

Extremal Mechanisms for Local Differential Privacy

Peter Kairouz

KAIROUZ2@ILLINOIS.EDU

*Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA*

Sewoong Oh

SWOH@ILLINOIS.EDU

*Department of Industrial and Enterprise Systems Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61820, USA*

Pramod Viswanath

PRAMODV@ILLINOIS.EDU

*Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA*

Editor: Mehryar Mohri

Abstract

Local differential privacy has recently surfaced as a strong measure of privacy in contexts where personal information remains private even from data analysts. Working in a setting where both the data providers and data analysts want to maximize the utility of statistical analyses performed on the released data, we study the fundamental trade-off between local differential privacy and utility. This trade-off is formulated as a constrained optimization problem: maximize utility subject to local differential privacy constraints. We introduce a combinatorial family of extremal privatization mechanisms, which we call staircase mechanisms, and show that it contains the optimal privatization mechanisms for a broad class of information theoretic utilities such as mutual information and f -divergences. We further prove that for any utility function and any privacy level, solving the privacy-utility maximization problem is equivalent to solving a finite-dimensional linear program, the outcome of which is the optimal staircase mechanism. However, solving this linear program can be computationally expensive since it has a number of variables that is exponential in the size of the alphabet the data lives in. To account for this, we show that two simple privatization mechanisms, the binary and randomized response mechanisms, are universally optimal in the low and high privacy regimes, and well approximate the intermediate regime.

Keywords: local differential privacy, privacy-preserving machine learning algorithms, information theoretic utilities, f -divergences, mutual information, statistical inference, hypothesis testing, estimation

1. Introduction

In statistical analyses involving data from individuals, there is an increasing tension between the need to share data and the need to protect sensitive information about the individuals. For example, users of social networking sites are increasingly cautious about their privacy, but still find it inevitable to agree to share their personal information in order to benefit

from customized services such as recommendations and personalized search (Acquisti, 2004; Acquisti and Grossklags, 2007). There is a certain utility in sharing data for both data providers and data analysts, but at the same time, individuals want *plausible deniability* when it comes to sensitive information.

For such applications, there is a natural core optimization problem to be solved. Assuming both the data providers and analysts want to maximize the utility of the released data, how can they do so while preserving the privacy of participating individuals? The formulation and study of a framework that addresses the fundamental tradeoff between utility and privacy is the focus of this paper.

1.1 Local differential privacy

The need for data privacy appears in two different contexts: the *local privacy* context, as in when individuals disclose their personal information (e.g., voluntarily on social network sites), and the *global privacy* context, as in when institutions release databases of information of several people or answer queries on such databases (e.g., US Government releases census data, companies like Netflix release proprietary data for others to test state of the art machine learning algorithms). In both contexts, privacy is achieved by *randomizing* the data before releasing it. We study the setting of local privacy, in which data providers do not trust the data collector (analyst). Local privacy dates back to Warner (1965), who proposed the *randomized response* method to provide plausible deniability for individuals responding to sensitive surveys.

A natural notion of privacy protection is making inference of information beyond what is released hard. *Differential privacy* has been proposed in the global privacy context to formally capture this notion of privacy (Dwork, 2006; Dwork et al., 2006b; Dwork and Lei, 2009). In a nutshell, differential privacy ensures that an adversary should not be able to reliably infer an individual’s record in a database, even with unbounded computational power and access to every other record in the database. Recently, Duchi et al. (2013) extended the notion of differential privacy to the local privacy context. Formally, consider a setting where there are n data providers each owning a data X_i defined on an input alphabet \mathcal{X} . The X_i ’s are independently sampled from some distribution P_ν parameterized by ν . A statistical privatization mechanism Q is a conditional distribution that maps $X_i \in \mathcal{X}$ stochastically to $Y_i \in \mathcal{Y}$, where \mathcal{Y} is an output alphabet possibly larger than \mathcal{X} . The Y_i ’s are referred to as the privatized (sanitized) views of X_i ’s. In a non-interactive setting, the same privatization mechanism Q is used locally by all individuals. This setting is shown in Figure 1 for the special case of $n = 2$. For some non-negative ε , we follow the definition of Duchi et al. (2013) and say that a mechanism Q is ε -*locally differentially private* if

$$\sup_{S \subset \mathcal{Y}, x, x' \in \mathcal{X}} \frac{Q(S|x)}{Q(S|x')} \leq e^\varepsilon, \quad (1)$$

where $Q(S|x) = \mathbb{P}(Y_i \in S | X_i = x)$ represents the privatization mechanism. This ensures that for small values of ε , given a privatized data Y_i , it is (almost) equally likely to have come from any data, i.e. x or x' . A small value of ε means that we require a high level of privacy and a large value corresponds to a low level of privacy. At one extreme, for $\varepsilon = 0$, the privatized output must be independent of the private data, and on the other extreme, for $\varepsilon = \infty$, the privatized output can be made equal to the private data.




Figure 1: Client i owns X_i sampled from P_ν . Each X_i is privatized by the same ε -locally differentially private mechanism Q . The data analyst only observes the privatized data (Y_i 's) and makes an inference on the statistics of the original distribution of the data.

1.2 Information theoretic utilities for statistical analysis

In analyses of statistical databases, the analyst is interested in the *statistics* of the data as opposed to individual records. Naturally, the utility should also be measured in terms of the distribution rather than sample quantities. Concretely, consider a client-server setting where each client with data X_i releases Y_i , a privatized version of the data, via a non-interactive ε -locally differentially private privatization mechanism Q . Assume all the clients use the same privatization mechanism Q , and each client's data is an i.i.d. sample from a distribution P_ν for some parameter ν . Given the privatized views $\{Y_i\}_{i=1}^n$, the data analyst would like to make inferences based on the induced marginal distribution

$$M_\nu(S) \equiv \sum_{x \in \mathcal{X}} Q(S|x) P_\nu(x), \quad (2)$$

for $S \subseteq \mathcal{Y}$. We consider a broad class of convex utility functions, and identify the class of optimal mechanisms, which we call *staircase mechanisms*, in Section 2. We apply this framework to two specific applications: (a) hypothesis testing where the utility is measured in f -divergences (Section 3) and (b) information preservation where the utility is measured in mutual information (Section 4).

In the binary hypothesis testing setting, $\nu \in \{0, 1\}$; therefore, X can be generated by one of two possible distributions P_0 and P_1 . The power to discriminate data generated from P_0 to data generated from P_1 depends on the ‘distance’ between the marginals M_0 and M_1 . To measure the ability of such statistical discrimination, our choice of utility of a particular privatization mechanism Q is an information theoretic quantity called Csiszár's f -divergence defined as

$$D_f(M_0||M_1) = \sum_{x \in \mathcal{X}} f\left(\frac{M_0(x)}{M_1(x)}\right) M_1(x), \quad (3)$$

for some convex function f such that $f(1) = 0$. The Kullback-Leibler (KL) divergence $D_{\text{kl}}(M_0||M_1)$ is a special case with $f(x) = x \log x$, and so is the total variation distance $\|M_0 - M_1\|_{\text{TV}}$ with $f(x) = (1/2)|x - 1|$. Such f -divergences capture the quality of statistical inference, such as minimax rates of statistical estimation or error exponents in hypothesis testing (Tsybakov and Zaiats, 2009; Cover and Thomas, 2012). As a motivating example, suppose a data analyst wants to test whether the data is generated from P_0 or P_1 based on privatized views Y_1, \dots, Y_n . According to Chernoff-Stein’s lemma, for a bounded type I error probability, the best type II error probability scales as $e^{-n D_{\text{kl}}(M_0||M_1)}$. Naturally, we are interested in finding a privatization mechanism Q that minimizes the probability of error by solving the following constraint maximization problem

$$\begin{aligned} & \underset{Q}{\text{maximize}} && D_{\text{kl}}(M_0||M_1) \\ & \text{subject to} && Q \in \mathcal{D}_\varepsilon \end{aligned}, \tag{4}$$

where \mathcal{D}_ε is the set of all ε -locally differentially private mechanisms satisfying (1).

In the information preservation setting, X is generated from an underlying distribution P . We are interested in quantifying how much information can be preserved when releasing a private view of the data. In other words, the data provider would like to release an ε -locally differentially private view Y of X that preserves the amount of information in X as much as possible. The utility in this case is measured by the mutual information between X and Y

$$I(X; Y) = \sum_x \sum_y P(x) Q(y|x) \log \left(\frac{Q(y|x)}{\sum_{l \in \mathcal{X}} P(l) Q(y|l)} \right). \tag{5}$$

Mutual information, as the name suggests, measures the mutual dependence between two random variables. It has been used as a criterion for feature selection and as a measure of similarity between two different clusterings of a data set, in addition to many other applications in signal processing and machine learning. To characterize the fundamental tradeoff between privacy and mutual information, we solve the following constrained maximization problem

$$\begin{aligned} & \underset{Q}{\text{maximize}} && I(X; Y) \\ & \text{subject to} && Q \in \mathcal{D}_\varepsilon \end{aligned}, \tag{6}$$

where \mathcal{D}_ε is the set of all ε -locally differentially private mechanisms satisfying (1).

Motivated by such applications in statistical analysis, our goal is to provide a general framework for finding optimal privatization mechanisms that maximize information theoretic utilities under local differential privacy. We demonstrate the power of our techniques in a very general setting that includes both hypothesis testing and information preservation.

1.3 Our contributions

We study the fundamental tradeoff between local differential privacy and a rich class of convex utility functions. This class of utilities includes several information theoretic quantities such as mutual information and f -divergences. The privacy-utility tradeoff is posed as a constrained maximization problem: maximize utility subject to local differential privacy constraints. This maximization problem is (a) nonlinear: the utility functions we consider

are convex in Q ; (b) non-standard: we are maximizing instead of minimizing a convex function; and (c) infinite dimensional: the space of all differentially private mechanisms is infinite dimensional. We show, in Theorem 2, that for all utility functions considered and any privacy level ε , a *finite* family of *extremal* mechanisms (a finite subset of the corner points of the space of differentially private mechanisms), which we call *staircase* mechanisms, contains the optimal privatization mechanism. We further prove, in Theorem 4, that solving the original privacy-utility problem is equivalent to solving a finite dimensional linear program, the outcome of which is the optimal mechanism. However, solving this linear program can be computationally expensive because it has $2^{|\mathcal{X}|}$ variables. To account for this, we show that two simple staircase mechanisms (the binary and randomized response mechanisms) are optimal in the high and low privacy regimes, respectively, and well approximate the intermediate regime. This contributes an important progress in the differential privacy area, where the privatization mechanisms have been few and almost no exact optimality results are known. As an application, we show that the effective sample size reduces from n to $\varepsilon^2 n$ under local differential privacy in the context of hypothesis testing.

We also study the fundamental tradeoff between utility and approximate differential privacy, a generalized notion of privacy that was first introduced in Dwork et al. (2006a). The techniques we develop for differential privacy do not generalize to approximate differential privacy. To account for this, we use an operational interpretation of approximate differential privacy (developed in Kairouz et al. (2014a)) to prove that a simple mechanism maximizes utility for all levels of privacy when the data is binary.

1.4 Related work

Our work is closely related to the recent work of Duchi et al. (2013) where an upper bound on $D_{\text{kl}}(M_0||M_1)$ was derived under the same local differential privacy setting. Precisely, Duchi et. al. proved that the KL-divergence maximization problem in (4) is at most $4(e^\varepsilon - 1)^2 \|P_1 - P_2\|_{TV}^2$. This bound was further used to provide a minimax bound on statistical estimation using information theoretic converse techniques such as Fano’s and Le Cam’s inequalities. Such tradeoffs also provide tools for comparing various notions of privacy (Barber and Duchi, 2014).

In a similar spirit, we are also interested in maximizing information theoretic quantities under local differential privacy. We generalize the results of Duchi et al. (2013), and provide stronger results in the sense that we (a) consider a broader class of information theoretic utilities; (b) provide explicit constructions for the optimal mechanisms; and (c) recover the existing result of (Duchi et al., 2013, Theorem 1) (with a stronger condition on ε).

Our work also provides a formal connection to an information theoretic notion of privacy called information leakage (Chatzikokolakis et al., 2010; Sankar et al., 2013). Given a privatization mechanism Q , the information leakage is measured by the mutual information between the private data X and the released output Y , i.e. $I(X; Y)$. Information leakage has been widely studied as a practical notion of privacy. However, connections to differential privacy have been studied only indirectly through comparisons to how much distortion is incurred under the two notions of privacy (Sarwate and Sankar, 2014; Wang et al., 2014a). We show that under ε -local differential privacy, $I(X; Y)$ is upper bounded by

$0.5\epsilon^2 \max_{A \subseteq \mathcal{X}} P(A)P(A^c) + O(\epsilon^3)$ for small ϵ . Moreover, we provide an explicit privatization mechanism that achieves this bound.

While there is a vast literature on differential privacy, exact optimality results are only known for a few cases. The typical recipe is to propose a differentially private mechanism inspired by the work of Dwork (2006); Dwork et al. (2006b); McSherry and Talwar (2007) and Hardt and Rothblum (2010), and then establish its near-optimality by comparing the achievable utility to a converse, for example in linear dynamical systems (Wang et al., 2014b), principal component analysis (Chaudhuri et al., 2012; Blocki et al., 2012; Hardt and Roth, 2012; Kapralov and Talwar, 2013), linear queries (Hardt and Talwar, 2010; Hardt et al., 2012), logistic regression (Chaudhuri and Monteleoni, 2008) and histogram release (Lei, 2011). In this paper, we take a different route and solve the utility maximization problem *exactly*.

Optimal differentially private mechanisms are known only in a few cases. Ghosh et al. (2012) showed that the geometric noise adding mechanism is optimal (under a Bayesian setting) for monotone utility functions under count queries (sensitivity one). This was generalized by Geng et. al. (for a worst-case input setting) who proposed a family of mechanisms and proved its optimality for monotone utility functions under queries with arbitrary sensitivity (Geng and Viswanath, 2012, 2013a,b). The family of optimal mechanisms was called *staircase mechanisms* because for any y and any neighboring x and x' , the ratio of $Q(y|x)$ to $Q(y|x')$ takes one of three possible values e^ϵ , $e^{-\epsilon}$, or 1. Since the optimal mechanisms we develop also have an identical property, we retain the same nomenclature.

1.5 Organization

The remainder of the paper is organized as follows. In Section 2, we introduce the family of staircase mechanisms, prove its optimality for a broad class of convex utility functions, and study its combinatorial structure. In Section 3, we study the problem of private hypothesis testing and prove that two staircase mechanisms, the binary and randomized response mechanisms, are optimal for KL-divergence in the high and low privacy regimes, respectively, and (nearly) optimal the intermediate regime. We show, in Section 4, similar results for mutual information. In Section 5, we study approximate local differential privacy, a more general notion of local privacy. Finally, we conclude this paper in Section 6 with a few interesting and nontrivial extensions.

2. Main Results

In this section, we first present a formal definition for staircase mechanisms and prove that they are the optimal solutions to optimization problems of the form (8). We then provide a combinatorial representation for staircase mechanisms that allows us to reduce the infinite dimensional nonlinear program of (8) to a finite dimensional linear program with $2^{|\mathcal{X}|}$ variables. For any given privacy level ϵ and utility function $U(\cdot)$, one can solve this linear program to obtain the optimal privatization mechanism, albeit with significant computational challenges since the number of variables scales exponentially in the alphabet size. To address this issue, we prove, in Sections 3 and 4, that two simple staircase mechanisms, which we call the binary mechanism and the randomized response mechanism, are optimal

in the high and low privacy regimes, respectively, and well approximate the intermediate regime.

2.1 Optimality of staircase mechanisms

For an input alphabet \mathcal{X} with $|\mathcal{X}| = k$, we represent the set of ε -locally differentially private mechanisms that lead to output alphabets \mathcal{Y} with $|\mathcal{Y}| = \ell$ by

$$\mathcal{D}_{\varepsilon,\ell} = \mathcal{Q}_{k \times \ell} \cap \left\{ Q : \forall x, x' \in \mathcal{X}, S \subseteq \mathcal{Y}, \left| \ln \frac{Q(S|x)}{Q(S|x')} \right| \leq \varepsilon \right\},$$

where $\mathcal{Q}_{k \times \ell}$ denotes the set of all $k \times \ell$ dimensional conditional distributions. The set of all ε -locally differentially private mechanisms is given by

$$\mathcal{D}_\varepsilon = \cup_{\ell \in \mathbb{N}} \mathcal{D}_{\varepsilon,\ell}. \tag{7}$$

The set of all conditional distributions acting on \mathcal{X} is given by $\mathcal{Q} = \cup_{\ell \in \mathbb{N}} \mathcal{Q}_{k,\ell}$.

We consider two types of utility functions, one for the hypothesis testing setup and another for the information preservation setup. In the hypothesis testing setup, the utility is a function of the privatization mechanism and two priors defined on the input alphabet. Namely, $U(P_0, P_1, Q) : \mathbb{S}^k \times \mathbb{S}^k \times \mathcal{Q} \rightarrow \mathbb{R}_+$, where P_0 and P_1 are positive priors defined on \mathcal{X} , and \mathbb{S}^k is the $(k - 1)$ -dimensional probability simplex. P_ν is said to be positive if $P_\nu(x) > 0$ for all $x \in \mathcal{X}$. In the information preservation setup, the utility is a function of the privatization mechanism and a prior defined on the input alphabet. Namely, $U(P, Q) : \mathbb{S}^k \times \mathcal{Q} \rightarrow \mathbb{R}_+$, where P is a positive prior defined on \mathcal{X} . For notational convenience, we will use $U(Q)$ to refer to both $U(P, Q)$ and $U(P_0, P_1, Q)$.

Definition 1 (Sublinear Functions) *A function $\mu(z) : \mathbb{R}^k \rightarrow \mathbb{R}$ is said to be sublinear if the following two conditions are met*

1. $\mu(\gamma z) = \gamma \mu(z)$ for all $\gamma \in \mathbb{R}_+$.
2. $\mu(z_1 + z_2) \leq \mu(z_1) + \mu(z_2)$ for all $z_1, z_2 \in \mathbb{R}^k$.

Let Q_y be the column of Q corresponding to $Q(y|\cdot)$ and μ be any sublinear function. We are interested in utilities that can be decomposed into a sum of sublinear functions. We study the *fundamental tradeoff between privacy and utility* by solving the following constrained maximization problem

$$\begin{aligned} & \underset{Q}{\text{maximize}} && U(Q) = \sum_{y \in \mathcal{Y}} \mu(Q_y) \\ & \text{subject to} && Q \in \mathcal{D}_\varepsilon \end{aligned} \tag{8}$$

This includes maximization over information theoretic quantities of interest in statistical estimation and hypothesis testing such as mutual information, total variation, KL-divergence, and χ^2 -divergence (Tsybakov and Zaiats, 2009). Since sub-linearity implies convexity, (8) is in general a complicated nonlinear program: we are maximizing (instead of minimizing) a convex function in Q ; further, the dimension of Q might be unbounded: the optimal

privatization mechanism Q^* might produce an infinite output alphabet \mathcal{Y} . The following theorem proves that one never needs an output alphabet larger than the input alphabet in order to achieve the maximum utility, and provides a combinatorial representation for the optimal solution.

Theorem 2 *For any sublinear function μ and any $\varepsilon \geq 0$, there exists an optimal mechanism Q^* maximizing the utility in (8) over all ε -locally differentially private mechanisms, such that*

- (a) *the output alphabet size is at most the input alphabet size, i.e. $|\mathcal{Y}| \leq |\mathcal{X}|$; and*
- (b) *for all $y \in \mathcal{Y}$, and $x, x' \in \mathcal{X}$*

$$\left| \ln \frac{Q^*(y|x)}{Q^*(y|x')} \right| \in \{0, \varepsilon\}. \tag{9}$$

The first claim of bounded alphabet size is more generally true for any general utility $U(Q)$ that is convex in Q (not necessarily decomposing into a sum of sublinear functions as in (8)). The second claim establishes that there is an optimal mechanism with an extremal structure; the absolute value of the log-likelihood ratios can only take one of the two extremal values: 0 or e^ε (see Figure 2 for example). We refer to such a mechanism as a staircase mechanism, and define the *family of staircase mechanisms* formally as

$$\mathcal{S}_\varepsilon \equiv \{Q \mid \text{satisfying (9)}\}.$$

For all choices of $U(Q) = \sum_{\mathcal{Y}} \mu(Q_y)$ and any $\varepsilon \geq 0$, Theorem 2 implies that the family of staircase mechanisms includes the optimal solutions to maximization problems of the form (8). Notice that staircase mechanisms are ε -locally differentially private, since any Q satisfying (9) implies that $Q(y|x)/Q(y|x') \leq e^\varepsilon$.

