MIT Open Access Articles

Low Influence, Utility, and Independence in Differential Privacy: A Curious Case of (3 2)

Massachusetts Institute of Technology

Low Influence, Utility, and Independence in Differential Privacy: A Curious Case of $\binom{3}{2}$

Rafael G. L. D'Oliveira, Salman Salamatian, Muriel Médard, Parastoo Sadeghi

RLE, Massachusetts Institute of Technology, USA

SEIT, University of New South Wales, Canberra, Australia

Emails: {rafaeld, salmansa, medard}@mit.edu, p.sadeghi@unsw.edu.au

Abstract

We study the relationship between randomized low influence functions and differentially private mechanisms. Our main aim is to formally determine whether differentially private mechanisms are low influence and whether low influence randomized functions can be differentially private. We show that differential privacy does not necessarily imply low influence in a formal sense. However, low influence implies *approximate* differential privacy. These results hold for both independent and non-independent randomized mechanisms, where an important instance of the former is the widely-used additive noise techniques in the differential privacy literature. Our study also reveals the interesting dynamics between utility, low influence, and independence of a differentially private mechanism. As the name of this paper suggests, we show that any two such features are simultaneously possible. However, in order to have a differentially private mechanism that has both utility and low influence, even under a very mild utility condition, one has to employ non-independent mechanisms.

I. INTRODUCTION

Since its inception in 2006, differential privacy [1] has emerged as one of the main frameworks to design, evaluate, and implement privacy-preserving data analysis methods (see e.g., [2], [3] for a survey of results). Some highlighted applications of differential privacy are Apple's large-scale private learning of users' preferences and behaviors [4], and the 2020 United States Census' privatization method to provide data privacy protection [5], each impacting hundreds of millions of individuals.

In the differential privacy framework, databases are mapped to randomized query outputs. The aim of such a randomized mechanism is to make the answer of a query (almost) statistically indistinguishable regardless of whether an individual participates in the database or not. In other words, even if an adversary had knowledge of all records in the database before the participation of a particular individual, they would still have severely limited capability in inferring the private record of the individual from the query's output. In the differential privacy literature, this is often referred to as ensuring individuals in a database have a *low influence* over the query's output, see for example [6], [7]. While this intuitive understanding appears consistent with the definition of differential privacy, in this paper, we argue it is prone to ignoring the core objective of data sharing: to provide utility. It turns out that even taking a mild utility constraint into account, the influence (statistical power) of an individual over the traditional class of independent randomized response DP mechanisms can be quite substantial in an absolute sense.

In this paper, we aim to shed light on whether differentially private mechanisms are low influence mechanisms (or vice versa) in a more formal sense.

A. Are differentially private mechanisms actually low influence?

We begin by revising the definition of differential privacy, formalizing the notion of low influence, and then examining a toy example to establish our key findings.

Setting: We consider a set $\mathcal{D}$ of datasets with a neighborhood relationship. This relationship is a symmetric relationship on $\mathcal{D}$, denoted by $d \sim d'$ whenever $d, d' \in \mathcal{D}$ are neighbors. In the differential privacy literature, a neighborhood is oftentimes defined as two datasets that differ in only one entry (corresponding to the response from one individual). A mechanism is a randomized function $\mathcal{M} : \mathcal{D} \to \mathcal{V}$, where $\mathcal{V}$ is referred to as the output space of the mechanism.

Differential Privacy and Low Influence Conditions:

Definition 1. A randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is (ϵ, δ)-differentially private if for any $d \sim d'$, we have $\Pr[\mathcal{M}(d) \in \mathcal{S}] \le e^\epsilon \Pr[\mathcal{M}(d') \in \mathcal{S}] + \delta$, for every $S \subseteq \mathcal{V}$.

In the original definition [1], $\delta = 0$. The case for $\delta > 0$ is a common relaxation or approximation [1], [8]. Usually, δ much smaller than 1 is desired. Throughout the paper, we will use the shorthand (ϵ, δ)-DP instead of both (ϵ, δ)-differential privacy or (ϵ, δ)-differentially private.

We propose the following adaptation of the notion of low influence from [9].[1]

Definition 2. A randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is ι-low influence (ι-LI for short) if for any $d \sim d'$, we have $\Pr[\mathcal{M}(d) \ne \mathcal{M}(d')] \le \iota$.

In other words, the output of a randomized mechanism does not change much, statistically speaking, for neighboring datasets. Indeed, if $\iota = 0$, then no two neighboring datasets can be statistically distinguished.

In short, the DP constraints on the probabilities are relative (modulo a constant factor δ), while the LI constraints are absolute. To see why such a definition matters and how it is different from the DP constraints, consider a dataset containing YES/NO private votes from $2n + 1$ individuals on a sensitive matter. Assume an adversary somehow knows n individuals voted YES and another n individuals voted NO. Hence, the majority is decided by the last individual, whose vote is unknown to the adversary. Assume a mildly useful (truthful) DP mechanism to approximate the majority function, where a 55% biased coin is tossed to determine whether to reveal the true majority outcome. In the DP framework, the deciding vote is said to be statistically indistinguishable by a rather small multiplicative factor $\epsilon = \ln(\frac{0.55}{0.45}) \approx 0.2$. However, the probability of giving the *same* YES or NO answer, even when the individual voted differently, is only $0.55 \times 0.45 \times 2 = 0.495$. In other words, $\iota = 0.505$ and more than half of the time the deciding vote influences the randomized outcome. In such randomized coin-toss models, insisting on even the slightest utility (slightest bias towards telling the truth) results in $\iota \ge 0.5$ – a 50% or more chance for an individual to influence the majority outcome in a borderline dataset.

