in\mathcal{G}_\alpha} \sum_{\nu\in\mathcal{V}} \left(\varphi_{\nu,i}(\gamma)\right)^2, \tag{8}$$

where $C_\alpha := 4/(e^{-\alpha} - 2(e^\alpha - 1))$.

By combining Proposition 1 with Lemma 1, it is possible to derive sharp lower bounds on arbitrary estimation procedures under α-local privacy. In particular, we may apply the bound (7) with Le Cam's method (6a), though lower bounds so obtained often lack dimension dependence we might hope to capture (see Section 3.2 of Duchi et al. [11] for more discussion of this issue). The bound (8), which (up to constants) implies the bound (7), allows more careful control via suitably constructed packing sets $\mathcal{V}$ and application of Fano's inequality (6b), since the supremum controls a more global view of the structure of $\mathcal{V}$. In the rest of this paper, we demonstrate this combination for probability estimation problems.

3 Multinomial Estimation under Local Privacy

In this section we return to the classical problem of avoiding answer bias in surveys, the original motivation for studying local privacy [27]. We provide a detailed study of estimation of a multinomial probability under α-local differential privacy, providing sharp upper and lower bounds for the minimax rate of convergence.

3.1 Minimax rates of convergence for multinomial estimation

Consider the probability simplex $\Delta_d := \left\{\theta \in \mathbb{R}^d \mid \theta \geq 0, \sum_{j=1}^d \theta_j = 1\right\}$ in $\mathbb{R}^d$. The multinomial estimation problem is defined as follows. Given a vector $\theta \in \Delta_d$, samples X are drawn i.i.d. from a multinomial with parameters θ, where $P_\theta(X = j) = \theta_j$ for $j \in \{1, \ldots, d\}$, and our goal is to estimate the probability vector θ. In one of the earliest evaluations of privacy, Warner [27] studied the Bernoulli variant of this problem, and proposed a simple privacy-preserving mechanism known as *randomized response*: for a given survey question, respondents provide a truthful answer with probability $p > 1/2$, and a lie with probability $1 - p$.

In our setting, we assume that the statistician sees random variables Z_i that are all α-locally private (1) for the corresponding samples X_i from the multinomial. In this case, we have the following result, which characterizes the minimax rate of estimation of a multinomial in terms of mean-squared error $\mathbb{E}[\|\widehat{\theta} - \theta\|_2^2]$.

Theorem 1. *There exist universal constants $0 < c_\ell \leq c_u < 5$ such that for all $\alpha \in [0, 1/4]$, the minimax rate for multinomial estimation satisfies the bounds*

$$c_\ell \max_{k \in \{1,\ldots,d\}} \min\left\{\frac{1}{k}, \frac{k \log \frac{d}{k}}{n\alpha^2}\right\} \;\leq\; \mathfrak{M}_n\left(\Delta_d, \|\cdot\|_2^2, \alpha\right) \;\leq\; c_u \min\left\{1, \frac{d}{n\alpha^2}\right\}. \tag{9}$$

We provide a proof of the lower bound in Theorem 1 in Section 5. Simple estimation strategies achieve the lower bound, and we believe exploring them is interesting, so we provide them in the next section.

Theorem 1 shows that providing local privacy can sometimes be quite detrimental to the quality of statistical estimators. Indeed, let us compare this rate to the classical rate in which there is no privacy. Then estimating θ via proportions (i.e., maximum likelihood), we have

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

On the other hand, an appropriate choice of k in Theorem 1 implies that

$$\min\left\{1, \frac{1}{\sqrt{n\alpha^2}}, \frac{d}{n\alpha^2}\right\} \lesssim \mathfrak{M}_n\left(\Delta_d, \|\cdot\|_2^2, \alpha\right) \lesssim \min\left\{1, \frac{d}{n\alpha^2}\right\}, \tag{10}$$

which we show in Section 5. Thus, for suitably large sample sizes n, the effect of providing differential privacy at a level α causes a reduction in the effective sample size of $n \mapsto n\alpha^2/d$.

3.2 Private multinomial estimation strategies

An interesting consequence of the lower bound in (9) is the following fact that we now demonstrate: Warner's classical randomized response mechanism [27] (with minor modification) achieves the optimal convergence rate. Thus, although it was not originally recognized as such, Warner's proposal is actually optimal in a strong sense. There are also other relatively simple estimation strategies that achieve convergence rate $d/n\alpha^2$; for instance, as we show, the Laplace perturbation approach proposed by Dwork et al. [15] is one such strategy. Nonetheless, its ease of use coupled with our optimality results provide support for randomized response as a preferred method for private estimation of population probabilities.

Let us now prove that these strategies attain the optimal rate of convergence. Since there is a bijection between multinomial samples $x \in \{1, \ldots, d\}$ and the d standard basis vectors $e_1, \ldots, e_d \in \mathbb{R}^d$, we abuse notation and represent samples x as either when designing estimation strategies.

Randomized response: In randomized response, we construct the private vector $Z \in \{0,1\}^d$ from a multinomial sample $x \in \{e_1, \ldots, e_d\}$ by sampling d coordinates independently via the procedure

$$[Z]_j = \begin{cases} x_j & \text{with probability } \frac{\exp(\alpha/2)}{1+\exp(\alpha/2)} \\ 1 - x_j & \text{with probability } \frac{1}{1+\exp(\alpha/2)}. \end{cases} \tag{11}$$

We claim that this channel (11) is α-differentially private: indeed, note that for any $x, x' \in \Delta_d$ and any vector $z \in \{0,1\}^d$ we have

$$\begin{aligned} \frac{Q(Z = z \mid x)}{Q(Z = z \mid x')} &= \frac{\left(\frac{\exp(\alpha/2)}{1+\exp(\alpha/2)}\right)^{\|z-x\|_1} \left(\frac{1}{1+\exp(\alpha/2)}\right)^{d-\|z-x\|_1}}{\left(\frac{\exp(\alpha/2)}{1+\exp(\alpha/2)}\right)^{\|z-x'\|_1} \left(\frac{1}{1+\exp(\alpha/2)}\right)^{d-\|z-x'\|_1}} \\ &= \exp\left(\frac{\alpha}{2}\left(\|z-x\|_1 - \|z-x'\|_1\right)\right) \in [\exp(-\alpha), \exp(\alpha)], \end{aligned}$$

where we used the triangle inequality to assert that $|\,\|z-x\|_1 - \|z-x'\|_1\,| \le \|x-x'\|_1 \le 2$. We can compute the expected value and variance of the random variables Z; indeed, by definition (11)

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

Since the Z are Bernoulli, we obtain the variance bound $\mathbb{E}[\|Z - \mathbb{E}[Z]\|_2^2] < d/4 + 1$. Recalling the definition of the projection Π_{Δ_d} onto the simplex, we arrive at the natural estimator

$$\widehat{\theta}_{\text{part}} := \frac{1}{n}\sum_{i=1}^n \left(Z_i - \mathbb{1}/(1+e^{\alpha/2})\right)\frac{e^{\alpha/2}+1}{e^{\alpha/2}-1} \quad \text{and} \quad \widehat{\theta} := \Pi_{\Delta_d}\left(\widehat{\theta}_{\text{part}}\right). \tag{12}$$

The projection of $\widehat{\theta}_{\text{part}}$ onto the probability simplex can be done in time linear in the dimension d of the problem [4], so the estimator (12) is efficiently computable. Since projections only decrease distance, vectors in the simplex are at most distance $\sqrt{2}$ apart, and $\mathbb{E}_\theta[\widehat{\theta}_{\text{part}}] = \theta$ by construction, we find that

$$\mathbb{E}\left[\|\widehat{\theta} - \theta\|_2^2\right] \le \min\left\{2, \mathbb{E}\left[\|\widehat{\theta}_{\text{part}} - \theta\|_2^2\right]\right\} \le \min\left\{2, \left(\frac{d}{4n} + \frac{1}{n}\right)\left(\frac{e^{\alpha/2}+1}{e^{\alpha/2}-1}\right)^2\right\} \lesssim \min\left\{1, \frac{d}{n\alpha^2}\right\}.$$