For global differential privacy, we can generalize the definition of staircase mechanisms to hold for all neighboring database queries x, x' (or equivalently within some sensitivity), and recover all known existing optimal mechanisms. Precisely, the geometric mechanism shown to be optimal in Ghosh et al. (2012), and the mechanisms shown to be optimal in Geng and Viswanath (2012, 2013a) (also called staircase mechanisms) are special cases of the staircase mechanisms defined above. We believe that the characterization of these extremal mechanisms and the analysis techniques developed in this paper can be of independent interest to researchers interested in optimal mechanisms for global privacy and more general utilities.

2.2 Combinatorial representation of the staircase mechanisms

Now that we know that staircase mechanisms are optimal, we can try to combinatorially search for the best staircase mechanism for an instance of the function μ and a fixed ε . To this end, we give a simple representation for all staircase mechanisms, exploiting the fact that they are scaled copies of a finite number of patterns.

Let $Q \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Y}|}$ be a staircase mechanism, and $k = |\mathcal{X}|$ denote the size of the input alphabet. Then, from the definition of staircase mechanisms, $Q(y|x)/Q(y|x') \in \{e^{-\varepsilon}, 1, e^\varepsilon\}$ and each column $Q(y|\cdot)$ must be proportional to one of the canonical staircase patterns we define next.




Figure 2: Examples of staircase mechanisms: the binary (left) and the randomized response (right) mechanisms.

Definition 3 (Staircase Pattern Matrix) Let b_j be the k -dimensional binary vector corresponding to the binary representation of j for $j \leq 2^k - 1$. A matrix $S^{(k)} \in \{1, e^\epsilon\}^{k \times 2^k}$ is called a staircase pattern matrix if the j -th column of $S^{(k)}$ is $S_j^{(k)} = (e^\epsilon - 1)b_{j-1} + \mathbf{1}$, for $j \in \{1, \dots, 2^k\}$. Each column of $S^{(k)}$ is a staircase pattern.

When $k = 3$, there are $2^k = 8$ staircase patterns and the staircase pattern matrix is given by

$$S^{(3)} = \begin{bmatrix} 1 & 1 & 1 & 1 & e^\epsilon & e^\epsilon & e^\epsilon & e^\epsilon \\ 1 & 1 & e^\epsilon & e^\epsilon & 1 & 1 & e^\epsilon & e^\epsilon \\ 1 & e^\epsilon & 1 & e^\epsilon & 1 & e^\epsilon & 1 & e^\epsilon \end{bmatrix}.$$

For all values of k , there are exactly 2^k such patterns, and any column $Q(y|\cdot)$ of Q , a staircase mechanism, is a scaled version of one of the columns of $S^{(k)}$. Using this pattern matrix, we can show that any staircase mechanism Q can be represented as

$$Q = S^{(k)}\Theta, \quad (10)$$

where $\Theta = \text{diag}(\theta)$ is a $2^k \times 2^k$ diagonal matrix and θ is a 2^k -dimensional vector representing the scaling of the columns of $S^{(k)}$. We can now formulate the problem of maximizing the utility as a linear program and prove their equivalence.

Theorem 4 *For any sublinear function μ and any $\varepsilon \geq 0$, the nonlinear program of (8) and the following linear program have the same optimal value*

$$\begin{aligned} \underset{\theta \in \mathbb{R}^{2^k}}{\text{maximize}} \quad & \sum_{j=1}^{2^k} \mu(S_j^{(k)})\theta_j = \mu^T \theta \\ \text{subject to} \quad & S^{(k)}\theta = \mathbb{1} \\ & \theta \geq 0, \end{aligned} \tag{11}$$

and the optimal solutions are related by (10).

Thus, the infinite dimensional nonlinear program of (8) is now reduced to a finite dimensional linear program. The constraints in (11) ensure that we get a valid probability matrix $Q = S^{(k)}\Theta$ with rows that sum to one. One could potentially solve this LP with 2^k variables but its computational complexity scales exponentially in the alphabet size $k = |\mathcal{X}|$. For practical values of k this might not always be possible. However, in the following sections, we prove that in the high privacy regime ($\varepsilon \leq \varepsilon^*$ for some positive ε^*), there is a single optimal mechanism, which we call the *binary mechanism*, which dominates over all other mechanisms in a very strong sense for all utility functions of practical interest.

In order to understand the above theorem, observe that both the objective function and differential privacy constraints are invariant under *permutations* (or relabelling) of the columns of a privatization mechanism Q . In other words, shuffling the columns of an ε -locally differentially private mechanism Q results in a new ε -locally differentially private mechanism Q' that achieves the same utility. Similarly, both the objective function and differential privacy constraints are invariant under *merging/splitting* of outputs with the same pattern. To be specific, consider a privatization mechanism Q and suppose that there exist two outputs y and y' that have the same pattern, i.e. $Q(y|\cdot) = C Q(y'|\cdot)$ for some positive constant C . Then, we can consider a new mechanism Q' by merging the two columns corresponding to y and y' . Let y'' denote this new output. It follows that Q' satisfies the differential privacy constraints and the resulting utility is also preserved. Precisely, using the fact that $Q(y|\cdot) = C Q(y'|\cdot)$, it follows that

$$\mu(Q_y) + \mu(Q_{y'}) = \mu((1+C)Q_y) = \mu(Q'_{y''}),$$

by the homogeneity property of μ . Therefore, we can naturally define equivalence classes for staircase mechanisms that are equivalent up to a permutation of columns and merging/splitting of columns with the same pattern:

$$[Q] = \{Q' \in \mathcal{S}_\varepsilon \mid \exists \text{ a sequence of permutations and merge/split of columns from } Q' \text{ to } Q\}.$$

To represent an equivalence class, we use a mechanism in $[Q]$ that is ordered and merged to match the patterns of the pattern matrix $S^{(k)}$. For any staircase mechanism Q , there exists a possibly different staircase mechanism $Q' \in [Q]$ such that $Q' = S^{(k)}\Theta$ for some diagonal matrix Θ with nonnegative entries. Therefore, to solve optimization problems of the form (8), we can restrict our attention to such representatives of equivalent classes. Further, for privatization mechanisms of the form $Q = S^{(k)}\Theta$, the objective function takes the form $\sum_j \mu(S_j^{(k)})\theta_j$, a simple linear function of Θ .

3. Hypothesis Testing

In this section, we study the fundamental tradeoff between local differential privacy and hypothesis testing. In this setting, there are n individuals each with data X_i sampled from a distribution P_ν for a fixed $\nu \in \{0, 1\}$. Let Q be a non-interactive privatization mechanism guaranteeing ε -local differential privacy. The output of the privatization mechanism Y_i is distributed according to the induced marginal M_ν defined in (2). With a slight abuse of notation, we will use M_ν and P_ν to represent both probability distributions and probability mass functions. The power to discriminate data sampled from P_0 to data sampled from P_1 depends on the ‘distance’ between the marginals M_0 and M_1 . To measure the ability of such statistical discrimination, our choice of utility of a privatization mechanism Q is an information theoretic quantity called Csiszár’s f -divergence defined as

$$D_f(M_0||M_1) = \sum_y M_1(y) f\left(\frac{M_0(y)}{M_1(y)}\right) = U(P_0, P_1, Q) = U(Q) , \quad (12)$$

for some convex function f such that $f(1) = 0$. The Kullback-Leibler (KL) divergence $D_{\text{kl}}(M_0||M_1)$ is a special case of f -divergence with $f(x) = x \log x$. The total variation distance $\|M_0 - M_1\|_{\text{TV}}$ is also special case with $f(x) = (1/2)|x - 1|$. Note that in general, the f -divergence is not necessarily a distance metric since it need not be symmetric or satisfy triangular inequality. We are interested in characterizing the optimal solution to

$$\begin{aligned} & \underset{Q}{\text{maximize}} && D_f(M_0||M_1) \\ & \text{subject to} && Q \in \mathcal{D}_\varepsilon \end{aligned} , \quad (13)$$

where \mathcal{D}_ε is the set of all ε -locally differentially private mechanisms defined in (7).

A motivating example for this choice of utility is the Neyman-Pearson hypothesis testing framework (Cover and Thomas, 2012). Given the privatized views $\{Y_i\}_{i=1}^n$, the data analyst wants to test whether they are generated from M_0 or M_1 . Let the null hypothesis be H_0 : Y_i ’s are generated from M_0 , and the alternative hypothesis H_1 : Y_i ’s are generated from M_1 . For a choice of rejection region $R \subseteq \mathcal{Y}^n$, the probability of false alarm (type I error) is $\alpha = M_0^n(R)$ and the probability of miss detection (type II error) is $\beta = M_1^n(\mathcal{Y}^n \setminus R)$. Let $\beta^{\alpha^*} = \min_{R \subseteq \mathcal{Y}^n, \alpha < \alpha^*} \beta$ denote the minimum type II error achievable while keeping the type I error rate at most α^* . According to Chernoff-Stein lemma (Cover and Thomas, 2012), we know that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \beta^{\alpha^*} = -D_{\text{kl}}(M_0||M_1) .$$

Suppose the analyst knows P_0 , P_1 , and Q . Then in order to achieve optimal asymptotic error rate, one would want to maximize the KL divergence of the induced marginals, over all ε -locally differentially private mechanisms Q . The results we present in this section (Theorems 5 and 8 to be precise) provide an explicit construction of optimal mechanisms in high and low privacy regimes. Using these optimality results, we prove a fundamental limit on the achievable error rates under differential privacy. Precisely, with data collected from an ε -locally differentially privatization mechanism, one cannot achieve an asymptotic

type II error smaller than

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \beta^{\alpha^*} \geq -\frac{(1+\delta)(e^\varepsilon - 1)^2}{(e^\varepsilon + 1)} \|P_0 - P_1\|_{\text{TV}}^2 \geq -\frac{(1+\delta)(e^\varepsilon - 1)^2}{2(e^\varepsilon + 1)} D_{\text{kl}}(P_0 \| P_1) ,$$

whenever $\varepsilon \leq \varepsilon^*$, where ε^* is dictated by Theorem 5 and $\delta > 0$ is some arbitrarily small but positive constant. In the equation above, the second inequality follows from Pinsker's inequality. Since $(e^\varepsilon - 1)^2 = O(\varepsilon^2)$ for small ε , the effective sample size is now reduced from n to $\varepsilon^2 n$. This is the price of privacy. In the low privacy regime where $\varepsilon \geq \varepsilon^*$, for ε^* dictated by Theorem 8, one cannot achieve an asymptotic type II error smaller than

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \beta^{\alpha^*} \geq -D_{\text{kl}}(P_0 \| P_1) + (1 - \delta)G(P_0, P_1)e^{-\varepsilon} .$$

3.1 Optimality of staircase mechanisms

From the definition of $D_f(M_0 \| M_1)$, we have that

$$D_f(M_0 \| M_1) = \sum_y (P_1^T Q_y) f(P_0^T Q_y / P_1^T Q_y) = \sum_y \mu(Q_y) ,$$

where $P_\nu^T Q_y = \sum_{\mathcal{X}} P_\nu(x) Q(y|x)$ and $\mu(Q_y) = (P_1^T Q_y) f(P_0^T Q_y / P_1^T Q_y)$. For any $\gamma > 0$,

$$\begin{aligned} \mu(\gamma Q_y) &= (P_1^T(\gamma Q_y)) f(P_0^T(\gamma Q_y) / P_1^T(\gamma Q_y)) \\ &= \gamma (P_1^T Q_y) f(P_0^T Q_y / P_1^T Q_y) \\ &= \gamma \mu(Q_y) . \end{aligned}$$

Moreover, since the function $\phi(z, t) = tf(\frac{z}{t})$ is convex in (z, t) for $0 \leq z, t \leq 1$, then μ is convex in Q_y . Convexity and homogeneity together imply sublinearity. Therefore, Theorems 2 and 4 apply to $D_f(M_0 \| M_1)$ and we have that staircases are optimal.

3.2 Optimality of the binary mechanism

For a given P_0 and P_1 , the *binary mechanism* is defined as a staircase mechanism with only two outputs $Y \in \{0, 1\}$ satisfying (see Figure 2)

$$Q(0|x) = \begin{cases} \frac{e^\varepsilon}{1 + e^\varepsilon} & \text{if } P_0(x) \geq P_1(x) , \\ \frac{1}{1 + e^\varepsilon} & \text{if } P_0(x) < P_1(x) . \end{cases} \quad Q(1|x) = \begin{cases} \frac{e^\varepsilon}{1 + e^\varepsilon} & \text{if } P_0(x) < P_1(x) , \\ \frac{1}{1 + e^\varepsilon} & \text{if } P_0(x) \geq P_1(x) . \end{cases} \quad (14)$$

Although this mechanism is extremely simple, perhaps surprisingly, we will establish that it is the optimal mechanism when a high level of privacy is required. Intuitively, the output should be very noisy in the high privacy regime, and we are better off sending just one bit of information that tells you whether your data is more likely to have come from P_0 or P_1 .

Theorem 5 *For any pair of distributions P_0 and P_1 , there exists a positive ε^* that depends on P_0 and P_1 such that for any f -divergences and any positive $\varepsilon \leq \varepsilon^*$, the binary mechanism maximizes the f -divergence between the induced marginals over all ε -locally differentially private mechanisms.*

This implies that in the high privacy regime, which is a typical setting studied in much of the differential privacy literature, the binary mechanism is universally optimal for all f -divergences. In particular this threshold ε^* is *universal*, in that it does not depend on the particular choice of which f -divergence we are maximizing. It is only a function of P_0 and P_1 . This is established by proving a very strong statistical dominance using Blackwell’s celebrated result on comparisons of statistical experiments Blackwell (1953). In a nutshell, we prove that any ε -locally differentially private mechanism can be simulated from the output of the binary mechanism for sufficiently small ε . Hence, the binary mechanism dominates over all other mechanisms and at the same time achieves the maximum divergence. A similar idea has been used previously in (Kairouz et al., 2013) to exactly characterize how much privacy degrades under composition attacks.

The optimality of binary mechanisms is not just for high privacy regimes. The next theorem shows that it is *the* optimal solution of (13) for all ε , when the objective function is the total variation distance: $D_f(M_0||M_1) = \|M_0 - M_1\|_{\text{TV}}$.

Theorem 6 *For any pair of distributions P_0 and P_1 , and any $\varepsilon \geq 0$, the binary mechanism maximizes the total variation distance between the induced marginals M_0 and M_1 among all ε -locally differentially private mechanisms.*

When maximizing the KL divergence between the induced marginals, we show that the binary mechanism still achieves good performance for $\varepsilon \leq C$ where C is a constant that does not depend on P_0 and P_1 . For the special case of KL divergence, let OPT denote the maximum value of (13) and BIN denote the KL divergence when the binary mechanism is used. The next theorem shows that

$$\text{BIN} \geq \frac{1}{2(e^\varepsilon + 1)^2} \text{OPT} .$$

Theorem 7 *For any ε and any pair of distributions P_0 and P_1 , the binary mechanism is an $1/(2(e^\varepsilon + 1)^2)$ approximation of the maximum KL divergence between the induced marginals M_0 and M_1 among all ε -locally differentially private mechanisms.*

Observe that $2(e^\varepsilon + 1)^2 \leq 32$ for $\varepsilon \leq 1$. Therefore, for any $\varepsilon \leq 1$, the simple binary mechanism is at most a constant factor away from the optimal mechanism.

3.3 Optimality of the randomized response mechanism

The *randomized response mechanism* (see Figure 2) is a staircase mechanism with $\mathcal{Y} = \mathcal{X}$ satisfying

$$Q(y|x) = \begin{cases} \frac{e^\varepsilon}{|\mathcal{X}| - 1 + e^\varepsilon} & \text{if } y = x , \\ \frac{1}{|\mathcal{X}| - 1 + e^\varepsilon} & \text{if } y \neq x . \end{cases} \tag{15}$$

In other words, the randomized response is a simple randomization over the same alphabet where the true data is released with probability $e^\varepsilon / (|\mathcal{X}| - 1 + e^\varepsilon)$. We view it as a multiple

choice generalization to the randomized response method proposed by Warner (1965). We now establish that for the special case of optimizing the KL divergence between the induced marginals, the randomized response mechanism is the optimal solution of (13) in the low privacy regime (i.e., $\varepsilon \geq \varepsilon^*$ for some threshold ε^* that depends on P_0 and P_1).

Theorem 8 *There exists a positive ε^* that depends on P_0 and P_1 such that for all P_0 and P_1 , and all $\varepsilon \geq \varepsilon^*$, the randomized response mechanism maximizes the KL divergence between the induced marginals among all ε -locally differentially private mechanisms.*

The randomized response mechanism is particularly important because it does not depend on P_0 or P_1 . Thus, even if the data providers and analysts do not have access to the priors, they can still use the randomized response mechanism to achieve the optimal (or near-optimal) utility in the moderate to low privacy regimes.

3.4 Numerical experiments

A typical approach for achieving ε -local differential privacy is to add geometric noise with appropriately chosen variance. For an input with alphabet size $|\mathcal{X}| = k$, this amounts to relabelling the inputs as integers $\{1, \dots, k\}$ and adding geometric noise, i.e., $Q(y|x) = ((1 - \varepsilon^{1/(k-1)})/(1 + \varepsilon^{1/(k-1)}))\varepsilon^{|y-x|/(k-1)}$ for $y \in \mathbb{Z}$. The output is then truncated at 1 and k to preserve the support.

For 100 instances of randomly chosen P_0 and P_1 defined over an input alphabet of size $|\mathcal{X}| = 6$, we compare the average performance of the binary, randomized response, and geometric mechanisms to the average performance of the optimal staircase mechanism for various values of ε . The optimal staircase mechanism is computed by solving the linear program in Equation (11) for each fixed pair (P_0, P_1) and ε . The left panel of Figure 3 shows the average performance measured by the normalized divergence $D_{\text{kl}}(M_0||M_1)/D_{\text{kl}}(P_0||P_1)$ for all 4 mechanisms. The average is taken over the 100 instances of P_0 and P_1 . In the low privacy (large ε) regime, the randomized response achieves optimal performance (which converges exponentially in ε to 1) as predicted. In the high privacy regime (small ε), the binary mechanism achieves optimal performance (which converges quadratically in ε to 0) as predicted. In all regimes, both the binary and randomized response mechanisms provide significant gains over the geometric mechanism.

To illustrate how much worse the binary and the randomized response mechanisms can be relative to the optimal staircase mechanism, we plot in the right panel of Figure 3 the divergence under each mechanism normalized by the divergence under the optimal mechanism. This is done for all 100 instances of P_0 and P_1 . In all instances, the binary mechanism is optimal for small ε and the randomized response mechanism is optimal for large ε . However, $D_{\text{kl}}(M_0||M_1)$ under the randomized response mechanism can be as bad as 10% of the optimal one (for small ε). Similarly, $D_{\text{kl}}(M_0||M_1)$ under the binary mechanism can be as bad as 25% of the optimal one (for large ε). To overcome this issue, we propose the following simple strategy: use the better among these two mechanisms. The performance of this strategy is illustrated in Figure 4. For various input alphabet size $|\mathcal{X}| \in \{3, 4, 6, 12\}$, we plot the performance of this mixed strategy for each value of ε and each of the 100 randomly generated instances of P_0 and P_1 . This mixed strategy achieves at least 70% for $|\mathcal{X}| = 6$ (and 55% for $|\mathcal{X}| = 12$) of the optimal divergence for all instances. Figure 4 shows




Figure 3: The binary and randomized response mechanisms are optimal in the high-privacy (small ϵ) and low-privacy (large ϵ) regimes, respectively, and improve over the geometric mechanism significantly (left). When the regimes are mismatched, $D_{\text{kl}}(M_0||M_1)$ under these mechanisms can be as bad as 10% of the optimal one (right).

that this mixed strategy is not too sensitive to the size of the alphabet k . Therefore, this strategy provides a good mechanism that can be readily used in practice for any value of ϵ .

