

Robust Privacy-Utility Tradeoffs Under Differential Privacy and Hamming Distortion

Kousha Kalantari^{1b}, *Student Member, IEEE*, Lalitha Sankar^{1b}, *Senior Member, IEEE*,
and Anand D. Sarwate^{1b}, *Senior Member, IEEE*

Abstract—A privacy-utility tradeoff is developed for an arbitrary set of finite-alphabet source distributions. Privacy is quantified using differential privacy (DP), and utility is quantified using expected Hamming distortion maximized over the set of distributions. The family of source distribution sets (source sets) is categorized into three classes, based on different levels of prior knowledge they capture. For source sets whose convex hull includes the uniform distribution, symmetric DP mechanisms are optimal. For source sets whose probability values have a fixed monotonic ordering, asymmetric DP mechanisms are optimal. For all other source sets, general upper and lower bounds on the optimal privacy leakage are developed and necessary and sufficient conditions for tightness are established. Differentially private leakage is an upper bound on mutual information leakage: the two criteria are compared analytically and numerically to illustrate the effect of adopting a stronger privacy criterion.

Index Terms—Differential privacy, Hamming distortion, information leakage, utility-privacy tradeoff.

I. INTRODUCTION

THE differential privacy (DP) framework offers strong guarantees on the risk of identifying an individual's presence in a database from public disclosures of functions of that database [3]. This metric has been applied to a variety of computational tasks where privacy guarantees are required, especially in theoretical computer science, databases, and machine learning. The monograph of Dwork and Roth [4] gives an in-depth treatment of the fundamentals; a short tutorial for signal processing applications introduces basic concepts

Manuscript received July 21, 2017; revised January 22, 2018 and April 5, 2018; accepted April 6, 2018. Date of publication April 30, 2018; date of current version May 23, 2018. The work of K. Kalantari and L. Sankar was supported by the National Science Foundation under Grant CIF-1422358. The work of A. D. Sarwate was supported in part by the National Science Foundation under Grant CCF-1453432 and Grant SaTC-1617849, in part by DARPA, and in part by SSC Pacific under Grant N66001-15-C-4070. This paper was presented in part at the 52nd Annual Allerton Conference [1] and in part at the 2016 IEEE International Symposium on Information Theory [2]. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Tobias Oechtering. (*Corresponding author: Kousha Kalantari.*)

K. Kalantari was with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287 USA. He is now with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: kousha.kalantari@asu.edu).

L. Sankar is with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: lalithasankar@asu.edu).

A. D. Sarwate is with the Department of Electrical and Computer Engineering, Rutgers, The State University of New Jersey, Piscataway, NJ 08854 USA (e-mail: anand.sarwate@rutgers.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIFS.2018.2831619

. Under differential privacy, randomizing the computation limits the privacy risk, or *leakage*, due to revealing the result of the computation. This randomization often incurs a significant penalty in terms of the usefulness of the published result: this is known as the privacy-utility tradeoff. Differential privacy is a property of the distribution of the computation's output conditioned on its input, which can be modeled information-theoretically as a noisy channel. In this paper we seek to understand the distortion properties of channels that guarantee differential privacy: this is the privacy-utility tradeoff for the task of publishing a differentially private approximation of the full dataset with utility quantified via a distortion measure.

The DP framework makes no modeling assumptions on the data distribution and gives distribution-independent privacy guarantees. However, in many applications a data publisher may know something *a priori* about the data distribution. For example, they may know that some elements of the alphabet have a higher probability, or may know something about the distribution up to the labeling of the alphabet.

There are many instances in which a data holder may be required to publish a version of the underlying data. To capture the data holder's knowledge, we assume they know the true distribution lies in an uncertainty set or *source set* of distributions but do not know the true distribution exactly. This knowledge could come from previously published population statistics, public data or estimation from the source itself, with the uncertainty set represented by confidence intervals. To match the spirit of DP models, we do not assume a Bayesian prior on the set of distributions.

In order to measure the effect of this uncertainty, we model utility as the maximum Hamming distortion over the entire source set. Many datasets contain categorical data for which Hamming distortion is a natural metric and previously studied privacy mechanisms using additive noise make less sense [6]: Hamming distortion captures whether the original data was altered or not. We show that the optimal mechanism guarantees the minimal leakage by effectively censoring low-probability symbols (which are hardest to protect).

A. Our Contributions

This paper extends our previous results on binary sources [2] to general finite alphabets under Hamming distortion. Larger alphabet sizes permit more complex structures for source sets. The following are our main contributions:

- 1) We categorize source distributions as belonging to one of three possible source classes as illustrated in Figure 1 for which different DP mechanisms are optimal.




Fig. 1. All three classes of source sets for $M = 3$.

- **Class I:** Source distribution sets whose convex hull includes the uniform distribution. For example, see the blue source set in Figure 1.

For Class I sources we show that the symmetric mechanism is optimal. Intuitively, this is because knowing that we have a class I source set does not give the data publisher any advantage compared to not knowing anything at all about the source distribution. Therefore, guaranteeing both utility and privacy for all distributions requires a symmetric privacy mapping.

- **Class II:** Source distribution sets that are not Class I, and have ordered probability values. That is, there is a permutation of the alphabet such that all distributions in the class have monotonically decreasing probability mass functions for this permutation. For example, in Figure 1, we have $P_1 \geq P_2 \geq P_3$.

We show how to exploit the ordering to characterize optimal non-symmetric mechanisms for the distribution. As the distortion increases, the optimal mechanism reduces the support size of the output set by mapping low-probability elements to high probability events. We can think of these low-probability events as outliers: since they are most informative (from a privacy perspective), they can be censored in the output to guarantee more privacy.

- **Class III:** All source distribution sets that cannot be classified as Class I or Class II. An example is depicted in Figure 1.

We first show that any arbitrary source set can be written as a disjoint union of Class II subsets, each having a different ordering, and use the characterization of Class II mechanisms to derive upper and lower bounds on the privacy leakage.

- 2) We show that the structure of the conditional probability (channel) matrix of optimal mechanisms depends on the location of what we call *critical pairs*: two elements in the same column with maximum ratio.
- 3) We show how the worst-case guarantee of differentially private leakage compares to average-case guarantee of mutual information (MI) leakage. Under standard DP, the mechanism is context-free in the sense that the guarantees do not rely on source distribution assumptions. This leakage upper bounds the context-aware MI leakage, whose guarantees depend on the source

distribution [7], [8]. This work shows how context-awareness can improve the utility of DP mechanisms and how the gap between the MI leakage and its DP leakage upper bound varies. To do this we study the min-max problems under DP and MI to derive bounds and compute numerical comparisons. To this end, we formulate the same min-max optimization problem proposed to compute DP mechanisms to determine optimal mechanisms in MI privacy guarantees. We formulate the same min-max optimization problem using mutual information as the privacy metric. Then, we show that for certain ranges of distortion we can obtain tight bounds and we present numerical comparison.

B. Related Work

There is a growing body of work on differential privacy (DP) the survey of which is beyond the scope of this paper; we refer the curious reader to the monograph [4]. However, comparing differential privacy to other statistical privacy models is a more recent area of study.

Mutual information has been proposed as a metric for privacy leakage [9]–[11] in a variety of settings including data communications, publishing, and mining. One of the earliest works comparing differential privacy and mutual information privacy is by Alvim *et al.* [12], [13]. Mutual information based privacy metrics have also been considered for data streaming applications [14]–[16]. Wang *et al.* [17] compare mutual information privacy with differential privacy under Hamming distortion. They also introduce a new privacy measure called identifiability and highlight its relationship to MI and DP. Building upon prior work [12], Issa *et al.* [18] have introduced maximal leakage (ML) as an information leakage measure for a guessing adversary; this measure can also be compared to DP. Cuff and Yu [19] present an equivalent definition of differential privacy using mutual information.

In this paper we consider utility metrics based on rate-distortion in Section III-A. A rate-distortion approach to mutual information privacy has been also considered by many researchers [7], [20]–[22]; this paper extends our prior work in this direction [23].

Local differential privacy (L-DP) [24]–[28] studies scenarios in which each data respondent, independently applies the same privacy mechanism. Recently, Kairouz *et al.* [29] determined the optimal local differential privacy mechanism for a class of utility functions that satisfy a sub-linearity property and show that the resulting L-DP mechanism has a *staircase property*, meaning that the ratio of any two conditional probabilities leading to same outputs is in $\{1, c, c^{-1}\}$, where c is some constant. However, while Hamming distortion is a sub-linear utility function, the worst case distortion over a source class is not. We show that the staircase property holds only for Class 1 source sets.

II. PRELIMINARIES AND PROBLEM SETUP

Let \mathcal{X} represent data value alphabet for each individual. For example, \mathcal{X} can represent a single attribute or a Cartesian product of other alphabet sets, i.e. $\mathcal{X} = \prod_{k=1}^K \mathcal{X}_k$, where

\mathcal{X}_k is the set of possible values for k^{th} attribute measured about an individual. Without loss of generality, we assume $\mathcal{X} = \{1, 2, 3, \dots, M\}$. A source distribution set (source set) \mathcal{P} is a subset of probability simplex on M atoms. We model the data of individuals as being drawn from a distribution $P_X \subset \mathcal{P}$ on \mathcal{X} , but our solutions can depend on \mathcal{P} and not the particular P_X is not known *a priori*.

Given an individual's data $X \in \mathcal{X}$, our goal is to find a conditional distribution $Q_{\hat{X}|X}$ that maps the input data X to an output data $\hat{X} \in \hat{\mathcal{X}}$. Our objective is to find a $Q_{\hat{X}|X}$ that is both *privacy preserving* and does not distort the data above some threshold. For any $P_X \in \mathcal{P}$, let $Q_{\hat{X}|X}$ and $P_{X,\hat{X}} = P_X Q_{\hat{X}|X}$ indicate the mechanism and the joint distribution of input and output data, respectively. When it is clear from context, we may drop the subscripts and simply use P and Q instead of P_X and $Q_{\hat{X}|X}$. We write P_i for $P_X(X = i)$ for $i \in \mathcal{X}$.

Let T be a permutation on $\{1, \dots, M\}$, such that $T(i)$ is the i^{th} element in the permuted sequence, and $T^{-1}(j)$ is the index of element j in the permuted sequence. Throughout the sequel, we also refer to the permuted version of a source distribution P as $T(P)$, which is a distribution \hat{P} such that for all i , we have $\hat{P}_{T(i)} = P_i$. Likewise, we write $T^{-1}(P)$ to denote the inverse of T applied to P . The notation $T(P)$ (respectively $T^{-1}(P)$) is the image of \mathcal{P} (respectively $T(P)$) when T (respectively T^{-1}) is applied to every $P \in \mathcal{P}$.

A. Classification of Source Distribution Sets

We divide the set of source distribution sets into three classes.

Definition 1: A source set \mathcal{P} is classified as follows:

- A source set \mathcal{P} is of **Class I** if its convex hull $\text{conv}(\mathcal{P})$ includes the uniform distribution $P_i = \frac{1}{M}$ for all $i \in \{1, 2, \dots, M\}$.
- A source set \mathcal{P} is of **Class II** if it is not Class I and there exists a single permutation $T(\cdot)$ such that for every distribution $P = (P_1, P_2, P_3, \dots, P_M) \in \mathcal{P}$, we have $P_{T(1)} \geq P_{T(2)} \geq P_{T(3)} \geq \dots \geq P_{T(M)}$.
- Any other source set is defined to be of **Class III**.

Remark 1: Without loss of generality, for Class II source sets we assume $P_1 \geq P_2 \geq \dots \geq P_M$.

We now provide some examples for source sets in Classes II and III defined above. These examples are used in Section IV to derive numerical comparisons between different source classes.

Example 1: In a Class II set, all source distributions have entries (as vectors) ordered in the same way. For example, a \mathcal{P} containing a single distribution, such as $\mathcal{P}_{\text{II}}^{(6)}$ shown in Table I for $M = 6$, is a Class II set. A line segment between two distributions with the same order is a Class II set. An example $\mathcal{P}_{\text{II}}^{(10)}$ for $M = 10$ is given in Table II: the line segment between the two rows is a Class II set.