Laplace perturbation: We now turn to the strategy of Dwork et al. [15], where we add Laplacian noise to the data. Let the vector $W \in \mathbb{R}^d$ have independent coordinates, each distributed as Laplace($\alpha/2$). Then $x + W$ is α-differentially private for $x \in \Delta_d$: the ratio of the densities $q(\cdot \mid x)$ and $q(\cdot \mid x')$ of $x + W$ and $x' + W$ is

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

By defining the private data $Z_i = X_i + W_i$, where $W_i \in \mathbb{R}^d$ are independent, we can define the partial estimator $\widehat{\theta}_{\text{part}} = \frac{1}{n}\sum_{i=1}^n Z_i$ and the projected estimator $\widehat{\theta} := \Pi_{\Delta_d}(\widehat{\theta}_{\text{part}})$, similar to our randomized response construction (12). Then by computing the variance of the noise samples W_i, it is clear that

$$\mathbb{E}\left[\|\widehat{\theta} - \theta\|_2^2\right] \le \min\left\{2, \mathbb{E}\left[\|\widehat{\theta}_{\text{part}} - \theta\|_2^2\right]\right\} \le \min\left\{2, \frac{1}{n} + 4\frac{d}{n\alpha^2}\right\} \lesssim \min\left\{1, \frac{d}{n\alpha^2}\right\}.$$

For small α, the Laplace perturbation approach has a somewhat sharper convergence rate in terms of constants than that of the randomized response estimator (12), so in some cases it may be preferred. Nonetheless, the simplicity of explaining the sampling procedure (11) may argue for its use in scenarios such as survey sampling.

4 Density Estimation under Local Privacy

In this section, we turn to studying a nonparametric statistical problem in which the effects of local differential privacy turn out to be somewhat more severe. We show that for the problem of density estimation, instead of just multiplicative loss in the effective sample size as in previous section (see also our paper [11]), imposing local differential privacy leads to a completely different convergence rate. This result holds even though we solve an estimation problem in which the function estimated and the samples themselves belong to compact spaces.

In more detail, we consider estimation of probability densities $f : \mathbb{R} \to \mathbb{R}_+$, $\int f(x)dx = 1$ and $f \ge 0$, defined on the real line, focusing on a standard family of densities of varying smoothness [e.g. 26]. Throughout this section, we let $\beta \in \mathbb{N}$ denote a fixed positive integer. Roughly, we consider densities that have bounded βth derivative, and we study density estimation using the squared L^2-norm $\|f\|_2^2 := \int f^2(x)dx$ as our metric; in formal terms, we impose these constraints in terms of Sobolev classes (e.g. [26, 16]). Let the countable collection of functions $\{\varphi_j\}_{j=1}^\infty$ be an orthonormal basis for $L^2([0,1])$. Then any function $f \in L^2([0,1])$ can be expanded as a sum $\sum_{j=1}^\infty \theta_j\varphi_j$ in terms of the basis coefficients $\theta_j := \int f(x)\varphi_j(x)dx$. By Parseval's theorem, we are guaranteed that $\{\theta_j\}_{j=1}^\infty \in \ell^2(\mathbb{N})$. The *Sobolev space* $\mathcal{F}_{\beta,C}$ is obtained by enforcing a particular decay rate on the coefficients θ:

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

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

If we choose the trignometric basis as our orthonormal basis, then membership in the class $\mathcal{F}_{\beta,C}$ corresponds to certain 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:

$$\varphi_0(t) = 1, \quad \varphi_{2j}(t) = \sqrt{2}\cos(2\pi jt), \quad \varphi_{2j+1}(t) = \sqrt{2}\sin(2\pi jt). \tag{13}$$

Now consider a β-times almost everywhere differentiable function f 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 it is known [26, Lemma A.3] that, uniformly for such f, there is a universal constant c such that that $f \in \mathcal{F}_{\beta,cC}$. Thus, Definition 1 (essentially) captures densities that have Lipschitz-continuous $(\beta-1)$th derivative. In the sequel, we write $\mathcal{F}_\beta$ when the bound C in $\mathcal{F}_{\beta,C}$ is $\mathcal{O}(1)$. It is well known [30, 29, 26] that the minimax risk for non-private estimation of densities in the class $\mathcal{F}_\beta$ scales as

$$\mathfrak{M}_n\left(\mathcal{F}_\beta, \|\cdot\|_2^2, \infty\right) \asymp n^{-\frac{2\beta}{2\beta+1}}. \tag{14}$$

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 (14) is *no longer attainable* when we require α-local differential privacy. In particular, we prove a lower bound that is substantially larger. In Sections 4.2 and 4.3, we show how to achieve this lower bound using histogram and orthogonal series estimators.

4.1 Lower bounds on density estimation

We begin by giving our main lower bound on the minimax rate of estimation of densities when samples from the density must be kept differentially private. We provide the proof of the following theorem in Section 6.1.

Theorem 2. *Consider the class of densities $\mathcal{F}_\beta$ defined using the trigonometric basis (13). For some $\alpha \in (0, 1/4]$, suppose Z_i are α-locally private (1) for the samples $X_i \in [0,1]$. There exists a constant $c > 0$, dependent only on β, such that*

$$\mathfrak{M}_n\left(\mathcal{F}_\beta, \|\cdot\|_2^2, \alpha\right) \geq c\left(n\alpha^2\right)^{-\frac{2\beta}{2\beta+2}}. \tag{15}$$

In comparison with the classical minimax rate (14), the lower bound (15) is substantially different, in that it involves a different polynomial exponent: namely, the exponent is $2\beta/(2\beta+1)$ in the classical case (14), while in the differentially private case (15), the exponent has been reduced to $2\beta/(2\beta+2)$. For example, when we estimate Lipschitz densities, we have $\beta = 1$, and the rate degrades from $n^{-2/3}$ to $n^{-1/2}$. Moreover, this degradation occurs even though our samples are drawn from a compact space and the set $\mathcal{F}_\beta$ is also compact.

Interestingly, no estimator based on Laplace (or exponential) perturbation of the samples X_i themselves can attain the rate of convergence (15). This can be established by connecting such a perturbation-based approach to the problem of nonparametric deconvolution. In their study of the deconvolution problem, Carroll and Hall [5] show that if samples 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 the samples $X + W$ and attain a rate of convergence better than $n^{-2\beta/(2\beta+2a+1)}$. Since the Laplace distribution's characteristic function has tails decaying as t^{-2}, no estimator based on perturbing the samples directly can attain a rate of convergence better than $n^{-2\beta/(2\beta+5)}$. If the lower bound (15) is attainable, we must then study privacy mechanisms that are not simply based on direct perturbation of the samples $\{X_i\}_{i=1}^n$.

4.2 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 ($\beta = 1$), for which 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. Any such vector defines a density estimate via the sum

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

where 1_E denotes the characteristic (indicator) function of the set E.

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

$$Z_i := \mathsf{e}_k(X_i) + W_i, \quad \text{for } i = 1, 2, \ldots, n, \tag{16}$$

where $\mathsf{e}_k(X_i) \in \Delta_k$ is a k-vector with the jth 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 $\{Z_i\}_{i=1}^n$ so-defined are α-locally differentially private for $\{X_i\}_{i=1}^n$.

Using these private variables, we then form the density estimate $\widehat{f} := f_{\widehat{\theta}} = \sum_{j=1}^k \widehat{\theta}_j 1_{\mathcal{X}_j}$ based on the vector

$$\widehat{\theta} := \Pi_k\left(\frac{k}{n}\sum_{i=1}^n Z_i\right), \tag{17}$$

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.

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

$$\mathbb{E}_f\left[\left\|\widehat{f} - f\right\|_2^2\right] \leq 5(\alpha^2 n)^{-\frac{1}{2}} + \sqrt{\alpha} n^{-3/4}. \tag{18}$$

For any fixed $\alpha > 0$, the first term in the bound (18) dominates, and the $\mathcal{O}((\alpha^2 n)^{-\frac{1}{2}})$ rate matches the minimax lower bound (15) in the case $\beta = 1$. Consequently, we have shown that the privatized histogram estimator is minimax-optimal for Lipschitz densities. This result provides the private analog of the classical result that histogram estimators are minimax-optimal (in the non-private setting) for Lipshitz densities. See Section 6.2 for a proof of Proposition 2. We remark in passing that a randomized response scheme parallel to that of Section 3.2 achieves the same rate of convergence; once again, randomized response is optimal.