3.5 Lower bounds

In this section, we provide converse results on the fundamental limit of differentially private mechanisms; these results follow from our main theorems and are of independent interest in other applications where lower bounds in statistical analysis are studied (Beimel et al., 2008; Hardt and Talwar, 2010; Chaudhuri and Hsu, 2012; De, 2012). For example, a bound similar to (16) was used to provide converse results on the sample complexity for statistical estimation with differentially private data in Duchi et al. (2013).

Corollary 9 *For any $\epsilon \geq 0$, let Q be any conditional distribution that guarantees ϵ -local differential privacy. Then, for any pair of distributions P_0, P_1 and any positive $\delta > 0$, there exists a positive ϵ^* that depends on P_0, P_1 and δ such that for any $\epsilon \leq \epsilon^*$ the induced marginals M_0 and M_1 satisfy the bound*

$$D_{\text{kl}}(M_0||M_1) + D_{\text{kl}}(M_1||M_0) \leq \frac{2(1+\delta)(e^\epsilon - 1)^2}{(e^\epsilon + 1)} \|P_0 - P_1\|_{\text{TV}}^2. \quad (16)$$

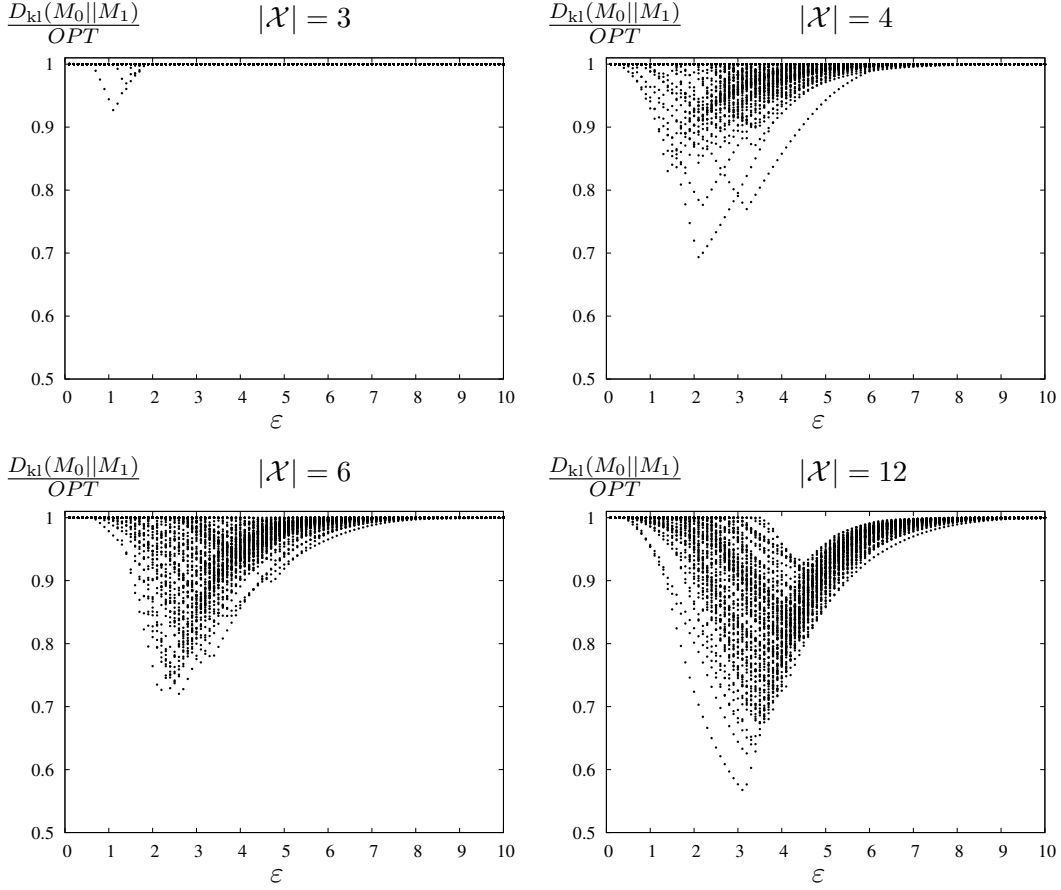


Figure 4: For varying input alphabet size $|\mathcal{X}| \in \{3, 4, 6, 12\}$, at least 55% of the optimal divergence can be achieved by taking the better one between the binary and the randomized response mechanisms.

This follows from Theorem 5 and observing that the binary mechanism achieves

$$\begin{aligned}
 & D_{\text{kl}}(M_0||M_1) \\
 &= \frac{(e^\varepsilon - 1)P_0(T) + 1}{e^\varepsilon + 1} \log \left(\frac{1 + (e^\varepsilon - 1)P_0(T)}{1 + (e^\varepsilon - 1)P_1(T)} \right) \\
 &\quad + \frac{(e^\varepsilon - 1)P_0(T^c) + 1}{e^\varepsilon + 1} \log \left(\frac{1 + (e^\varepsilon - 1)P_0(T^c)}{1 + (e^\varepsilon - 1)P_1(T^c)} \right) \\
 &= \frac{(e^\varepsilon - 1)^2}{e^\varepsilon + 1} (P_0(T) - P_1(T)) + O(\varepsilon^3) \\
 &= \frac{(e^\varepsilon - 1)^2}{e^\varepsilon + 1} \|P_0 - P_1\|_{\text{TV}}^2 + O(\varepsilon^3),
 \end{aligned}$$

where $T \subseteq \mathcal{X}$ is the set of x such that $P_0(x) \geq P_1(x)$. Compared to (Duchi et al., 2013, Theorem 1), we recover their bound of $4(e^\varepsilon - 1)^2 \|P_0 - P_1\|_{\text{TV}}^2$ with a smaller constant.

We want to note that Duchi et al.’s bound holds for all values of ε and uses a different technique of bounding the KL divergence directly, however no achievable mechanism has been provided. We instead provide an explicit mechanism, that is optimal in high privacy regime.

Similarly, in the low privacy regime, we can show the following converse result.

Corollary 10 *For any $\varepsilon \geq 0$, let Q be any conditional distribution that guarantees ε -local differential privacy. Then, for any pair of distributions P_0 and P_1 and any positive $\delta > 0$, there exists a positive ε^* that depends on P_0 and P_1 and δ such that for any $\varepsilon \geq \varepsilon^*$ the induced marginals M_0 and M_1 satisfy the bound*

$$D_{\text{kl}}(M_0||M_1) + D_{\text{kl}}(M_1||M_0) \leq D_{\text{kl}}(P_0||P_1) - (1 - \delta)G(P_0, P_1)e^{-\varepsilon}. \quad (17)$$

where $G(P_0, P_1) = \sum_{\mathcal{X}} (1 - P_0(x)) \log(P_1(x)/P_0(x))$.

This follows directly from Theorem 8 and observing that the randomized response mechanism achieves $D_{\text{kl}}(M_0||M_1) = D_{\text{kl}}(P_0||P_1) - G(P_0, P_1)e^{-\varepsilon} + O(e^{-2\varepsilon})$.

Figure 5 illustrates the gap between the divergence achieved by the geometric mechanism described in the previous section and the optimal mechanisms (the binary mechanism for the high privacy regime and the randomized response mechanism for the low privacy regime). For each instance of the 100 randomly generated P_0 and P_1 defined over input alphabets of size $k = 6$, we plot the resulting divergence $D_{\text{kl}}(M_0||M_1)$ as a function of $\|P_0 - P_1\|_{\text{TV}}$ for $\varepsilon = 0.1$, and as a function of $D_{\text{kl}}(P_0||P_1)$ for $\varepsilon = 10$. The binary and the randomized response mechanisms exhibit the scaling predicted by Equation (16) and (17), respectively. Similarly, for total variation, we can get the following converse result.

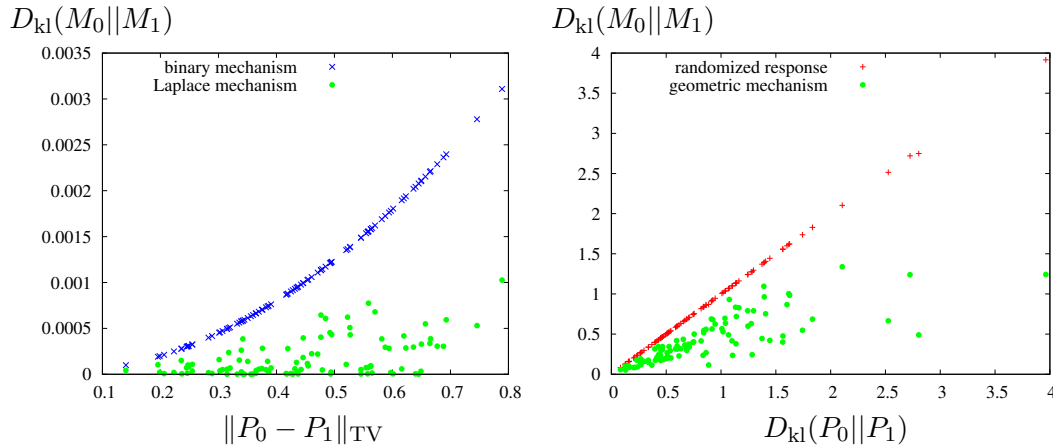


Figure 5: For small $\varepsilon = 0.1$ (left) the binary mechanism achieves the optimal KL divergence, which scales as Equation (16). For large $\varepsilon = 10$ (right) the randomized response achieves the optimal KL divergence, which scales as Equation (17). Both mechanisms improve significantly over the geometric mechanism.

Corollary 11 *For any $\varepsilon \geq 0$, let Q be any conditional distribution that guarantees ε -local differential privacy. Then, for any pair of distributions P_0 and P_1 , the induced marginals M_0 and M_1 satisfy the bound $\|M_0 - M_1\|_{\text{TV}} \leq ((e^\varepsilon - 1)/(e^\varepsilon + 1)) \|P_0 - P_1\|_{\text{TV}}$, and equality is achieved by the binary mechanism.*

This follows from Theorem 6 and explicitly computing the total variation achieved by the binary mechanism.

4. Information Preservation

In this section, we study the fundamental tradeoff between local privacy and mutual information. Consider a random variable X distributed according to P . The information content of X is captured by information theoretic quantity called entropy

$$H(X) = - \sum_x P(x) \log P(x).$$

We are interested in releasing a differentially private version of X represented by Y . The random variable Y should preserve the information content of X as much as possible while meeting the local differential privacy constraints. Similar to the hypothesis testing setting, we will show that a variant of the binary mechanism is optimal in the high privacy regime and that the randomized response mechanism is optimal in the low privacy regime.

Let Q be a non-interactive privatization mechanism guaranteeing ε -local differential privacy. The output of the privatization mechanism Y is distributed according to the induced marginal M given by

$$M(S) = \sum_{x \in \mathcal{X}} Q(S|x)P(x),$$

for $S \subseteq \mathcal{Y}$. With a slight abuse of notation, we will use M and P to represent both probability distributions and probability mass functions. The information content in Y about X is captured by the well celebrated information theoretic quantity called mutual information. The mutual information between X and Y is given by

$$I(X;Y) = \sum_x \sum_y P(x) Q(y|x) \log \left(\frac{Q(y|x)}{\sum_{l \in \mathcal{X}} P(l) Q(y|l)} \right) = U(P, Q) = U(Q).$$

Notice that $I(X;Y) \leq H(X)$ and $I(X;Y)$ is convex in Q (Cover and Thomas, 2012). To preserve the information context of X , we wish to choose a privatization mechanism Q such that the mutual information between X and Y is maximized subject to differential privacy constraints. In other words, we are interested in characterizing the optimal solution to

$$\begin{aligned} & \underset{Q}{\text{maximize}} && I(X;Y) \\ & \text{subject to} && Q \in \mathcal{D}_\varepsilon \end{aligned}, \tag{18}$$

where \mathcal{D}_ε is the set of all ε -locally differentially private mechanisms defined in (7). The above mutual information maximization problem can be thought of as a conditional entropy minimization problem since $I(X;Y) = H(X) - H(X|Y)$.

4.1 Optimality of staircase mechanisms

From the definition of $I(X; Y)$, we have that

$$I(X; Y) = \sum_y \sum_x P(x) Q(y|x) \log \left(\frac{Q(y|x)}{P^T Q_y} \right) = \sum_y \mu(Q_y),$$

where $P^T Q_y = \sum_x P(x) Q(y|x)$ and $\mu(Q_y) = \sum_x P(x) Q(y|x) \log(Q(y|x)/P^T Q_y)$. Notice that $\mu(\gamma Q_y) = \gamma \mu(Q_y)$, and by the log-sum inequality, μ is convex. Convexity and homogeneity together imply sublinearity. Therefore, Theorems 2 and 4 apply to $I(X; Y)$ and we have that staircase mechanisms are optimal.

4.2 Optimality of the binary mechanism

For a given P , the *binary mechanism for mutual information* is a staircase mechanism with only two outputs $Y \in \{0, 1\}$ (see Figure 2)

$$Q(0|x) = \begin{cases} \frac{e^\varepsilon}{1 + e^\varepsilon} & \text{if } x \in T, \\ \frac{1}{1 + e^\varepsilon} & \text{if } x \notin T, \end{cases} \quad Q(1|x) = \begin{cases} \frac{e^\varepsilon}{1 + e^\varepsilon} & \text{if } x \notin T, \\ \frac{1}{1 + e^\varepsilon} & \text{if } x \in T, \end{cases} \quad (19)$$

with

$$T \in \arg \min_{A \subseteq \mathcal{X}} \left| P(A) - \frac{1}{2} \right|. \quad (20)$$

Observe that there are multiple valid choices for T . Indeed, for any minimizing set T , T^c is also a minimizing set since $|P(T) - 1/2| = |P(T^c) - 1/2|$. When there are multiple pairs, any pair (T, T^c) can be chosen to define the binary mechanism. All valid binary mechanisms are equivalent from a utility maximization perspective.

In what follows, we will establish that this simple mechanism is optimal in the high privacy regime. Intuitively, in the high privacy regime, we cannot release more than one bit of information, and hence, the input alphabet is reduced to a binary output alphabet. In this case, it makes sense to partition the original alphabet \mathcal{X} in a way that preserves the information content of X as much as possible. Indeed, our choice of T in Equation (20) maximizes the information contained in the released bit because $T \in \arg \max_{A \subseteq \mathcal{X}} |P(A) - 1/2| = \arg \max_{A \subseteq \mathcal{X}} (-P(A) \log P(A) - P(A^c) \log P(A^c))$ (see Section 9.1 for a proof).

Theorem 12 *For any distribution P , there exists a positive ε^* that depends on P such that for any positive $\varepsilon \leq \varepsilon^*$, the binary mechanism maximizes the mutual information between the input and the output of a privatization mechanism over all ε -locally differentially private mechanisms.*

This implies that in the high privacy regime, the binary mechanism is the optimal solution for (18).

Next, we show that the binary mechanism achieves near-optimal performance for all (\mathcal{X}, P) and $\varepsilon \leq 1$ even when $\varepsilon^* < 1$. Let OPT denote the maximum value of (18) and BIN

denote the mutual information achieved by the binary mechanism given in (19). The next theorem shows that

$$\text{BIN} \geq \frac{1}{1 + e^\varepsilon} \text{OPT}.$$

Theorem 13 *For any $\varepsilon \leq 1$ and any distribution P , the binary mechanism is an $(1 + e^\varepsilon)$ -approximation of the maximum mutual information between the input and the output of a privatization mechanism among all ε -locally differentially private mechanisms.*

Note that $1 + e^\varepsilon \leq 4$ for $\varepsilon \leq 1$ which is a commonly studied regime in differential privacy applications. Therefore, we can always use the simple binary mechanism and the resulting mutual information is at most a constant factor away from the optimal value.

4.3 Optimality of the randomized response mechanism

In the low privacy regime ($\varepsilon \geq \varepsilon^*$), the *randomized response mechanism* defined in(15) is optimal.

Theorem 14 *There exists a positive ε^* that depends on P such that for any distribution P and all $\varepsilon \geq \varepsilon^*$, the randomized response mechanism maximizes the mutual information between the input and the output of a privatization mechanism over all ε -locally differentially private mechanisms.*

Observe that the randomized response is not a function of P . Therefore, it can be used even when the distribution P is unknown.

4.4 Numerical experiments

For 100 instances of randomly chosen P defined over an input alphabet of size $|\mathcal{X}| = 6$, we compare the average performance of the binary, randomized response, and the geometric mechanisms to the average performance of the optimal mechanism. We plot (in Figure 6, left) the average performance measured by the normalized mutual information $I(X;Y)/H(X)$ for all 4 mechanisms. The average is taken over the 100 random instances of P . In the low privacy (large ε) regime, the randomized response achieves optimal performance as predicted, which converges to one. In the high privacy regime (small ε), the binary mechanism achieves optimal performance as predicted. In all regimes, both mechanisms significantly improve over the geometric mechanism. To illustrate how much worse the binary and randomized response mechanisms can be (relative to the optimal staircase mechanism), we plot (in Figure 6, right) the mutual information under each mechanism normalized by the mutual information under the optimal staircase mechanism. This is done for all 100 instances of P . In all instances, the binary mechanism is optimal for small ε and the randomized response mechanism is optimal for large ε . However, $I(X;Y)$ under the randomized response mechanism can be as bad as 35% of the optimal one (for small ε). Similarly, $I(X;Y)$ under the binary mechanism can be as bad as 40% of the optimal one (for large ε).

For $|\mathcal{X}| \in \{3, 4, 6, 12\}$, we plot (in Figure 7) the performance of the better between the binary and randomized response mechanisms normalized by the optimal mechanism for

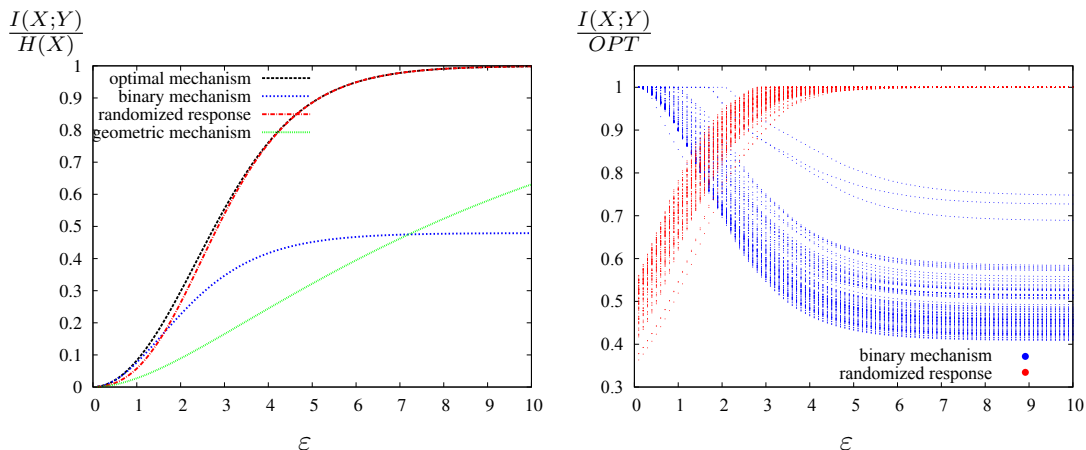


Figure 6: The binary and randomized response mechanisms are optimal in the high-privacy (small ε) and low-privacy (large ε) regimes, respectively, and improve over the geometric mechanism significantly (left). When the regimes are mismatched, $I(X; Y)$ under these mechanisms can each be as bad as 35% of the optimal one (right).

all 100 randomly generated instances of P . This mixed strategy achieves at least 75% for $|\mathcal{X}| = 6$ (and 65% for $|\mathcal{X}| = 12$) of the optimal mutual information for all instances of P . Moreover, it is not sensitive to the size of the alphabet $|\mathcal{X}|$.

4.5 Lower bounds

In this section, we provide converse results on the fundamental limit of locally differentially private mechanisms when utility is measured via mutual information.

Corollary 15 *For any $\varepsilon \geq 0$, let Q be any conditional distribution that guarantees ε -local differential privacy. Then, for any distribution P and any positive $\delta > 0$, there exists a positive ε^* that depends on P and δ such that for any $\varepsilon \leq \varepsilon^*$ the following bound holds*

$$I(X; Y) \leq (1 + \delta) \frac{1}{2} P(T) P(T^c) \varepsilon^2, \quad (21)$$

where T is defined in (20).