It may appear that the only way to achieve lower than 50% influence is through being statistically not truthful in one realization of the majority. For example, assume the mechanism declares YES as majority 91% and 89% of the time when the actual majority is YES and NO, respectively. At the cost of rendering the mechanism almost useless (giving the wrong answer with probability $0.5 \times (0.09 + 0.89) = 0.49$ assuming equally-likely votes), the probability of giving the same answer is raised to $0.91 \times 0.89 + 0.09 \times 0.11 = 0.8198$, meaning $\iota = 0.1802$, while still achieving $\epsilon = \max\{\ln(\frac{0.11}{0.09}), \ln(\frac{0.91}{0.89})\} \approx 0.2$ according to the DP framework.

---

[1]This notion is widely used in the fields of social choice theory and in the analysis of Boolean functions [10].

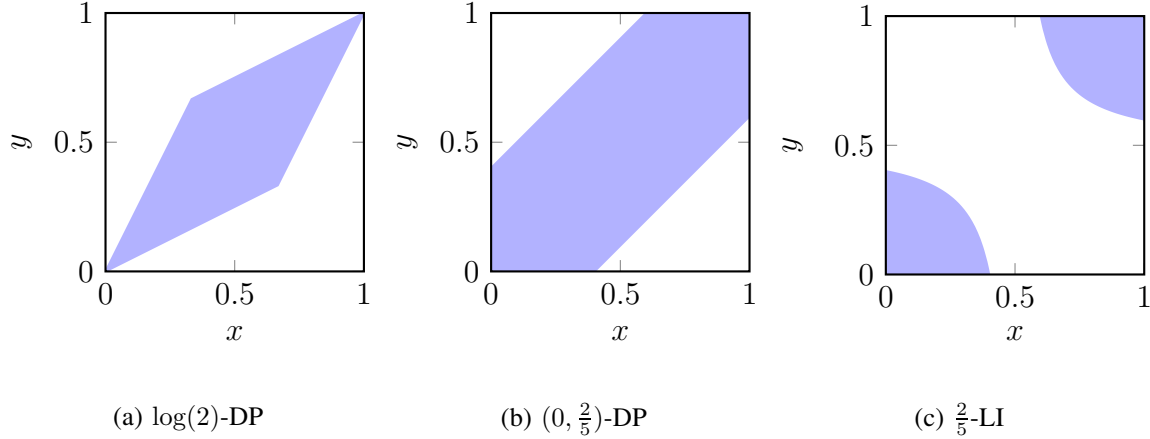(a) $\log(2)$-DP   (b) $(0, \frac{2}{5})$-DP   (c) $\frac{2}{5}$-LI

Fig. 1: Figure (a) and (b) illustrate the (ϵ, δ)-DP regions, and Figure (c) shows the ι-LI region for independent mechanisms, where $x = \Pr[\mathcal{M}(d_1) = 1]$ and $y = \Pr[\mathcal{M}(d_2) = 1]$.

We show in this paper that there is a systematic way to address the above issue through generalizing the randomized response model from independent across datasets to *joint* across datasets. To see how this can be fixed in the above scenario, refer to Example 2.

On Mechanism Independence: In studying the low influence conditions it will be important whether the mechanism is independent or not. A mechanism is independent if the random variables $\{\mathcal{M}(d_i) : d_i \in \mathcal{D}\}$ are mutually independent.

To the best of the authors' knowledge, the differential privacy framework has been solely based on independent mechanisms. Indeed, one of the most common ways of constructing differentially private schemes for continuous-valued queries is to add a noise variable N_u, such as Laplace or Gaussian noise [1], [3] to the true query output $f(d) = u$ to obtain $\mathcal{M}(d) = u + N_u$. For discrete queries taking values over a finite field $\mathbb{F}_q$, this construction is equivalent to simulating noise over a discrete memoryless channel with transition probability $\Pr[\mathcal{M}(d) = v | f(d) = u]$.

Even if the noise statistics are query-output dependent, as is proposed in [11] (see also [12] for more discussion), such schemes are still independent mechanisms. This is because the random variables $\{N_u\}$ are mutually independent and

$$\Pr[\mathcal{M}(d_i) = v_i, \cdots, \mathcal{M}(d_j) = v_j | f(d_i) = u_i, \cdots, f(d_j) = u_j]$$

decomposes into the product $\prod_{\ell=i}^{j} \Pr[\mathcal{M}(d_\ell) = v_\ell | f(d_\ell) = u_\ell]$ for any $i, j \in \{1, \cdots, |\mathcal{D}|\}, i \leq j$. In this paper, we will consider both independent and general (i.e., non-independent) mechanisms on $\mathcal{D}$. By a non-independent mechanism, we mean that the random variables $\{\mathcal{M}(d_i) : d_i \in \mathcal{D}\}$ are in general dependent upon each other.