4.3 Achievability by orthogonal projection estimators

For higher degrees of smoothness ($\beta > 1$), histogram estimators no longer achieve optimal rates in the classical setting. 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 (15) for all orders of smoothness $\beta \geq 1$.

Recall the elliptical Sobolev space (Definition 1), in which a function f is represented in terms of its basis expansion $f = \sum_{j=1}^{\infty} \theta_j \varphi_j$, where $\theta_j = \int f(x)\varphi_j(x)dx$. This representation underlies the classical method of orthonormal series estimation: given a data set, $\{X_1, X_2, \ldots, X_n\}$, drawn i.i.d. according to a density $f \in L^2([0,1])$, we first compute the empirical basis coefficients

$$\widehat{\theta}_j = \frac{1}{n}\sum_{i=1}^{n} \varphi_j(X_i) \quad \text{and then set} \quad \widehat{f} = \sum_{j=1}^{k} \widehat{\theta}_j \varphi_j, \tag{19}$$

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 [16, 26].

In the setting of local privacy, 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 $Z_i = (Z_{i,1}, \ldots, Z_{i,k})$ with the property that $\mathbb{E}[Z_{i,j} \mid X_i] = \varphi_j(X_i)$ for each $j = 1, 2, \ldots, k$. We assume the basis functions are uniformly bounded; i.e., there exists a constant $B_0 = \sup_j \sup_x |\varphi_j(x)| < \infty$. (This boundedness condition holds for many standard bases, including the trigonometric basis underlying the classical Sobolev classes and the Walsh basis.) For a fixed number B strictly larger than B_0 (to be specified momentarily), consider the following scheme:

Sampling strategy Given a vector $\tau \in [-B_0, B_0]^k$, construct $\widetilde{\tau} \in \{-B_0, B_0\}^k$ with coordinates $\widetilde{\tau}_j$ sampled independently from $\{-B_0, B_0\}$ with probabilities $\frac{1}{2} - \frac{\tau_j}{2B_0}$ and $\frac{1}{2} + \frac{\tau_j}{2B_0}$. Sample T from a Bernoulli$(e^\alpha/(e^\alpha + 1))$ distribution. Then choose $Z \in \{-B, B\}^k$ via

$$Z \sim \begin{cases} \text{Uniform on } \big\{z \in \{-B, B\}^k : \langle z, \widetilde{\tau}\rangle > 0\big\} & \text{if } T = 1 \\ \text{Uniform on } \big\{z \in \{-B, B\}^k : \langle z, \widetilde{\tau}\rangle \leq 0\big\} & \text{if } T = 0. \end{cases} \tag{20}$$

By inspection, Z is α-differentially private for any initial vector in the box $[-B_0, B_0]^k$, and moreover, the samples (20) are efficiently computable (for example by rejection sampling).

Starting from the vector $\tau \in \mathbb{R}^k$, $\tau_j = \varphi_j(X_i)$, in the above sampling strategy, iteration of expectation yields

$$\mathbb{E}[[Z]_j \mid X = x] = c_k \frac{B}{B_0\sqrt{k}} \left(\frac{e^\alpha}{e^\alpha + 1} - \frac{1}{e^\alpha + 1} \right) \varphi_j(x) = c_k \frac{B}{B_0\sqrt{k}} \frac{e^\alpha - 1}{e^\alpha + 1} \varphi_j(x), \tag{21}$$

for a constant c_k that may depend on k but is $\mathcal{O}(1)$ and bounded away from 0. Consequently, to attain the unbiasedness condition $\mathbb{E}[[Z_i]_j \mid X_i] = \varphi_j(X_i)$, it suffices to take $B = \mathcal{O}(B_0\sqrt{k}/\alpha)$.

The full sampling and inferential scheme are as follows: given a data point X_i, we sample Z_i according to the strategy (20), where we start from the vector $\tau = [\varphi_j(X_i)]_{j=1}^k$ and use the bound $B = B_0\sqrt{k}(e^\alpha + 1)/c_k(e^\alpha - 1)$, where the constant c_k is as in the expression (21). Defining the density estimator

$$\widehat{f} := \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{k} Z_{i,j}\varphi_j, \tag{22}$$

we obtain the following proposition.

Proposition 3. *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 the estimator (22) with $k = (n\alpha^2)^{1/(2\beta+2)}$ satisfies*

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

for any f in the Sobolev space $\mathcal{F}_{\beta,C}$.

See Section 6.3 for a proof.

Propositions 2 and 3 make clear that the minimax lower bound (15) is sharp, as claimed. We have thus attained a variant of the known minimax density estimation rates (14), but with a polynomially worse sample complexity as soon as we require local differential privacy.

Before concluding our exposition, we make a few remarks on other potential density estimators. Our orthogonal-series estimator (22) (and sampling scheme (21)), while similar in spirit to that proposed by Wasserman and Zhou [28, Sec. 6], is different in that it is locally private and requires a different noise strategy to obtain both α-local privacy and optimal convergence rate. 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. [28]) 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 (22), except that we use the Laplacian perturbed Z_i. However, this estimator suffers the following drawback (see section 6.4):

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*

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

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.

5 Proof of Theorem 1

At a high level, our proof can be split into three steps, the first of which is relatively standard, while the second two exploit specific aspects of the local privacy set-up:

(1) The first step is a standard reduction, based on Lemma 1, 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 the set $\mathcal{V} \subset \mathbb{R}^d$ such that each pair is δ-separated and the resulting set is as large as possible (a maximal δ-packing). In addition, our arguments require that, for a random variable V uniformly distributed over $\mathcal{V}$, the covariance $\text{Cov}(V)$ has relatively small operator norm.

(3) The final step is to apply Proposition 1 in order to control the mutual information associated with the testing problem. To do so, it is necessary to show that controlling the supremum subsets of $L^\infty(\mathcal{X})$ in the bound (8) can be reduced to bounding the operator norm of $\mathrm{Cov}(V)$.

We have already described the reduction of Step 1 in Section 2.1. Accordingly, we turn to the second step.

Constructing a good packing: The following result on the binary hypercube $H_d := \{0,1\}^d$ underlies our construction:

Lemma 2. *There exist universal constants $c_1, c_2 \in (0,\infty)$ such that for each $k \in \{1,\ldots,d\}$, there is a set $\mathcal{V} \subset H_d$ with the following properties:*

(i) Any $\nu \in \mathcal{V}$ has exactly k non-zero entries.

(ii) For any $\nu, \nu' \in \mathcal{V}$ with $\nu \neq \nu'$, the ℓ_1-norm is lower bounded as $\|\nu - \nu'\|_1 \geq \max\{\lfloor k/4 \rfloor, 1\}$.

(iii) The set $\mathcal{V}$ has cardinality $\mathrm{card}(\mathcal{V}) \geq (d/k)^{c_1 k}$.

(iv) For a random vector V uniformly distributed over $\mathcal{V}$, we have

$$\mathrm{Cov}(V) \preceq c_2 \frac{k}{d} I_{d\times d}.$$

The proof of Lemma 2 is based on the probabilistic method [1]: we show that a certain randomized procedure generates such a packing with strictly positive probability. Along the way, we use matrix Bernstein inequalities [23] and some approximation-theoretic ideas developed by Kühn [22]. We provide details in Appendix A.