This follows from Theorem 12 (optimality of the binary mechanism) and observing that the binary mechanism achieves

$$\begin{aligned} I(X; Y) &= \frac{1}{e^\varepsilon + 1} \left\{ P(T) e^\varepsilon \log \frac{e^\varepsilon}{P(T^c) + e^\varepsilon P(T)} + P(T^c) \log \frac{1}{P(T^c) + e^\varepsilon P(T)} \right\} \\ &\quad + \frac{1}{e^\varepsilon + 1} \left\{ P(T^c) e^\varepsilon \log \frac{e^\varepsilon}{P(T) + e^\varepsilon P(T^c)} + P(T) \log \frac{1}{P(T) + e^\varepsilon P(T^c)} \right\} \\ &= \frac{1}{2} P(T) P(T^c) \varepsilon^2 + O(\varepsilon^3). \end{aligned}$$




Figure 7: For varying input alphabet size $|\mathcal{X}| \in \{3, 4, 6, 12\}$, at least 65% of the maximum $I(X;Y)$ can be achieved by taking the better one between the binary and the randomized response mechanisms.

Similarly, in the low privacy regime, we can show the following converse result.

Corollary 16 *For any $\varepsilon \geq 0$, let Q be any conditional distribution that guarantees ε -local differential privacy. Then, for any distribution P and any positive $\delta > 0$, there exists a positive ε^* that depends on P and δ such that for any $\varepsilon \geq \varepsilon^*$ the following bound holds*

$$I(X;Y) \leq H(X) - (1 - \delta)(k - 1)\varepsilon e^{-\varepsilon}.$$

This follows directly from Theorem 14 (optimality of the randomized response mechanism) and observing that the randomized response mechanism achieves

$$I(X;Y) = H(X) - (k - 1)\varepsilon e^{-\varepsilon} + O(e^{-2\varepsilon}). \quad (22)$$

Figure 8 illustrates the gap between the mutual information achieved under the geometric and the optimal mechanisms (the binary mechanism for the high privacy regime and




Figure 8: For $\varepsilon = 0.1$ (left) the binary mechanism achieves the maximum $I(X;Y)$, which scales as Equation (21). For $\varepsilon = 10$ (right) the randomized response mechanism achieves the optimal mutual information, which scales as Equation (22).

the randomized response mechanism for the low privacy regime). For each instance of the 100 randomly generated P over an input of size $k = 6$, we plot the resulting mutual information $I(X;Y)$ as a function of $P(T)P(T^c)$ for $\varepsilon = 0.1$, and as a function of $H(X)$ for $\varepsilon = 10$. The binary and the randomized response mechanisms exhibit the scaling predicted by Equations (21) and (22), respectively.

5. Generalizations to approximate differential privacy

In this section, we generalize the results of the previous sections in the following ways.

1. We consider the class of utility functions that obey the data processing inequality. Consider the composition of two privatization mechanisms $QW = Q \circ W$ where the output of the first mechanism Q is applied to another mechanism W . We say that a utility function $U(\cdot)$ obeys the data processing inequality if the following inequality holds for all Q and W

$$U(QW) \leq U(Q).$$

The following proposition, proved in Section 10, shows that the class of utilities obeying the data processing inequality includes all the utility functions we studied in Section 2.

Proposition 17 *Any utility function that can be written in the form of $U(Q) = \sum_{\mathcal{Y}} \mu(Q_y)$, where μ is any sublinear function, obeys the data processing inequality.*

2. We consider (ε, δ) -differential privacy which generalizes the notion of ε -differential privacy. (ε, δ) -differential privacy is commonly referred to as approximate differential

privacy and it was first introduced in Dwork et al. (2006a). For the release of a random variable $X \in \mathcal{X}$, we say that a mechanism Q is (ε, δ) -locally differentially private if

$$Q(S|x) \leq e^\varepsilon Q(S|x') + \delta, \quad (23)$$

for all $S \subseteq \mathcal{Y}$ and all $x, x' \in \mathcal{X}$. Note that ε -local differential privacy is a special case of (ε, δ) -local differential privacy where $\delta = 0$.

3. We prove that the *quaternary mechanism*, defined in Equation (24), is optimal for any ε and any δ . This is different from the treatment conducted in the previous sections where we proved the optimality of the binary (randomized response) mechanism for sufficiently small (large) ε and $\delta = 0$.

The treatment in this section, even though more general than the one in previous sections in the ways described above, holds only for binary alphabets (i.e., $|\mathcal{X}| = 2$). Finding optimal privatization mechanisms under (ε, δ) -local differential privacy for larger input alphabets (i.e., $|\mathcal{X}| > 2$) is an interesting open question. Unlike ε -local differential privacy, the privacy constraints under (ε, δ) -local differential privacy no longer decompose into separate constraints on each output y . This makes it difficult to generalize the techniques developed in previous sections of this paper. However, for the special case of binary input alphabets, we can prove the optimality of one mechanism for all values of (ε, δ) and all utility functions that obey the data processing inequality.

For a binary random variable $X \in \mathcal{X} = \{0, 1\}$, the *quaternary mechanism* maps X to a quaternary random variable $Y \in \mathcal{Y} = \{0, 1, 2, 3\}$ and is defined as

$$Q_{\text{QT}}(0|x) = \begin{cases} \delta & \text{if } x = 0, \\ 0 & \text{if } x = 1. \end{cases} \quad Q_{\text{QT}}(1|x) = \begin{cases} 0 & \text{if } x = 0, \\ \delta & \text{if } x = 1. \end{cases} \quad (24)$$

$$Q_{\text{QT}}(2|x) = \begin{cases} (1 - \delta) \frac{1}{1+e^\varepsilon} & \text{if } x = 0, \\ (1 - \delta) \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } x = 1. \end{cases} \quad Q_{\text{QT}}(3|x) = \begin{cases} (1 - \delta) \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } x = 0, \\ (1 - \delta) \frac{1}{1+e^\varepsilon} & \text{if } x = 1. \end{cases}$$

In other words, the quaternary mechanism passes X unchanged with probability δ and applies the binary mechanism (defined in previous sections) with probability $1 - \delta$. The main result of this section can be stated formally as follows.

Theorem 18 *If $|\mathcal{X}| = 2$, then for any ε , any δ , and any utility $U(Q)$ that obeys the data processing inequality, the quaternary mechanism maximizes $U(Q)$ subject to $Q \in \mathcal{D}_{(\varepsilon, \delta)}$, the set of all (ε, δ) -locally differentially private mechanism.*

The proof of Theorem 18 depends on an *operational definition* of differential privacy which we describe next. Consider a privatization mechanism Q that maps $X \in \{0, 1\}$ stochastically to $Y \in \mathcal{Y}$. Given Y , construct a binary hypothesis test on whether $X = 0$ or $X = 1$. Any binary hypothesis test is completely described by a, possibly randomized, decision rule $\hat{X} : Y \rightarrow \{0, 1\}$. The two types of error associated with \hat{X} are *false alarm*: $\hat{X} = 1$ when $X = 0$, and *miss detection*: $\hat{X} = 0$ when $X = 1$. The probability of false alarm is given by $P_{\text{FA}} = \mathbb{P}(\hat{X} = 1|X = 0)$ while the probability of miss detection is given




Figure 9: The quaternary mechanism

by $P_{\text{MD}} = \mathbb{P}(\hat{X} = 0 | X = 1)$. For a fixed Q , the convex hull of all pairs $(P_{\text{MD}}, P_{\text{FA}})$ for all decision rules \hat{X} defines a two-dimensional *error region* where P_{MD} is plotted against P_{FA} . For example, the quaternary mechanism given in Figure 9a has an error region $\mathcal{R}_{Q_{\text{QT}}}$ shown in Figure 9b.

It turns out that (ε, δ) -local differential privacy imposes the following conditions on the error region of all (ε, δ) -locally differentially private mechanisms

$$P_{\text{FA}} + e^\varepsilon P_{\text{MD}} \geq 1 - \delta, \quad \text{and} \quad e^\varepsilon P_{\text{FA}} + P_{\text{MD}} \geq 1 - \delta,$$

for any decision rule \hat{X} . These two conditions define an error region $\mathcal{R}_{\varepsilon, \delta}$ shown in Figure 9b. Interestingly, the next theorem shows that the converse result is also true.

Theorem 19 *A mechanism Q is (ε, δ) -locally differentially private if and only if $\mathcal{R}_Q \subseteq \mathcal{R}_{\varepsilon, \delta}$.*

The proof of the above theorem can be found in Kairouz et al. (2013). Notice that it is no coincidence that $\mathcal{R}_{Q_{\text{QT}}} = \mathcal{R}_{\varepsilon, \delta}$. This property will be essential in proving the optimality of the quaternary mechanism.

Theorem 19 allows us to benefit from the data processing inequality (DPI) and its converse, which follows from a celebrated result by Blackwell (1953). These inequalities, while simple by themselves, lead to surprisingly strong technical results. Indeed, there is a long line of such a tradition in the information theory literature (see Chapter 17 of Cover and Thomas (2012)). Consider two privatization mechanisms, $Q^{(1)}$ and $Q^{(2)}$. Let Y and Z denote the output of the mechanisms $Q^{(1)}$ and $Q^{(2)}$, respectively. We say that $Q^{(1)}$ dominates $Q^{(2)}$ if there exists a coupling of Y and Z such that $X-Y-Z$ forms a Markov chain. In other words, we say $Q^{(1)}$ dominates $Q^{(2)}$ if there exists a stochastic mapping Q such that $Q^{(2)} = Q^{(1)} \circ Q$.

Theorem 20 *A mechanism $Q^{(1)}$ dominates a mechanism $Q^{(2)}$ if and only if $\mathcal{R}_{Q^{(2)}} \subseteq \mathcal{R}_{Q^{(1)}}$.*

The proof of the above theorem can be found in Blackwell (1953). Observe that by Theorems 20 and 19, and the fact that $\mathcal{R}_{Q_{QT}} = \mathcal{R}_{\varepsilon, \delta}$, the quaternary mechanism dominates any other differentially private mechanism. In other words, for any differentially private mechanism Q , there exists a stochastic mapping W such that $Q = W \circ Q_{QT}$. Therefore, for any (ε, δ) and any utility function $U(\cdot)$ obeying the data processing inequality, we have that $U(Q) \leq U(Q_{QT})$. This finishes the proof of Theorem 18.

6. Discussion

In this paper, we have considered a broad class of convex utility functions and assumed a setting where individuals cannot collaborate (communicate with each other) before releasing their data. It turns out that the techniques developed in this work can be generalized to find optimal privatization mechanisms in a setting where different individuals can collaborate interactively and each individual can be an analyst (Kairouz et al., 2014b).

Binary hypothesis testing and information preservation are two canonical problems with a wide range of applications. However, there are a number of non-trivial and interesting extensions to our work.

Correlation among data. In some scenarios the X_i 's could be correlated (e.g., when different individuals observe different functions of the same random variable). In this case, the data analyst is interested in inferring whether the data was generated from P_0^n or P_1^n , where P_ν^n is one of two possible joint priors on X_1, \dots, X_n . This is a challenging problem because knowing X_i reveals information about X_j , $j \neq i$. Therefore, the utility maximization problems for different individuals are coupled in this setting.

Robust and m -ary hypothesis testing. In some cases the data analyst need not have access to P_0 and P_1 , but rather to two classes of prior distribution P_{θ_0} and P_{θ_1} for $\theta_0 \in \Lambda_0$ and $\theta_1 \in \Lambda_1$. Such problems are studied under the rubric of universal hypothesis testing and robust hypothesis testing. One possible direction is to select the privatization mechanism that maximizes the worst case utility: $Q^* = \arg \max_{Q \in \mathcal{D}_\varepsilon} \min_{\theta_0 \in \Lambda_0, \theta_1 \in \Lambda_1} D_f(M_{\theta_0} || M_{\theta_1})$, where M_{θ_ν} is the induced marginal under P_{θ_ν} .

The more general problem of private m -ary hypothesis testing is also an interesting but challenging one. In this setting, the X_i 's can follow one of m distributions P_0, P_1, \dots, P_{m-1} . Consequently, the Y_i 's can follow one of m distributions M_0, M_1, \dots, M_{m-1} . In this case, the utility can be defined as the average f -divergence between any two distributions: $1/(m(m-1)) \sum_{i \neq j} D_f(M_i || M_j)$, or the worst case one: $\min_{i \neq j} D_f(M_i || M_j)$.

Non-exchangeable utility functions. The utility studied in this paper was measured by functions that are exchangeable, i.e. the utility does not depend on the naming (labelling) or topology of the private and privatized data (X and Y). This made sense for statistical learning applications that depend on information theoretic quantities such as f -divergences and mutual information. However, in some other applications, the utility might be defined over $\mathcal{X} \cup \mathcal{Y}$ in a metric space, where there exists a natural measure of distance (or distortion) between the data points. In this case, we can formulate the problem as a distortion

minimization one

$$\text{minimize}_{Q \in \mathcal{D}_\varepsilon} \sum_{x,y} d(x,y) P(x) Q(y|x),$$

where $d(x,y)$ is some distortion metric. Wang et al. (2014a) studied this problem, and showed that the mechanism $Q(y|x) \propto e^{\varepsilon(1-d(x,y))}/(k-1+e^\varepsilon)$ achieves near optimal performance when ε is large enough, which is the low privacy regime. Notice that when Hamming distance is used, $d(x,y) = \mathbb{I}(x \neq y)$, this recovers the randomized response mechanism exactly. This provides a starting point for generalizing the search for optimal mechanisms under non-exchangeable utility functions.

7. Proof of Theorems 2 and 4

We start the proof with a few definitions, a lemma, and a general result that applies to any convex utility function that obeys a mild assumption.

Recall that for an input alphabet \mathcal{X} with $|\mathcal{X}| = k$, we represent the set of ε -locally differentially private mechanisms that lead to output alphabets \mathcal{Y} with $|\mathcal{Y}| = \ell$ by $\mathcal{D}_{\varepsilon,\ell}$. The set of all ε -locally differentially private mechanisms is given by $\mathcal{D}_\varepsilon = \cup_{\ell \in \mathbb{N}} \mathcal{D}_{\varepsilon,\ell}$. A utility function $U(Q)$ is convex in Q if $U(\lambda Q^{(1)} + (1-\lambda)Q^{(2)}) \leq \lambda U(Q^{(1)}) + (1-\lambda)U(Q^{(2)})$ for any $\lambda \in (0,1)$. Convex utility functions are ubiquitous in information theory and statistics.

Assumption 1 *If a $k \times \ell$ privatization mechanism $Q^{(1)} \in \mathcal{D}_{\varepsilon,\ell}$ is obtained by deleting an all-zero column of a $k \times \ell + 1$ privatization mechanism $Q^{(2)} \in \mathcal{D}_{\varepsilon,\ell+1}$, then $U(Q^{(1)}) = U(Q^{(2)})$.*

Naturally, one would expect that if we delete the zero columns of a privatization mechanism $Q^{(2)}$ to obtain a new privatization mechanism $Q^{(1)}$, we would still get the same utility. This is because a “reasonable” utility function should not depend on output alphabets with zero probability.

Theorem 21 *If $U(Q)$ is a convex utility function that satisfies Assumption 1, then the following holds*

$$\max_{Q \in \mathcal{D}_\varepsilon} U(Q) = \max_{Q \in \cup_{\ell=1}^k \mathcal{D}_{\varepsilon,\ell}} U(Q).$$

This theorem implies that among all ε -locally differentially private mechanisms, we only need to consider those that lead to output alphabets of size $\ell \leq k$. In other words, enlarging the input alphabet cannot further maximize the utility. The proof of Theorem 21 is given in Section 7.1.

Lemma 22 *A $k \times \ell$ conditional distribution Q is ε -locally differentially private if and only if it can be written as $Q = S\Theta$, where S is a $k \times \ell$ matrix with $S_{ij} \in [1, e^\varepsilon]$ and $\Theta = \text{diag}(\theta_1, \dots, \theta_\ell)$ with its diagonal entries in \mathbb{R}_+ .*

The proof of Lemma 22 is provided in Section 7.2. With the above results, we are now ready to prove Theorems 2 and 4. By Lemma 22, for any $Q \in \mathcal{D}_{\varepsilon,\ell}$ we have that $Q_j = \theta_j S_j$.

Suppose $U(Q) = \sum_{j \in [\ell]} \mu(Q_j)$, where μ is a sublinear function. Since μ is sublinear, it is convex and $\mu(\theta_j S_j) = \theta_j \mu(S_j)$. $U(Q)$ is convex in Q because

$$\begin{aligned} U\left(\lambda Q^{(1)} + (1-\lambda)Q^{(2)}\right) &= \sum_{j \in [\ell]} \mu\left(\lambda \theta_j^{(1)} S_j^{(1)} + (1-\lambda) \theta_j^{(2)} S_j^{(2)}\right) \\ &\leq \sum_{j \in [\ell]} \lambda \mu\left(\theta_j^{(1)} S_j^{(1)}\right) + (1-\lambda) \mu\left(\theta_j^{(2)} S_j^{(2)}\right) \\ &= \lambda U\left(Q^{(1)}\right) + (1-\lambda) U\left(Q^{(2)}\right), \end{aligned} \quad (25)$$

for any $\lambda \in (0, 1)$. Furthermore, $U(Q)$ satisfies Assumption 1 because $\mu(Q_j) = 0$ whenever $\theta_j = 0$. Let $Q^* = S^* \Theta^* \in \arg \max_{Q \in \cup_{\ell=1}^k \mathcal{D}_{\varepsilon, \ell}} U(Q)$ and note that by Theorem 21, $U(Q^*) = \max_{Q \in \mathcal{D}_{\varepsilon}} U(Q)$. Suppose that Q^* is of dimensions $k \times \ell$, where $\ell \leq k$. Each of the ℓ columns of Q^* can be expressed as a convex combination of the columns of $S^{(k)}$, the staircase pattern matrix. This is because the 2^k columns of $S^{(k)}$ are the corner points of the cube $[1, e^{\varepsilon}]^k$ and each $S_j^* \in [1, e^{\varepsilon}]^k$. Therefore, $S_j^* = \sum_{i=1}^{2^k} \lambda_{ij} S_i^{(k)}$, where $\lambda_{ij} \geq 0$ for all i and j , and $\sum_{i=1}^{2^k} \lambda_{ij} = 1$ for all j . Create the 2^k -dimensional vector $\tilde{\theta}$ such that $\tilde{\theta}_i = \sum_{j=1}^{\ell} \lambda_{ij} \theta_j^*$ and let $\tilde{Q} = S^{(k)} \tilde{\Theta}$.

$$\begin{aligned} U(Q^*) - U(\tilde{Q}) &= \sum_{j=1}^{\ell} \mu(S_j^*) \theta_j^* - \sum_{i=1}^{2^k} \mu\left(\left(\sum_{j=1}^{\ell} \lambda_{ij} \theta_j^*\right) S_i^{(k)}\right) \\ &= \sum_{j=1}^{\ell} \mu\left(\sum_{i=1}^{2^k} \lambda_{ij} S_i^{(k)}\right) \theta_j^* - \sum_{i=1}^{2^k} \sum_{j=1}^{\ell} \lambda_{ij} \theta_j^* \mu\left(S_i^{(k)}\right) \\ &= \sum_{j=1}^{\ell} \theta_j^* \left\{ \mu\left(\sum_{i=1}^{2^k} \lambda_{ij} S_i^{(k)}\right) - \sum_{i=1}^{2^k} \lambda_{ij} \mu\left(S_i^{(k)}\right) \right\} \\ &\leq 0, \end{aligned} \quad (26)$$

by the convexity of $\mu(z)$ and the non-negativity of θ_j^* 's. Moreover, observe that since $S^{(k)} \tilde{\theta} = \mathbf{1}$, $\tilde{\theta}$ is a valid choice for the linear program of (11). This implies that

$$\max_{S^{(k)} \theta = \mathbf{1}, \theta \geq 0} \sum_{j=1}^{2^k} \mu\left(S_j^{(k)}\right) \theta_j \geq U(\tilde{Q}) \geq U(Q^*) = \max_{Q \in \mathcal{D}_{\varepsilon}} U(Q)$$

On the other hand, for any $\tilde{Q} = S^{(k)} \tilde{\Theta}$, where $\tilde{\theta}$ is valid for the linear program of (11), we have that $\tilde{Q} \in \mathcal{D}_{\varepsilon, 2^k} \subset \mathcal{D}_{\varepsilon}$ and therefore, $\max_{S^{(k)} \theta = \mathbf{1}, \theta \geq 0} \sum_{j=1}^{2^k} \mu\left(S_j^{(k)}\right) \theta_j \leq \max_{Q \in \mathcal{D}} U(Q)$.