Example 2: One way to generate a Class III source set is by taking the unions of Class II sets and (some of) their permutations. Examples based on $\mathcal{P}_{\text{II}}^{(6)}$ and $\mathcal{P}_{\text{II}}^{(10)}$ are shown in Tables III and IV. Specifically, we created three additional distributions for each Class II distribution considered by permuting the entries in three different ways to obtain additional

TABLE I
 $\mathcal{P}_{\text{II}}^{(6)}$ WITH $M = 6$

P_1	P_2	P_3	P_4	P_5	P_6
0.7	0.15	0.06	0.04	0.03	0.02

TABLE II
 $\mathcal{P}_{\text{II}}^{(10)}$ WITH $M = 10$

P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9	P_{10}
0.3	0.2	0.15	0.08	0.07	0.06	0.05	0.04	0.03	0.02
0.35	0.16	0.12	0.10	0.09	0.09	0.05	0.02	0.01	0.01

TABLE III
 $\mathcal{P}_{\text{III}:a}^{(6)}$, $\mathcal{P}_{\text{III}:b}^{(6)}$, AND $\mathcal{P}_{\text{III}:c}^{(6)}$ WITH $M = 6$

		P_1	P_2	P_3	P_4	P_5	P_6	
c	b	a	0.7	0.15	0.06	0.04	0.03	0.02
			0.15	0.7	0.06	0.04	0.03	0.02
			0.06	0.15	0.7	0.04	0.03	0.02
			0.04	0.15	0.06	0.7	0.03	0.02

TABLE IV
 $\mathcal{P}_{\text{III}:a}^{(10)}$, $\mathcal{P}_{\text{III}:b}^{(10)}$, AND $\mathcal{P}_{\text{III}:c}^{(10)}$ WITH $M = 10$

		P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9	P_{10}	
c	b	a	0.3	0.2	0.15	0.08	0.07	0.06	0.05	0.04	0.03	0.02
			0.35	0.16	0.12	0.10	0.09	0.09	0.05	0.02	0.01	0.01
			0.2	0.3	0.15	0.08	0.07	0.06	0.05	0.04	0.03	0.02
			0.16	0.35	0.12	0.10	0.09	0.09	0.05	0.02	0.01	0.01
			0.15	0.2	0.3	0.08	0.07	0.06	0.05	0.04	0.03	0.02
			0.12	0.16	0.35	0.10	0.09	0.09	0.05	0.02	0.01	0.01
			0.08	0.2	0.15	0.3	0.07	0.06	0.05	0.04	0.03	0.02
			0.10	0.16	0.12	0.35	0.09	0.09	0.05	0.02	0.01	0.01

distributions. We label these resulting sets as $\mathcal{P}_{\text{III}:a}^{(i)}$, $\mathcal{P}_{\text{III}:b}^{(i)}$, and $\mathcal{P}_{\text{III}:c}^{(i)}$, where $i \in \{6, 10\}$.

Thus, for example, to generate $\mathcal{P}_{\text{III}:a}^{(6)}$ we consider both the original $\mathcal{P}_{\text{II}}^{(6)}$ and a permuted version of it obtained by swapping the first and second entries of $\mathcal{P}_{\text{II}}^{(6)}$. To obtain $\mathcal{P}_{\text{III}:b}^{(6)}$, we permute the first and third entries of $\mathcal{P}_{\text{II}}^{(6)}$ and add this to the set $\mathcal{P}_{\text{III}:a}^{(6)}$. Finally, the set $\mathcal{P}_{\text{III}:c}^{(6)}$ is obtained by adding to the set $\mathcal{P}_{\text{III}:b}^{(6)}$ the new distribution obtained by permuting the first and fourth entries of $\mathcal{P}_{\text{II}}^{(6)}$. One can similarly construct the sets $\mathcal{P}_{\text{III}:a}^{(10)}$, $\mathcal{P}_{\text{III}:b}^{(10)}$, and $\mathcal{P}_{\text{III}:c}^{(10)}$ by permuting the first and second, first and third, and first and fourth entries of $\mathcal{P}_{\text{II}}^{(10)}$, respectively. These sets are highlighted in Tables III and IV with their entries denoted by the groupings a, b, and c. Clearly, for both the $M = 6$ and $M = 10$ cases, the number of entries increases from (a) to (c) reflecting less and less structured sets.

B. Distortion Measure

We measure the distortion between X and \hat{X} using Hamming distortion, i.e. $d(x, y) = 1$ if $x \neq y$, and $d(x, y) = 0$ if $x = y$.

Since Hamming distortion imposes a penalty when the published data is different from the original, it suffices to limit our search for optimal mechanisms to those with an output

support set at most equal to M . We formally prove this in Section III. Hamming distortion is particularly meaningful for categorical data in which there may be no natural metric: any difference captures a semantic difference.

We show later that the output alphabet size is at most M as well, depending on the distortion, so $\mathcal{X} = \hat{\mathcal{X}}$. The average distortion is then given as $\mathbb{E}_{X, \hat{X}}[d(X, \hat{X})]$. To indicate the dependence of the average distortion on the source distribution and the mechanism we write $\mathbb{E}_{P_X, Q_{\hat{X}|X}}[d(X, \hat{X})]$. For Hamming distortion, the average distortion is $\sum_{i=1}^M P_i(1 - Q(i|i))$. Thus, we can simplify the Q matrix by defining $D_i = 1 - Q(i|i)$ for all i . Henceforth, it suffices to consider mechanisms $Q(j|i)$ with the following form:

$$[Q_{\hat{X}|X}]_{ij} = \begin{bmatrix} 1 - D_1 & Q(2|1) & \dots & Q(M|1) \\ Q(1|2) & 1 - D_2 & \dots & Q(M|2) \\ \vdots & \vdots & \ddots & \vdots \\ Q(1|M) & Q(2|M) & \dots & 1 - D_M \end{bmatrix}.$$

The sub-matrix of Q induced by rows from i_{\min} to i_{\max} and columns from j_{\min} to j_{\max} is written as $Q(j_{\min} : j_{\max} | i_{\min} : i_{\max})$.

Definition 2: A mechanism $Q_{\hat{X}|X}$, or equivalently its corresponding distortion set $\{D_i\}_{i=1}^M$, is called (\mathcal{P}, D) -valid if it satisfies the average distortion constraint for every $P_X \in \mathcal{P}$. The set of all (\mathcal{P}, D) -valid mechanisms is

$$\mathcal{Q}(\mathcal{P}, D) \triangleq \left\{ Q_{\hat{X}|X} : \mathbb{E}[d(X, \hat{X})] \leq D, \quad \forall P_X \in \mathcal{P} \right\}.$$

C. Local Differential Privacy

We use the same model for local differential privacy as Kairouz *et al.* [29]. We borrow the formalization by Kasiviswanathan *et al.* [26], which was based on the randomized response mechanism of Warner [24] and Evfimievski *et al.* [25]. It is stronger than non-local privacy [30] and implies local ϵ -differential privacy [29], but allows for columns of Q to be all-0:

Definition 3: A mechanism $Q_{\hat{X}|X}$ is ϵ -differentially private (ϵ -DP) if

$$Q(\hat{x}|x_1) \leq e^\epsilon Q(\hat{x}|x_2) \quad \text{for all } x_1, x_2 \in \mathcal{X}, \hat{x} \in \hat{\mathcal{X}}, \quad (1)$$

and

$$\epsilon_{\text{DP}}(Q_{\hat{X}|X}) \triangleq \min \left\{ \epsilon : Q(\hat{x}|x_1) \leq e^\epsilon Q(\hat{x}|x_2) \right. \\ \left. \text{for all } x_1, x_2 \in \mathcal{X}, \hat{x} \in \hat{\mathcal{X}} \right\}. \quad (2)$$

Remark 2: Note that the privacy parameter ϵ does not depend on the source class \mathcal{P} .

Remark 3: For a finite $\epsilon > 0$, an ϵ -differentially private mechanism is such that every column has either all non-zero or all zero entries, i.e. there cannot be a zero and a non-zero entry in the same column. Thus, any mechanism achieving a finite $\epsilon_{\text{DP}}(\cdot)$ can have $M - k$ non-zero columns and k all-zero columns for some integer $0 \leq k \leq M - 1$. Also note that $\epsilon_{\text{DP}}(Q_{\hat{X}|X}) \geq 0$ for any $Q_{\hat{X}|X}$.

Note that $D = 0$ (perfect utility) implies that $\hat{X} = X$, i.e. the optimal mechanism is an identity matrix Q with $\epsilon_{\text{DP}}(Q) = \infty$. Thus, we focus only on $D > 0$ in the sequel.

Lemma 1: $\epsilon_{\text{DP}}(\cdot)$ is quasi-convex in $Q_{\hat{X}|X}$, where quasi-convexity is defined according to [31, Sec. 3.4.1]. Equivalently [31, Sec. 3.4.2], for all Q_1, Q_2 , and $\lambda \in [0, 1]$ we have

$$\epsilon_{\text{DP}}(\lambda Q_1 + (1 - \lambda) Q_2) \leq \max\{\epsilon_{\text{DP}}(Q_1), \epsilon_{\text{DP}}(Q_2)\}. \quad (3)$$

The proof is given in Section V-A.

From Definitions 2 and 3, the minimal achievable ϵ -DP for a given distribution set under Hamming distortion is defined as follows.

Definition 4: For a source distribution set \mathcal{P} , and a distortion D , where $0 < D \leq 1$, let

$$\epsilon_{\text{DP}}^*(\mathcal{P}, D) \triangleq \min_{Q_{\hat{X}|X} \in \mathcal{Q}(\mathcal{P}, D)} \epsilon_{\text{DP}}(Q_{\hat{X}|X}). \quad (4)$$

Also denote the set of all $Q \in \mathcal{Q}(\mathcal{P}, D)$ that achieve (4) by $Q^*(\mathcal{P}, D)$.

D. Worst-Case Distortion Is Not Sub-Linear

The worst-case distortion for a mechanism Q is $\max_{P \in \mathcal{P}} \sum_{i=1}^M P_i D_i$. Since $\epsilon_{\text{DP}}^*(\mathcal{P}, D)$ is decreasing in D , instead of minimizing leakage for a limited worst-case distortion, we can minimize worst-case distortion for limited leakage. Hence, one can formulate the optimization problem in (4) as

$$\min_{Q \in \mathcal{Q}_\epsilon} \max_{P \in \mathcal{P}} \sum_{i=1}^M P_i D_i = \max_{Q \in \mathcal{Q}_\epsilon} U(Q), \quad (5)$$

where utility $U(Q) = -\max_{P \in \mathcal{P}} \sum_{i=1}^M P_i D_i$, and \mathcal{Q}_ϵ is the set of all ϵ -DP mechanisms.

Kairouz *et al.* [29] show the optimality of staircase mechanisms for U satisfying $U(\gamma Q) = \gamma U(Q)$ and $U(Q_1 + Q_2) \leq U(Q_1) + U(Q_2)$; they call such U *sub-linear*. We show by example that our utility is not sub-linear. Let $M = 2$ and consider two mechanisms Q_1 and Q_2 with distortions $D^{(1)} = \{1, 0\}$ and $D^{(2)} = \{0, 1\}$ respectively, as well as a source set $\mathcal{P} = \{P^{(1)} = (1, 0), P^{(2)} = (0, 1)\}$. Then:

$$-\max_{P \in \mathcal{P}} \sum_{i=1}^M P_i (D_i^{(1)} + D_i^{(2)}) \\ > -\max_{P \in \mathcal{P}} \sum_{i=1}^M P_i D_i^{(1)} - \max_{P \in \mathcal{P}} \sum_{i=1}^M P_i D_i^{(2)}. \quad (6)$$

III. MAIN RESULTS

In the prior work Sarwate and Sankar [1] conjecture that the optimal differentially private (DP) mechanism for a discrete source of alphabet size M and distortion level D is

$$Q_D(j|i) \triangleq \begin{cases} 1 - D, & i = j \\ \frac{D}{M-1}, & i \neq j. \end{cases} \quad (7)$$

In the following, we show that the achievable scheme in (7) is tight for Class I source classes; for the Class II source sets,

we exactly characterize $\epsilon_{\text{DP}}^*(\mathcal{P}, D)$ and show that it matches to that of (7) for well-defined subsets of $D \in (0, 1]$, specifically for high and low utility regimes. Finally, we characterize the optimal leakage for Class III source sets of any other form. Note that, $\epsilon_{\text{DP}}^*(\mathcal{P}, D) = 0$ is achievable for $D \geq \frac{M-1}{M}$ by $Q(j|i) = \frac{1}{M}$, $1 \leq i \leq M$, $1 \leq j \leq M$, for any source set of any class.