We now construct a suitable packing of the the unit simplex Δ_d. Given an integer $k \in \{1,\ldots,d\}$, consider the packing $\mathcal{V} \subset \{0,1\}^d$ given by Lemma 2. For a fixed $\delta \in [0,1]$, consider the following family of vectors in $\mathbb{R}^d$

$$\theta_\nu := \frac{\delta}{k}\nu + \frac{(1-\delta)}{d}\mathbb{1}, \quad \text{for each } \nu \in \mathcal{V}.$$

By inspection, each of these vectors belongs to the d-variate probability simplex (i.e., satisfies $\langle \mathbb{1}, \theta_\nu \rangle = 1$ and $\theta_\nu \geq 0$). Moreover, since the vector $\nu - \nu'$ can have at most $2k$ non-zero entries, we have $\|\nu - \nu'\|_1 \leq \sqrt{2k}\|\nu - \nu'\|_2$. Combined with property (ii), we conclude that for universal constants $c, c' > 0$

$$\left\|\nu - \nu'\right\|_2 \geq \frac{\|\nu - \nu'\|_1}{\sqrt{2k}} \geq c' k \frac{1}{\sqrt{2k}} = c\sqrt{k}.$$

By the definition of θ_ν, we then have for a universal constant c that

$$\|\theta_\nu - \theta_{\nu'}\|_2^2 = \frac{\delta^2}{k^2}\left\|\nu - \nu'\right\|_2^2 \geq c\frac{\delta^2}{k}. \tag{23}$$

Upper bounding the mutual information: Our next step is to upper bound the mutual information $I(Z_1, \ldots, Z_n; V)$. Recall the definition of the linear functionals φ_ν from Proposition 1. Since $\mathcal{X} = \{1, 2, \ldots, d\}$, any element of $L^\infty(\mathcal{X})$ may be identified with a vector $\gamma \in \mathbb{R}^d$. Following this identification, we have

$$\varphi_\nu(\gamma) = \sum_{j=1}^{d} \theta_{\nu,j}\gamma_j - \frac{1}{|\mathcal{V}|}\sum_{\nu' \in \mathcal{V}}\sum_{j=1}^{d} \theta_{\nu',j}\gamma_j = \frac{\delta}{k}\left\langle \gamma, \nu - \mathbb{E}[V]\right\rangle,$$

where V is a random variable distributed uniformly over $\mathcal{V}$. As a consequence, we have

$$\frac{1}{|\mathcal{V}|}\sum_{\nu \in \mathcal{V}} (\varphi_\nu(\gamma))^2 = \frac{\delta^2}{k^2}\gamma^\top \operatorname{Cov}(V)\gamma \leq c_2 \frac{\delta^2}{dk}\|\gamma\|_2^2,$$

where the final inequality follows from Lemma 2(iv). For any $\gamma \in \mathcal{G}_\alpha$, we have the upper bound $\|\gamma\|_2^2 \leq d(e^\alpha - e^{-\alpha})^2/4$, whence

$$\sup_{\gamma \in \mathcal{G}_\alpha} \frac{1}{|\mathcal{V}|}\sum_{\nu \in \mathcal{V}} (\varphi_\nu(\gamma))^2 \leq c\frac{(e^\alpha - e^{-\alpha})^2\delta^2}{k},$$

for some universal constant c. Consequently, by applying the information inequality (8), we conclude that there is a universal constant constant C such that

$$I(Z_1, \ldots, Z_n; V) \leq C\frac{n\alpha^2\delta^2}{k} \qquad \text{for all } \alpha \leq 1/4. \tag{24}$$

Applying testing inequalities: The final step is to lower bound the testing error. Since the vectors $\{\theta_\nu, \nu \in \mathcal{V}\}$ are $c/\sqrt{k}$-separated in the ℓ_2-norm (23) and Lemma 2 implies $\operatorname{card}(\mathcal{V}) \geq (d/k)^{c_2 k}$ for a constant c_2, Fano's inequality (6b) implies

$$\mathfrak{M}_n\left(\Delta_d, \|\cdot\|_2^2, \alpha\right) \geq c_0\frac{\delta^2}{k}\left(1 - \frac{c_1 n\delta^2\alpha^2/k + \log 2}{c_2 k\log(d/k)}\right), \tag{25}$$

for universal constants c_0, c_1, c_2. We split the remainder of our analysis into cases, depending on the values of (k, d).

Case 1: First, suppose that (k, d) are large enough to guarantee that

$$c_2 k\log(d/k) \geq 3\log 2. \tag{26}$$

In this case, if we set

$$\delta^2 = \min\left\{1, \frac{c_2 k^2}{2c_1 n\alpha^2}\log\frac{d}{k}\right\},$$

then we have

$$\begin{aligned} 1 - \frac{c_1 n\delta^2\alpha^2/k + \log 2}{c_2 k\log(d/k)} &\geq 1 - \frac{\frac{c_1 n\alpha^2}{k}\frac{c_2 k^2}{2c_1 n\alpha^2}\log\frac{d}{k} + \log 2}{c_2 k\log\frac{d}{k}} \\ &= 1 - \frac{\frac{1}{2}c_2 k\log\frac{d}{k} + \log 2}{c_2 k\log\frac{d}{k}} \geq 1 - \frac{1}{2} - \frac{1}{3} = \frac{1}{6}. \end{aligned}$$

Combined with the Fano lower bound (25), this yields the claim (9) under condition (26).

Case 2: Alternatively, when d is small enough that condition (26) is violated, we instead apply Le Cam's inequality (6a) to a two-point hypothesis. For our purposes, it suffices to consider the case $d = 2$, since for the purposes of lower bounds, any higher-dimensional problem is at least as hard as this case. Define the two vectors

$$\theta_1 = \frac{1+\delta}{2} e_1 + \frac{1-\delta}{2} e_2 \quad \text{and} \quad \theta_2 = \frac{1-\delta}{2} e_1 + \frac{1+\delta}{2} e_2.$$

By construction, each of these vectors belongs to the probability simplex in $\mathbb{R}^2$, and moreover, we have $\|\theta_1 - \theta_2\|_2^2 = 2\delta^2$. Letting P_j denote the multinomial distribution defined by θ_j, we also have $\|P_1 - P_2\|_{\text{TV}} = \|\theta_1 - \theta_2\|_1 / 2 = \delta$.

In terms of the marginal measures M_ν^n defined in equation (4), Pinsker's inequality (e.g. [26, Lemma 2.5]) implies that

$$\|M_1^n - M_2^n\|_{\text{TV}} \leq \sqrt{D_{\text{kl}}\left(M_1^n \| M_2^n\right)/2}.$$

Combined with Le Cam's inequality (6a) and the upper bound on KL divergences from Proposition 1, we find that the minimax risk is lower bounded as

$$\mathfrak{M}_n(\Delta_d, \|\cdot\|_2^2, \alpha) \geq \frac{\delta^2}{2}\left(\frac{1}{2} - \frac{1}{2}\sqrt{2(e^\alpha - 1)^2 n \, \|P_1 - P_2\|_{\text{TV}}^2}\right),$$

where P_i denotes the multinomial probability associated with the vector θ_i. Since $\|P_1 - P_2\|_{\text{TV}} = \delta$ by construction and $e^\alpha - 1 \leq (5/4)\alpha$ for all $\alpha \in [0, 1/4]$, we have

$$\frac{\delta^2}{2}\left(\frac{1}{2} - \frac{1}{2}\sqrt{25 n \alpha^2 \delta^2 / 8}\right) = \frac{\delta^2}{2}\left(\frac{1}{2} - \frac{5}{4\sqrt{2}}\sqrt{n}\alpha\delta\right).$$

Choosing $\delta = \min\{1, \sqrt{2}/(5\sqrt{n\alpha^2})\}$ guarantees that $\frac{1}{2} - 5\sqrt{n}\alpha\delta/4\sqrt{2} \geq 1/4$, and hence

$$\mathfrak{M}_n(\Delta_d, \|\cdot\|_2^2, \alpha) \geq \frac{1}{8}\min\left\{1, \frac{2}{25 n \alpha^2}\right\},$$

which completes the proof of Theorem 1.

Proof of inequality (10) We conclude by proving inequality (10). We distinguish three cases:

$$\text{(i) } n\alpha^2 < \log d, \quad \text{(ii) } \log d \leq n\alpha^2 \leq \frac{1}{2}d^2, \quad \text{(iii) } n\alpha^2 \geq \frac{1}{2}d^2.$$

In case (i), by taking $k = 1$ in the lower bound (9) we obtain the lower bound 1. In case (ii), we set $k = \sqrt{n\alpha^2} \in [\sqrt{\log d}, d/\sqrt{2}]$, and we obtain

$$\min\left\{\frac{1}{k}, \frac{k \log \frac{d}{k}}{n\alpha^2}\right\} = \min\left\{\frac{1}{\sqrt{n\alpha^2}}, \frac{\log d}{\sqrt{n\alpha^2}} - \frac{\frac{1}{2}\log(n\alpha^2)}{\sqrt{n\alpha^2}}\right\} \geq \frac{\log d}{\sqrt{n\alpha^2}} - \frac{\frac{1}{2}(2\log d + \log\frac{1}{2})}{\sqrt{n\alpha^2}} = \frac{\log 2}{2\sqrt{n\alpha^2}}.$$

In the final case (iii), choosing $k = d/2$ yields the bound $\min\{2/d, d\log 2/(n\alpha^2)\} \geq d\log 2/(n\alpha^2)$.