Thus, $\max_{S^{(k)} \theta = \mathbf{1}, \theta \geq 0} \sum_{j=1}^{2^k} \mu\left(S_j^{(k)}\right) \theta_j = \max_{Q \in \mathcal{D}} U(Q)$. This proves Theorem 4.

The polytope given by $S^{(k)} \theta = \mathbf{1}$ and $\theta \geq 0$ is a closed and bounded one. Thus, the linear program of (11) is bounded and has a solution, say θ^* , at one of the corner points of the polytope. Since there are k equality constraints given by $S^{(k)} \theta = \mathbf{1}$ and 2^k inequality constraints given by $\theta \geq 0$, any corner point, including θ^* , cannot have more than k non-zero entries. Form \tilde{S} by deleting the columns of $S^{(k)}$ corresponding to zero entries of θ^* .

Similarly, form $\tilde{\theta}$ by deleting the zero entries of θ^* and let $\tilde{Q} = \tilde{S}\tilde{\Theta}$, where $\tilde{\Theta} = \text{diag}\tilde{\theta}$. It is easy to verify that $U(\tilde{Q}) = U(Q^*) = \mu^T \theta^*$; hence, \tilde{Q} solves linear program of (11). Moreover, \tilde{Q} has at most k columns and $\tilde{S}_{ij} = \{1, e^\varepsilon\}$. Therefore, \tilde{Q} is a staircase mechanism of dimension $k \times \ell$, where $\ell \leq k$.

7.1 Proof of Theorem 21

We start the proof of Theorem 21 with an important lemma the proof of which is presented in Section 7.3.

Lemma 23 *The set of all $k \times \ell$, ε -locally differentially private mechanisms $\mathcal{D}_{\varepsilon,\ell}$ forms a closed and bounded polytope in $\mathbb{R}_+^{k\ell}$. Moreover, let Q be a corner point of the polytope formed by $\mathcal{D}_{\varepsilon,\ell}$, then Q has at most k non-zero columns.*

Fix an $\ell > k$. Since $U(Q)$ is convex in Q , it suffices to consider the corner points of $\mathcal{D}_{\varepsilon,\ell}$ when maximizing $U(Q)$ subject to $Q \in \mathcal{D}_{\varepsilon,\ell}$. By Lemma 23, any $Q^{(1)}$, a $k \times \ell$ corner point of $\mathcal{D}_{\varepsilon,\ell}$, has at most k non-zero columns. Therefore, the privatization mechanism $Q^{(2)}$, obtained by deleting the all-zero columns of $Q^{(1)}$, has at most k columns. Notice that $Q^{(2)} \in \cup_{i=1}^k \mathcal{D}_{\varepsilon,i}$. Since $U(Q)$ satisfies Assumption 1, we have that $U(Q^{(1)}) = U(Q^{(2)})$ and therefore, it suffices to consider $Q \in \cup_{i=1}^k \mathcal{D}_{\varepsilon,i}$ when maximizing $U(Q)$ subject to $Q \in \mathcal{D}_{\varepsilon,\ell}$. Thus,

$$\begin{aligned} \sup_{Q \in \mathcal{D}_\varepsilon} U(Q) &= \sup_{\ell \in \mathbb{N}} \left\{ \max_{Q \in \mathcal{D}_{\varepsilon,\ell}} U(Q) \right\} \\ &= \sup_{\ell \in \mathbb{N}} \left\{ \max_{Q \in \cup_{i=1}^k \mathcal{D}_{\varepsilon,i}} U(Q) \right\} \\ &= \max_{Q \in \cup_{i=1}^k \mathcal{D}_{\varepsilon,i}} U(Q), \end{aligned} \tag{27}$$

which finishes the proof.

7.2 Proof of Lemma 22

Claim 1 *Let $Q \in \mathcal{D}_{\varepsilon,\ell}$. If $Q_{ij} = 0$ for some $j \in \{1, \dots, \ell\}$ then $Q_{ij} = 0$ for all $i \in \{1, \dots, k\}$.*

Proof Assume that $Q_{i_1 j} = 0$ and $Q_{i_2 j} \neq 0$ for some $i_1, i_2 \in \{1, \dots, k\}$. It is obvious that $q(y_j | x_{i_2}) \leq q(y_j | x_{i_1}) e^\varepsilon$ is not satisfied. Therefore, Q is not a locally differentially private mechanism. \blacksquare

It is easy to check that any $k \times \ell$ stochastic matrix $Q = S\Theta$, where Θ is a diagonal matrix with non-negative entries and S is a $k \times \ell$ matrix with $S_{ij} \in [1, e^\varepsilon]$, satisfies the local differential privacy constraints. Thus, $Q \in \mathcal{D}_{\varepsilon,\ell}$. On the other hand, assume that $Q \in \mathcal{D}_{\varepsilon,\ell}$. If $Q_{ij} = 0$ for some j then by Claim 1 we have that $Q_{ij} = 0$ for all i and therefore, we can set $\theta_j = 0$ and $S_{ij} = 1$ for all i . If $Q_{ij} > 0$ then by Claim 1 we have that $Q_{ij} > 0$ for all i . In this case, let $\theta_j = \min_i Q_{ij}$ and observe that $\theta_j > 0$ since $Q_{ij} > 0$ for all i . Let $S_{ij} = Q_{ij}/\theta_j$, then it is clear (from the definition of θ_j) that $S_{ij} \geq 1$. On the other hand, from the differential privacy constraints, we have that $Q_{ij} \leq Q_{kj} e^\varepsilon$ for all k and thus,

$Q_{ij} \leq \min_k Q_{kj} e^\varepsilon$ which proves that $S_{ij} = Q_{ij} / \min_k Q_{kj} \leq e^\varepsilon$. This establishes that any $Q \in \mathcal{D}_{\varepsilon, \ell}$ can be written as $Q = S\Theta$, where Θ is a diagonal matrix with non-negative entries and S is a $k \times \ell$ matrix with $S_{ij} \in [1, e^\varepsilon]$.

7.3 Proof of Lemma 23

We start by showing that $\mathcal{D}_{\varepsilon, \ell}$ forms a closed and bounded polytope in $\mathbb{R}_+^{k\ell}$. We are interested in studying the corner points of the polytope formed by $\mathcal{D}_{\varepsilon, \ell}$ because convex utility functions are maximized at one of these corner points whenever the space of privatization mechanisms is restricted to $\mathcal{D}_{\varepsilon, \ell}$.

Claim 2 *A privatization mechanism $Q \in \mathcal{D}_{\varepsilon, \ell}$ if and only if for all $x, x' \in \mathcal{X}$ and all $y \in \mathcal{Y}$ we have that $Q(y|x) \leq Q(y|x') e^\varepsilon$.*

Proof By definition, Q is differentially private if for all $x, x' \in \mathcal{X}$ and all $B \subseteq \mathcal{Y}$ we have that $Q(B|x) \leq Q(B|x') e^\varepsilon$. By choosing $B = \{y\}$ for some $y \in \mathcal{Y}$ the first direction of the above lemma is proven. In order to prove the other direction, assume that for all $x, x' \in \mathcal{X}$ and all $y \in \mathcal{Y}$ we have that $Q(y|x) \leq Q(y|x') e^\varepsilon$. Then for any $B \subseteq \mathcal{Y}$, the following holds

$$\begin{aligned} \sum_{y \in B} Q(y|x) &\leq \sum_{y \in B} Q(y|x') e^\varepsilon \\ \Leftrightarrow Q(B|x) &\leq Q(B|x') e^\varepsilon. \end{aligned} \tag{28}$$

■

Let $Q \in \mathcal{D}_{\varepsilon, \ell}$, then by Claim 2, it is easy to see that Q must satisfy $\ell k(k-1)$ inequalities of the form $Q(y|x) \leq Q(y|x') e^\varepsilon$. These inequalities can be compactly represented by

$$\tilde{A}q \leq 0,$$

where $q = [Q(y_1|x_1), \dots, Q(y_1|x_k), \dots, Q(y_\ell|x_1), \dots, Q(y_\ell|x_k)]^T$ and \tilde{A} is a $\ell k(k-1) \times k\ell$ matrix that contains all the local differential privacy linear constraints. Observe that there is a one-to-one mapping between Q and q . Here is an example for the case when $k = \ell = 2$

$$\underbrace{\begin{bmatrix} 1 & -e^\varepsilon & 0 & 0 \\ -e^\varepsilon & 1 & 0 & 0 \\ 0 & 0 & 1 & -e^\varepsilon \\ 0 & 0 & -e^\varepsilon & 1 \end{bmatrix}}_{\tilde{A}} \begin{bmatrix} Q(y_1|x_1) \\ Q(y_1|x_2) \\ Q(y_2|x_1) \\ Q(y_2|x_2) \end{bmatrix} \leq 0.$$

Moreover, since Q is a row stochastic matrix (a conditional distribution) it satisfies $Q\mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ represents the all ones vector of appropriate dimensions. This condition can be rewritten as

$$Bq = \mathbf{1},$$

where B is a $k \times k\ell$ binary matrix. For the case when $k = \ell = 2$, we have that

$$\underbrace{\begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}}_B \begin{bmatrix} Q(y_1|x_1) \\ Q(y_1|x_2) \\ Q(y_2|x_1) \\ Q(y_2|x_2) \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Finally, observe that $Q(y|x) \geq 0$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. These constraints can be incorporated as follows. Let $A = \begin{bmatrix} \tilde{A}^T & -I_{\ell k} \end{bmatrix}^T$, where $I_{\ell k}$ is the $\ell k \times \ell k$ identity matrix, then $Aq \leq 0$. To summarize, $Q \in \mathcal{D}_{\varepsilon, \ell}$ if and only if

$$\begin{aligned} Aq &\leq 0 \\ Bq &= \mathbf{1}. \end{aligned} \tag{29}$$

Therefore, the set of all $k \times \ell$, ε -locally differentially private mechanisms $\mathcal{D}_{\varepsilon, \ell}$ forms a convex polytope in $\mathbb{R}_+^{k\ell}$.

We now proceed to proving that if Q is a corner point of the polytope formed by $\mathcal{D}_{\varepsilon, \ell}$, then Q has at most k non-zero columns. This claim is obvious for all $k \times \ell$ privatization mechanisms with $\ell \leq k$. Therefore, we restrict our attention to the case where $\ell > k$. Let A_j be the matrix including all the inequality constraints imposed on the j^{th} column of Q . Observe that the rows of A_j form a subset of the rows of A , defined in (29), and recall that there are $k(k-1)$ differential privacy and k non-negativity inequality constraints imposed on the j^{th} column of Q . Therefore, A_j is a $k^2 \times k$ matrix and we have that $A_j Q_j \leq 0$, where Q_j represents the j^{th} column of Q . By Claim 1, we know that Q_j is either equal to zero or contains non-zero entries.

Claim 3 *In what follows, the term linearly independent inequality constraints refers to linear independent rows of A_j .*

- If $Q_j = 0$, then k linearly independent inequality constraints are achieved with equality.
- If $Q_j \neq 0$, then at most $k-1$ linearly independent inequality constraints can be achieved with equality.

Proof In fact, the number of linearly independent inequality constraints (achieved or not) cannot exceed k because A_j has a rank less than or equal to k . If $Q_j = 0$, then the k non-negativity inequality constraints are achieved with equality and it is easy to see that they are all linearly independent (in fact, they form an orthonormal basis to \mathbb{R}^k). This proves the first part of the claim. We now establish the second part of the claim by showing that if $Q_j \neq 0$, we cannot have k linearly independent inequality constraints achieved with equality. Assume that $Q_j \neq 0$ and let \tilde{A}_j be the matrix including the largest collection of linearly independent rows of A_j corresponding to the inequality constraints that are achieved with equality. In other words, $\tilde{A}_j Q_j = 0$. If \tilde{A}_j contains k rows, then its rank is equal to k . However, this implies that $Q_j = 0$, a contradiction. Therefore, at most $k-1$ linearly independent inequality constraints can be achieved with equality when $Q_j \neq 0$. ■

Suppose that Q is a corner point of $\mathcal{D}_{\varepsilon, \ell}$ and out of its ℓ columns, $\ell_{>0}$ are non-zero and $\ell_{=0}$ ($\ell_{=0} = \ell - \ell_{>0}$) are zero. Moreover, assume that the number of non-zero columns of Q is larger than k (i.e., $\ell_{>0} > k$). In this case, from Claim 3, we can see that Q achieves at most $\ell_{>0}(k-1) + (\ell - \ell_{>0})k$ linearly independent inequality constraints with equality. Furthermore, at most k additional linearly independent equality constraints (linearly independent rows of the matrix B defined in (29)) can be met by Q . Therefore, the total number of linearly independent constraints that Q achieves with equality is at most $\ell_{>0}(k-1) + (\ell - \ell_{>0})k + k = -\ell_{>0} + (\ell + 1)k < \ell k$, where the last strict inequality follows from the fact that $\ell_{>0} > k$. This implies that Q cannot be a corner point of $\mathcal{D}_{\varepsilon, \ell}$. Therefore, any corner point of $\mathcal{D}_{\varepsilon, \ell}$ must have at most k non-zero columns.

8. Proofs for Hypothesis Testing

8.1 Proof of Theorem 5

Let $T = \{x : P_0(x) \geq P_1(x)\}$. In other words, $P_0(T) - P_1(T) = \max_{A \subseteq \mathcal{X}} P_0(A) - P_1(A)$. Recall that for a given P_0 and P_1 , the binary mechanism is defined as a staircase mechanism with only two outputs $y \in \{0, 1\}$ satisfying

$$Q(0|x) = \begin{cases} \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } P_0(x) \geq P_1(x), \\ \frac{1}{1+e^\varepsilon} & \text{if } P_0(x) < P_1(x). \end{cases} \quad Q(1|x) = \begin{cases} \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } P_0(x) < P_1(x), \\ \frac{1}{1+e^\varepsilon} & \text{if } P_0(x) \geq P_1(x). \end{cases} \quad (30)$$

Lemma 24 *For any pair of distributions P_0 and P_1 , there exists a positive ε^* that depends on P_0 and P_1 such that for all $y \in \mathcal{Y}$, all $\ell \in \mathbb{N}$, and all $Q \in \mathcal{D}_{\varepsilon, \ell}$ with $\varepsilon \leq \varepsilon^*$, we have that*

$$\frac{(e^\varepsilon - 1)P_0(T^c) + 1}{(e^\varepsilon - 1)P_1(T^c) + 1} \leq \frac{M_0(y)}{M_1(y)} \leq \frac{(e^\varepsilon - 1)P_0(T) + 1}{(e^\varepsilon - 1)P_1(T) + 1}.$$

Moreover, the above upper and lower bounds are achieved by the binary mechanism given in (30).

Observe that because $P_0(T) \geq P_1(T)$ and $P_0(T^c) \leq P_1(T^c)$, the direction of the above inequalities makes sense.

Let \tilde{M}_ν be the induced marginal for the binary mechanism when P_ν is the original distribution. Following the analysis techniques developed in Kairouz et al. (2013), we define hypothesis testing region $R(\tilde{M}_0, \tilde{M}_1)$ as the convex hull of all achievable probabilities of missed detection and false alarm, when testing whether $\nu = 0$ or $\nu = 1$ based on Y_{bin} distributed as \tilde{M}_ν :

$$R(\tilde{M}_0, \tilde{M}_1) \equiv \text{conv}\left(\{(\tilde{M}_1(S), \tilde{M}_0(S^c)) : \forall S \subseteq \mathcal{Y}\}\right),$$

where $S \in \mathcal{Y}$ is the accept region for hypothesis $\nu = 0$. For the binary mechanism, this ends up being a very simple triangular region. The slopes defining the two sides of the triangular region are: $-\max_S \tilde{M}_0(S)/\tilde{M}_1(S) = -((e^\varepsilon - 1)P_0(T) + 1)/((e^\varepsilon - 1)P_1(T) + 1)$ and $-\min_S \tilde{M}_0(S^c)/\tilde{M}_1(S^c) = -((e^\varepsilon - 1)P_0(T^c) + 1)/((e^\varepsilon - 1)P_1(T^c) + 1)$.

For any other mechanism satisfying the ε -local differential privacy for $\varepsilon \leq \varepsilon^*$, the above lemma implies that for any choice of the rejection region S , the slopes satisfy




Figure 10: Hypothesis testing regions for two mechanisms.

$-M_0(S)/M_1(S) \geq -((e^\varepsilon - 1)P_0(T) + 1)/((e^\varepsilon - 1)P_1(T) + 1)$ and $-M_0(S^c)/M_1(S^c) \leq -((e^\varepsilon - 1)P_0(T^c) + 1)/((e^\varepsilon - 1)P_1(T^c) + 1)$. In the hypothesis testing region, this implies that

$$R(M_0, M_1) \subseteq R(\tilde{M}_0, \tilde{M}_1) ,$$

as in the following Figure 10.

From Theorem 2.5 of Kairouz et al. (2013), we know that this implies a certain Markov property. Precisely, let Y_{bin} denote the output of the binary mechanism, and Y_{dp} denote the output of any ε -local differentially private mechanism. Then, it follows that there exists a coupling of Y_{bin} and Y_{dp} such that they form a Markov chain: $\nu - Y_{\text{bin}} - Y_{\text{dp}}$, where ν is the hypothesis on P_ν whether the data was generated from $\nu = 0$ or $\nu = 1$. Then, it follows from the data processing inequality of f -divergences that

$$D_f(\tilde{M}_0, \tilde{M}_1) \geq D_f(M_0, M_1) .$$

It follows that there is no algorithm with larger f -divergence than the binary mechanism.

8.2 Proof of Lemma 24

We start by showing that the binary mechanism achieves the upper and lower bounds presented in the statement of the lemma. Let M_0^B and M_1^B denote the induced marginals

under the binary mechanism given in (30). For $\nu \in \{0, 1\}$, we have that

$$\begin{aligned} M_\nu^B(0) &= \sum_{x \in \mathcal{X}} P_0(x) Q(0|x) = \frac{1}{e^\varepsilon + 1} ((e^\varepsilon - 1) P_\nu(T) + 1) \\ M_\nu^B(1) &= \sum_{x \in \mathcal{X}} P_0(x) Q(1|x) = \frac{1}{e^\varepsilon + 1} ((e^\varepsilon - 1) P_\nu(T^c) + 1). \end{aligned} \quad (31)$$

Computing $M_0^B(0)/M_1^B(0)$ and $M_0^B(1)/M_1^B(1)$ we see that the binary mechanism achieves the upper and lower bounds for all values of ε .