Lemma 2: Under Hamming distortion, the minimal leakage of \mathcal{P} is the same as the minimal leakage of the convex hull of \mathcal{P} .

The proof is given in Section V-C. Hence, we can always assume \mathcal{P} is convex.

The following lemma shows that it suffices to limit our search for optimal mechanisms to only those with output support set sizes at most equal to M . Therefore, we focus on only such mechanisms throughout the rest of this article

Lemma 3: For a source set \mathcal{P} , there exists an optimal mechanism with an output alphabet $\hat{\mathcal{X}}$ satisfying $|\hat{\mathcal{X}}| \leq M$. The proof is given in Section V-B.

Lemma 4: For a source set \mathcal{P} of Class II (without loss of generality let $P_1 \geq P_2 \geq \dots \geq P_M$ for any $P \in \mathcal{P}$, there exists an optimal mechanism whose corresponding set of $\{D_i\}$ satisfy $D_1 \leq D_2 \leq \dots \leq D_M$.

The proof is given in Section V-D.

Theorem 1: For any source set \mathcal{P} of Class I, we have

$$\epsilon_{\text{DP}}^*(\mathcal{P}, D) = \begin{cases} \log(M-1) \frac{1-D}{D}, & D \in [0, \frac{M-1}{M}), \\ 0, & D \in [\frac{M-1}{M}, 1]. \end{cases} \quad (8)$$

For a full proof see Section V-E. Since the source set includes the uniform point, i.e. the worst distribution, there is no choice other than applying a symmetric mechanism as if there is no knowledge available.

We now proceed to Class II source sets, where there is a known order on the probability of each outcome. For such source sets, we use a coloring argument on the entries of Q to prove specific properties that hold for any optimal mechanism. This, in turn, helps us to reduce the dimension of the feasible space and derive the optimal leakage in terms of a minimization over only M diagonal entries of the mechanism. This is formally stated in the next Theorem.

Since utility is a statistical quantity, the statistical knowledge about the source class can be exploited to obtain a better mechanism than the symmetric one. As the distortion increases, i.e. lower utility is allowed, the size of the output space can decrease. Conversely, for increasing utility, i.e. decreasing distortion, the output space cannot be smaller than a certain size. This leads to a collection of distortion thresholds $D^{(k)}$ at which an additional decrease in output size becomes optimal.

In addition to this observation, we also use the properties of ϵ -DP, and in fact properties of any optimal mechanism, to reduce the dimension of the feasible space from M^2 to just M entries.

Theorem 2: For a Class II source set \mathcal{P} with ordered statistics:

- (a) There is no (\mathcal{P}, D) -valid mechanism with k or more all-zero columns for $D < D^{(k)}$, where $D^0 \triangleq 0$ and

$$D^{(k)} \triangleq \max_{P \in \mathcal{P}} \sum_{i=M-k+1}^M P_i, \quad 1 \leq k \leq M. \quad (9)$$

- (b) The optimal leakage is

$$\epsilon_{\text{DP}}^*(\mathcal{P}, D) = \begin{cases} \log \left((M-1) \frac{1-D}{D} \right), & 0 < D < D^{(1)}, \\ \min_{l \in \{0, 1, \dots, k\}} \epsilon_{\text{DP:II}}^{(l)*}(\mathcal{P}, D), & D^{(k)} \leq D < D^{(k+1)}, \\ 0, & D^{(M-1)} \leq D \leq 1, \end{cases} \quad (10)$$

where $\epsilon_{\text{DP:II}}^{(l)*}(\mathcal{P}, D)$ is the minimum leakage achievable over all (\mathcal{P}, D) -valid mechanisms with exactly l columns with all zero elements and $M-l$ columns with positive elements, formally defined as

$$\epsilon_{\text{DP:II}}^{(l)*}(\mathcal{P}, D) \triangleq \min_{\{D_i\}_{i=1}^{M-l}} \log(M-1-l) \frac{1 - \frac{\sum_{i=2}^{M-l} D_i}{M-1-l}}{D_1} \quad \text{subject to} \begin{cases} \sum_{i=1}^{M-l} P_i D_i \leq D - D^{(l)}, & \forall P \in \mathcal{P}, \\ \sum_{i=1}^{M-l} D_i \leq M-1-l, \\ D_i \in [0, 1], & \forall 1 \leq i \leq M-l, \end{cases} \quad (11)$$

and the subscript DP:II in (10) and (11) denote the Class II source set.


For a detailed proof see Section V-F. Note that each column of a mechanism Q with finite $\epsilon_{\text{DP}}(Q)$ has elements that are either all zero, or all positive. The proof hinges on the fact that for a Class II source set, where we have a complete knowledge on the order of the source distribution probabilities, only mechanisms with specific number of all-zero columns can be feasible for a given distortion D . This limitation, together with some properties of the $\epsilon_{\text{DP}}(\cdot)$ function, result in a specific structure imposed on the optimal mechanism. Therefore, the dimension of the variable space that we need to optimize over reduces to only M instead of M^2 .

We now consider the Class III source sets. We show that the optimal mechanism for Class III source sets can be obtained from results for Class II source sets. As a first step to presenting the main result for Class III source sets, we introduce the following notation and definitions.

Let \mathcal{S}_0 be the set of all Class II distributions with decreasingly ordered probabilities. Formally,

$$\mathcal{S}_0 \triangleq \{P : P_1 \geq P_2 \geq \dots \geq P_M\}. \quad (12)$$

Note that the simplex of distributions can be partitioned into $M!$ such ordered subsets, one for each permutation of $\{1, \dots, M\}$, and thus, there are a total of $M! - 1$ other subregions similar to \mathcal{S}_0 . For example, for $M = 3$, as shown in Figure 2, the simplex is a union of six disjoint ordered sets. More generally, a subset \mathcal{P} of the simplex is a union of


 Fig. 2. A Class III source set \mathcal{P} .

distributions P that lie in one or more ordered partitions. For a source distribution P belonging to any one of these partitions, there exists a corresponding folding permutation T such that $P_{T(1)} \geq P_{T(2)} \geq \dots \geq P_{T(M)}$, or equivalently $T(P) \in \mathcal{S}_0$. Specifically, any Class III source set \mathcal{P} can be written as a disjoint union of Class II source sets using what we call *folding permutations*.

Definition 5: Given a Class III source set \mathcal{P} , its *folding permutation set* $\mathcal{T}_{\mathcal{P}}$ is the set of all permutations T , for which there exists at least one $P \in \mathcal{P}$ with $P_{T(1)} \geq P_{T(2)} \geq \dots \geq P_{T(M)}$. Then, for each $T \in \mathcal{T}_{\mathcal{P}}$ define

$$\mathcal{P}|_T \triangleq \{P \in \mathcal{P} : P_{T(1)} \geq P_{T(2)} \geq \dots \geq P_{T(M)}\}. \quad (13)$$

Thus, a Class III source set \mathcal{P} is a union of Class II source sets, i.e. $\mathcal{P} = \bigcup_{T \in \mathcal{T}_{\mathcal{P}}} \mathcal{P}|_T$. For example, the source set \mathcal{P} in Fig. 3 lies in three partitions, with corresponding folding permutations T_1, T_2 , and T_3 . Thus, $\mathcal{P}|_{T_i}$ is the intersection of \mathcal{P} with the partition whose folding permutation is T_i , $i = 1, 2, 3$, such that $\mathcal{P} = \mathcal{P}|_{T_1} \cup \mathcal{P}|_{T_2} \cup \mathcal{P}|_{T_3}$.

Without loss of generality, we only focus on those Class III source sets \mathcal{P} that have a non-empty intersection with \mathcal{S}_0 . This is due to the fact that for any other Class III source set, the optimal mechanism can be found using a similar analysis with appropriate change of indices. We now show that for any such Class III source set \mathcal{P} , the optimal leakage can be bounded using the result in Theorem 2. We do so by mapping each $\mathcal{P}|_T$ into \mathcal{S}_0 , using its corresponding permutation.

Definition 6: For any permutation function $T \in \mathcal{T}_{\mathcal{P}}$, we define a *folded equivalent* of $\mathcal{P}|_T$ as its mapped image to \mathcal{S}_0 , defined as

$$\overline{\mathcal{P}|_T} \triangleq \{\bar{P} \in \mathcal{S}_0 : \exists P \in \mathcal{P}|_T \text{ s.t. } \bar{P} = T(P)\}. \quad (14)$$


Furthermore, let

$$\mathcal{P}^{\cap} \triangleq \bigcap_{T \in \mathcal{T}_{\mathcal{P}}} \overline{\mathcal{P}|_T}, \quad \mathcal{P}^{\cup} \triangleq \bigcup_{T \in \mathcal{T}_{\mathcal{P}}} \overline{\mathcal{P}|_T}. \quad (15)$$

Clearly, $\mathcal{P}^{\cap} \subseteq \mathcal{P}^{\cup} \subseteq \mathcal{S}_0$, and thus, \mathcal{P}^{\cap} and \mathcal{P}^{\cup} are Class II source sets. This is depicted in Figure 3.

We now proceed to our main result for Class III source sets.

Theorem 3: Let \mathcal{P} be a Class III source set, such that \mathcal{P}^{\cap}


 Fig. 3. a Class III source set \mathcal{P} .

and \mathcal{P}^{\cup} are non-empty. Then

$$\epsilon_{\text{DP:III}}^*(\mathcal{P}^{\cap}, D, \mathcal{T}_{\mathcal{P}}) \leq \epsilon_{\text{DP}}^*(\mathcal{P}, D) \leq \epsilon_{\text{DP:III}}^*(\mathcal{P}^{\cup}, D, \mathcal{T}_{\mathcal{P}}), \quad (16)$$

where for any Class II source set \mathcal{P}_{II} and a folding permutation set \mathcal{T} we have:

$$\epsilon_{\text{DP:III}}^*(\mathcal{P}_{\text{II}}, D, \mathcal{T}) \triangleq \begin{cases} \log \left((M-1) \frac{1-D}{D} \right), & 0 < D < D^{(1)}, \\ \min_{l \in \{0, 1, \dots, M\}} \epsilon_{\text{DP:III}}^{(l)*}(\mathcal{P}_{\text{II}}, D, \mathcal{T}), & D^{(1)} \leq D < \frac{M-1}{M}, \\ 0, & \frac{M-1}{M} \leq D \leq 1, \end{cases} \quad (17)$$

with

$$\epsilon_{\text{DP:III}}^{(k)*}(\mathcal{P}_{\text{II}}, D, \mathcal{T}) \triangleq \min_{\{D_i\}_{i=1}^{M-k}} \log(M-1-k) \frac{1 - \sum_{i=2}^{M-k} D_i}{D_1}$$

subject to

$$\begin{cases} \sum_{i=1}^{M-k} P_i D_i \leq D - D^{(k)}, & \forall P \in \mathcal{P}_{\text{II}}, \\ \sum_{i=1}^{M-k} D_i \leq M-1-k, \\ D_{T(i)} = D_i, & \forall T \in \mathcal{T}, \\ 1 \leq i \leq M, \\ D_i \in [0, 1], & \forall 1 \leq i \leq M. \end{cases} \quad (18)$$

See Section V-G for a detailed proof. For any Class III source set, one can determine \mathcal{P}^{\cap} and \mathcal{P}^{\cup} located inside \mathcal{S}_0 , as shown in Figure 3. The bound results from focusing on \mathcal{P}^{\cap} and \mathcal{P}^{\cup} , and mapping them back using the inverses of all permutations in $\mathcal{T}_{\mathcal{P}}$. The union of all these mapped sets forms \mathcal{P}^{LB} and \mathcal{P}^{UB} , which is contained in and contains \mathcal{P} , respectively. However, \mathcal{P}^{LB} and \mathcal{P}^{UB} have this specific property that their corresponding leakage can be calculated from applying Theorem 2 on \mathcal{P}^{\cap} and \mathcal{P}^{\cup} , with an additional constraint of $D_i = D_{T(i)}$, for all $1 \leq i \leq M$ and $T \in \mathcal{T}_{\mathcal{P}}$.