As alluded to previously, at the surface level, the low influence and differential privacy conditions may seem synonymous to each other. Our first key observation is that this is not the case. As it turns out, the low influence and differential privacy conditions do not generally imply one another. However, low influence does imply the approximate form of differential privacy, i.e., for $\delta > 0$. We illustrate this in the following example.

Example 1. Consider the case where $\mathcal{D}$ consists of two neighboring datasets $d_1 \sim d_2$, and the output space is binary, i.e. $\mathcal{V} = \{1, 2\}$. For the sake of brevity, let $x = \Pr[\mathcal{M}(d_1) = 1]$ and
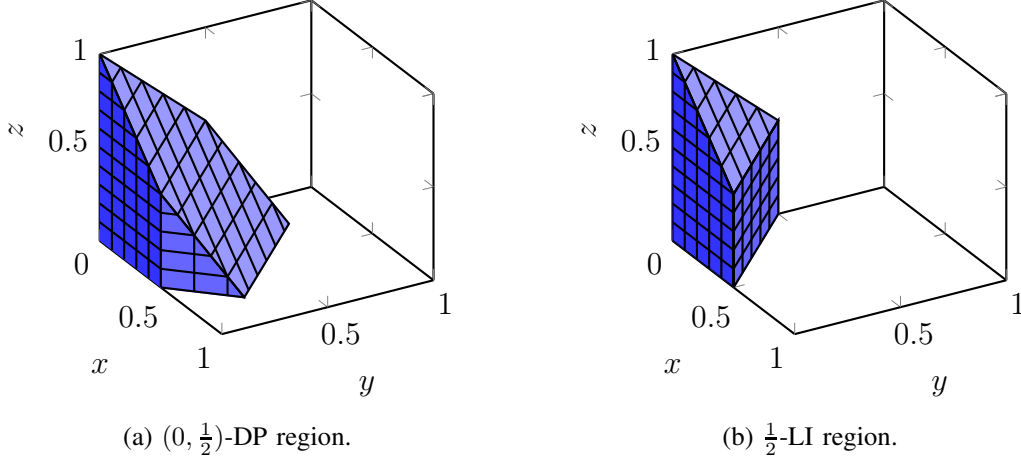
(a) $(0, \frac{1}{2})$-DP region.

(b) $\frac{1}{2}$-LI region.

Fig. 2: (ϵ, δ)-DP and ι-LI regions for non-independent mechanisms for $\delta = \iota = 1/2$.

$y = \Pr[\mathcal{M}(d_2) = 1]$. Thus, the full set of (ϵ, δ)-DP conditions are given by

$$x \le e^\epsilon y + \delta, \quad y \le e^\epsilon x + \delta, \quad 1 - x \le e^\epsilon(1 - y) + \delta, \quad \text{and} \quad 1 - y \le e^\epsilon(1 - x) + \delta.$$

In Figs. 1a and 1b, we show the set of points $(x, y) \in \mathbb{R}^2$ that satisfy these constraints, for $\epsilon = \log(2), \delta = 0$ and for $\epsilon = 0, \delta = \frac{2}{5}$, respectively.

On the other hand, an independent mechanism is ι-LI if and only if x and y satisfy the following hyperbolic constraint.

$$\begin{aligned}
\iota &\ge \Pr[\mathcal{M}(d_1) \ne \mathcal{M}(d_2)] = 1 - \Pr[\mathcal{M}(d_1) = \mathcal{M}(d_2)] \\
&= 1 - \Pr[\mathcal{M}(d_1) = 1, \mathcal{M}(d_2) = 1] - \Pr[\mathcal{M}(d_1) = 2, \mathcal{M}(d_2) = 2] \\
&= 1 - \Pr[\mathcal{M}(d_1) = 1] \Pr[\mathcal{M}(d_2) = 1] - \Pr[\mathcal{M}(d_1) = 2] \Pr[\mathcal{M}(d_2) = 2] \\
&= x + y - 2xy.
\end{aligned}$$

In Fig. 1c, we show the set of points $(x, y) \in \mathbb{R}^2$ that satisfy this constraint for $\iota = \frac{2}{5}$.

Fig. 1 illustrates how, when $\delta = 0$, the (ϵ, δ)-DP region and the ι-LI region are distinct and neither can be embedded into the other. However, an ι-LI region can be embedded into an (ϵ, ι)-DP region for any $\epsilon \ge 0$, meaning that ι-LI implies (ϵ, ι)-DP for independent mechanisms.