As in Lemma 22, for any $\ell \in \mathbb{N}$, $Q \in \mathcal{D}_{\varepsilon, \ell}$ can be represented as $Q = S\Theta$, where $S \in [1, e^\varepsilon]^{k \times \ell}$ and $\Theta = \text{diag}(\theta_1, \dots, \theta_\ell)$ with its diagonal entries in \mathbb{R}_+ . We now show that for any $\ell \in \mathbb{N}$ and any $Q \in \mathcal{D}_{\varepsilon, \ell}$, the following upper bound holds

$$\frac{M_0(y)}{M_1(y)} = \frac{\sum_{i \in [k]} P_0(x_i) S_{ij}}{\sum_{i \in [k]} P_1(x_i) S_{ij}} \leq \frac{(e^\varepsilon - 1) P_0(T) + 1}{(e^\varepsilon - 1) P_1(T) + 1},$$

for all $y \in \mathcal{Y}$ and sufficiently small ε . The above expression can be alternatively written as

$$\begin{aligned} (e^\varepsilon - 1) (P_0(T) - P_1(T)) + (e^\varepsilon - 1) \sum_{i \in [k]} (S_{ij} - 1) (P_0(T) P_1(x_i) - P_1(T) P_0(x_i)) \\ - \sum_{i \in [k]} (S_{ij} - 1) (P_0(x_i) - P_1(x_i)) \geq 0, \end{aligned} \quad (32)$$

where $S_j \in [1, e^\varepsilon]^k$. Equation (32) is linear in S_j and is therefore minimized (and maximized) at the corner points of $[1, e^\varepsilon]^{k \times \ell}$, a cube in $\mathbb{R}_+^{k \times \ell}$. The corner points of this cube are given by the staircase patterns: $S_j \in \{1, e^\varepsilon\}^k$. To begin with, let S_j be a staircase pattern with $T_j = \{x_i : S_{ij} = e^\varepsilon\} \neq T$. Then Equation (32) is equivalent to

$$\begin{aligned} (e^\varepsilon - 1) \{(P_0(T) - P_1(T)) - (P_0(T_j) - P_1(T_j))\} \\ + (e^\varepsilon - 1)^2 \{P_0(T) P_1(T_j) - P_1(T) P_0(T_j)\} \geq 0. \end{aligned} \quad (33)$$

Using the fact that $P_0(T) - P_1(T) > P_0(T_j) - P_1(T_j)$ for all $T_j \neq T$, the inequality in (32) holds true for all ε whenever $P_0(T) P_1(T_j) \geq P_1(T) P_0(T_j)$. If $P_0(T) P_1(T_j) < P_1(T) P_0(T_j)$, then the inequality in (32) holds true for all $\varepsilon \leq \varepsilon(j)$, where

$$\varepsilon(j) = \log \left(\frac{(P_0(T) - P_1(T)) - (P_0(T_j) - P_1(T_j))}{P_1(T) P_0(T_j) - P_0(T) P_1(T_j)} + 1 \right) > 0.$$

On the other hand, it is easy to verify that when $T_j = T$, we have that

$$\begin{aligned} (e^\varepsilon - 1) \{(P_0(T) - P_1(T)) - (P_0(T_j) - P_1(T_j))\} \\ + (e^\varepsilon - 1) (P_0(T) P_1(T_j) - P_1(T) P_0(T_j)) = 0, \end{aligned} \quad (34)$$

for all ε . In this case, set $\varepsilon(j) = 0$ and $\varepsilon_1 = \min_{j \in [2^k]} \varepsilon(j)$. Therefore, for any $\ell \in \mathbb{N}$ and any $Q \in \mathcal{D}_{\varepsilon, \ell}$, the upper bound in the statement of the lemma holds for all $\varepsilon \leq \varepsilon_1$.

We now show that for any $\ell \in \mathbb{N}$ and any $Q \in \mathcal{D}_{\varepsilon, \ell}$, the following lower bound holds

$$\frac{(e^\varepsilon - 1) P_0(T^c) + 1}{(e^\varepsilon - 1) P_1(T^c) + 1} \leq \frac{M_0(y)}{M_1(y)} = \frac{\sum_{i \in [k]} P_0(x_i) S_{ij}}{\sum_{i \in [k]} P_1(x_i) S_{ij}},$$

for all $y \in \mathcal{Y}$ and sufficiently small ε . The above expression can be alternatively written as

$$\begin{aligned} (e^\varepsilon - 1) (P_0(T) - P_1(T)) + (e^\varepsilon - 1) \sum_{i \in [k]} (S_{ij} - 1) (P_0(T) P_1(x_i) - P_1(T) P_0(x_i)) \\ + e^\varepsilon \sum_{i \in [k]} (S_{ij} - 1) (P_0(x_i) - P_1(x_i)) \geq 0, \end{aligned} \quad (35)$$

where $S_j \in [1, e^\varepsilon]^k$. Equation (35) is linear in S_j and is therefore minimized at the corner points of $[1, e^\varepsilon]^k$, a cube in \mathbb{R}_+^k . The corner points of this cube are given by staircase patterns: $S_j \in \{1, e^\varepsilon\}^k$. To begin with, let S_j be a staircase pattern with $T_j = \{x_i : S_{ij} = e^\varepsilon\} \neq T^c$, then Equation (35) is equivalent to

$$\begin{aligned} (e^\varepsilon - 1) \{(P_0(T) - P_1(T)) + e^\varepsilon (P_0(T_j) - P_1(T_j))\} \\ + (e^\varepsilon - 1)^2 \{P_0(T) P_1(T_j) - P_1(T) P_0(T_j)\} \geq 0. \end{aligned} \quad (36)$$

Using the fact that $P_0(T) - P_1(T) > P_1(T_j) - P_0(T_j)$ for all $T_j \neq T^c$, then for sufficiently small ε , Equation (35) can be written as

$$\varepsilon \{(P_0(T) - P_1(T)) - (P_1(T_j) - P_0(T_j))\} + O(\varepsilon^2) > 0.$$

This proves that there exists a positive $\varepsilon(j)$ such that the left hand side of Equation (36) is positive for all $\varepsilon \leq \varepsilon(j)$. On the other hand, it is easy to verify that when $T_j = T^c$, we have that

$$\begin{aligned} (e^\varepsilon - 1) \{(P_0(T) - P_1(T)) + e^\varepsilon (P_0(T_j) - P_1(T_j))\} \\ + (e^\varepsilon - 1) (P_0(T) P_1(T_j) - P_1(T) P_0(T_j)) = 0, \end{aligned} \quad (37)$$

for all ε . In this case, let $\varepsilon_1(j) = 0$ and let $\varepsilon_2 = \min_{j \in [2^k]} \varepsilon(j)$. Therefore, for any $\ell \in \mathbb{N}$ and any $Q \in \mathcal{D}_{\varepsilon, \ell}$, the lower bound in the statement of the lemma holds for all $\varepsilon \leq \varepsilon_2$. To conclude, let $\varepsilon^* = \min(\varepsilon_1, \varepsilon_2)$. Then both, the upper and lower bounds, hold for all $\varepsilon \leq \varepsilon^*$.

8.3 Proof of Theorem 6

The total variation (TV) distance $\|M_0 - M_1\|_{\text{TV}}$ is a special case of f -divergence $D_f(M_0 \| M_1)$ with $f(x) = \frac{1}{2}|x - 1|$. Therefore, by Theorem 4, we have that

$$\begin{aligned} \max_{Q \in \mathcal{D}_\varepsilon} \|M_0 - M_1\|_{\text{TV}} = \underset{\theta}{\text{maximize}} \quad & \mu^T \theta \\ \text{subject to} \quad & S^{(k)} \theta = \mathbf{1} \\ & \theta \geq 0, \end{aligned} \quad (38)$$

where $\mu_j = \mu \left(S_j^{(k)} \right) = \frac{1}{2} \left| \sum_{i \in [k]} (P_0(x_i) - P_1(x_i)) S_{ij}^{(k)} \right|$ for $j \in \{1, \dots, 2^k\}$ and $S^{(k)}$ is the $k \times 2^k$ staircase pattern matrix given in Definition 3.

The polytope given by $S^{(k)}\theta = \mathbf{1}$ and $\theta \geq 0$ is a closed and bounded one. Thus, there is no duality gap and solving the above linear program is equivalent to solving its dual

$$\begin{aligned} & \underset{\alpha}{\text{minimize}} && \mathbf{1}^T \alpha \\ & \text{subject to} && S^{(k)T} \alpha \geq \mu. \end{aligned} \tag{39}$$

Note that any satisfiable solution α^* to (39) provides an upper bound to (38) since $\max \mu^T \theta = \min \mathbf{1}^T \alpha \leq \mathbf{1}^T \alpha^*$. Let $T = \{x : P_0(x) \geq P_1(x)\}$ and $T_j = \{x_i : S_{ij}^{(k)} = e^\varepsilon\}$ for $j \in [2^k]$. Consider the following choice of dual variable

$$\alpha_i^* = \frac{1}{2} \frac{e^\varepsilon - 1}{e^\varepsilon + 1} |P_0(x_i) - P_1(x_i)|,$$

for $i \in [k]$. Observe that

$$\begin{aligned} \mathbf{1}^T \alpha^* &= \frac{1}{2} \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \sum_{i \in [k]} |P_0(x_i) - P_1(x_i)| \\ &= \frac{1}{2} \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \|P_0 - P_1\|_1 \\ &= \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \|P_0 - P_1\|_{\text{TV}}. \end{aligned} \tag{40}$$

We claim that α^* is a feasible dual variable for all values of ε . In order to prove that α^* is a feasible dual variable, we show that $S^{(k)T} \alpha^* - \mu_j \geq 0$ for all $j \in [2^k]$ and all ε . For all $j \in [2^k]$, we have that

$$\begin{aligned} g_j &= 2 \left(S^{(k)T} \alpha^* - \mu_j \right) \\ &= \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \sum_{i \in [k]} |P_0(x_i) - P_1(x_i)| S_{ij}^{(k)} - \left| \sum_{i \in [k]} (P_0(x_i) - P_1(x_i)) S_{ij}^{(k)} \right| \\ &= \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \left\{ \sum_{x_i \in T} (P_0(x_i) - P_1(x_i)) S_{ij}^{(k)} + \sum_{x_i \in T^c} (P_1(x_i) - P_0(x_i)) S_{ij}^{(k)} \right\} \\ &\quad - \left| \sum_{x_i \in T} (P_0(x_i) - P_1(x_i)) S_{ij}^{(k)} - \sum_{x_i \in T^c} (P_1(x_i) - P_0(x_i)) S_{ij}^{(k)} \right|. \end{aligned} \tag{41}$$

Notice that we have arranged the equation such that all the summands are non-negative. Without loss of generality, we will assume that

$$\sum_{x_i \in T} (P_0(x_i) - P_1(x_i)) S_{ij}^{(k)} \geq \sum_{x_i \in T^c} (P_1(x_i) - P_0(x_i)) S_{ij}^{(k)}.$$

From the equality $\sum_{x_i \in T} (P_0(x_i) - P_1(x_i)) = \sum_{x_i \in T^c} (P_1(x_i) - P_0(x_i))$ and the fact that $S_{ij}^{(k)} \in \{1, e^\varepsilon\}$ for all i and j , it follows that

$$e^{-\varepsilon} \sum_{x_i \in T} (P_0(x_i) - P_1(x_i)) S_{ij}^{(k)} \leq \sum_{x_i \in T^c} (P_1(x_i) - P_0(x_i)) S_{ij}^{(k)}. \tag{42}$$

This is true because the right-hand side is minimized when the $S_{ij}^{(k)}$'s for $x_i \in T^c$ are all equal to 1 and the left-hand side is maximized when the $S_{ij}^{(k)}$'s for $x_i \in T$ are all equal to e^ε . Now, (41) can be written as

$$g_j = \frac{1}{e^\varepsilon + 1} \left\{ -2 \sum_{x_i \in T} (P_0(x_i) - P_1(x_i)) S_{ij}^{(k)} + 2e^\varepsilon \sum_{x_i \in T^c} (P_1(x_i) - P_0(x_i)) S_{ij}^{(k)} \right\} \geq 0,$$

where the last inequality follows from (42).

This establishes the satisfiability of α^* for all ε which, in turn, shows that (40) is indeed an upper bound to the primal problem. It remains to show that this upper bound can be achieved via the binary mechanism. To this extent, recall that for a given P_0 and P_1 , the binary mechanism is defined as a staircase mechanism with only two outputs $y \in \{0, 1\}$ satisfying

$$Q(0|x) = \begin{cases} \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } P_0(x) \geq P_1(x), \\ \frac{1}{1+e^\varepsilon} & \text{if } P_0(x) < P_1(x). \end{cases} \quad Q(1|x) = \begin{cases} \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } P_0(x) < P_1(x), \\ \frac{1}{1+e^\varepsilon} & \text{if } P_0(x) \geq P_1(x). \end{cases} \quad (43)$$

Computing the TV distance between M_0 and M_1 under (43), we get that

$$\|M_0 - M_1\|_{\text{TV}} = \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \|P_0 - P_1\|_{\text{TV}}.$$

Hence, the binary mechanism in (43) achieves the upper bound in (40). This proves the optimality of the binary mechanism for all ε .

8.4 Proof of Theorem 8

The Kullback-Leibler (KL) divergence $D_{\text{kl}}(M_0||M_1)$ is a special f -divergence $D_f(M_0||M_1)$ with $f(x) = x \log x$. Therefore, by Theorem 4, we have that

$$\begin{aligned} \max_{Q \in \mathcal{D}_\varepsilon} D_{\text{kl}}(M_0||M_1) &= \underset{\theta}{\text{maximize}} && \mu^T \theta \\ &\text{subject to} && S^{(k)} \theta = \mathbf{1} \\ &&& \theta \geq 0, \end{aligned} \quad (44)$$

where $\mu_j = \mu \left(S_j^{(k)} \right) = \sum_{i \in [k]} P_0(x_i) S_{ij}^{(k)} \log \left(\frac{\sum_{i \in [k]} P_0(x_i) S_{ij}^{(k)}}{\sum_{i \in [k]} P_1(x_i) S_{ij}^{(k)}} \right)$ for $j \in \{1, \dots, 2^k\}$ and $S^{(k)}$ is the $k \times 2^k$ staircase pattern matrix given in Definition 3.

The polytope given by $S^{(k)} \theta = \mathbf{1}$ and $\theta \geq 0$ is a closed and bounded one. Thus, there is no duality gap and solving the above linear program is equivalent to solving its dual

$$\begin{aligned} \underset{\alpha}{\text{minimize}} & \quad \mathbf{1}^T \alpha \\ \text{subject to} & \quad S^{(k)T} \alpha \geq \mu. \end{aligned} \quad (45)$$

Note that any satisfiable solution α^* to (45) provides an upper bound to (44) since $\max \mu^T \theta = \min \mathbb{1}^T \alpha \leq \mathbb{1}^T \alpha^*$. Let $T = \{x : P_0(x) \geq P_1(x)\}$ and $T_j = \{x_i : S_{ij}^{(k)} = e^\varepsilon\}$ for $j \in [2^k]$. Set $j_i = \{j : T_j = x_i\}$ for $i \in [k]$, and consider the following choice of dual variable

$$\alpha_i^* = \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ (e^\varepsilon + k - 2) \mu \left(S_{j_i}^{(k)} \right) - \sum_{l \in [k], l \neq i} \mu \left(S_{j_l}^{(k)} \right) \right\},$$

for $i \in [k]$. Observe that since $T_{j_i} = x_i$ we have that $P_\nu(T_{j_i}) = P_\nu(x_i)$ and since

$$\begin{aligned} \mu_j &= \sum_{i \in [k]} P_0(x_i) S_{ij}^{(k)} \log \left(\frac{\sum_{i \in [k]} P_0(x_i) S_{ij}^{(k)}}{\sum_{i \in [k]} P_1(x_i) S_{ij}^{(k)}} \right) \\ &= (P_0(T_j)(e^\varepsilon - 1) + 1) \log \frac{(P_0(T_j)(e^\varepsilon - 1) + 1)}{(P_1(T_j)(e^\varepsilon - 1) + 1)} \end{aligned} \quad (46)$$

we have that

$$\begin{aligned} \mathbb{1}^T \alpha^* &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \sum_{i \in [k]} \left\{ (e^\varepsilon + k - 2) \mu \left(S_{j_i}^{(k)} \right) - \sum_{l \in [k], l \neq i} \mu \left(S_{j_l}^{(k)} \right) \right\} \\ &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ (e^\varepsilon + k - 2) \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) - \sum_{i \in [k]} \sum_{l \in [k], l \neq i} \mu \left(S_{j_l}^{(k)} \right) \right\} \\ &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ (e^\varepsilon + k - 2) \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) - (k - 1) \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) \right\} \\ &= \frac{1}{(e^\varepsilon + k - 1)} \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) \\ &= \frac{1}{(e^\varepsilon + k - 1)} \sum_{i \in [k]} (P_0(x_i)(e^\varepsilon - 1) + 1) \log \frac{(P_0(x_i)(e^\varepsilon - 1) + 1)}{(P_1(x_i)(e^\varepsilon - 1) + 1)}. \end{aligned} \quad (47)$$

We claim that α^* is a feasible dual variable for sufficiently large ε . In order to prove that α^* is a feasible dual variable, we show that $S^{(k)T} \alpha^* - \mu_j \geq 0$ for all $j \in [2^k]$ for all $\varepsilon \geq \varepsilon^*$, where ε^* is a positive quantity that depends on the priors P_0 and P_1 . Using the facts that

$$\begin{aligned} \log(a + e^\varepsilon b) &= \varepsilon + \log b + O(e^{-\varepsilon}) \\ \frac{1}{e^\varepsilon + k - 1} &= e^{-\varepsilon} + O(e^{-2\varepsilon}), \end{aligned} \quad (48)$$

for large ε , we get that

$$\begin{aligned} \mu_j &= (P_0(T_j)(e^\varepsilon - 1) + 1) \log \frac{(P_0(T_j)(e^\varepsilon - 1) + 1)}{(P_1(T_j)(e^\varepsilon - 1) + 1)} \\ &= \left(P_0(T_j) \log \frac{P_0(T_j)}{P_1(T_j)} \right) e^\varepsilon + (1 - P_0(T_j)) \log \frac{P_0(T_j)}{P_1(T_j)} + O(e^{-\varepsilon}). \end{aligned} \quad (49)$$

On the other hand,

$$\begin{aligned}
 S_j^{(k)T} \alpha^* &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ \sum_{i \in [k]} S_{ij}^{(k)} (e^\varepsilon + k - 2) \left(P_0(x_i) \log \frac{P_0(x_i)}{P_1(x_i)} e^\varepsilon + O(1) \right) \right\} \\
 &\quad - \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ \sum_{i \in [k]} \sum_{l \in [k], l \neq i} S_{ij}^{(k)} \left(P_0(x_l) \log \frac{P_0(x_l)}{P_1(x_l)} e^\varepsilon + O(1) \right) \right\} \\
 &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left(\left(\sum_{x_i \in T_j} P_0(x_i) \log \frac{P_0(x_i)}{P_1(x_i)} \right) e^{3\varepsilon} + O(e^{2\varepsilon}) \right) \\
 &= \left(\sum_{x_i \in T_j} P_0(x_i) \log \frac{P_0(x_i)}{P_1(x_i)} \right) e^\varepsilon + O(1). \tag{50}
 \end{aligned}$$

Assume, to begin with, that $j \neq \{j_1, j_2, \dots, j_k\}$. Then

$$S_j^{(k)T} \alpha^* - \mu_j = \left(P_0(T_j) \log \frac{P_0(T_j)}{P_1(T_j)} - \sum_{x_i \in T_j} P_0(x_i) \log \frac{P_0(x_i)}{P_1(x_i)} \right) e^\varepsilon + O(1).$$

Notice that for $j \neq \{j_1, j_2, \dots, j_k\}$, $P_0(T_j) \log \frac{P_0(T_j)}{P_1(T_j)} > \sum_{x_i \in T_j} P_0(x_i) \log \frac{P_0(x_i)}{P_1(x_i)}$ by the log-sum inequality. Therefore, there exists a $\varepsilon(j) > 0$ such that $S_j^{(k)T} \alpha^* - \mu_j \geq 0$ for all $\varepsilon \geq \varepsilon(j)$. If $j \in \{j_1, j_2, \dots, j_k\}$, it is not hard to check that $S_j^{(k)T} \alpha^* - \mu_j = 0$ for all ε . In this case, set $\varepsilon(j) = 0$. This establishes the satisfiability of α^* for all $\varepsilon \geq \varepsilon^* = \max_{j \in [2^k]} \varepsilon(j)$. The satisfiability of α^* , in turn, shows that (47) is indeed an upper bound to the primal problem. It remains to show that this upper bound can be achieved via the randomized response. To this extent, recall that the randomized response is given by

$$Q(y|x) = \begin{cases} \frac{e^\varepsilon}{|\mathcal{X}| - 1 + e^\varepsilon} & \text{if } y = x, \\ \frac{1}{|\mathcal{X}| - 1 + e^\varepsilon} & \text{if } y \neq x. \end{cases} \tag{51}$$

Computing the KL divergence between M_0 and M_1 under (51), we get that

$$D_{\text{kl}}(M_0 || M_1) = \frac{1}{(e^\varepsilon + k - 1)} \sum_{i \in [k]} (P_0(x_i) (e^\varepsilon - 1) + 1) \log \frac{(P_0(x_i) (e^\varepsilon - 1) + 1)}{(P_1(x_i) (e^\varepsilon - 1) + 1)}.$$

Hence, the randomized response in (51) achieves the upper bound in (47). This proves the optimality of the randomized response for all $\varepsilon \geq \varepsilon^*$.