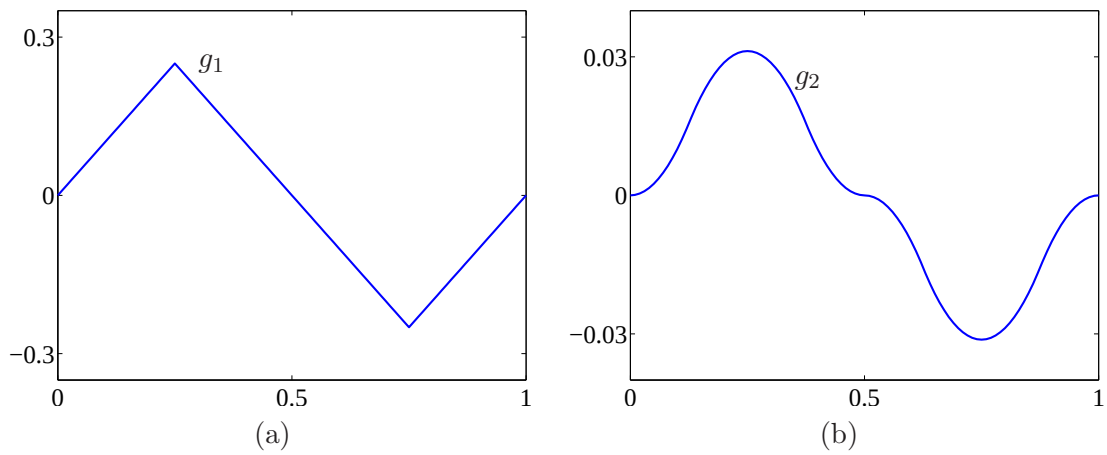

Figure 1. Panel (a): illustration of 1-Lipschitz continuous bump function g_1 used to pack $\mathcal{F}_\beta$ when $\beta = 1$. Panel (b): bump function g_2 with $|g_2''(x)| \leq 1$ used to pack $\mathcal{F}_\beta$ when $\beta = 2$.

6 Proofs of Density Estimation Results

In this section, we provide the proofs of the results stated in Section 4 on density estimation. We defer the proofs of more technical results to the appendices. Throughout all proofs, we use c to denote a universal constant whose value may change from line to line.

6.1 Proof of Theorem 2

As with our previous proof, the argument follows the general outline described at the beginning of Section 5. We remark that our proof is based on a local packing technique, a more classical approach than the metric entropy approach developed by Yang and Barron [29]. We do so because in the setting of local differential privacy we do not expect that global results on metric entropy will be generally useful; rather, we must carefully construct our packing set to control the mutual information, relating the geometry of the packing to the actual information communicated. In comparison with our proof of Theorem 1, 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 Proposition 1, is not immediately obvious. In all cases, we use the trigonometric basis to prove our lower bounds, so we may work directly with smooth density functions f.

Constructing a good packing: We begin by describing the collection of functions we use to prove our lower bound. Our construction and identification of density functions by vectors is essentially standard [26], but we specify some necessary conditions that we use later. First, let g_β be a function defined on $[0, 1]$ satisfying the following properties:

(a) The function g_β is β-times differentiable, and

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

(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

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

(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

$$G_1(A) = \int_{A\cap[0,1/2]} g_\beta(x)dx \quad \text{and} \quad G_2(A) = -\int_{A\cap[1/2,1]} g_\beta(x)dx. \tag{27}$$

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

The functions g_β are smooth "bumps" that we use as pieces in our general construction; see Figure 1 for an illustration of such functions in the cases $\beta = 1$ and $\beta = 2$

Fix a positive integer k (to be specified momentarily). Our proof makes use of the following result from our previous paper [11, Lemma 7]:

Lemma 3 (Re-stated from the paper [11])**.** *There exists a packing $\mathcal{V}$ of size at least $\exp(c_0 k)$ of the hypercube $\{-1, 1\}^k$ such that*

$$\left\|\nu - \nu'\right\|_1 \geq c_1 k \quad \textit{for all } \nu \neq \nu', \textit{ and} \quad \frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}} \nu\nu^\top \preceq c_2 I_{k\times k},$$

where (c_0, c_1, c_2) are universal positive constants.

We now make use of this packing of the hypercube in order to construct a packing of our density class. For each $j \in \{1, \ldots, k\}$, define the function

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

Based on this definition, we define the family of densities

$$\left\{ f_\nu := 1 + \sum_{j=1}^k \nu_j g_{\beta,j} \quad \text{for } \nu \in \mathcal{V} \right\} \subseteq \mathcal{F}_\beta. \tag{28}$$

It is a standard fact [30, 26] that for any $\nu \in \mathcal{V}$, the function f_ν is β-times differentiable, satisfies $|f^{(\beta)}(x)| \leq 1$ for all x, and $\|f_\nu - f_{\nu'}\|_2^2 \geq ck^{-2\beta}$. Consequently, the class (28) is a $(ck^{-\beta})$-packing of $\mathcal{F}_\beta$ with cardinality at least $\exp(c_0 k)$.

Controlling the operator norm of the packing: Having constructed a suitable packing of the space $\mathcal{F}_\beta$, we now turn to bounding the mutual information associated with a certain multi-way hypothesis testing problem. Suppose that an index V is drawn uniformly at random from $\mathcal{V}$, and conditional on $V = \nu$, the data points X_i are drawn i.i.d. according to the density f_ν. The data $\{X_1, \ldots, X_n\}$ are then passed through an α-locally private distribution Q, yielding the perturbed quantities $\{Z_1, \ldots, Z_n\}$. The following lemma bounds the mutual information between the random index V and the outputs Z_i.

Lemma 4. *There exists a universal constant c such that for any α-locally private (1) conditional distribution Q, the mutual information is upper bounded as*

$$I(Z_1, \ldots, Z_n; V) \leq n \frac{c\alpha^2}{k^{2\beta+1}}.$$

The proof of this claim is fairly involved, so we defer it to Appendix B. We remark, however, that standard mutual information bounds [30, 26] show $I(Z_1, \ldots, Z_n; V) \lesssim n/k^{2\beta}$; our bound is thus essentially a factor of the "dimension" k tighter.

Applying testing inequalities: The remainder of the proof is an application of Fano's inequality. In particular, we apply Lemma 1 with our $k^{-\beta}$ packing of $\mathcal{F}_\beta$ in $\|\cdot\|_2$ of size $\exp(c_0 k)$, and we find that for any α-locally private channel Q, there are universal constants c_0, c_1, c_2 such that

$$\mathfrak{M}_n\left(\mathcal{F}_\beta, \|\cdot\|_2^2, Q\right) \geq \frac{c_0}{k^{2\beta}}\left(1 - \frac{I(Z_{1:n}; V) + \log 2}{c_1 k}\right) \geq \frac{c_0}{k^{2\beta}}\left(1 - \frac{c_2 n\alpha^2 k^{-2\beta-1} + \log 2}{c_1 k}\right).$$

Choosing $k_{n,\alpha,\beta} = \left(2c_2 n\alpha^2\right)^{\frac{1}{2\beta+2}}$ ensures that the quantity inside the parentheses is a strictly positive constant. As a consequence, there are universal constants $c, c' > 0$ such that

$$\mathfrak{M}_n\left(\mathcal{F}_\beta, \|\cdot\|_2^2, \alpha\right) \geq \frac{c}{k_{n,\alpha,\beta}^{2\beta}} = c'\left(n\alpha^2\right)^{-\frac{2\beta}{2\beta+2}}$$

as claimed.