This is also true for non-independent mechanisms. Indeed, let $x = \Pr[\mathcal{M}(d_1) = 1, \mathcal{M}(d_2) = 2]$, $y = \Pr[\mathcal{M}(d_1) = 2, \mathcal{M}(d_2) = 1]$, and $z = \Pr[\mathcal{M}(d_1) = 1, \mathcal{M}(d_2) = 1]$, subject to being inside the probability simplex: $0 \le x + y + z \le 1$. To ensure ι-LI, one must have

$$\begin{aligned}
\iota &\ge \Pr[\mathcal{M}(d_1) \ne \mathcal{M}(d_2)] = 1 - \Pr[\mathcal{M}(d_1) = \mathcal{M}(d_2)] \\
&= 1 - \Pr[\mathcal{M}(d_1) = 1, \mathcal{M}(d_2) = 1] - \Pr[\mathcal{M}(d_1) = 2, \mathcal{M}(d_2) = 2] \\
&= 1 - z - (1 - x - y - z) = x + y.
\end{aligned}$$

For brevity, we detail only two of the (ϵ, δ)-DP conditions below

$$x + z \le e^\epsilon(y + z) + \delta \quad \text{and} \quad 1 - x - z \le e^\epsilon(1 - y - z) + \delta.$$

Fig. 2 shows the set of points $(x, y, z) \in \mathbb{R}^3$, which satisfy $(0, \frac{1}{2})$-DP and $\frac{1}{2}$-LI conditions (and also belong to the probability simplex). We can see from the figure that the low influence region, Fig. 2b, is contained in the differential privacy region, Fig. 2a.

(a) $\frac{2}{5}$-LI region
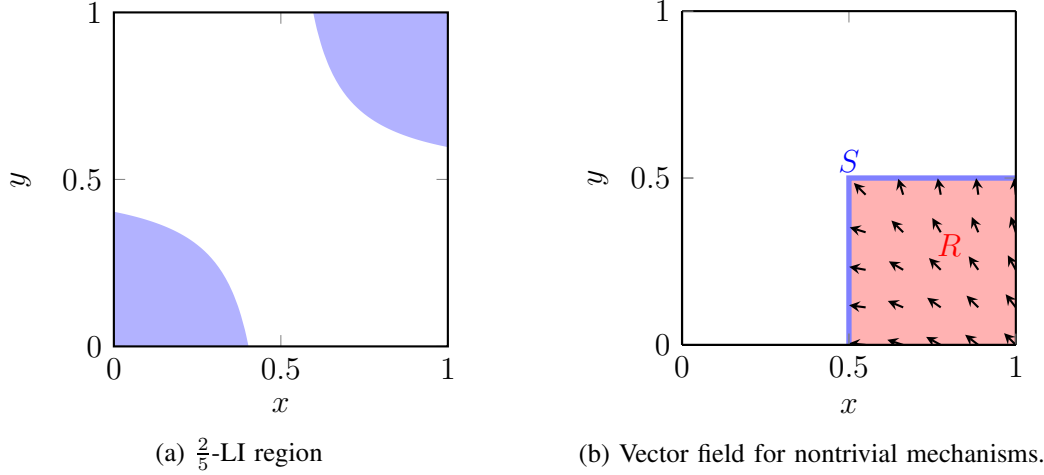

(b) Vector field for nontrivial mechanisms.

Fig. 3: Figure (a) illustrates how an independent mechanism can only have arbitrarily low influence if it is trivial, i.e. x and y are both close to 0 or close to 1. Figure (b) shows that the influence for a nontrivial independent mechanism, i.e. belonging to the red region R, is lower bounded by the influence in the blue boundary region S, denoted by $I(S)$. We show that $I(S) = \frac{1}{2}$.

B. What about the relationship between low influence and utility?

In application, randomized mechanisms are generally used to approximate some function of interest. This approximation is often measured through some notion of utility, which is highly dependent on the application. It is then natural to ask if ι-LI mechanisms can be useful, i.e. can a mechanism have utility while being low influence. It turns out that the answer to this heavily depends on whether the mechanism is independent, operating on single datasets, or not, jointly operating on the space of datasets. To formalize this, we introduce a universal notion of utility without being too restrictive. Roughly speaking, we say a mechanism is *nontrivial* if the most likely output of the mechanism shows some variability across the datasets. Let us continue with the toy example.

Example 1 Continued: We say the binary mechanism operating on $d_1 \sim d_2$ is nontrivial if outputting $v = 0$ is more likely when the database is d_1 and outputting $v = 1$ is more likely when the database is d_2 (or vice versa).[2]

Looking again at independent mechanisms, and setting $x = \Pr[\mathcal{M}(d_1) = 1]$ and $y = \Pr[\mathcal{M}(d_2) = 1]$, a mechanism is nontrivial if $x \geq 1 - x$ and $y \leq 1 - y$ (or vice versa). Without loss of generality, let us assume that $x \geq \frac{1}{2}$ and $y \leq \frac{1}{2}$.[3] It turns out that minimizing the influence function of the independent mechanism $I(x, y) = x + y - 2xy$ subject to $x \geq \frac{1}{2}$ and $y \leq \frac{1}{2}$ leads to $\min I(x, y) = \frac{1}{2}$. Thus, if an independent mechanism is ι-LI, then $\iota \geq \frac{1}{2}$. This result is illustrated in Figure 3. In Section III, we show that this is also true for non-binary mechanisms. *Thus, nontrivial independent mechanisms cannot be arbitrarily low influence.*

However, if we consider general mechanisms, we can achieve both low influence $\iota < \frac{1}{2}$ and a nontrivial utility. Let $x = \Pr[\mathcal{M}(d_1) = 1, \mathcal{M}(d_2) = 2]$, $y = \Pr[\mathcal{M}(d_1) = 2, \mathcal{M}(d_2) = 1]$, and

---

[2] If the most likely query output is $v = 0$ (or $v = 1$) regardless of d_1 or d_2, then there will be no need to query and a fixed guess will always give the maximum likelihood outcome. Recall the voting example, where in two marginal neighboring datasets, being nontrivial means the randomized response must give the true majority answer with a probability more than 0.5.