8.5 Proof of Theorem 7

We start the proof with a fundamental bound on the symmetrized KL divergence between the M_0 and M_1 .

Lemma 25 *For any $\varepsilon \geq 0$, let Q be any conditional distribution that guarantees ε differential privacy. Then for any pair of distributions P_0 and P_1 , the induced marginals M_0 and M_1 must satisfy the bound*

$$D_{\text{kl}}(M_0||M_1) + D_{\text{kl}}(M_1||M_0) \leq 4(e^\varepsilon - 1)^2 \|P_0 - P_1\|_{\text{TV}}^2.$$

The above lemma appears as Theorem 1 in Duchi et al. (2013). By Lemma 25, we have that

$$\text{OPT} = \max_{Q \in \mathcal{D}_\varepsilon} D_{\text{kl}}(M_0||M_1) \leq 4(e^\varepsilon - 1)^2 \|P_0 - P_1\|_{\text{TV}}^2. \quad (52)$$

Let M_0^B and M_1^B be the marginals obtained by using the binary mechanism given in (14). By Corollary 11, we have that $\|M_0^B - M_1^B\|_{\text{TV}} = \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \|P_0 - P_1\|_{\text{TV}}$. Consequently, by applying Pinsker's inequality to the KL divergence between M_0^B and M_1^B we get that

$$\begin{aligned} \text{BIN} &= D_{\text{kl}}(M_0^B||M_1^B) \\ &\geq 2\|M_0^B - M_1^B\|_{\text{TV}}^2 \\ &= 2\left(\frac{e^\varepsilon - 1}{e^\varepsilon + 1}\right)^2 \|P_0 - P_1\|_{\text{TV}}^2. \end{aligned} \quad (53)$$

Combining (52) and (53) we get that $\text{BIN} \geq \frac{1}{2(e^\varepsilon + 1)^2} \text{OPT}$ which was to be shown.

9. Proofs for Information Preservation

9.1 Proof of Theorem 12

By Theorem 4, we have that

$$\begin{aligned} \max_{Q \in \mathcal{D}_\varepsilon} I(X; Y) &= \underset{\theta}{\text{maximize}} && \mu^T \theta \\ &\text{subject to} && S^{(k)} \theta = \mathbf{1} \\ &&& \theta \geq 0, \end{aligned} \quad (54)$$

where $\mu_j = \mu\left(S_j^{(k)}\right) = \sum_{i \in [k]} P(x_i) S_{ij}^{(k)} \log\left(\frac{S_{ij}^{(k)}}{\sum_{i \in [k]} P(x_i) S_{ij}^{(k)}}\right)$ for $j \in \{1, \dots, 2^k\}$ and $S^{(k)}$ is the $k \times 2^k$ staircase pattern matrix given in Definition 3. The polytope given by $S^{(k)} \theta = \mathbf{1}$ and $\theta \geq 0$ is a closed and bounded one. Thus, there is no duality gap and solving the above linear program is equivalent to solving its dual

$$\begin{aligned} &\underset{\alpha}{\text{minimize}} && \mathbf{1}^T \alpha \\ &\text{subject to} && S^{(k)T} \alpha \geq \mu. \end{aligned} \quad (55)$$

Note that any satisfiable solution α^* to (55) provides an upper bound to (54) since $\max \mu^T \theta = \min \mathbf{1}^T \alpha \leq \mathbf{1}^T \alpha^*$. Let $T_j = \{x_i : S_{ij}^{(k)} = e^\varepsilon\}$ and set $j_1 = \{j : T_j = T\}$ and $j_2 = \{j : T_j = T^c\}$. Consider the following choice of dual variable

$$\alpha_i^* = \frac{1}{(e^\varepsilon + 1)(e^\varepsilon - 1)} \begin{cases} \frac{e^\varepsilon \mu(S_{j_1}^{(k)}) - \mu(S_{j_2}^{(k)})}{|T|} & \forall i \in T \\ \frac{e^\varepsilon \mu(S_{j_2}^{(k)}) - \mu(S_{j_1}^{(k)})}{|T^c|} & \forall i \in T^c \end{cases}.$$

Observe that since $T_{j_1} = T$, $T_{j_2} = T^c$, and

$$\mu_j = P(T_j) e^\varepsilon \log \frac{e^\varepsilon}{P(T_j^c) + e^\varepsilon P(T_j)} + P(T_j^c) \log \frac{1}{P(T_j^c) + e^\varepsilon P(T_j)}, \quad (56)$$

we have that

$$\begin{aligned} \mathbb{1}^T \alpha^* &= \frac{1}{(e^\varepsilon + 1)(e^\varepsilon - 1)} \left\{ \sum_{i \in T} \frac{1}{|T|} \left(e^\varepsilon \mu(S_{j_1}^{(k)}) - \mu(S_{j_2}^{(k)}) \right) \right. \\ &= \left. + \sum_{i \in T^c} \frac{1}{|T^c|} \left(e^\varepsilon \mu(S_{j_2}^{(k)}) - \mu(S_{j_1}^{(k)}) \right) \right\} \\ &= \frac{1}{(e^\varepsilon + 1)} \left(\mu(S_{j_1}^{(k)}) + \mu(S_{j_1}^{(k)}) \right) \\ &= \frac{1}{e^\varepsilon + 1} \left\{ P(T) e^\varepsilon \log \frac{e^\varepsilon}{P(T^c) + e^\varepsilon P(T)} + P(T^c) \log \frac{1}{P(T^c) + e^\varepsilon P(T)} \right\} + \\ &= \frac{1}{e^\varepsilon + 1} \left\{ P(T^c) e^\varepsilon \log \frac{e^\varepsilon}{P(T) + e^\varepsilon P(T^c)} + P(T) \log \frac{1}{P(T) + e^\varepsilon P(T^c)} \right\}. \quad (57) \end{aligned}$$

We claim that α^* is a feasible dual variable for sufficiently small ε . In order to prove that α^* is a feasible dual variable, we show that $(S^{(k)T} \alpha^*)_j - \mu_j \geq 0$ for all $j \in \{1, \dots, 2^k\}$ and all $\varepsilon \leq \varepsilon^*$, where ε^* is a positive quantity that depends on P . Using the following facts

$$\begin{aligned} e^\varepsilon &= 1 + \varepsilon + \frac{1}{2}\varepsilon^2 + O(\varepsilon^3) \\ \log(a + e^\varepsilon b) &= b\varepsilon + \frac{b(1-b)}{2}\varepsilon^2 + O(\varepsilon^3) \\ \frac{1}{1 + e^\varepsilon} &= \frac{1}{2} - \frac{1}{4}\varepsilon + O(\varepsilon^2), \quad (58) \end{aligned}$$

for small ε , we get that

$$\begin{aligned} \mu_j &= P(T_j) e^\varepsilon \log \frac{e^\varepsilon}{P(T_j^c) + e^\varepsilon P(T_j)} + P(T_j^c) \log \frac{1}{P(T_j^c) + e^\varepsilon P(T_j)} \\ &= P(T_j) e^\varepsilon \varepsilon - (P(T_j) (e^\varepsilon - 1) + 1) \log(P(T_j) (e^\varepsilon - 1) + 1) \\ &= \frac{1}{2} P(T_j) P(T_j^c) \varepsilon^2 + O(\varepsilon^3). \quad (59) \end{aligned}$$

On the other hand,

$$\begin{aligned}
 \left(S^{(k)T} \alpha^* \right)_j &= S_j^{(k)T} \alpha^* \\
 &= \frac{1}{(e^\varepsilon + 1)(e^\varepsilon - 1)} \left\{ \sum_{i \in T} \frac{S_{ij}^{(k)}}{|T|} \left(e^\varepsilon \mu(S_{j_1}^{(k)}) - \mu(S_{j_2}^{(k)}) \right) \right. \\
 &\quad \left. + \sum_{i \in T^c} \frac{S_{ij}^{(k)}}{|T^c|} \left(e^\varepsilon \mu(S_{j_2}^{(k)}) - \mu(S_{j_1}^{(k)}) \right) \right\} \\
 &= \frac{1}{(e^\varepsilon + 1)(e^\varepsilon - 1)} \left(e^\varepsilon \mu(S_{j_1}^{(k)}) - \mu(S_{j_2}^{(k)}) \right) \left(\frac{|T_j \cap T|}{|T|} e^\varepsilon + \frac{|T_j^c \cap T|}{|T|} \right) \\
 &\quad + \frac{1}{(e^\varepsilon + 1)(e^\varepsilon - 1)} \left(e^\varepsilon \mu(S_{j_2}^{(k)}) - \mu(S_{j_1}^{(k)}) \right) \left(\frac{|T_j \cap T^c|}{|T^c|} e^\varepsilon + \frac{|T_j^c \cap T^c|}{|T^c|} \right) \\
 &= \frac{1}{(e^\varepsilon + 1)} \left(\frac{1}{2} P(T) P(T^c) \varepsilon^2 + O(\varepsilon^3) \right) \left\{ \frac{|T_j \cap T^c|}{|T^c|} + \frac{|T_j^c \cap T^c|}{|T^c|} \right. \\
 &\quad \left. + \frac{|T_j \cap T|}{|T|} + \frac{|T_j^c \cap T|}{|T|} + O(\varepsilon) \right\} \\
 &= \frac{1}{2} P(T) P(T^c) \varepsilon^2 + O(\varepsilon^3), \tag{60}
 \end{aligned}$$

where we have used the facts that $T_{j_1} = T$, $T_{j_2} = T^c$, and

$$\begin{aligned}
 \mu(S_{j_1}^{(k)}) &= \frac{1}{2} P(T) P(T^c) \varepsilon^2 + O(\varepsilon^3) \\
 \mu(S_{j_2}^{(k)}) &= \frac{1}{2} P(T) P(T^c) \varepsilon^2 + O(\varepsilon^3). \tag{61}
 \end{aligned}$$

Let $f(z) = |z - \frac{1}{2}|$, $g(z) = -z \log z - (1 - z) \log(1 - z)$, and $h(z) = z(1 - z)$ for $0 \leq z \leq 1$. On the one hand, g and h are monotonically increasing over $0 \leq z \leq \frac{1}{2}$ and monotonically decreasing over $\frac{1}{2} \leq z \leq 1$ but on the other hand, f is monotonically decreasing over $0 \leq z \leq \frac{1}{2}$ and monotonically increasing over $\frac{1}{2} \leq z \leq 1$. Therefore,

$$\begin{aligned}
 T \in \arg \min_{A \subseteq \mathcal{X}} \left| P(A) - \frac{1}{2} \right| &\Leftrightarrow T \in \arg \max_{A \subseteq \mathcal{X}} -P(A) \log P(A) - P(A^c) \log P(A^c) \\
 &\Leftrightarrow T \in \arg \max_{A \subseteq \mathcal{X}} P(A) P(A^c). \tag{62}
 \end{aligned}$$

Since the set T was chosen so that it maximizes $P(T) P(T^c)$, we have that $P(T) P(T^c) \geq P(T_j) P(T_j^c)$ for all $j \in \{1, \dots, 2^k\}$. Assume, to begin with, that $j \neq \{j_1, j_2\}$. Then by the uniqueness of the maximizer assumption stated in the theorem, we have that $P(T) P(T^c) > P(T_j) P(T_j^c)$.

$$\left(S^T \alpha^* \right)_j - \mu_j = \frac{1}{2} (P(T) P(T^c) - P(T_j) P(T_j^c)) \varepsilon^2 + O(\varepsilon^3),$$

and thus, there exists an ε^* that depends on P such that $\left(S^{(k)T} \alpha^* \right)_j - \mu_j \geq 0$ for all $\varepsilon \leq \varepsilon^*$.

If $j = \{j_1, j_2\}$, it is not hard to check that $\left(S^{(k)T} \alpha^* \right)_j - \mu_j = 0$ for all ε . This establishes

the satisfiability of α^* for all $\varepsilon \leq \varepsilon^*$ which proves an upper bound on the primal problem (given in (57)). It remains to show that the upper bound can be indeed achieved via the binary mechanism. To this extent, recall that the binary mechanism is given by

$$Q(0|x) = \begin{cases} \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } x \in T, \\ \frac{1}{1+e^\varepsilon} & \text{if } x \notin T. \end{cases} \quad Q(1|x) = \begin{cases} \frac{e^\varepsilon}{1+e^\varepsilon} & \text{if } x \notin T, \\ \frac{1}{1+e^\varepsilon} & \text{if } x \in T. \end{cases} \quad (63)$$

Computing the $I(X;Y)$ under (63), we get that

$$I(X;Y) = \frac{1}{e^\varepsilon + 1} \left\{ P(T) e^\varepsilon \log \frac{e^\varepsilon}{P(T^c) + e^\varepsilon P(T)} + P(T^c) \log \frac{1}{P(T^c) + e^\varepsilon P(T)} \right\} + \frac{1}{e^\varepsilon + 1} \left\{ P(T^c) e^\varepsilon \log \frac{e^\varepsilon}{P(T) + e^\varepsilon P(T^c)} + P(T) \log \frac{1}{P(T) + e^\varepsilon P(T^c)} \right\}. \quad (64)$$

Hence, the binary mechanism in (63) achieves the upper bound in (57). This proves the optimality of the binary mechanism for all $\varepsilon \leq \varepsilon^*$.

9.2 Proof of Theorem 13

We start by proving an upper bound on $\max_{Q \in \mathcal{D}_\varepsilon} I(X;Y)$ which is tight for $\varepsilon \leq 1$. Recall that by Theorem 4, we have that

$$\begin{aligned} \text{OPT} = \max_{Q \in \mathcal{D}_\varepsilon} I(X;Y) &= \underset{\theta}{\text{maximize}} \quad \sum_{j=1}^{2^k} \mu_j \theta_j \\ &\text{subject to} \quad S^{(k)} \theta = \mathbf{1} \\ &\quad \theta \geq 0, \end{aligned}$$

where

$$\begin{aligned} \mu_j &= \mu(S_j^{(k)}) \\ &= \sum_{i \in [k]} P(x_i) S_{ij}^{(k)} \log \left(\frac{S_{ij}^{(k)}}{\sum_{i \in [k]} P(x_i) S_{ij}^{(k)}} \right) \\ &= P(T_j) e^\varepsilon \varepsilon - (P(T_j) (e^\varepsilon - 1) + 1) \log (P(T_j) (e^\varepsilon - 1) + 1), \end{aligned} \quad (65)$$

$T_j = \{i : S_{ij}^{(k)} = e^\varepsilon\}$, and $S^{(k)}$ is the $k \times 2^k$ staircase pattern matrix given in Definition 3.

Lemma 26 *For all distributions P and all ε , the following bound holds*

$$\text{OPT} = \max_{Q \in \mathcal{D}_\varepsilon} I(X;Y) \leq \left(\max_j \mu_j \right) \frac{k}{e^\varepsilon + k - 1}.$$

The proof of this lemma is given in Section 9.3. In what follows, we will make the dependency of μ_j on $P(T_j)$ and ε explicit by writing $\mu_j(P(T_j), \varepsilon)$ for μ_j . From the proof of Theorem 12, we have that the partition set T defined in (20) is given by $T \in \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$. It is easy to check that the binary mechanism given in (19) achieves the following utility

$$\text{BIN} = \frac{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)}{e^\varepsilon + 1}.$$

Lemma 27 *For all distributions P and all $\varepsilon \leq 1$, the following bound holds*

$$\frac{\max_j \mu_j}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} \leq 1.$$

The proof of the above lemma is given in Section 9.4. Combining the results of lemmas 26 and 27 we get that

$$\begin{aligned} \frac{\text{OPT}}{\text{BIN}} &\leq \frac{\max_j \mu_j}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} \frac{k}{e^\varepsilon + k - 1} (e^\varepsilon + 1) \\ &\leq \frac{k}{e^\varepsilon + k - 1} (e^\varepsilon + 1) \\ &\leq e^\varepsilon + 1, \end{aligned} \tag{66}$$

for all $\varepsilon \leq 1$. This concludes the proof.

9.3 Proof of Lemma 26

To begin with, since $S_1^{(k)} = \mathbf{1} = \frac{1}{e^\varepsilon} S_{2^k}^{(k)}$ and μ is homogenous, we have that $\theta_1 \mu_1 + \theta_{2^k} \mu_{2^k} = (\frac{1}{e^\varepsilon} \theta_1 + \theta_{2^k}) \mu_{2^k}$. Therefore, the following two maximization problems are equivalent

$$\begin{aligned} \underset{\theta}{\text{maximize}} \quad & \sum_{j=1}^{2^k} \mu_j \theta_j & \underset{\theta}{\text{maximize}} \quad & \sum_{j=1}^{2^k-1} \tilde{\mu}_j \theta_j \\ \text{subject to} \quad & S^{(k)} \theta = \mathbf{1} & \text{subject to} \quad & \tilde{S}^{(k)} \theta = \mathbf{1} \\ & \theta \geq 0 & & \theta \geq 0, \end{aligned}$$

where $\tilde{\mu}_j = \mu_{j+1}$ and $\tilde{S}^{(k)}$ is obtained by deleting the first column of $S^{(k)}$. Moreover, using the fact that $\max_{j \in [2^k-1]} \tilde{\mu}_j \leq \max_{j \in [2^k]} \mu_j$ and weak duality, we get that

$$\begin{aligned} \underset{\theta}{\text{maximize}} \quad & \tilde{\mu}^T \theta & \leq & \left(\max_{j \in [2^k-1]} \tilde{\mu}_j \right) \underset{\theta}{\text{maximize}} \quad \mathbf{1}^T \theta \\ \text{subject to} \quad & \tilde{S}^{(k)} \theta = \mathbf{1} & & \text{subject to} \quad \tilde{S}^{(k)} \theta = \mathbf{1} \\ & \theta \geq 0 & & \theta \geq 0 \\ & & \leq & \left(\max_{j \in [2^k]} \mu_j \right) \underset{\alpha}{\text{minimize}} \quad \mathbf{1}^T \alpha \\ & & & \text{subject to} \quad \tilde{S}^{(k)T} \alpha \geq \mathbf{1}. \end{aligned} \tag{67}$$

Consider the following choice of dual variable $\alpha_i^* = \frac{1}{e^\varepsilon + k - 1}$. We claim that α^* is satisfiable. This can be easily verified by noting that

$$\left(\tilde{S}^{(k)T} \alpha^* \right)_j = \tilde{S}_j^{(k)T} \alpha^* = \frac{|T_j| e^\varepsilon + (k - |T_j|)}{e^\varepsilon + k - 1} = \frac{|T_j| (e^\varepsilon - 1) + k}{e^\varepsilon + k - 1} \geq 1$$

where the last inequality holds since $|T_j| \geq 1$ (this is true because we have deleted the first column of $S^{(k)}$). Therefore, $\text{OPT} \leq (\max_j \mu_j) \mathbf{1}^T \alpha^* = (\max_j \mu_j) \frac{k}{e^\varepsilon + k - 1}$ which was to be shown.

9.4 Proof of Lemma 27

Let $\mu(z, \varepsilon)$ be the function obtained by replacing $P(T_j)$ by the continuous variable $z \in [0, 1]$ in $\mu_j(P(T_j), \varepsilon)$. Taking the derivative of $\mu(z, \varepsilon)$ with respect to z we get

$$\mu'(z, \varepsilon) = e^\varepsilon z - (e^\varepsilon - 1) - (e^\varepsilon - 1) \log(z(e^\varepsilon - 1) + 1).$$

Observe that $\mu'(z, \varepsilon) > 0$ for all

$$z < z^*(\varepsilon) = \frac{1}{e^\varepsilon - 1} \left(e^{\left\{ \frac{e^\varepsilon \varepsilon}{e^\varepsilon - 1} - 1 \right\}} - 1 \right),$$

$\mu'(z, \varepsilon) < 0$ for all $z > z^*(\varepsilon)$, and $\mu'(z, \varepsilon) = 0$ for $z = z^*(\varepsilon)$. Combining this with the fact that $\mu(0, \varepsilon) = \mu(1, \varepsilon) = 0$ we get that $\mu(z, \varepsilon) \geq 0$ for all $z \in [0, 1]$ and for any fixed ε , $\mu(z, \varepsilon)$ is maximized at $z^*(\varepsilon)$.

Set $x^* \in \arg \max_{x \in \mathcal{X}} P(x)$ and fix an $\varepsilon \leq 1$. We will treat the following three cases separately.

Case 1: $P(x^*) \in [1 - z^*(\varepsilon), 1]$.