6.2 Proof of Proposition 2

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 $\widehat{\theta}$ for the unprojected version. Using this notation, we have

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

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

$$\begin{aligned}\mathbb{E}\left[\left\|\widehat{f} - f\right\|_2^2\right] &\leq \sum_{j=1}^k \mathbb{E}_f\left[\int_{\frac{j-1}{k}}^{\frac{j}{k}}\left(f(x) - \frac{k}{n}\sum_{i=1}^n [\mathsf{e}_k(X_i)]_j\right)^2 dx\right] + \sum_{j=1}^k \int_{\frac{j-1}{k}}^{\frac{j}{k}} \mathbb{E}\left[\left(\frac{k}{n}\sum_{i=1}^n W_{ij}\right)^2\right] \\ &= \sum_{j=1}^k \int_{\frac{j-1}{k}}^{\frac{j}{k}} \mathbb{E}_f\left[\left(f(x) - \frac{k}{n}\sum_{i=1}^n [\mathsf{e}_k(X_i)]_j\right)^2\right] dx + k\,\frac{1}{k}\,\frac{4k^2}{n\alpha^2}. \end{aligned} \tag{29}$$

We bound the error term inside the expectation (29). Defining $p_j := \mathbb{P}_f(X \in \mathcal{X}_j) = \int_{\mathcal{X}_j} f(x)dx$, we have

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

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 (29), 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.

6.3 Proof of Proposition 3

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 (21). 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.

6.4 Proof of Observation 1

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 $f \geq 0$.)

7 Discussion

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 bounds in Proposition 1, shows that applying differentially private sampling schemes essentially acts as a contraction on distributions, and we think that such results may be more generally applicable. In this paper particularly, we showed how to apply our divergence bounds to obtain sharp bounds on the convergence rate for certain nonparametric problems in addition to standard finite-dimensional settings. With our earlier paper [11], we have developed a set of techniques that show 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 Proposition 1 provides is strong.

By providing (to our knowledge, the first) sharp convergence rates for many standard statistical inference procedures under local differential privacy, we have developed and explored some tools that may be used to better understand privacy-preserving statistical inference and estimation procedures. 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. Formally identifying this tradeoff in other statistical problems should allow us to better understand the costs and benefits of privacy; we believe we have laid some of the groundwork for doing so.

Acknowledgments

We thank Guy Rothblum for very helpful discussions. JCD was supported by a Facebook Graduate Fellowship and an NDSEG fellowship. Our work was supported in part by the U.S. Army Research Laboratory, U.S. Army Research Office under grant number W911NF-11-1-0391, and Office of Naval Research MURI grant N00014-11-1-0688.

A Proof of Lemma 2

In the regime $k \in (d/2, d]$, the statement of lemma follows Lemma 7 in Duchi et al. [11]; consequently, we prove the claim for $k \le d/2$. If $k \in \{1, 2, 3, 4\}$, taking $\mathcal{V} = \{\nu \in H_d : \|\nu\|_1 = k\}$ implies that $\|\nu - \nu'\|_1 \ge k/4$ for $\nu \ne \nu'$, $\mathrm{card}(\mathcal{V}) = \binom{d}{k} \ge (d/k)^{ck}$ for some constant $c > 0$, and

$$\mathrm{Cov}(V) = \left(\frac{k}{d} - \frac{k^2}{d^2}\right) I_{d\times d},$$

from which the claim follows.

Accordingly, we focus on $k \in (4, d/2]$. To further simplify the analysis, we claim it suffices to establish the claim in the case that $k/4$ is integral (i.e., $k \in 4\mathbb{N}$). Indeed, assume that the result holds for all such integers. Given some $k \notin 4\mathbb{N}$, we may consider a packing $\mathcal{V}'$ of the binary hypercube $H_{d'}$ with $d' = d - (k - 4\lfloor k/4 \rfloor)$ and $\|\nu\|_1 = k' = 4\lfloor k/4 \rfloor$ for $\nu \in \mathcal{V}'$. By assumption, there is a packing $\mathcal{V}'$ of $H_{d'}$ satisfying the lemma. Now to each vector $\nu \in \mathcal{V}'$, we concatenate the $(k - 4\lfloor k/4 \rfloor)$-vector $\mathbb{1}$, which gives $[\nu^\top \ \mathbb{1}^\top]^\top \in \{0,1\}^d$ and $\left\|[\nu^\top \ \mathbb{1}^\top]\right\|_1 = k$. This concatenation does not increase $\mathrm{Cov}(V)$—the last $k - 4\lfloor k/4 \rfloor$ coordinates have covariance zero—and the rest of the terms in items (i)–(iv) incur only constant factor changes.

It remains to prove the claim for $k \in 4\mathbb{N}$ over the range $\{5, \ldots, \lfloor d/2 \rfloor\}$. To ease notation, we let $\ell = k/4$ belong to the interval $[2, d/8]$. Our proof is based on the probabilistic method [1]: we propose a random construction of a packing, and show that it satisfies the desired properties with strictly positive probability. Our random construction is straightforward: letting $H_d = \{0,1\}^d$ denote the Boolean hypercube, we sample K i.i.d. random vectors U_i from the uniform distribution over the set

$$S_\ell := \{\nu \in H_d \mid \|\nu\|_1 := 4\ell\}. \tag{30}$$

We claim that for $K = (d/(6\ell))^{3\ell/2}$, the resulting random set $\mathcal{U}_K := \{U_1, \ldots, U_K\}$ satisfies the claimed properties with non-zero probability. We say that $\mathcal{U}_K$ is ℓ-separated if $\|U_i - U_j\|_1 > \ell$ for all $i \ne j$, and we use $\mathrm{Cov}(\mathcal{U}_K)$ to denote the covariance of a random vector V drawn uniformly at random from $\mathcal{U}_K$. Our proof is based on the following two tail bounds, which we prove shortly: for a universal constant $c < \infty$,

$$\mathbb{P}\left[\mathcal{U}_K \text{ is not } \ell\text{-separated}\right] \le \binom{K}{2}\left(6\,\frac{\ell}{d}\right)^{3\ell}, \qquad \text{and} \tag{31a}$$

$$\mathbb{P}\left[\lambda_{\max}\big(\mathrm{Cov}(\mathcal{U}_K)\big) \ge t\right] \le d\exp\left(-\frac{Kt^2}{3c\max\{\ell, \ell^3/d\} + ct\ell}\right) \qquad \text{for all } t > 0. \tag{31b}$$

For the moment, let us assume the validity of these bounds and use them to complete the proof. By the union bound, we have

$$\mathbb{P}\left(\mathcal{U}_K \text{ is not } \ell\text{-separated or } \mathrm{Cov}(\mathcal{U}_K) \not\preceq tI\right) \le \binom{K}{2}\left(6\,\frac{\ell}{d}\right)^{3\ell} + d\exp\left(-\frac{Kt^2}{3c\max\{\ell, \ell^3/d\} + ct\ell}\right).$$

By choosing $t = C\ell/d$ and recalling that $K = (d/(6\ell))^{3\ell/2}$, we obtain the bound

$$\frac{1}{2} + d\exp\left(-C^2\frac{\ell^2(d/(6\ell))^{3\ell/2}}{3c\max\{d^2\ell, d\ell^3\} + Ccd\ell^2}\right).$$

If $\ell \geq \ell^3/d$, the second term can be easily seen to be less than $\frac{1}{2}$ for suitably large constants C, so assume that $\ell \leq \ell^3/d$. Then we have, where c is a constant whose value may change from inequality to inequality,

$$\frac{\ell^2(d/(6\ell))^{3\ell/2}}{3c\max\{d^2\ell, d\ell^3\} + Cd\ell^2} = \frac{\ell^2 d^{3\ell/2}}{(6\ell)^{3\ell/2}(3cd\ell^3 + Cd\ell^2)} \geq c\frac{d^{3\ell/2}}{(6\ell)^{3\ell/2}d\ell} \geq \frac{c}{d\ell}\left(\frac{d}{6\ell}\right)^{3\ell/2}.$$

For suitably large d and any $\ell \geq 2$, the final term is greater than $c' \log d$ for some constant $c' > 0$, which implies that with appropriate choice of the constant C earlier, we have the bound

$$\mathbb{P}\left(\mathcal{U}_K \text{ is not } \ell\text{-separated or } \operatorname{Cov}(\mathcal{U}_K) \not\preceq tI\right) < 1.$$

Consequently, recalling that $k = 4\ell$ by definition, a packing as described in the statement of the lemma must exist.