[3] To satisfy the nontriviality condition, we additionally assume x and y cannot simultaneously be equal to $\frac{1}{2}$.
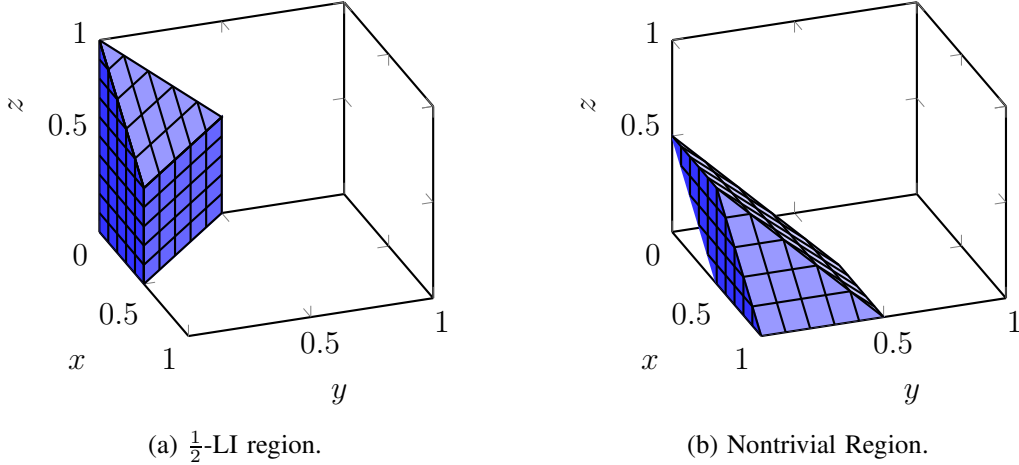
(a) $\frac{1}{2}$-LI region.

(b) Nontrivial Region.

Fig. 4: Nontrivial non-independent mechanisms can have arbitrarily low influence.

$z = \Pr[\mathcal{M}(d_1) = 1, \mathcal{M}(d_2) = 1]$, subject to the probability simplex: $0 \leq x + y + z \leq 1$. Fig. 4 illustrates how the nontrivial region specified by $x + z > \frac{1}{2}$ and $y + z < \frac{1}{2}$ (Fig. 4b) intersects with the $\frac{1}{2}$-LI region specified by $x + y \leq \frac{1}{2}$ (Fig. 4a). This holds true in general, i.e., for any influence $\iota > 0$.

C. Bringing the message home

The toy example illustrates the interesting relationships and tensions between influence, differential privacy, utility and the independence property for randomized mechanisms. To understand the landscape of randomized mechanisms, we consider general mechanisms, beyond independent mechanisms, and characterize precisely how each notion relates to another. In particular, we show the following.

1) Independent nontrivial mechanisms cannot be ι-LI, for $\iota < \frac{1}{2}$. In particular, additive noise mechanisms often used for differential privacy cannot be simultaneously low influence and useful.
2) Low influence mechanisms are differentially private. More specifically, ι-LI finite output mechanisms are $(0, \delta)$-DP, for $\delta = \iota(|\mathcal{V}| - 1)$.
3) There exist non-independent mechanisms which are simultaneously (ϵ, δ)-DP, ι-LI, and nontrivial.

Together, these results form a more complete taxonomy of randomized mechanisms and their properties. An important consequence is that, while independent mechanisms may be sufficient for differential privacy, they fall short when one wants to design mechanisms that satisfy additional properties. The low influence condition is one instance that exemplifies the opportunities that arise from studying general non-independent randomized mechanisms. In other words, when restricting randomized mechanisms to satisfy additional conditions, e.g. low-influence, fairness, or composition rules, it might be necessary to move away from independent mechanisms to achieve the best performance.

Additionally, it turns out that it is possible to characterize the space of general joint mechanisms, at least for mechanisms with finite inputs and outputs. While this space has a larger dimension due to the increased number of degrees of freedom, we show that this increased dimensionality may be a blessing, rather than a curse. For example, we shall show that the ι-LI
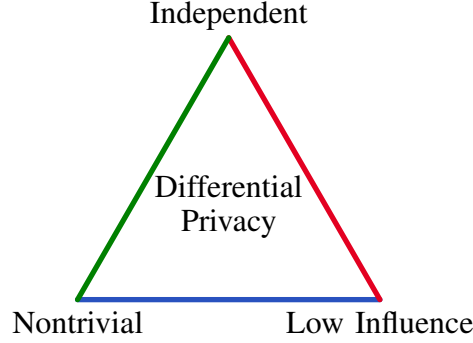
Fig. 5: A differentially private mechanism can satisfy any two of the properties, but not all three.

property can be expressed as a linear constraint over the space of general mechanisms – a fact that is no longer true if one restricts oneself to independent mechanisms.