Claim 4 *Let $T = \{x^*\}$. Then $\{T, T^c\} = \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$ and $\max_{A \subseteq \mathcal{X}} \mu(P(A), \varepsilon) = \max(\mu(P(T), \varepsilon), \mu(P(T^c), \varepsilon))$.*

Proof Observe that $z^*(\varepsilon) \leq \frac{1}{2}$ for all ε and $T^c = \mathcal{X} \setminus \{x^*\}$. The function $f(z) = z(1 - z)$ decreases over the range $[\frac{1}{2}, 1] \supseteq [1 - z^*(\varepsilon), 1]$. Thus, for all $A \supset T$, $P(T)P(T^c) > P(A)P(A^c)$ because $P(T) \geq 1 - z^*(\varepsilon)$. This proves that $T \in \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$ and for all $A \supset T$, $A \notin \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$. Using a similar approach, we can show that $T^c \in \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$ and for all $A \subset T^c$, $A \notin \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$. Therefore, $\{T, T^c\} = \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$. This proves the first part of the claim. The function $\mu(z, \varepsilon)$ increases over the range $[0, z^*(\varepsilon)]$. Thus, for all $A \subseteq T^c$, $\mu(P(A), \varepsilon) \leq \mu(P(T^c), \varepsilon)$ because $P(T^c) \leq z^*(\varepsilon)$. On the other hand, note that $\mu(z, \varepsilon)$ decreases over the range $[z^*(\varepsilon), 1]$ which includes the range $[1 - z^*(\varepsilon), 1]$. Thus, for all A such that $A \supseteq T$, $\mu(P(A), \varepsilon) \leq \mu(P(T), \varepsilon)$ because $P(T) \geq 1 - z^*(\varepsilon)$. This proves that $\max(\mu(P(T), \varepsilon), \mu(P(T^c), \varepsilon)) = \max_{A \subseteq \mathcal{X}} \mu(P(A), \varepsilon)$. \blacksquare

Using the above claim, we can conclude that the partition set T defined in (20) is equal to $\{x^*\}$ and

$$\begin{aligned} \frac{\max_j \mu_j}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} &= \frac{\max_{A \subseteq \mathcal{X}} \mu(P(A), \varepsilon)}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} \\ &\leq \frac{\max_{A \subseteq \mathcal{X}} \mu(P(A), \varepsilon)}{\max(\mu(P(T), \varepsilon), \mu(P(T^c), \varepsilon))} \\ &= 1. \end{aligned} \tag{68}$$

Case 2: $P(x^*) \in [\frac{1}{2}, 1 - z^*(\varepsilon)]$. Using the first part of the proof of Claim 4, we can show that if $T = \{x^*\}$, then $\{T, T^c\} = \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)$. Therefore, the partition set T

defined in (20) is equal to $\{x^*\}$ and

$$\begin{aligned} \frac{\max_j \mu_j}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} &= \frac{\max_{A \subseteq \mathcal{X}} \mu(P(A), \varepsilon)}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} \\ &\leq \frac{\mu(z^*(\varepsilon), \varepsilon)}{\mu(P(x^*), \varepsilon) + \mu(1 - P(x^*), \varepsilon)} \\ &\leq 1. \end{aligned} \tag{69}$$

Case 3: $P(x^*) \in [0, \frac{1}{2}]$.

Claim 5 *There exists a set $A \subset \mathcal{X}$ such that $\frac{1}{2} - P(x^*) \leq P(A) \leq \frac{1}{2}$.*

Proof Without loss of generality, assume that the sequence $P(x_i)$, $i \in [k]$, is sorted in increasing order. Let $l^* = \min\{l : \sum_{i=1}^l P(x_i) \geq \frac{1}{2}\}$. From the definition of l^* , $P(\{x_1, \dots, x_{l^*-1}\}) < \frac{1}{2}$ and $P(\{x_1, \dots, x_{l^*}\}) \geq \frac{1}{2}$. Further,

$$P(\{x_1, \dots, x_{l^*-1}\}) = P(\{x_1, \dots, x_{l^*}\}) - P(x_{l^*})$$

and since $x^* \in \arg \max_{x \in \mathcal{X}} P(x)$, $P(x_{l^*}) \leq P(x^*)$. Therefore, if $A = \{x_1, \dots, x_{l^*-1}\}$, then $\frac{1}{2} - P(x^*) \leq P(A) \leq \frac{1}{2}$. \blacksquare

Let $P(T) = \min\{P(B) : B \in \arg \max_{A \subseteq \mathcal{X}} P(A)P(A^c)\}$. We claim that $\frac{1}{4} \leq P(T) \leq \frac{1}{2}$. The upper bound on $P(T)$ follows immediately from its definition. To prove the lower bound on $P(T)$, consider the set A given in Claim 5 and observe that

$$\begin{aligned} P(T) &\geq \max(P(x^*), P(A)) \\ &\geq \max(P(x^*), \frac{1}{2} - P(x^*)) \\ &\geq \frac{1}{4}. \end{aligned} \tag{70}$$

All the inequalities follow from Claim 5 and the fact that $P(x^*) \in [0, \frac{1}{2}]$.

Since $\frac{1}{4} \leq P(T) \leq \frac{1}{2}$, we have that $\frac{1}{2} \leq P(T^c) \leq \frac{3}{4}$. Moreover, the function $\mu(z, \varepsilon)$ decreases over the range $[z^*(\varepsilon), 1] \supset [\frac{1}{2}, \frac{3}{4}]$ and increases over the range $[\frac{1}{4}, z^*(\varepsilon)]$. Therefore, $\mu(P(T^c), \varepsilon) \geq \mu(\frac{3}{4}, \varepsilon)$ and $\mu(P(T), \varepsilon) \geq \min(\mu(\frac{1}{2}, \varepsilon), \mu(\frac{1}{4}, \varepsilon))$. Putting it all together, we have that

$$\begin{aligned} \frac{\max_j \mu_j}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} &= \frac{\max_{A \subseteq \mathcal{X}} \mu(P(A), \varepsilon)}{\mu(P(T), \varepsilon) + \mu(P(T^c), \varepsilon)} \\ &\leq \frac{\mu(z^*(\varepsilon), \varepsilon)}{\min(\mu(\frac{1}{2}, \varepsilon), \mu(\frac{1}{4}, \varepsilon)) + \mu(\frac{3}{4}, \varepsilon)} \\ &\leq 1. \end{aligned} \tag{71}$$

9.5 Proof of Theorem 14

By Theorem 4, we have that

$$\begin{aligned} \max_{Q \in \mathcal{D}_\varepsilon} I(X; Y) &= \underset{\theta}{\text{maximize}} && \mu^T \theta \\ &\text{subject to} && S^{(k)} \theta = \mathbf{1} \\ &&& \theta \geq 0, \end{aligned} \tag{72}$$

where $\mu_j = \mu \left(S_j^{(k)} \right) = \sum_{i \in [k]} P(x_i) S_{ij}^{(k)} \log \left(\frac{S_{ij}^{(k)}}{\sum_{i \in [k]} P(x_i) S_{ij}^{(k)}} \right)$ for $j \in \{1, \dots, 2^k\}$ and $S^{(k)}$ is the $k \times 2^k$ staircase pattern matrix given in Definition 3. The polytope given by $S^{(k)}\theta = \mathbf{1}$ and $\theta \geq 0$ is a closed and bounded one. Thus, there is no duality gap and solving the above linear program is equivalent to solving its dual

$$\begin{aligned} & \underset{\alpha}{\text{minimize}} && \mathbf{1}^T \alpha \\ & \text{subject to} && S^{(k)T} \alpha \geq \mu. \end{aligned} \tag{73}$$

Note that any satisfiable solution α^* to (73) provides an upper bound to (72) since $\max \mu^T \theta = \min \mathbf{1}^T \alpha \leq \mathbf{1}^T \alpha^*$. Let $T_j = \{x_i : S_{ij}^{(k)} = e^\varepsilon\}$ and set $j_i = \{j : T_j = i\}$ for $i \in \{1, \dots, k\}$. Consider the following choice of dual variable

$$\alpha_i^* = \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ (e^\varepsilon + k - 2) \mu \left(S_{j_i}^{(k)} \right) - \sum_{l \in [k], l \neq i} \mu \left(S_{j_l}^{(k)} \right) \right\},$$

for $i \in \{1, \dots, k\}$. Observe that since $T_{j_i} = i$ we have that $P(T_{j_i}) = P(x_i)$ and since

$$\mu_j = P(T_j) e^\varepsilon \log \frac{e^\varepsilon}{P(T_j^c) + e^\varepsilon P(T_j)} + P(T_j^c) \log \frac{1}{P(T_j^c) + e^\varepsilon P(T_j)}, \tag{74}$$

we have that

$$\begin{aligned} \mathbf{1}^T \alpha^* &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \sum_{i \in [k]} \left\{ (e^\varepsilon + k - 2) \mu \left(S_{j_i}^{(k)} \right) - \sum_{l \in [k], l \neq i} \mu \left(S_{j_l}^{(k)} \right) \right\} \\ &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ (e^\varepsilon + k - 2) \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) - \sum_{i \in [k]} \sum_{l \in [k], l \neq i} \mu \left(S_{j_l}^{(k)} \right) \right\} \\ &= \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ (e^\varepsilon + k - 2) \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) - (k - 1) \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) \right\} \\ &= \frac{1}{(e^\varepsilon + k - 1)} \sum_{i \in [k]} \mu \left(S_{j_i}^{(k)} \right) \\ &= \frac{1}{(e^\varepsilon + k - 1)} \sum_{i \in [k]} \left\{ P(x_i) e^\varepsilon \log \frac{e^\varepsilon}{P(x_i)(e^\varepsilon - 1) + 1} \right. \\ &\quad \left. + (1 - P(x_i)) \log \frac{1}{P(x_i)(e^\varepsilon - 1) + 1} \right\}. \end{aligned} \tag{75}$$

We claim that α^* is a feasible dual variable for sufficiently large ε . In order to prove that α^* is a feasible dual variable, we show that $\left(S^{(k)T} \alpha^* \right)_j - \mu_j \geq 0$ for all $j \in \{1, \dots, 2^k\}$ and all $\varepsilon \geq \varepsilon^*$, where ε^* is a positive quantity that depends on P . Using the fact that

$$\log(a + e^\varepsilon b) = \varepsilon + \log b + O(e^{-\varepsilon}),$$

for large ε , we get that

$$\begin{aligned}
 \mu_j &= P(T_j) e^\varepsilon \log \frac{e^\varepsilon}{P(T_j^c) + e^\varepsilon P(T_j)} + P(T_j^c) \log \frac{1}{P(T_j^c) + e^\varepsilon P(T_j)} \\
 &= P(T_j) e^\varepsilon \varepsilon - (P(T_j) (e^\varepsilon - 1) + 1) \log (P(T_j) (e^\varepsilon - 1) + 1) \\
 &= P(T_j) e^\varepsilon \varepsilon - (P(T_j) (e^\varepsilon - 1) + 1) (\varepsilon + \log P(T_j) + O(e^{-\varepsilon})) \\
 &= -(P(T_j) \log P(T_j)) e^\varepsilon + O(\varepsilon).
 \end{aligned} \tag{76}$$

On the other hand,

$$\begin{aligned}
 \left(S^{(k)T} \alpha^* \right)_j &= S_j^{(k)T} \alpha^* \\
 &= e^\varepsilon \sum_{i \in T_j} \alpha_i^* + \sum_{i \in T_j^c} \alpha_i^* \\
 &= -\frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ \sum_{i \in [k]} S_{ij}^{(k)} (e^\varepsilon + k - 2) (P(x_i) \log P(x_i)) e^\varepsilon \right. \\
 &\quad \left. + O(\varepsilon) \right\} \\
 &\quad + \frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left\{ \sum_{i \in [k]} \sum_{l \in [k], l \neq i} S_{ij}^{(k)} ((P(x_l) \log P(x_l)) e^\varepsilon + O(\varepsilon)) \right\} \\
 &= -\frac{1}{(e^\varepsilon - 1)(e^\varepsilon + k - 1)} \left(\left(\sum_{i \in T_j} P(x_i) \log P(x_i) \right) e^{3\varepsilon} + O(e^{2\varepsilon} \varepsilon) \right) \\
 &= -\left(\sum_{i \in T_j} P(x_i) \log P(x_i) \right) e^\varepsilon + O(\varepsilon).
 \end{aligned} \tag{77}$$

Assume, to begin with, that $j \neq \{j_1, j_2, \dots, j_k\}$. Then

$$\left(S^{(k)T} \alpha^* \right)_j - \mu_j = \left(P(T_j) \log P(T_j) - \sum_{i \in T_j} P(x_i) \log P(x_i) \right) e^\varepsilon + O(\varepsilon).$$

Notice that for $j \neq \{j_1, j_2, \dots, j_k\}$, $P(T_j) \log P(T_j) > \sum_{i \in T_j} P(x_i) \log P(x_i)$. Therefore, there exists an $\varepsilon^* > 0$ such that $\left(S^{(k)T} \alpha^* \right)_j - \mu_j \geq 0$ for all $\varepsilon \geq \varepsilon^*$. If $j \in \{j_1, j_2, \dots, j_k\}$, it is not hard to check that $\left(S^{(k)T} \alpha^* \right)_j - \mu_j = 0$ for all ε . This establishes the satisfiability of α^* for all $\varepsilon \geq \varepsilon^*$. It remains to show that the upper bound can be indeed achieved via the randomized response mechanism. To this extent, recall that the randomized response is given by

$$Q(y|x) = \begin{cases} \frac{e^\varepsilon}{|\mathcal{X}| - 1 + e^\varepsilon} & \text{if } y = x, \\ \frac{1}{|\mathcal{X}| - 1 + e^\varepsilon} & \text{if } y \neq x. \end{cases} \tag{78}$$

Computing the $I(X;Y)$ under (78), we get that

$$\begin{aligned}
 I(X;Y) = \frac{1}{e^\varepsilon + k - 1} \sum_{i \in [k]} \left\{ P(x_i) e^\varepsilon \log \frac{e^\varepsilon}{P(x_i)(e^\varepsilon - 1) + 1} \right. \\
 \left. + (1 - P(x_i)) \log \frac{1}{P(x_i)(e^\varepsilon - 1) + 1} \right\}. \tag{79}
 \end{aligned}$$

Hence, the randomized response mechanism achieves the upper bound (75). This proves the optimality of the randomized response for all $\varepsilon \geq \varepsilon^*$.

10. Proof of Proposition 17

Let $U(Q)$ be a utility mechanism of the form $U(Q) = \sum_{\mathcal{Y}} \mu(Q_y)$, where μ is a sublinear function. Consider a stochastic mapping W of dimensions $\ell \times m$ and let QW be the stochastic mapping obtained by first applying Q to $X \in \mathcal{X}$ to obtain $Y \in \mathcal{Y}$ and then applying W to Y to obtain $Z \in \mathcal{Z}$.

$$\begin{aligned}
 U(QW) &= \sum_{\mathcal{Z}} \mu((QW)_z) \\
 &= \sum_{\mathcal{Z}} \mu \left(\sum_{\mathcal{Y}} Q_y W_{y,z} \right) \\
 &\leq \sum_{\mathcal{Y}, \mathcal{Z}} W_{y,z} \mu(Q_y) \\
 &= \sum_{\mathcal{Y}} \mu(Q_y) \\
 &= U(Q), \tag{80}
 \end{aligned}$$

where the inequality follows from sublinearity and the second to last equality follows from the row stochastic property of W . Therefore, $U(Q)$ obeys the data processing inequality.

Acknowledgments

This research is supported in part by NSF CISE award CCF-1422278, NSF SaTC award CNS-1527754, NSF CMMI award MES-1450848 and NSF ENG award ECCS-1232257.

References

- A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 21–29. ACM, 2004.
- A. Acquisti and J. Grossklags. What can behavioral economics teach us about privacy. *Digital Privacy*, page 329, 2007.

- R. F. Barber and J. C. Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.
- A. Beimel, K. Nissim, and E. Omri. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology-CRYPTO 2008*, pages 451–468. Springer, 2008.
- D. Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.
- J. Blocki, A. Blum, A. Datta, and O. Sheffet. The Johnson-Lindenstrauss transform itself preserves differential privacy. In *Foundations of Computer Science, 2012 IEEE 53rd Annual Symposium on*, pages 410–419. IEEE, 2012.
- K. Chatzikokolakis, T. Chothia, and A. Guha. Statistical measurement of information leakage. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 390–404. Springer, 2010.
- K. Chaudhuri and D. Hsu. Convergence rates for differentially private statistical estimation. *arXiv preprint arXiv:1206.6395*, 2012.
- K. Chaudhuri and C. Monteleoni. Privacy-preserving logistic regression. In *NIPS*, volume 8, pages 289–296, 2008.
- K. Chaudhuri, A. D. Sarwate, and K. Sinha. Near-optimal differentially private principal components. In *NIPS*, pages 998–1006, 2012.
- T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- A. De. Lower bounds in differential privacy. In *Theory of Cryptography*, pages 321–338. Springer, 2012.
- J. C Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. In *Foundations of Computer Science, 2013 IEEE 54th Annual Symposium on*, pages 429–438. IEEE, 2013.
- C. Dwork. Differential privacy. In *Automata, languages and programming*, pages 1–12. Springer, 2006.
- C. Dwork and J. Lei. Differential privacy and robust statistics. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 371–380. ACM, 2009.
- 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*, pages 486–503. Springer, 2006a.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pages 265–284. Springer, 2006b.
- Q. Geng and P. Viswanath. The optimal mechanism in differential privacy. *arXiv preprint arXiv:1212.1186*, 2012.

- Q. Geng and P. Viswanath. The optimal mechanism in (ϵ, δ) -differential privacy. *arXiv preprint arXiv:1305.1330*, 2013a.
- Q. Geng and P. Viswanath. The optimal mechanism in differential privacy: Multidimensional setting. *arXiv preprint arXiv:1312.0655*, 2013b.
- A. Ghosh, T. Roughgarden, and M. Sundararajan. Universally utility-maximizing privacy mechanisms. *SIAM Journal on Computing*, 41(6):1673–1693, 2012.
- M. Hardt and A. Roth. Beating randomized response on incoherent matrices. In *Proceedings of the 44th symposium on Theory of Computing*, pages 1255–1268. ACM, 2012.
- M. Hardt and G. N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *Foundations of Computer Science, 2010 51st Annual IEEE Symposium on*, pages 61–70. IEEE, 2010.
- M. Hardt and K. Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 705–714. ACM, 2010.
- M. Hardt, K. Ligett, and F. McSherry. A simple and practical algorithm for differentially private data release. In *NIPS*, pages 2348–2356, 2012.
- P. Kairouz, S. Oh, and P. Viswanath. The composition theorem for differential privacy. *arXiv preprint arXiv:1311.0776*, 2013.
- P. Kairouz, S. Oh, and P. Viswanath. Extremal mechanisms for local differential privacy. *arXiv preprint arXiv:1407.1338*, 2014a.
- P. Kairouz, S. Oh, and P. Viswanath. Differentially private multi-party computation: Optimality of non-interactive randomized response. *arXiv preprint arXiv:1407.1546*, 2014b.
- M. Kapralov and K. Talwar. On differentially private low rank approximation. In *SODA*, volume 5, page 1. SIAM, 2013.
- J. Lei. Differentially private m-estimators. In *NIPS*, pages 361–369, 2011.
- F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007.
- L. Sankar, S. R. Rajagopalan, and H. V. Poor. Utility-privacy tradeoffs in databases: An information-theoretic approach. *Information Forensics and Security, IEEE Transactions on*, 8(6):838–852, 2013.
- A. D. Sarwate and L. Sankar. A rate-distortion perspective on local differential privacy. In *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*, pages 903–908. IEEE, 2014.
- A. B. Tsybakov and V. Zaiats. *Introduction to nonparametric estimation*, volume 11. Springer, 2009.

- W. Wang, L. Ying, and J. Zhang. On the relation between identifiability, differential privacy and mutual-information privacy. *arXiv preprint arXiv:1402.3757*, 2014a.
- Y. Wang, Z. Huang, S. Mitra, and G.E. Dullerud. Entropy-minimizing mechanism for differential privacy of discrete-time linear feedback systems. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 2130–2135, Dec 2014b. doi: 10.1109/CDC.2014.7039713.
- S. L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.