It remains to prove the tail bounds (31a) and (31b). Beginning with the former bound, define the set

$$N(\nu, \ell) := \left\{\nu' \in H_d \mid \left\|\nu - \nu'\right\|_1 \leq \ell\right\}.$$

Recalling the definition (30) S_ℓ, let U_i and U_j be sampled independently and uniformly at random from S_ℓ. Then

$$\mathbb{P}\left(\|U_i - U_j\|_1 \leq \ell\right) \leq \frac{\operatorname{card}(N(\nu, \ell))}{\operatorname{card}(S_\ell)}.$$

Note that $N(\nu, \ell)$ can be constructed by choosing an arbitrary subset $J \subset \{1, \ldots, d\}$ of size ℓ, and then setting $\nu'_j = \nu_j$ for $j \notin J$ and ν_j arbitrarily otherwise; consequently, its cardinality is upper bounded as $\operatorname{card}(N(\nu, \ell)) \leq \binom{d}{\ell}2^\ell$. Since $\operatorname{card}(S_\ell) = \binom{d}{4\ell}$, we find that

$$\frac{\operatorname{card}(N(\nu, \ell))}{\operatorname{card}(S_\ell)} = \frac{\binom{d}{\ell}2^\ell}{\binom{d}{4\ell}} = \frac{2^\ell (d-4\ell)!(4\ell)!}{(d-\ell)!\ell!} = 2^\ell \prod_{j=1}^{3\ell} \frac{\ell + j}{d - 4\ell + j} \leq 2^\ell \left(\frac{4\ell}{d-\ell}\right)^{3\ell},$$

where the final inequality follows because the function $x \mapsto h(x) = \frac{\ell + x}{d - 4\ell + x}$ is increasing for $x > 0$. Since $\ell \leq d/8$ by assumption, we arrive at the upper bound

$$\frac{\operatorname{card}(N(\nu, \ell))}{\operatorname{card}(S_\ell)} \leq 2^\ell \left(\frac{4\ell}{d-\ell}\right)^{3\ell} = \left(\frac{4 \cdot 2^{1/3}\ell}{d-\ell}\right)^{3\ell} \leq \left(6\,\frac{\ell}{d}\right)^{3\ell}.$$

Since we have to compare $\binom{K}{2}$ such pairs over the set $\mathcal{U}_K$, the claim (31a) follows from the union bound.

We now turn to establishing the claim (31b), for which we make use of matrix Bernstein inequalities. Letting U be drawn uniformly at random from S_ℓ, we have

$$\mathbb{E}[UU^\top] = \beta_{\ell,d}\mathbb{1}\mathbb{1}^\top + \left(\frac{4\ell}{d} - \beta_{\ell,d}\right) I_{d\times d}.$$

where $\beta_{\ell,d} := \binom{4\ell}{2}\binom{d}{2}^{-1}$. Consequently, the $d \times d$ random matrix

$$A := UU^\top - \beta_{\ell,d}\mathbb{1}\mathbb{1}^\top - \left(\frac{4\ell}{d} - \beta_{\ell,d}\right) I_{d\times d}.$$

is centered ($\mathbb{E}[A] = 0$), and by definition of our construction, $\mathrm{Cov}(\mathcal{U}_K) = \frac{1}{K}\sum_{i=1}^K A_i$, where the random matrices $\{A_i\}_{i=1}^K$ are drawn i.i.d.

In order to apply a matrix Bernstein inequality, it remains to bound the operator norm (maximum singular value) of A and its variance. The operator norm of A is upper bounded as

$$|||A||| \leq \left|\left|\left|UU^\top - (4\ell/d - \beta_{\ell,d})I\right|\right|\right| + \beta_{\ell,d}\left|\left|\left|\mathbb{1}\mathbb{1}^\top\right|\right|\right| = 4\ell - \frac{4\ell}{d} + \beta_{\ell,d} + d\beta_{\ell,d} \leq 5\ell.$$

Moreover, we claim that there is a universal positive constant c such that

$$\left|\left|\left|\mathbb{E}[A^2]\right|\right|\right| \leq c\max\{\ell, \ell^3/d\}. \tag{32}$$

To establish this claim, we begin by computing

$$\begin{aligned}\mathbb{E}[A^2] &= \mathbb{E}[UU^\top UU^\top] - \left(\left(\frac{4\ell}{d} - \beta_{\ell,d}\right)I_{d\times d} + \beta_{\ell,d}\mathbb{1}\mathbb{1}^\top\right)^2 \\ &= 4\ell\left(\left(\frac{4\ell}{d} - \beta_{\ell,d}\right)I_{d\times d} + \beta_{\ell,d}\mathbb{1}\mathbb{1}^\top\right) - \left(\left(\frac{4\ell}{d} - \beta_{\ell,d}\right)I_{d\times d} + \beta_{\ell,d}\mathbb{1}\mathbb{1}^\top\right)^2.\end{aligned}$$

Consequently, if we define the constants,

$$a_{\ell,d} := \left(4\ell - \frac{4\ell}{d} + \beta_{\ell,d}\right) \quad \text{and} \quad b_{\ell,d} := \left(4\ell\beta_{\ell,d} - \frac{8\ell\beta_{\ell,d}}{d} + 2\beta_{\ell,d}^2 - d\beta_{\ell,d}^2\right),$$

then $\mathbb{E}[A^2] = a_{\ell,d}I_{d\times d} + b_{\ell,d}\mathbb{1}\mathbb{1}^\top$. It is easy to see that $|a_{\ell,d}| \leq 4\ell$ and that $|b_{\ell,d}| \leq c'\frac{\ell^3}{d^2}$ for some universal constant c', from which the intermediate claim (32) follows. With these pieces in place, the claimed tail bound (31b) follows a matrix Bernstein inequality (e.g., [23, Corollary 5.2]), applied to the quantity $\mathrm{Cov}(\mathcal{U}_K) = \frac{1}{K}\sum_{i=1}^K A_i$.

B Proof of Lemma 4

This result relies on inequality (8) from Proposition 1, along with a careful argument to understand the extreme points of $\gamma \in L^\infty([0,1])$ that we use when applying the proposition. First, we take a packing $\mathcal{V}$ as guaranteed by Lemma 3 and consider densities f_ν for $\nu \in \mathcal{V}$. Overall, our first step is to show for the purposes of applying inequality (8), 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 use the packing set $\mathcal{V}$ from Lemma 3 to provide a bound on the correlation of any $\gamma \in L^\infty([0,1])$ with the densities f_ν, 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. Recalling the linear functionals φ_ν in Proposition 1, we have $\varphi_\nu : L^\infty([0,1]) \to \mathbb{R}$ defined via

$$\varphi_\nu(\gamma) = \sum_{i=1}^{2k}\int_{D_i}\gamma(x)(f_\nu(x) - \overline{f}(x))dx, = \sum_{i=1}^{2k}\int_{D_i}\gamma(x)(f_\nu(x) - 1 - (\overline{f}(x) - 1))dx,$$

where $\overline{f} = (1/|\mathcal{V}|)\sum_{\nu\in\mathcal{V}} f_\nu$. Expanding the square, we find that since $\overline{f}$ is the average, we have

$$\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\varphi_\nu(\gamma)^2 \le \frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\left(\sum_{i=1}^{2k}\int_{D_i}\gamma(x)(f_\nu(x)-1)dx\right)^2.$$

Since the set $\mathcal{G}_\alpha$ from Proposition 1 is compact, convex, and Hausdorff, the Krein-Milman theorem [24, Proposition 1.2] guarantees that it is equal to the convex hull of its extreme points; moreover, since the functionals $\gamma \mapsto \varphi_\nu^2(\gamma)$ are convex, the supremum in Proposition 1 must be attained at the extreme points of $\mathcal{G}_\alpha$. As a consequence, when applying the information bound

$$I(Z_1,\ldots,Z_n;V) \le nC_\alpha\frac{1}{|\mathcal{V}|}\sup_{\gamma\in\mathcal{G}_\alpha}\sum_{\nu\in\mathcal{V}}\varphi_\nu^2(\gamma), \tag{33}$$

we can restrict our attention to $\gamma \in L^\infty([0,1])$ for which $\gamma(x) \in \{e^{-\alpha}-e^{\alpha}, e^{\alpha}-e^{-\alpha}\}/2$.