The relationship between independence, differential privacy, triviality, and low influence is represented graphically in Figure 5. The green edge signifies independent and useful mechanisms that are prevalent in differential privacy literature (but cannot be arbitrarily low influence). The red edge signifies the subset of independent differentially private mechanisms that can be arbitrarily low influence, but are trivial (hence, not useful). The results of this paper advocate for further studying non-independent differentially private mechanisms (existing on the blue edge) that can be jointly low influence and useful. For example, non-independent schemes could be useful for differentially private voting schemes. As we showed in the first numerical example, if a private voting scheme is not low influence, then either a single voter changing their vote can change the outcome significantly, or the noise in the mechanism is playing a significant role in the outcome, both undesirable properties.

D. Notation and Problem Setup

We use $[k]$ to denote the set $\{1, 2, \cdots, k\}$. We denote by $\mathcal{D}$ the space of databases and often enumerate them as $\mathcal{D} = \{d_1, \ldots, d_{|\mathcal{D}|}\}$. We consider a symmetric neighborhood relationship in $\mathcal{D}$ where $d_i, d_j \in \mathcal{D}$ are said to be neighbors if $d_i \sim d_j$. When clear from context, we use the shorthand notation d, d' to refer to two generic neighboring datasets. In many applications, $\mathcal{D} = \mathbb{F}_2^n$ and two databases $d, d' \in \mathbb{F}_2^n$ are neighbors if they differ in one element.

To avoid degenerate cases, in this paper, we only consider spaces of databases with at least two datasets and which are connected, i.e. for any two d and d' in $\mathcal{D}$, there exists a finite sequence $d = d_1, \ldots, d_k = d'$, such that $d_i \sim d_{i+1}$ for $i \in [1 : k - 1]$.

We also consider a finite output space $\mathcal{V}$ which corresponds to the space over which the output of the queries lie. A few remarks are in order about finite output mechanisms, which are relatively less studied in the differential privacy literature compared to continuous-output or countable discrete-output mechanisms. Adding continuous noise to an inherently discrete-valued query (such as a counting query) and post processing may make the output less reliable. With the adoption of differential privacy framework for the US Census, there is emerging motivation to either understand existing discrete privacy-preserving mechanisms through the lens of differential privacy [13] or propose new discrete and finite-valued mechanisms. In this context, the Geometric distribution has been studied in [14]–[16] as the discrete counterpart to the Laplace distribution. The Differential privacy properties of the discrete Gaussian distribution was recently studied

Symbol	Definition
$\mathcal{M}$	A random mechanism.
$\mathcal{D}$	Set of datasets.
$\mathcal{V}$	Output set of the mechanism.
$d \in \mathcal{D}$	A dataset.
$v \in \mathcal{V}$	An output value.
(ϵ, δ)	Differential privacy parameters.
ι	Low influence parameter.
x, y, z	Coordinates of the Euclidean space $\mathbb{R}^3$.
$\mathfrak{I}$	Set of independent mechanisms.
$\mathfrak{M}$	Set of general mechanisms.

TABLE I: List of the main symbols used with their descriptions.

in [17]. However, the support of the noise is the entirety of $\mathbb{Z}$, which may lead to undesirable outcomes. For example, according to a recent article in the New York Times [18], the population of a county in the 2010 US Census was over-reported by a factor of almost 8. Post-processing (truncating or folding the probability mass function) can be used to limit the range of a discrete mechanism [19]. However, such an approach may not be tailored to the problem at hand. A recent work [20] has reported an optimized discrete distribution for counting queries with a finite range as low as $|\mathcal{V}| = 8$ that can outperform the discrete Gaussian distribution in terms of privacy-utility tradeoff.

In Table I we list the main symbols used throughout this paper together with their description.

II. A CLASSIFICATION OF RANDOMIZED MECHANISMS

A. Independent versus Joint Mechanisms

Randomized mechanisms are a key notion in differential privacy. In the literature, these have usually been defined as we do in Definition 3. We, however, refer to them as independent mechanisms. The reason for this is that, in Definition 4, we present a more general notion of mechanism, which we call joint mechanisms. These more general mechanisms are needed for the results in Sections IV and V.

We begin by defining independent mechanisms. To do this, we need the notion of a probability simplex. Given a finite set $\mathcal{V}$, the probability simplex over $\mathcal{V}$ is the set

$$\Delta(\mathcal{V}) = \left\{ \lambda = (\lambda_1, \cdots, \lambda_{|\mathcal{V}|}) \in \mathbb{R}^{|\mathcal{V}|} : \lambda_i \geq 0 \quad \text{for all } i, \text{ and } \sum_{i=1}^{|\mathcal{V}|} \lambda_i = 1 \right\}.$$

Definition 3. An independent mechanism $\mathcal{M}$ with domain $\mathcal{D}$ and range $\mathcal{V}$ is associated with a mapping $g : \mathcal{D} \rightarrow \Delta(\mathcal{V})$. On input $d \in \mathcal{D}$, the randomized mechanism outputs $v \in \mathcal{V}$ with probability $\Pr[\mathcal{M}(d) = v]$. We denote the set of independent mechanisms by $\mathfrak{I}$.