Remark 4: In the special case where $\mathcal{P}^{\cap} = \emptyset$, $\epsilon_{\text{DP}}^*(\mathcal{P}, D)$ can be simply lower bounded by $\epsilon_{\text{DP}}^*(\mathcal{P}^{\cup} \cap \mathcal{S}_0, D)$.

Remark 5: Observe that $\epsilon_{\text{DP:II}}^{(k)*}$ in (11) and $\epsilon_{\text{DP:III}}^{(k)*}$ in (18) differ in an additional constraint. This comes from the fact that the image of \mathcal{P}^{LB} and \mathcal{P}^{UB} in each Class II partition

is similar, and therefore, forces some distortion values to be equal.

Corollary 1: For $\mathcal{P}^\cap = \mathcal{P}^\cup$, we have

$$\epsilon_{\text{DP}}^*(\mathcal{P}^\cap, D) = \epsilon_{\text{DP}}^*(\mathcal{P}, D) = \epsilon_{\text{DP}}^*(\mathcal{P}^\cup, D). \quad (19)$$

Remark 6: If $\mathcal{P}^\cap = \mathcal{P}^\cup$, then the upper and lower bound match, and the minimal leakage is equal to that of \mathcal{P}^\cap obtained by Theorem 2, with the additional constraint $D_i = D_{T(i)}$, for all $1 \leq i \leq M$ and $T \in \mathcal{T}_P$.

Finally, note that the solutions provided in Theorem 2 and Theorem 3 are found by solving a linear program for fixed D_1 . This simplifies the optimization considerably for large M : a naïve exhaustive search over $\Theta(M^2)$ options is reduced to $\Theta(M)$. These formulations are exploited in Section IV to provide intuition comparing Class I, II, and III source sets.

A. Information Theoretic Leakage

Another metric used for leakage is the mutual information between the original and released data, often referred to as “mutual information (MI) leakage”. Unlike DP leakage that provides worst case guarantees, MI leakage provides average case guarantees for all entries of a dataset by taking the statistics of the data into account. Another difference between the two is the fact that for any given mechanism, the MI leakage is not only a function of the mechanism $P_{\hat{X}|X}$, but also it is dependent on the specific data distribution P_X . For known source distributions, mutual information leakage is studied in [7], [20], and [23], where both asymptotic and non-asymptotic results are derived. However, for the case wherein the source distribution is not known precisely, but some knowledge of source distribution is available, then the worst-case MI leakage of any mechanism Q is defined in [1] as:

$$\epsilon_{\text{IT}}(Q) = \max_{P \in \mathcal{P}} I(P; Q), \quad (20)$$

such that the minimal mutual information leakage is defined as

$$\epsilon_{\text{IT}}^*(\mathcal{P}, D) = \min_{Q \in \mathcal{Q}(\mathcal{P}, D)} \epsilon_{\text{IT}}(Q) = \min_{Q \in \mathcal{Q}(\mathcal{P}, D)} \max_{P \in \mathcal{P}} I(P; Q). \quad (21)$$

Note that it is not in general straightforward to get analytical closed form results for $\epsilon_{\text{IT}}^*(\mathcal{P}, D)$ for any \mathcal{P} and a desired utility function. However, we can characterize its general behavior and use that to make comparisons with DP. Since any mechanism Q that is (\mathcal{P}, D) -valid for two source distributions P_1 and P_2 , is also valid for any convex combination of P_1 and P_2 as well, any source distribution set \mathcal{P} can be replaced with its convex hull without loss of generality. As a result, the set of valid mechanisms $\mathcal{Q}(\mathcal{P}, D)$ is also convex. Also note that both \mathcal{P} and $\mathcal{Q}(\mathcal{P}, D)$ are compact, i.e. closed and bounded, and mutual information is convex in conditional distribution and convex in source distribution. Therefore, according to the minimax theorem [32] we can conclude that the minimax inequality holds as equality and we have:

$$\begin{aligned} \epsilon_{\text{IT}}^*(\mathcal{P}, D) &= \min_{Q \in \mathcal{Q}(\mathcal{P}, D)} \max_{P \in \mathcal{P}} I(P; Q) \\ &= \max_{P \in \mathcal{P}} \min_{Q \in \mathcal{Q}(\mathcal{P}, D)} I(P; Q). \end{aligned} \quad (22)$$

We stress that MI leakage and DP leakage reflect two very different privacy sensitivity models; in particular, DP leakage is always an upper bound on MI leakage. Therefore, for a common utility function, and a given source set, it is worthwhile to compare their performance. To this end, we present some analytical results under MI leakage for source classes I and II.

Lemma 5: For any Class I source set \mathcal{P} , we have

$$\begin{aligned} \epsilon_{\text{IT}}^*(\mathcal{P}, D) &= \begin{cases} \log M - H(D) - D \log(M-1), & D < \frac{M-1}{M}, \\ 0, & D \geq \frac{M-1}{M}. \end{cases} \end{aligned} \quad (23)$$

Proof: We first show that $\epsilon_{\text{IT}}^*(\mathcal{P}, D) = 0$, if $D \geq \frac{M-1}{M}$. Consider the mechanism which maps every letter of the input alphabet to the first letter. This mechanism results in a distortion of $\frac{M-1}{M}$ and achieves $\epsilon_{\text{IT}}^*(\mathcal{P}, D) = 0$, because the resulting output distribution is totally independent of the input.

We now proceed to the case where $D < \frac{M-1}{M}$. Recall that \mathcal{P} is of Class I and includes the uniform point. Since MI leakage is a concave function of $P \in \mathcal{P}$, for any given mechanism Q the resulting worst case MI leakage is the one corresponding to the uniform source distribution.

The resulting leakage can be lower bounded as

$$\begin{aligned} I(X; \hat{X}) &= H(X) - H(X|\hat{X}) \\ &\geq \log M - H(D) - D \log(M-1), \end{aligned} \quad (24)$$

where (24) follows from Fano’s inequality. The lower bound in (24) can be achieved by the following mechanism:

$$Q(j|i) = \begin{cases} 1 - D, & i = j, \\ \frac{D}{M-1}, & i \neq j. \end{cases} \quad (25)$$

Lemma 6: For any Class II source set \mathcal{P} , $\epsilon_{\text{IT}}^*(\mathcal{P}, D) = 0$ iff $D \geq D^{(M-1)}$.

Proof: Let $D \geq D^{(M-1)}$ and $P^* \in \mathcal{P}$ be the distribution achieving the maximum in definition of $D^{(M-1)}$ in (9). Consider a mechanism that maps every input independently to the output element of P^* with the highest probability. One can verify that the resulting distortion is $D^{(M-1)} \leq D$. Furthermore, one can also verify that for this mechanism $I(X; \hat{X}) = 0$.

We now show that no mechanism can achieve $I(X; \hat{X}) = 0$, if $D < D^{(M-1)}$. Without loss of generality, let $P_1 \geq P_2 \geq \dots \geq P_M$. Assume to the contrary that there exists a \hat{Q} achieving $\epsilon_{\text{IT}}^*(\mathcal{P}, D) = 0$ for some $D < D^{(M-1)}$. Since $I(X; \hat{X}) = 0$, $P(\hat{x}|x) = p(\hat{x})$ for all x . This in turn result in a distortion at least equal to $\sum_{i=2}^M P_i^*$, which is equal to $D^{(M-1)}$, and thus, \hat{Q} cannot be (\mathcal{P}, D) -valid. ■

IV. ILLUSTRATION OF RESULTS

We now illustrate our result by first giving examples of DP leakage for class I,II, and III sources as well as comparisons




Fig. 4. DP leakage-distortion tradeoff for Class I, Class II, and Class III source sets with $M = 6$.




Fig. 5. DP leakage-distortion tradeoff for Class I, Class II, and Class III source sets with $M = 10$.

between DP and MI leakage. The central question motivating this work is how partial knowledge of the source distribution can be exploited to improve privacy-utility tradeoffs. We revisit our examples from Section II-A to illustrate our theoretical results.

We first illustrate the reduction in leakage obtained when one goes beyond Class I source knowledge to Class II. Figures 4 and 5 show the minimal DP leakage for $\mathcal{P}_{II}^{(6)}$ with $M = 6$ and $\mathcal{P}_{II}^{(10)}$ with $M = 10$, respectively. When compared against Class I, for low distortion requirements there is no benefit to source knowledge, but in regimes where a moderate level of distortion is tolerable, the data publisher can significantly decrease the privacy leakage by taking advantage of the source set structure.

Yet another comparison made here is between Class II and Class III source sets. Specifically, we expect the leakage guarantees to diminish for Class III which is less structured than Class II and indeed we observe this behavior. In fact, since the distortion guarantee for a Class III set also holds for its convex hull, eventually this hull will contain the uniform distribution and the tradeoff will correspond to the Class I leakage.

Finally, since DP is distribution-agnostic, the MI leakage is always upper bounded by the DP leakage. In Figures 6 and 7




Fig. 6. Differential Privacy vs Information Theoretic Leakage for Class I source sets and $M = 10$.




Fig. 7. Differential Privacy vs Information Theoretic Leakage for $\mathcal{P}_{II}^{(10)}$ in Table II.

we compare the MI and DP leakages for Class I and Class II sets. The MI leakage we use is the source-aware worst-case mutual information.

These plots clearly show the convexity of the MI leakage and nonconvexity of the DP leakage as a function of the distortion constraint D . The bounds only coincide at perfect privacy, where the output is independent of the input, as indicated by Lemmas 5 and 6.

V. PROOFS

A. Proof of Lemma 1: Quasi-Convexity of $\epsilon_{DP}(\cdot)$ in Q

Proof: Based on the definition of quasi-convexity in [31, Sec. 3.4], it suffices to show that all the sub-level sets of the function $\epsilon_{DP}(\cdot)$ are convex, i.e. if two different mechanisms Q_1 and Q_2 are ϵ -differentially private mechanisms for some finite ϵ , then their convex combination $Q_\theta = \theta Q_1 + (1 - \theta)Q_2$, $0 < \theta < 1$, is also ϵ -differentially private. Let x_1 and x_2 be two arbitrary input elements, and let \hat{x} be an arbitrary output element. We have

$$Q_\theta(\hat{x}|x_1) = \theta Q_1(\hat{x}|x_1) + (1 - \theta)Q_2(\hat{x}|x_1) \quad (26a)$$

$$\leq \theta Q_1(\hat{x}|x_2)e^\epsilon + (1 - \theta)Q_2(\hat{x}|x_2)e^\epsilon \quad (26b)$$

$$= e^\epsilon Q_\theta(\hat{x}|x_2). \quad (26c)$$

Therefore, Q_θ is also ϵ -differentially private, and thus $\epsilon_{\text{DP}}(\cdot)$ is a quasi-convex function. ■

B. Proof of Lemma 3

Proof: We first show that for any optimal mechanism P with output support set of size N , where $N > M + 1$, there exists an optimal mechanism Q with output support set of size $N - 1$. It suffices to build Q from P by merging the last two columns of P , i.e. adding them element-wise to make one single column. One can verify that $\epsilon_{\text{DP}}(Q) \leq \epsilon_{\text{DP}}(P)$ due to quasi-convexity of $\epsilon_{\text{DP}}(\cdot)$ shown in Lemma 1. Note that the resulting distortion is exactly identical in both Q and P since their diagonal elements are equal. ■

We now show that for an optimal mechanism P with output support set of size $M + 1$, we can construct an optimal mechanism Q with output support set of size M . Take columns M and $M + 1$ of P , and merge them similar to the previous part. One can similarly verify $\epsilon_{\text{DP}}(Q) \leq \epsilon_{\text{DP}}(P)$. We now check distortion feasibility. Note that $P_{M,M} = 1 - D_M$, and therefore once $P_{M,M+1}$ is added to it to obtain $Q_{M,M}$, the updated D_M does not increase, and therefore the total distortion under Q is at most equal to that of P . This holds for any distribution point in \mathcal{P} .