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). Using μ to denote Lebesgue measure, define the shorthand $\kappa = (e^\alpha - e^{-\alpha})/2$. Fix $i \in [2k]$, and assume for the sake of contradiction that there exist sets $B_i, C_i \subset D_i$ such that $\gamma(B_i) = \{\kappa\}$ and $\gamma(C_i) = \{-\kappa\}$, while $\mu(B_i) > 0$ and $\mu(C_i) > 0$.[1] We will construct vectors γ_1 and $\gamma_2 \in \mathcal{G}_\alpha$ and a value $\lambda \in (0,1)$ such that

$$\int_{D_i}\gamma(x)(f_\nu(x)-1)dx = \lambda\int_{D_i}\gamma_1(x)(f_\nu(x)-1)dx + (1-\lambda)\int_{D_i}\gamma_2(x)(f_\nu(x)-1)dx$$

simultaneously for all $\nu \in \mathcal{V}$, 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 $\gamma_1(D_i) = \{\kappa\}$ and $\gamma_2(D_i) = \{-\kappa\}$, otherwise setting $\gamma_1(x) = \gamma_2(x) = \gamma(x)$ for $x \notin D_i$. We define

$$\lambda := \frac{\int_{B_i}(f_\nu(x)-1)dx}{\int_{D_i}(f_\nu(x)-1)dx} \quad \text{so} \quad 1-\lambda = \frac{\int_{C_i}(f_\nu(x)-1)dx}{\int_{D_i}(f_\nu(x)-1)dx}.$$

By the construction of the function g_β, the function $f_\nu - 1$ does not change signs on D_i, and the absolute continuity conditions on g_β specified in equation (27) guarantee $1 > \lambda > 0$, since $\mu(B_i) > 0$ and $\mu(C_i) > 0$. Moreover, the quantity λ is constant for all ν by the construction of the f_ν, since $B_i \subset D_i$ and $C_i \subset D_i$. We thus find that for any $\nu \in \mathcal{V}$,

$$\begin{aligned}
&\int_{D_i}\gamma(x)(f_\nu(x)-1)dx = \int_{B_i}\gamma_1(x)(f_\nu(x)-1)dx + \int_{C_i}\gamma_2(x)(f_\nu(x)-1)dx \\
&= \kappa\int_{B_i}(f_\nu(x)-1)dx - \kappa\int_{C_i}(f_\nu(x)-1)dx = \kappa\lambda\int_{D_i}(f_\nu(x)-1)dx - \kappa(1-\lambda)\int_{D_i}(f_\nu(x)-1)dx \\
&= \lambda\int\gamma_1(x)(f_\nu(x)-1)dx + (1-\lambda)\int\gamma_2(x)(f_\nu(x)-1)dx.
\end{aligned}$$

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

By linearity and the strong convexity of the function $x \mapsto x^2$, then, we find that

$$\sum_{\nu \in \mathcal{V}} \left(\sum_{i=1}^{2k} \int_{D_i} \gamma(x)(f_\nu(x) - 1)dx \right)^2$$
$$< \lambda \sum_{\nu \in \mathcal{V}} \left(\sum_{i=1}^{2k} \int_{D_i} \gamma_1(x)(f_\nu(x) - 1)dx \right)^2 + (1-\lambda) \sum_{\nu \in \mathcal{V}} \left(\sum_{i=1}^{2k} \int_{D_i} \gamma_2(x)(f_\nu(x) - 1)dx \right)^2.$$

Thus one of the densities γ_i, $i \in \{1, 2\}$ 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 (33) 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 (33) can be reduced to a finite-dimensional problem over the subset

$$\mathcal{G}_{\alpha,2k} := \left\{ u \in \mathbb{R}^{2k} \mid \|u\|_\infty \le \frac{e^\alpha - e^{-\alpha}}{2} \right\}$$

of $\mathbb{R}^{2k}$. In terms of this subset, we have the upper bound

$$\frac{|\mathcal{V}|}{C_\alpha n} I(Z_1, \ldots, Z_n; V) \le \sup_{\gamma \in \mathcal{G}_\alpha} \sum_{\nu \in \mathcal{V}} \varphi_\nu(\gamma)^2 \le \sup_{\gamma \in \mathcal{G}_{\alpha,2k}} \sum_{\nu \in \mathcal{V}} \left(\sum_{i=1}^{2k} \gamma_i \int_{D_i} (f_\nu(x) - 1)dx \right)^2.$$

By construction of the f_ν and g_β, we have the equality

$$\int_{D_i} (f_\nu(x) - 1)dx = (-1)^{i+1} \nu_i \int_0^{\frac{1}{2k}} g_{\beta,1}(x)dx = (-1)^{i+1} \nu_i \int_0^{\frac{1}{2k}} \frac{1}{k^\beta} g(kx)dx = (-1)^{i+1} \nu_i \frac{c_{1/2}}{k^{\beta+1}},$$

which implies that

$$\begin{aligned} \frac{|\mathcal{V}|}{C_\alpha n} I(Z_1, \ldots, Z_n; V) &\le \sup_{\gamma \in \mathcal{G}_{\alpha,2k}} \sum_{\nu \in \mathcal{V}} \left(\frac{c_{1/2}}{k^{\beta+1}} \gamma^\top \left(\begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \nu \right) \right)^2 \\ &= \frac{c_{1/2}^2}{k^{2\beta+2}} \sup_{\gamma \in \mathcal{G}_{\alpha,2k}} \gamma^\top \left(\sum_{\nu \in \mathcal{V}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \nu\nu^\top \otimes \begin{bmatrix} 1 \\ -1 \end{bmatrix}^\top \right) \gamma, \end{aligned} \tag{34}$$

where $\otimes$ denotes the Kronecker product. By our construction of the packing $\mathcal{V}$ of $\{-1, 1\}^k$, there exists a constant c such that $(1/|\mathcal{V}|) \sum_{\nu \in \mathcal{V}} \nu\nu^\top \preceq c I_{k \times k}$. Moreover, observe that the mapping

$$A \mapsto \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes A \otimes \begin{bmatrix} 1 \\ -1 \end{bmatrix}^\top \quad \text{satisfies} \quad \begin{bmatrix} x \\ y \end{bmatrix}^\top \left(\begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes A \otimes \begin{bmatrix} 1 \\ -1 \end{bmatrix}^\top \right) \begin{bmatrix} x \\ y \end{bmatrix} = (x-y)^\top A (x-y),$$

whence it is operator monotone ($A \succeq B$ implies $(x-y)^\top A(x-y) \ge (x-y)^\top B(x-y)$). Consequently, by linearity of the Kronecker product $\otimes$ and Lemma 3, there is a universal constant c such that

$$\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \nu\nu^\top \otimes \begin{bmatrix} 1 \\ -1 \end{bmatrix}^\top = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \left(\frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} \nu\nu^\top \right) \otimes \begin{bmatrix} 1 \\ -1 \end{bmatrix}^\top \preceq c I_{2k \times 2k}.$$

Combining this bound with our inequality (34), we obtain

$$I(Z_1, \ldots, Z_n; V) \leq n \frac{c}{k^{2\beta+2}} \sup_{\gamma \in \mathcal{G}_{\alpha,2k}} \gamma^\top I \gamma = n \frac{c(e^\alpha - e^{-\alpha})^2 k}{2k^{2\beta+2}}$$

for some universal numerical constant c. Since $\alpha \in [0, 1/4]$, we have $(e^\alpha - e^{-\alpha})^2 \leq c'\alpha^2$, which completes the proof.

References

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

[2] 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.

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

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

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

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

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

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

[9] A. De. Lower bounds in differential privacy. In *Proceedings of the Ninth Theory of Cryptography Conference*, 2012.

[10] 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.

[11] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. *arXiv:1302.3203 [math.ST]*, 2013. URL http://arxiv.org/abs/1302.3203.

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

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

[14] C. Dwork. Differential privacy: a survey of results. In *Theory and Applications of Models of Computation*, volume 4978 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2008.

[15] 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.

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

[17] 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.

[18] I. P. Fellegi. On the question of statistical confidentiality. *Journal of the American Statistical Association*, 67(337):7–18, 1972.

[19] 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.

[20] 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.

[21] 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.

[22] T. Kühn. A lower estimate for entropy numbers. *Journal of Approximation Theory*, 110: 120–124, 2001.

[23] 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.

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

[25] 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.

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

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

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

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

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