An independent mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is solely determined by the marginal probabilities $M_{ij} \doteq \Pr[\mathcal{M}(d_i) = v_j]$, which can be arranged in a $|\mathcal{D}| \times |\mathcal{V}|$ stochastic matrix. In other words,

$$\mathfrak{I} \cong \left\{ M \in \mathbb{R}^{|\mathcal{D}| \times |\mathcal{V}|} : M_{ij} \geq 0, \quad (i, j) \in [|\mathcal{D}|] \times [|\mathcal{V}|], \quad \sum_{j=1}^{|\mathcal{V}|} M_{ij} = 1, \quad i \in [|\mathcal{D}|] \right\}.$$

Because of this relation, we often identify $\mathfrak{I}$ with its isomorphic subset of $\mathbb{R}^{|\mathcal{D}| \times |\mathcal{V}|}$.

Note that independent mechanisms contain both (less common) *query-output-dependent* mechanisms, such as [11] (see also [12] for more discussion) and *query-output-independent* mechanisms that are much more commonplace in the differential privacy literature. For discrete queries taking values over a finite field $\mathbb{F}_q$, we say an (additive) independent mechanism operating on the true query output $f(d)$ is query-output-independent if for all $u, v \in \mathbb{F}_q$ and $d \in \mathcal{D}$, we have $\Pr[\mathcal{M}(d) = v | f(d) = u] = \Pr[N = v - u]$, where N is an independent noise variable whose distribution does not depend on u. We say an (additive) independent mechanism operating on the true query output $f(d)$ is query-output-dependent if for all $u, v \in \mathbb{F}_q$ and $d \in \mathcal{D}$, we have $\Pr[\mathcal{M}(d) = v | f(d) = u] = \Pr[N_u = v - u]$, where N_u is an independent noise variable whose distribution depends on u. We refer to both such mechanisms as independent because, as for all independent mechanisms, the random variables $\{\mathcal{M}(d) : d \in \mathcal{D}\}$ are mutually independent.

This is not the case for the joint mechanisms we now define.

Definition 4. A joint randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is defined by the joint probabilities $\Pr[\mathcal{M}(d_1) = v_1, \ldots, \mathcal{M}(d_{|\mathcal{D}|}) = v_{|\mathcal{D}|}]$ for every $(v_1, \ldots, v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}$. We denote the set of joint mechanisms by $\mathfrak{M}$.

A joint mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is solely determined by the joint probabilities of the form $P_{v_1, \ldots, v_{|\mathcal{D}|}} \doteq \Pr[\mathcal{M}(d_1) = v_1, \ldots, \mathcal{M}(d_{|\mathcal{D}|}) = v_{|\mathcal{D}|}]$. Thus,

$$\mathfrak{M} \cong \left\{ P_{v_1, \ldots, v_{|\mathcal{D}|}} \in \mathbb{R}^{|\mathcal{V}|^{|\mathcal{D}|}} : P_{v_1, \ldots, v_{|\mathcal{D}|}} \geq 0, \quad \text{for every } (v_1, \ldots, v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|} \text{ and} \right.$$

$$\left. \sum_{(v_1, \ldots, v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}} P_{v_1, \ldots, v_{|\mathcal{D}|}} = 1 \right\}. \quad (1)$$

Because of this relation, we often identify $\mathfrak{M}$ with its isomorphic subset of $\mathbb{R}^{|\mathcal{V}|^{|\mathcal{D}|}}$.

Clearly the set of joint mechanisms contains the class of independent mechanisms as a special case. That is, for an independent mechanism and for every $(v_1, \ldots, v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}$, we have

$$\Pr[\mathcal{M}(d_1) = v_1, \ldots, \mathcal{M}(d_{|\mathcal{D}|}) = v_{|\mathcal{D}|}] = \prod_{i=1}^{|\mathcal{D}|} \Pr[\mathcal{M}(d_i) = v_i].$$

We have already identified the space of independent mechanisms $\mathfrak{I}$ with a subset of $\mathbb{R}^{|\mathcal{D}| \times |\mathcal{V}|}$, and the space of general mechanisms $\mathfrak{M}$ with a subset of $\mathbb{R}^{|\mathcal{V}|^{|\mathcal{D}|}}$. Since an independent mechanism is a special case of joint mechanism, it is clear that there must exist a subset of $\mathfrak{M}$ isomorphic to $\mathfrak{I}$. We characterize this subset in the following proposition.