C. Proof of Lemma 2: Convexity of \mathcal{P}

Proof: This is due to the fact that any (\mathcal{P}, D) -valid mechanism should be also valid for any \bar{P} that is a convex combination of distributions in \mathcal{P} . More formally, suppose that a \bar{P} is in the form of $\sum_{i=1}^r \theta_i P^{(i)}$, where $P^{(i)} \in \mathcal{P}$, for all $i \in \{1, 2, \dots, r\}$. Then, for a (\mathcal{P}, D) -valid mechanism Q with distortion set $\{D_i\}$ we have

$$\sum_{i=1}^M \bar{P}_i D_i = \sum_{i=1}^M \sum_{j=1}^r \theta_j P_i^{(j)} D_i \quad (27a)$$

$$= \sum_{j=1}^r \theta_j \sum_{i=1}^M P_i^{(j)} D_i \quad (27b)$$

$$\leq \sum_{j=1}^r \theta_j D = D. \quad (27c)$$

Thus, Q is also a valid mechanism for \bar{P} and any \mathcal{P} can be extended to its convex hull without loss of generality. ■

D. Proof of Lemma 4

Without loss of generality, let $P_1 \geq P_2 \geq \dots \geq P_M$ for any $P \in \mathcal{P}$. We now show that there exists an optimal mechanism with Q with $D_1 \leq D_2 \leq \dots \leq D_M$. Let $\hat{Q}_{\hat{X}|X}$ be some optimal mechanism such that $\hat{D}_{T(1)} \leq \hat{D}_{T(2)} \leq \dots \leq \hat{D}_{T(M)}$, for some permutation T . Then, let $Q_{\hat{X}|X}^* = \hat{Q}_{T(\hat{X})|T(X)}$. Clearly, we have $D_1^* \leq D_2^* \leq \dots \leq D_M^*$ and $\epsilon_{\text{DP}}(\hat{Q}) = \epsilon_{\text{DP}}(Q^*)$. Finally, $Q_{\hat{X}|X}^*$ is a (\mathcal{P}, D) -valid mechanism, because:

$$\sum_{i=1}^M P_i D_i^* \leq \sum_{i=1}^M P_i D_{T^{-1}(i)}^* = \sum_{i=1}^M P_i \hat{D}_i \leq D. \quad (28)$$

E. Proof of Theorem 1: Class I Source Sets

We now determine the optimal mechanism for Class I and show that it is indeed the conjectured mechanism in [1]. From Lemma 2, we know that we can replace \mathcal{P} with $\text{conv}(\mathcal{P})$ without loss of generality, and henceforth our results hold for $\text{conv}(\mathcal{P})$. We begin by assuming to the contrary that there exists a (\mathcal{P}, D) -valid mechanism $Q_{\hat{X}|X}$ with lower risk guarantees than conjectured in [1], i.e. $\epsilon_{\text{DP}}(Q) < \log(M-1) \frac{1-D}{D}$ for $0 < D < \frac{M-1}{M}$. For any Q with $\epsilon_{\text{DP}}(Q) < \epsilon_{\text{DP}}(Q_D)$, we require that $Q_{\hat{X}|X}(j|i) > e^{\epsilon_{\text{DP}}(Q_D)} Q_{\hat{X}|X}(j|j)$, for at least one pair (i, j) , $i \neq j$. Thus, by summing over all columns in $Q_{\hat{X}|X}$ and recalling that $e^{\epsilon_{\text{DP}}(Q_D)} = \frac{1-D}{D}(M-1)$, we have

$$M = \sum_{i=1}^M \sum_{j=1}^M Q(j|i) \quad (29a)$$

$$> \sum_{j=1}^M \left[Q(j|j) + \sum_{i \neq j} \frac{Q(i|i)}{e^{\epsilon_{\text{DP}}(Q_D)}} \right] \quad (29b)$$

$$= \sum_{j=1}^M \left[(1 - D_j) + \frac{(M-1) - \sum_{i \neq j} D_i}{e^{\epsilon_{\text{DP}}(Q_D)}} \right] \quad (29c)$$

$$= M - \sum_{j=1}^M D_j + \frac{M(M-1)}{e^{\epsilon_{\text{DP}}(Q_D)}} - \frac{M-1}{e^{\epsilon_{\text{DP}}(Q_D)}} \sum_{j=1}^M D_j \quad (29d)$$

$$= \left(\frac{M-1}{e^{\epsilon_{\text{DP}}(Q_D)}} + 1 \right) \left(M - \sum_{j=1}^M D_j \right) \quad (29e)$$

$$= \left(\frac{1}{1-D} \right) \left(M - \sum_{j=1}^M D_j \right). \quad (29f)$$

Therefore

$$\sum_{j=1}^M (1 - Q(j|j)) = \sum_{j=1}^M D_j > MD. \quad (30)$$

This, however, contradicts satisfying the distortion constraint for the uniform distribution.

F. Proof of Theorem 2: Class II Source Sets

We now prove Theorem 2, which exactly characterizes $\epsilon_{\text{DP}}^*(\mathcal{P}, D)$ for the Class II source sets as introduced in Definition 1. Recall that we defined distortion levels $D^{(k)}$ in (9) such that for any k , $D^{(k)}$ corresponds to the case wherein at most $k - 1$ letters of the input are suppressed and the output alphabet size is at least $M - k$, if $D < D^{(k)}$. On the other hand, for $D \geq D^{(k)}$, the output alphabet size may be suppressed by k or more elements. Through the following lemma, we first prove that perfect privacy, i.e. zero leakage, can be achieved if and only if $D \geq D^{(M-1)}$.

Lemma 7: $\epsilon_{\text{DP}}^*(\mathcal{P}, D) = 0$ if and only if $D \geq D^{(M-1)}$.

Proof: We first prove the converse and show that $\epsilon_{\text{DP}}^*(\mathcal{P}, D) = 0$ only if $D \geq D^{(M-1)}$. Let Q be a (\mathcal{P}, D) -valid mechanism with $\epsilon_{\text{DP}}(Q) = 0$. This implies all elements of the i^{th} column have the same value, namely a_i , where $0 \leq a_i \leq 1$ and $\sum_{i=1}^M a_i = 1$. Hence, the corresponding

distortion values for Q are $D_i = 1 - a_i, \forall i$, where $0 \leq D_i \leq 1$ and $\sum_{i=1}^M D_i = M - 1$. Also, recall that for any distribution P , we have $P_1 \geq P_2 \geq \dots \geq P_M$. Therefore, by replacing D_1 with zero and $D_i, i > 1$ with one, we can further lower bound the distortion as $\sum_{i=1}^M P_i D_i \geq \sum_{i=2}^M P_i$. Note that Q is a (\mathcal{P}, D) -valid mechanism, and therefore $\sum_{i=2}^M P_i \leq D$. Taking the maximum over all $P \in \mathcal{P}$ gives $D^{(M-1)} = \max_{P \in \mathcal{P}} \sum_{i=2}^M P_i \leq D$.

For proving the achievability, i.e. $\epsilon_{\text{DP}}^*(\mathcal{P}, D) = 0$ for $D \geq D^{(M-1)}$, consider the mechanism with zero elements everywhere except the first column where all entries are 1, i.e. $Q(i|j) = 0$, if $i > 1$, and $Q(i|j) = 1$, if $i = 1$. This mechanism achieves $\epsilon_{\text{DP}}^*(\mathcal{P}, D) = 0$ and the distortion is bounded by

$$\max_{P \in \mathcal{P}} \sum_{i=1}^M P_i D_i = \max_{P \in \mathcal{P}} \sum_{i=2}^M P_i = D^{(M-1)} \leq D. \quad (31)$$

We now restrict ourselves to $0 \leq D \leq D^{(M-1)}$, and in the following collection of lemmas we prove structural conditions on the optimal mechanisms for Class II sources. We first describe the need for different distortion levels, and then provide achievability and converse proofs. In particular, as the distortion increases there are specific distortion values at which the support of output is allowed to shrink more. The following lemma captures this observation precisely.

Lemma 8: For a $k \in \{1, 2, \dots, M\}$ and $D < D^{(k)}$, no (\mathcal{P}, D) -valid mechanism can have an output support size of less than or equal to $(M - k)$.

Proof: For any $P \in \mathcal{P}$, any mechanism with k or more all-zero columns results in an average distortion $\sum_{i=1}^M P_i D_i$, which is strictly greater than $\sum_{i=M-k+1}^M P_i$ because at least k elements in the set $\{D_i\}_{i=1}^M$ are equal to one. Hence, for $D < D^{(k)}$, no mechanism with k or more all-zero columns can be (\mathcal{P}, D) -valid. ■

Recall that without loss of generality, we can assume a given Class II source set has the ordering $P_1 \geq P_2 \geq \dots \geq P_M$, for any $P \in \mathcal{P}$. Then, based on Lemma 4, there exists an optimal mechanism with $D_1 \geq D_2 \geq \dots \geq D_M$.

Using these lemmas, we now present a converse proof by exploiting the definition of differential privacy. We provide a sequence of properties that any optimal mechanism must satisfy. We can therefore obtain a lower bound on the leakage by minimizing parameters of those properties. Then, we present an achievable scheme by providing a mechanism that achieves the minimum value given by the converse.

1) *Converse for Theorem 2:* We now prove a lower bound on $\epsilon_{\text{DP}}^*(\mathcal{P}, D)$ for $0 < D < D^{(M-1)}$. We first define *critical pairs* in a matrix and then introduce a matrix coloring scheme to prove specific properties of the optimal mechanism. We illustrate this definition and the properties using Figure 8.

Definition 7: For a mechanism $Q_{\hat{X}|X}$ with $\epsilon_{\text{DP}}(Q_{\hat{X}|X}) > 0$, a *critical pair* in $Q_{\hat{X}|X}$ is a pair of elements $\{Q(k|i), Q(k|j)\}$ in a non-zero column, such that $Q_{\hat{X}|X}(k|i) = \exp(\epsilon_{\text{DP}}(Q_{\hat{X}|X})) Q_{\hat{X}|X}(k|j)$. Note that there exists at least one critical pair, but in general if there are multiple critical pairs in different columns of a

Fig. 8. For Class II source sets, optimal mechanisms with k all-zero columns have at most one non-black diagonal element, while all other non-zero off diagonal elements are either red or white.

matrix Q , they may have different values. However, their ratio needs to be equal to $\exp(\epsilon_{\text{DP}}(Q))$. Furthermore, note that not all columns may have a critical pair. However, the maximal ratio of two elements in any column is at most $\exp(\epsilon_{\text{DP}}(Q))$.

We color the entries of non-zero columns of any given matrix Q black, white or red as follows:

- An element is colored black if it is the larger element in a critical pair.
- An element is colored red if it is the smaller element in a critical pair.
- All other elements are colored white.

Remark 7: Note that if a black element is decreased (or a red is increased), either the $\epsilon_{\text{DP}}(\cdot)$ of the matrix has to decrease, or that element can no longer be black (red).

Our proof involves manipulating the elements of Q while maintaining it as a valid mechanism: any change in Q that neither increases a black element nor decreases a red element keeps $\epsilon_{\text{DP}}(Q_{\hat{X}|X})$ at most equal to its previous value.

Definition 8: For any $0 \leq k \leq M - 1$, let $Q^{(k)}(\mathcal{P}, D)$ be the set of all (\mathcal{P}, D) -valid mechanisms with k all-zero columns. Also let $Q^{*(k)}(\mathcal{P}, D) \subset Q^{(k)}(\mathcal{P}, D)$ be the set of mechanisms with the smallest $\epsilon_{\text{DP}}(\cdot)$ among all mechanisms in $Q^{(k)}(\mathcal{P}, D)$.

Recall from Section II that $Q(j_{\min} : j_{\max} | i_{\min} : i_{\max})$ is the sub-matrix of Q induced by rows from i_{\min} to i_{\max} and columns from j_{\min} to j_{\max} . This notation is used extensively throughout the following lemmas and propositions.

Lemma 9: For a given $0 < D < D^{(M-1)}$ and $0 \leq k \leq M - 1$, there exists a $Q_{\hat{X}|X}^k \in Q^{*(k)}(\mathcal{P}, D)$ with the distortions $D_1 \leq D_2 \leq \dots \leq D_M$ such that

$$\epsilon_{\text{DP}}(Q_{\hat{X}|X}^k) = \log(M - 1 - k) \frac{1 - \frac{\sum_{i=2}^{M-k} D_i}{M-1-k}}{D_1}. \quad (32)$$

Proof: Denote those mechanisms $Q^{*(k)}(\mathcal{P}, D)$ that have the smallest $\sum_{i=1}^M D_i$ by $Q_{\text{MS}}^{*(k)}(\mathcal{P}, D)$, where MS in the subscript stands for “Minimum Sum”. Through five sequential claims, we now show that there exists a $Q_{\hat{X}|X} \in Q_{\text{MS}}^{*(k)}(\mathcal{P}, D)$ with a specific color structure as shown in Figure 8 and $\epsilon_{\text{DP}}(Q_{\hat{X}|X})$ is given by (32).

For any mechanism $Q_{\hat{X}|X}$ with k all-zero columns (or in other words, with k of its D_i s being equal to 1), without loss of generality we can assume that columns $M-k+1$ to M are the all-zero ones due to Lemma 4. Let $Q(1 : (M-k)|1 : (M-k))$ and $Q(1 : (M-k)|(M-k+1) : M)$ be sub-matrices of any given Q . Then there exists a $Q_{\hat{X}|X} \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$ that satisfies the following claims sequentially. In other words, each claim i states that there exists at least one $Q_{\hat{X}|X} \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$ among all the ones satisfying the previous claims $\{1, \dots, i-1\}$, such that the statement of claim i is true.

Claim 1: No row is all black (or all red) in $Q(1 : (M-k)|1 : M)$.

Proof: Assume the contrary that the i^{th} row in $Q(1 : (M-k)|1 : M)$ is all black for any $Q_{\hat{X}|X} \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$. Consider the ratio $\epsilon_{\text{DP}}(Q) = \log \frac{Q(1|i)}{Q(1|j)}$, where $Q(1|j)$ is the red element associated with $Q(1|j)$ which is black, i.e. $Q(1|j)$ and $Q(1|i)$ form critical pairs. Also note that for any other $2 < l \leq M-k$, $Q(l|i) \geq Q(l|j)$. Now, if $\epsilon_{\text{DP}}(Q) > 0$, we have

$$1 = Q(1|i) + \sum_{l=2}^{M-k} Q(l|i) > Q(1|j) + \sum_{l=2}^{M-k} Q(l|j) = 1, \quad (33)$$

which is a contradiction. Otherwise, if $\epsilon_{\text{DP}}(Q) = 0$, by Lemma 7 $D \geq D^{(M-1)}$, which is also a contradiction to our assumption of $0 < D < D^{(M-1)}$ in the beginning of converse proof. The proof for no row with all red elements is similar. ■

Claim 2: All the off-diagonal elements of a row in $Q(1 : (M-k)|1 : M)$ have the same color.

Proof: Take an arbitrary mechanism $Q_{\hat{X}|X} \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$ that $Q(1 : (M-k)|1 : M)$ satisfies Claim 1. Fix a row i , and let the number of off-diagonal elements in row i of $Q(1 : (M-k)|1 : M)$ with each of the colors black, white and red be n_B , n_W and n_R respectively, where $n_B + n_W + n_R = M-1-k$. If only one of n_R , n_W or n_B is non-zero, then the claim is satisfied. We split the remaining scenarios into two cases and show that in each it is sufficient to make all off-diagonal elements of the i^{th} row white.

- If $n_R, n_B > 0$, then there exists an arbitrarily small $\delta > 0$ such that each of the off-diagonal black elements can be decreased by $\frac{\delta}{n_B}$, and each of the off-diagonal red elements can be increased by $\frac{\delta}{n_R}$. Consider a Q' which is identical to Q everywhere except for a segment of the i -th row $Q'(1 : (M-k)|i)$. For $1 \leq j \leq M-k$, $j \neq i$, let $Q'(j|i) = Q(j|i) - \frac{\delta}{n_B}$ if $Q(j|i)$ is black, $Q'(j|i) = Q(j|i) + \frac{\delta}{n_R}$ if $Q(j|i)$ is red, and $Q'(j|i) = Q(j|i)$ if $Q(j|i)$ is white. Note that it does not matter if $n_W = 0$ or not, because the white off-diagonal elements $Q'(1 : (M-k)|i)$ are not changed.
- If $n_R = 0, n_W, n_B > 0$ (or $n_B = 0, n_W, n_R > 0$), then there exists an arbitrarily small $\delta > 0$ such that each of off-diagonal white elements can be increased (or decreased) by $\frac{\delta}{n_W}$, and each of off-diagonal black elements (or red elements) can be decreased (or increased) by $\frac{\delta}{n_B}$ (or $\frac{\delta}{n_R}$). Consider a Q' which is identical to Q everywhere except for $Q'(1 : (M-k)|i)$. For $1 \leq$

$j \leq M-k$, $j \neq i$, let $Q'(j|i) = Q(j|i) - \frac{\delta}{n_B}$ if $Q(j|i)$ is black and $Q'(j|i) = Q(j|i) + \frac{\delta}{n_W}$ if $Q(j|i)$ is white (or $Q'(j|i) = Q(j|i) + \frac{\delta}{n_R}$ if $Q(j|i)$ is red and $Q'(j|i) = Q(j|i) - \frac{\delta}{n_W}$ if $Q(j|i)$ is white), and $Q'(j|i) = Q(j|i)$ elsewhere.

In both cases Q' has the same off-diagonal row sum and the same set of $\{D_i\}_{i=1}^M$ as Q , but for sufficiently small δ , Q' is still a valid row stochastic matrix and none of the elements in $Q'(1 : (M-k)|1 : M)$ become 1 or 0. Thus, Q' would still be a (\mathcal{P}, D) -valid mechanism. Besides, all off-diagonal elements in row i of Q' , $i \notin \{a_1, a_2, \dots, a_k\}$, are white: If not we would have a smaller $\epsilon_{\text{DP}}(Q'_{\hat{X}|X})$ due to Remark 7, which contradicts our first assumption that $Q_{\hat{X}|X} \in \mathcal{Q}^{k*}(\mathcal{P}, D)$. In this construction, all off-diagonal elements of $Q'(1 : (M-k)|1 : M)$ in rows other than i are colored the same as Q without affecting the average distortion, while keeping $\epsilon_{\text{DP}}(Q') \leq \epsilon_{\text{DP}}(Q)$ and thus $Q' \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$. This operation can be done for each row i repeatedly, to get the final Q' to satisfy the claim. ■

Remark 8: As a result of Claims 1 and 2, all the off-diagonal elements in $Q(1 : (M-k)|(M-k+1) : M)$ are white.

Claim 3: If a diagonal element in $Q(1 : (M-k)|1 : (M-k))$ is not black then all off-diagonal elements of $Q(1 : (M-k)|1 : (M-k))$ in the same row are red.

Proof: Take an arbitrary $Q_{\hat{X}|X} \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$ satisfying Claims 1 and 2, and for some $1 \leq i \leq M-k$ suppose $Q(i|i)$ is red or white. By Claim 2, all elements in the set $\{Q(j|i) : j \neq i, 1 \leq j \leq M-k\}$ have the same color. Assume to the contrary that they are not all red, so they are all black or white. Consider a Q' which is equal to Q , except in $Q'(1 : (M-k)|i)$ where $Q'(i|i) = Q(i|i) + \delta$ and $Q'(j|i) = Q(j|i) - \frac{\delta}{M-k-1}$ for $j \neq i, 1 \leq j \leq M-k$. For sufficiently small $\delta > 0$ this is also a (\mathcal{P}, D) -valid mechanism. Although $\epsilon_{\text{DP}}(Q'_{\hat{X}|X})$ remains unchanged due to Remark 7, we have

$$\sum_{i=1}^M D'_i = \sum_{i=1}^M (1 - Q'(i|i)) < \sum_{i=1}^M (1 - Q(i|i)) = \sum_{i=1}^M D_i, \quad (34)$$

where D'_i is the distortion of i^{th} element under Q' . This clearly contradicts the assumption that $Q \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$. ■

Claim 4: There is at most one non-black element on the diagonal of $Q(1 : (M-k)|1 : (M-k))$.

Proof: Take an arbitrary $Q_{\hat{X}|X} \in \mathcal{Q}_{\text{MS}}^{*(k)}(\mathcal{P}, D)$ satisfying all previous claims. Assume the contrary that there are at least two non-black diagonal elements $Q(i|i)$ and $Q(j|j)$, ($i \neq j, i, k \leq M-k$). Thus, $Q(i|j)$ is red by Claim 3, which implies that there exists a $k \neq j$ where $Q(i|k)$ is black, because there should be a black element for each red element in a column. However, $k \neq i$ because we already know that $Q(i|i)$ is non-black. Thus, $Q(i|k)$ has to be an off-diagonal black element in $Q(1 : (M-k)|1 : M)$, which is contradictory to our first assumption of Q satisfying all previous claims, including Claim 1 and 2. Therefore, at most one diagonal element is non-black. ■

Claim 5: The only possible non-black element along the diagonal of $Q(1 : (M - k)|1 : (M - k))$ is the one corresponding to the smallest, or one of the smallest D_i s.

Proof: Take an arbitrary $Q_{\hat{X}|X} \in Q_{MS}^{*(k)}(\mathcal{P}, D)$ satisfying all previous claims. Let the diagonal element in row i , $1 \leq i \leq M - k$, be non-black. We show that for any $j \neq i$, $1 \leq j \leq M - k$, $\hat{D}_j \geq \hat{D}_i$. By Claim 3, we know that any other off-diagonal entry in row i of $Q(1 : (M - k)|1 : (M - k))$, including $Q(j|i)$ is red. We also know that $Q(j|j)$ is black and other entries in row j of $Q(1 : (M - k)|1 : (M - k))$, including $Q(i|j)$, are either all red or all white due to Claims 1 and 2. Thus, for all $1 \leq k \leq M - k$ other than i or j we have $Q(k|j) \geq Q(k|i)$ because $Q(k|j)$ is either red or white, and $Q(k|i)$ is red, where both of them are in the same column. Since any row has to sum up to one, summing over rows i and j results in

$$Q(j|j) + Q(i|j) \leq Q(i|i) + Q(j|i). \quad (35)$$

Besides, since $Q(j|j)$ is black, $Q(j|i)$ is either red or white, $Q(i|j)$ is red, and $Q(i|i)$ is either red or white, we have

$$\frac{Q(j|j)}{Q(j|i)} > \frac{Q(i|i)}{Q(i|j)}. \quad (36)$$

We now show that

$$1 - \hat{D}_j = Q(j|j) \leq Q(i|i) = 1 - \hat{D}_i. \quad (37)$$

Assume the contrary that $Q(j|j) > Q(i|i)$. Thus, by (35) we have

$$0 < Q(j|j) - Q(i|i) \leq Q(j|i) - Q(i|j), \quad (38)$$

which means $\frac{Q(j|j) - Q(i|i)}{Q(j|i) - Q(i|j)} \leq 1$. However, (36) shows that

$$\frac{Q(j|j) - Q(i|i)}{Q(j|i) - Q(i|j)} \geq \frac{Q(j|j)}{Q(j|i)} > \frac{Q(i|i)}{Q(i|j)} \geq 0, \quad (39)$$

which means

$$Q(j|j) - Q(i|i) > Q(j|i) - Q(i|j) \quad (40)$$

because $e^{\epsilon_{DP}(Q)} = \frac{Q(j|j)}{Q(j|i)} > 1$, which contradicts (38). ■

The claims above imply that for D_1 as one of the smallest $\{D_i\}_{i=1}^{M-k}$, all the other diagonal elements are black, and the non-zero elements in the row corresponding to D_1 are all red. This implies that for each red $Q(j|1)$, $j \neq 1$, there exists a diagonal element $1 - D_j$ in the same column j which is $e^{\epsilon_{DP}(Q)}$ times bigger than the red element $Q(j|1)$. The proof is completed by summing over row 1 entries and solving for ϵ_{DP} . ■

Let $0 \leq k \leq M - 1$. Lemma 9 provides a formula for the optimal $\epsilon_{DP}(\cdot)$ among (\mathcal{P}, D) -valid mechanisms with k all-zero columns, in terms of their corresponding distortion values $\{D_i\}_{i=1}^{M-k}$. Besides, no mechanism with at least k all-zero columns can be (\mathcal{P}, D) -valid for $D < D^{(k)}$ due to Lemma 8. Thus, for any k and $D^{(k)} \leq D \leq D^{(k+1)}$, a lower bound on $\epsilon_{DP}^*(\mathcal{P}, D)$ can be derived by taking the minimum over $0 \leq l \leq k$ and all (\mathcal{P}, D) -valid sets of $\{D_i\}_{i=1}^M$ that satisfy

$(M - 1 - k) \frac{1 - \sum_{i=2}^{M-k} D_i}{D_1} \geq 1$, or equivalently $\sum_{i=1}^{M-k} D_i \leq M - 1$. Moreover, for a mechanism with k all-zero columns we have

$D_i = 1$ for $i > M - k$, and thus $\{D_i\}_{i=1}^M$ can be (\mathcal{P}, D) -valid if and only if $\{D_i\}_{i=1}^{M-k}$ is $(\mathcal{P}, D - D^{(k)})$ -valid, i.e. $\sum_{i=1}^{M-k} P_i D_i \leq D - D^{(k)}$. This result in (11) and completes the proof of the lower bound in Theorem 2.

We now proceed to the special case where $D < D^{(1)}$. For proving $\epsilon_{DP}^*(\mathcal{P}, D) \geq \log(M - 1) \frac{1-D}{D}$, it suffices to show that $\epsilon_{DP}^{*(0)}(\mathcal{P}, D)$ is greater than or equal to $\log(M - 1) \frac{1-D}{D}$. We need the following Lemma.

Lemma 10: Let $\{a_i\}_{i=1}^n$, $\{b_i\}_{i=1}^n$, and $\{a'_i\}_{i=1}^n$ be a collection of real numbers between 0 and 1, such that

$$\sum_{i=1}^n a_i = \sum_{i=1}^n a'_i, \quad (41)$$

and $b_1 \leq b_2 \leq \dots \leq b_n$. If $a'_1 \geq a_1$ and $a'_i \leq a_i$ for $i = 2, 3, \dots, n$, then

$$\sum_{i=1}^n a'_i b_i \leq \sum_{i=1}^n a_i b_i. \quad (42)$$

Then, assume the contrary that for some $D < D^{(1)}$, all mechanisms in $Q^*(\mathcal{P}, D)$ achieve a strictly smaller $\epsilon_{DP}(\cdot)$ than $\log(M - 1) \frac{1-D}{D}$. From Lemma 9 we know that there exists an optimal mechanism $Q_{\hat{X}|X}$ with the set of distortions $\{D_i\}_{i=1}^M$, such that $\epsilon_{DP}(Q)$ is given by (32), and without loss of generality $D_1 \leq D_2 \leq \dots \leq D_M$ due to Lemma 4. Hence, the contrary assumption is that there exists an optimal $Q_{\hat{X}|X}$ such that

$$\epsilon_{DP}^*(Q, D) = \epsilon_{DP}(Q_{\hat{X}|X}) = \log(M - 1) \frac{1 - \frac{\sum_{i=2}^M D_i}{M-1}}{D_1} \quad (43a)$$

$$< \log(M - 1) \frac{1 - D}{D}. \quad (43b)$$

Thus:

$$(1 - D)D_1 + \frac{D}{M - 1} \sum_{i=2}^M D_i > D. \quad (44)$$

On the other hand, since $Q_{\hat{X}|X}$ is supposed to satisfy the distortion constraint for any $P \in \mathcal{P}$, including $P^* = \operatorname{argmax}_{P \in \mathcal{P}} P_M$, we have

$$\sum_{i=1}^M P_i^* D_i = P_1^* D_1 + \sum_{i=2}^M P_i^* D_i \quad (45a)$$

$$= (1 - D)D_1 + \frac{D}{M - 1} \sum_{i=2}^M D_i \leq D, \quad (45b)$$

where (45b) is in view of Lemma 10 and the fact that $D \leq D^{(1)}$ implies $P_i^* \geq \frac{D}{M-1}$ for $i = 2, 3, \dots, M$ and $P_1^* \leq 1 - D$. Obviously (45b) contradicts (44). Thus, for $0 \leq D \leq D^{(1)}$, we have

$$\epsilon_{DP}^*(Q, D) \geq \log(M - 1) \frac{1 - D}{D}. \quad (46)$$

2) *Achievability for Theorem 2:* First, we show that for $0 \leq D < D^{(1)}$, the optimal leakage in (10) is achievable. Consider the following mechanism

$$Q(j|i) = \begin{cases} 1 - D, & i = j, \\ D, & i \neq j. \end{cases} \quad (47)$$

Observe that Q is (\mathcal{P}, D) -valid and $\epsilon_{\text{DP}}(Q) = \log(M-1) \frac{1-D}{D}$. Therefore, the lower bound in (46) is tight and $\epsilon_{\text{DP}}^*(\mathcal{P}, D) = \log(M-1) \frac{1-D}{D}$ for $0 \leq D < D^{(1)}$.

We now prove that the lower bound in (10) is achievable for $D^{(1)} \leq D < D^{(M-1)}$. To this end, we construct the following mechanism. For any given $0 \leq k \leq M-1$ and the optimal set $\{D_i^*\}_{i=1}^{M-1}$ in (11), consider the following mechanism.

$$Q^{(k)*}(j|i) = \begin{cases} 1 - D_i^*, & i = j \leq M - k, \\ D_i \frac{1 - D_j^*}{\sum_{l \neq i} 1 - D_l^*}, & i \neq j, i, j < M - k, \\ Q^{(k)*}(j|M - k), & i > M - k, j \leq M - k, \\ 0, & j > M - k. \end{cases} \quad (48)$$

We now verify that $\epsilon_{\text{DP}}(Q^{(k)*}) = (M-1-k) \frac{1 - \frac{\sum_{i=2}^{M-k} D_i^*}{M-1-k}}{D_1^*}$. Since each of the last k rows in the above matrix are equal to the $(M-k)^{\text{th}}$ row, and the last k columns are all equal to zero, it suffices to check the $\epsilon_{\text{DP}}(\cdot)$ for the square matrix formed by the first $M-k$ rows and columns. The ratio of any two elements in the same column in $Q^{(k)*}$ belongs to the set $\{c_1, c_2, \dots, c_{M-k}\}$, where

$$c_i = (M-1-k) \frac{1 - \frac{\sum_{j \neq i} D_j}{M-1-k}}{D_i}. \quad (49)$$

From Lemma 4 we have that $D_1 \leq \dots \leq D_{M-k}$, which in turn implies that for any $0 \leq j \leq M-k$:

$$(M-1-k) \frac{1 - \frac{\sum_{i=2}^{M-k} D_i}{M-1-k}}{D_1} \geq (M-1-k) \frac{1 - \frac{\sum_{i \neq j} D_i}{M-1-k}}{D_j}. \quad (50)$$

Therefore, $\epsilon_{\text{DP}}(Q^{(k)*}) = (M-1-k) \frac{1 - \frac{\sum_{i=2}^{M-k} D_i}{M-1-k}}{D_1}$.

G. Proof of Theorem 3: Class III Source Sets

Proof: Recall that a Class III source set \mathcal{P} can be written as a union of Class II source sets as $\mathcal{P} = \bigcup_{T \in \mathcal{T}_{\mathcal{P}}} \mathcal{P}|_T$. Furthermore, each of these partitions $\mathcal{P}|_T$ can be mapped to \mathcal{S}_0 with the appropriate permutation to get $\overline{\mathcal{P}}|_T$. We write the intersection and union of the mapped partitions as \mathcal{P}^\cap and \mathcal{P}^\cup , respectively. Since \mathcal{P}^\cap and \mathcal{P}^\cup are Class II source sets, we can compute the optimal leakage and the corresponding mechanism for these two sets. Moreover, mapping \mathcal{P}^\cup and \mathcal{P}^\cap back into the original partitions results in two sets \mathcal{P}^{UB} and \mathcal{P}^{LB} that contain and are contained in \mathcal{P} , respectively.

Formally, let

$$\mathcal{P}^{\text{LB}} = \bigcup_{T \in \mathcal{T}_{\mathcal{P}}} T^{-1}(\mathcal{P}^\cap), \quad (51a)$$

$$\mathcal{P}^{\text{UB}} = \bigcup_{T \in \mathcal{T}_{\mathcal{P}}} T^{-1}(\mathcal{P}^\cup). \quad (51b)$$




Fig. 9. A source set \mathcal{P} and its folded versions.

For the \mathcal{P} shown in Figure 2 and the corresponding \mathcal{P}^\cup and \mathcal{P}^\cap in Figure 3, Figure 9 illustrates the \mathcal{P}^{LB} and \mathcal{P}^{UB} .

Recall that for two sets \mathcal{P}_1 and \mathcal{P}_2 , if $\mathcal{P}_1 \subseteq \mathcal{P}_2$, then $\epsilon_{\text{DP}}^*(\mathcal{P}_1, D) \geq \epsilon_{\text{DP}}^*(\mathcal{P}_2, D)$. Thus, it suffices to show the following:

- 1) $\mathcal{P}^{\text{LB}} \subseteq \mathcal{P} \subseteq \mathcal{P}^{\text{UB}}$, and
- 2) The optimal DP leakage for \mathcal{P}^{LB} and \mathcal{P}^{UB} are given by

$$\epsilon_{\text{DP}}^*(\mathcal{P}^{\text{LB}}, D) = \epsilon_{\text{DP:III}}^*(\mathcal{P}^\cap, D, \mathcal{T}_{\mathcal{P}}), \quad (52)$$

$$\epsilon_{\text{DP}}^*(\mathcal{P}^{\text{UB}}, D) = \epsilon_{\text{DP:III}}^*(\mathcal{P}^\cup, D, \mathcal{T}_{\mathcal{P}}), \quad (53)$$

where $\epsilon_{\text{DP:III}}^*(\cdot)$ is defined in (17).

Proof of (i): For each $T \in \mathcal{T}_{\mathcal{P}}$, since $\mathcal{P}^\cap \subseteq \overline{\mathcal{P}}|_T$ we have $T^{-1}(\mathcal{P}^\cap) \subseteq \mathcal{P}|_T$. After taking union over all $T \in \mathcal{T}_{\mathcal{P}}$, we have $\mathcal{P}^{\text{LB}} \subseteq \mathcal{P}$. One can immediately show that $\mathcal{P} \subseteq \mathcal{P}^{\text{UB}}$.

Proof of (ii): We first prove (52). A similar argument proves (53). Recall that $\mathcal{P}^{\text{LB}} = \bigcup_{T \in \mathcal{T}_{\mathcal{P}}} T^{-1}(\mathcal{P}^\cap)$, and thus, for any $P \in \mathcal{P}^\cap$ and $T \in \mathcal{T}_{\mathcal{P}}$ we have $T^{-1}(P) \in \mathcal{P}^{\text{LB}}$. This means that for any given $(\mathcal{P}^{\text{LB}}, D)$ -valid mechanism $Q_{\hat{X}|X}$, $Q_{T(\hat{X})|T(X)}$ is also $(\mathcal{P}^{\text{LB}}, D)$ -valid and $\epsilon_{\text{DP}}(Q_{\hat{X}|X}) = \epsilon_{\text{DP}}(Q_{T(\hat{X})|T(X)})$. Since $\epsilon_{\text{DP}}(Q)$ is a quasi-convex function of Q due to Lemma 1, there exists an optimal mechanism achieving $\epsilon_{\text{DP}}^*(\mathcal{P}^{\text{LB}}, D)$ for which

$$D_{T(i)} = D_i, \quad \text{for any } T \in \mathcal{T}_{\mathcal{P}}, i = 1, \dots, M. \quad (54)$$

Hence, it suffices to search over only those (\mathcal{P}^\cap, D) -valid mechanism that satisfy (54) in order to find $\epsilon_{\text{DP}}^{(k)*}(\mathcal{P}^{\text{LB}}, D)$. Furthermore, since \mathcal{P}^\cap is a Class II source set, we can use the results from Theorem 2. We now show $\epsilon_{\text{DP}}^*(\mathcal{P}^{\text{LB}}, D) = \epsilon_{\text{DP:III}}^*(\mathcal{P}^\cap, D, \mathcal{T}_{\mathcal{P}})$.

First consider the case where $D \geq \frac{M-1}{M}$. Clearly, choosing $Q(i|j) = \frac{1}{M}$ achieves $\epsilon_{\text{DP}}^*(\mathcal{P}^{\text{LB}}, D) = 0$, while the distortion constraint is also satisfied.

For $D < \frac{M-1}{M}$, similar to the proof for Theorem 2, we first restrict the set of mechanisms to those that have a fixed number k of all-zero columns, $k = 0, 1, \dots, M-1$. For any such k , the optimal leakage is given by (18), where the third constraint is a result of (54). Note that (18) results from the addition of the constraint in (54) to the constraints in (11) for a Class II source set. The optimal $\epsilon_{\text{DP}}^*(\mathcal{P}^{\text{LB}}, D)$ is then the minimum of $\epsilon_{\text{DP:III}}^{(k)*}(\mathcal{P}^\cap, D, \mathcal{T}_{\mathcal{P}})$ over all k , resulting in (17).

Finally, for $D < D^{(1)}$, analogous to Theorem 2, we can still show that the optimal mechanism is symmetric. Recall that for a Class II source set \mathcal{P}^\cap and $D < D^{(1)}$, the optimal mechanism achieving $\epsilon_{\text{DP}}^*(\mathcal{P}^\cap, D)$ is symmetric, and thus, does not

violate (54). Hence, we have

$$\epsilon_{\text{DP:III}}^*(\mathcal{P}^\cap, D, \mathcal{T}_P) = \epsilon_{\text{DP}}^*(\mathcal{P}^\cap, D) = \log \left((M-1) \frac{1-D}{D} \right).$$

Note that in contrast to Theorem 2, we no longer have distortion thresholds $D^{(2)}, D^{(3)}, \dots, D^{(M-2)}$, where in each of them only mechanisms with specific number of all-zero columns are allowed. This is due to the constraint in (54), which may not allow a gradual shrinking of output support set.

VI. CONCLUSION

In this paper, we have quantified the privacy-utility tradeoffs for a dataset under different assumptions on distribution knowledge (classes) and for Hamming distortion using differential privacy as the leakage metric. The guarantees we can make under differential privacy are stronger than those under mutual information-based measures of privacy leakage: DP leakage is lower bounded by MI leakage. We divide source sets into three classes. For Class I the optimal mechanism is symmetric. For Class II achieving optimal leakage involves reducing the output space as the distortion increases. For Class III sets we can use Class II results to develop upper and lower bounds on the leakage.

Our results show that symmetric distortion, such as randomized response [24], is optimal when very little is known about the source distribution or when the distortion requirement is very strict. In cases where the source distribution is partially known, data publishers can take advantage of this to tailor a local privacy mechanism to guarantee lower privacy leakage for the same distortion, or lower distortion for the same privacy leakage. These gains can be significant if quite a lot is known about the source, such as Class II sources, and degrades as less and less information is known. These results imply that domain knowledge or public data should be used when designing mechanisms for publishing private data.

There are several interesting questions which we leave for future work. Obviously, a full characterization of Class III sources would be welcome, but the techniques here should extend directly to general discrete distortion measures (linear and nonlinear). Extensions to continuous source distributions may be trickier, but perhaps a starting point would be distributions with bounded support. Finally, understanding the implications of this simple model to categorical and hierarchically categorized data would help build insight into designing practical source-aware data release mechanisms.

REFERENCES

- [1] A. Sarwate and L. Sankar, "A rate-distortion perspective on local differential privacy," in *Proc. 52nd Annu. Allerton Conf. Commun., Control Comput. (Allerton)*, Sep. 2014, pp. 903–908.
- [2] K. Kalantari, L. Sankar, and A. D. Sarwate, "Optimal differential privacy mechanisms under Hamming distortion for structured source classes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2016, pp. 2069–2073.
- [3] C. Dwork, F. McSherry, K. Nissim, and A. D. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. 3rd Theory Cryptograph. Conf. (TCC)*, vol. 3876. New York, NY, USA, Mar. 2006, pp. 265–284.
- [4] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, nos. 3–4, pp. 211–407, 2014.
- [5] A. D. Sarwate and K. Chaudhuri, "Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data," *IEEE Signal Process. Mag.*, vol. 30, no. 5, pp. 86–94, Sep. 2013.
- [6] Q. Geng and P. Viswanath, "Optimal noise adding mechanisms for approximate differential privacy," *IEEE Trans. Inf. Theory*, vol. 62, no. 2, pp. 952–969, Feb. 2016.
- [7] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoffs in databases: An information-theoretic approach," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 6, pp. 838–852, Jun. 2013.
- [8] F. du P. Calmon and N. Fawaz, "Privacy against statistical inference," in *Proc. IEEE 50th Annu. Allerton Conf. Commun. Control Comput. (Allerton)*, Sep. 2012, pp. 1401–1408.
- [9] D. Rebollo-Monedero, J. Forne, and J. Domingo-Ferrer, "From t -closeness-like privacy to postrandomization via information theory," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 11, pp. 1623–1636, Nov. 2010.
- [10] D. Agrawal and C. C. Aggarwal, "On the design and quantification of privacy preserving data mining algorithms," in *Proc. 20th ACM SIGMOD-SIGACT-SIGART Symp. Principles Database Syst.*, New York, NY, USA, 2001, pp. 247–255.
- [11] H. Yamamoto, "A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers (Corresp.)," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 6, pp. 918–923, Nov. 1983.
- [12] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, and C. Palamidessi, "On the relation between differential privacy and quantitative information flow," in *Automata, Languages and Programming*, L. Aceto, M. Henzinger, and J. Sgall, Eds. Berlin, Germany: Springer, 2011, pp. 60–76.
- [13] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, P. Degano, and C. Palamidessi, "Differential privacy: On the trade-off between utility and information leakage," in *Formal Aspects Security Trust*, G. Barthe, A. Datta, and S. Etalle, Eds. Berlin, Germany: Springer, 2012, pp. 39–54.
- [14] Y. H. Liu, S.-H. Lee, and A. Khisti, "Information-theoretic privacy in smart metering systems using cascaded rechargeable batteries," *IEEE Signal Process. Lett.*, vol. 24, no. 3, pp. 314–318, Mar. 2017.
- [15] S. Han, U. Topcu, and G. J. Pappas, "Event-based information-theoretic privacy: A case study of smart meters," in *Proc. Amer. Control Conf.*, Jul. 2016, pp. 2074–2079.
- [16] O. Javidbakht and P. Venkatasubramanian, "Differential privacy in networked data collection," in *Proc. Annu. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2016, pp. 117–122.
- [17] W. Wang, L. Ying, and J. Zhang, "On the relation between identifiability, differential privacy, and mutual-information privacy," *IEEE Trans. Inf. Theory*, vol. 62, no. 9, pp. 5018–5029, Sep. 2016.
- [18] I. Issa, S. Kamath, and A. B. Wagner, "An operational measure of information leakage," in *Proc. Annu. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2016, pp. 234–239.
- [19] P. Cuff and L. Yu, "Differential privacy as a mutual information constraint," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, New York, NY, USA, 2016, pp. 43–54.
- [20] F. du Pin Calmon, A. Makhdoumi, and M. Médard, "Fundamental limits of perfect privacy," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2015, pp. 1796–1800.
- [21] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Information extraction under privacy constraints," *Information*, vol. 7, no. 1, p. 15, Mar. 2016.
- [22] Y. O. Basciftci, Y. Wang, and P. Ishwar, "On privacy-utility tradeoffs for constrained data release mechanisms," in *Proc. Inf. Theory Appl. Workshop (ITA)*, Jan. 2016, pp. 1–6.
- [23] K. Kalantari, O. Kosut, and L. Sankar, "On the fine asymptotics of information theoretic privacy," in *Proc. 54th Annu. Allerton Conf. Commun., Control Comput. (Allerton)*, Sep. 2016, pp. 532–539.
- [24] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *J. Amer. Statist. Assoc.*, vol. 60, no. 309, pp. 63–69, 1965.
- [25] A. Evfimievski, J. Gehrke, and R. Srikant, "Limiting privacy breaches in privacy preserving data mining," in *Proc. 22nd ACM SIGMOD-SIGACT-SIGART Symp. Principles Database Syst. (PODS)*, New York, NY, USA, 2003, pp. 211–222.
- [26] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, "What can we learn privately?" *SIAM J. Comput.*, vol. 40, no. 3, pp. 793–826, Jun. 2011. [Online]. Available: <http://dx.doi.org/10.1137/090756090>

- [27] J. C. Duchi, M. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Proc. IEEE 54th Annu. Symp. Found. Comput. Sci. (FOCS)*, Oct. 2013, pp. 429–438.
- [28] A. Smith, A. Thakurta, and J. Upadhyay, "Is interaction necessary for distributed private learning?" in *Proc. IEEE Symp. Secur. Privacy*, May 2017, pp. 58–77.
- [29] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," *J. Mach. Learn. Res.*, vol. 17, no. 17, pp. 1–51, 2016.
- [30] C. Dwork, "Differential privacy," in *Encyclopedia Cryptography Security*, 2nd ed., H. C. A. van Tilborg and S. Jajodia, Eds. Boston, MA, USA: Springer, 2011, pp. 338–340.
- [31] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge Univ. Press, 2004.
- [32] D.-Z. Du and P. M. Pardalos, Eds., *Minimax and Applications (Nonconvex Optimization and Its Applications)*, 2nd ed. Dordrecht, The Netherlands: Springer, 1995.



Kousha Kalantari received the B.Sc. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, and the M.Sc. degree in electrical engineering from Arizona State University, Tempe, AZ, USA, in 2013 and 2016, respectively. He is currently pursuing the Ph.D. degree with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA, USA. His research interests include wireless communications, information theory, and statistical learning.



Lalitha Sankar (S'02–M'07–SM'15) received the B.Tech. degree from IIT Bombay, Mumbai, India, the M.S. degree from the University of Maryland at College Park, and the Ph.D. degree from Rutgers University. She was an Associate Research Scholar with Princeton University. She is currently an Assistant Professor with the Electrical, Computer, and Energy Engineering Department, Arizona State University. Her research interests include information privacy and cybersecurity in distributed and cyber-physical systems. For her doctoral work, she received the 2007–2008 Electrical Engineering Academic Achievement Award from Rutgers University. She was also a recipient of the IEEE Globecom 2011 Best Paper Award for her work on privacy of side-information in multi-user data systems and the National Science Foundation CAREER Award in 2014.



Anand D. Sarwate (S'99–M'09–SM'14) received the B.S. degrees in electrical engineering and computer science and mathematics from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2002, and the M.S. and Ph.D. degrees in electrical engineering from the Department of Electrical Engineering and Computer Sciences (EECS), University of California at Berkeley (U.C. Berkeley), Berkeley, CA, USA. He was a Post-Doctoral Researcher with the University of California, San Diego, CA, USA, from 2008 to 2011. He was a Research Assistant Professor with the Toyota Technological Institute at Chicago from 2011 to 2013. He has been an Assistant Professor with the Department of Electrical and Computer Engineering, The State University of New Jersey, New Brunswick, NJ, USA, since 2014. His research interests include information theory, machine learning, signal processing, optimization, and privacy and security. He is a member of Phi Beta Kappa and Eta Kappa Nu. He received the National Defense Science and Engineering Graduate Fellowship from 2002 to 2005. He received the NSF CAREER Award in 2015, the Samuel Silver Memorial Scholarship Award, and the Demetri Angelakos Memorial Award from the EECS Department at U.C. Berkeley.