Proposition 1. *It holds that,*

$$\mathfrak{I} \cong \left\{ P_{v_1,\ldots,v_{|\mathcal{D}|}} \in \mathfrak{M} : P_{v_1,\ldots,v_{|\mathcal{D}|}} = \prod_{i=1}^{|\mathcal{D}|} \sum_{(w_1,\ldots,w_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|} : w_i = v_i} P_{w_1,\ldots,w_{|\mathcal{D}|}}, \right.$$

$$\left. \text{for every } (v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|} \right\}.$$

In other words, $\mathfrak{I}$ is a polynomial of degree $|\mathcal{D}|$ in $\mathbb{R}^{|\mathcal{V}|^{|\mathcal{D}|}}$.

Proof. See the Appendix. □

Not to confuse the previous identifications, we refer to the space in the above proposition as $\mathfrak{I}_{\mathfrak{M}}$. The statement of the proposition is then equivalent to $\mathfrak{I} \cong \mathfrak{I}_{\mathfrak{M}} \subset \mathfrak{M}$. Throughout the paper, when we refer to a mechanism without specification, we mean a joint mechanism as these include independent ones.

We finish this section by showing a different characterization of randomized mechanisms and giving an explicit example of a non-independent mechanism.

Proposition 2. *A randomized mechanism is equivalent to a probability distribution on the space of all functions $h : \mathcal{D} \to \mathcal{V}$.*

Proof. See the Appendix. □

This proposition gives an intuitive way of constructing randomized mechanisms by randomly choosing a function $h : \mathcal{D} \to \mathcal{V}$, as we do in the following example.

Example 2. Consider the case where $\mathcal{D}$ consists of two neighboring datasets $d_1 \sim d_2$, and the output space is binary, i.e., $\mathcal{V} = \{1, 2\}$. Then, the set of all functions $h : \mathcal{D} \to \mathcal{V}$ is given by $\mathcal{V}^{\mathcal{D}} = \{f_{11}, f_{12}, f_{21}, f_{22}\}$, where $f_{ij} : \mathcal{D} \to \mathcal{V}$ is such that

$$f_{ij}(d_k) = \begin{cases} i & \text{if } k = 1, \\ j & \text{if } k = 2. \end{cases}$$

By proposition 2, a randomized mechanism is equivalent to a probability distribution over $\mathcal{V}^{\mathcal{D}}$. We denote this distribution by a 2×2 matrix P such that P_{ij} is the probability of choosing the function f_{ij}. We define a random variable F to denote this random function, thus $P_{ij} = \Pr[F = f_{ij}]$. The random mechanism is then obtained by randomly choosing a function $f_{ij} \in \mathcal{V}^{\mathcal{D}}$ according to P and evaluating it on whatever dataset one has. Thus,

$$\Pr[\mathcal{M}(d_1) = i, \mathcal{M}(d_2) = j] = \Pr[F(d_1) = i, F(d_2) = j] = \Pr[F = f_{ij}] = P_{ij}.$$

Thus, the mechanism $\mathcal{M}$ is independent if and only if $P_{ij} = (P_{i1} + P_{i2})(P_{1j} + P_{2j})$.

Revisiting the first numerical example in the introduction, let us denote by d_1, d_2 two datasets where the last individual voted YES (outcome 1) and NO (outcome 2), respectively. We use the following probability matrix for choosing f_{ij}:

$$P = \begin{bmatrix} 0.45 & 0.1 \\ 0 & 0.45 \end{bmatrix} \tag{2}$$

It is clear that this mechanism is not independent and the probability of giving different answers due to the vote of the last individual is capped at $P(\mathcal{M}(d_1) \neq \mathcal{M}(d_2)) = P(\mathcal{M}(d_1) =$

$1, \mathcal{M}(d_2) = 2) = P(F = f_{12}) = P_{12} = 0.1$, while the probability of being truthful in both YES and NO majority cases is $P(\mathcal{M}(d_1) = 1) = P(F = f_{11}) + P(F = f_{12}) = P(\mathcal{M}(d_2) = 2) = P(F = f_{12}) + P(F = f_{22}) = 0.55$, resulting in the same $\epsilon \approx 0.2$ as before. Note that we do not have to have a symmetric randomized response matrix to ensure a mechanism which does not favor one majority outcome over another.

B. Differential Privacy, Low Influence, and Nontrivial Mechanisms

In this section we define different families of mechanisms.

Recall that differential privacy was defined in Definition 1. We also use the following notion.

Definition 5. A mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is (ϵ, δ)-value differentially private, or (ϵ, δ)-VDP in short, if for any two neighboring $d \sim d'$, we have $\Pr[\mathcal{M}(d) = v] \leq e^\epsilon \Pr[\mathcal{M}(d') = v] + \delta$, for every $v \in \mathcal{V}$.

A mechanism which is value differentially private is also differentially private.

Proposition 3. Let the mechanism $\mathcal{M}$ be (ϵ, δ)-VDP. Then, $\mathcal{M}$ is $(\epsilon, (|\mathcal{V}| - 1)\delta)$-DP.

Proof. Since $\mathcal{M}$ is (ϵ, δ)-VDP, it follows that $\Pr[\mathcal{M}(d) = v] \leq e^\epsilon \Pr[\mathcal{M}(d') = v] + \delta$, for every neighboring datasets $d \sim d'$. But then, for every neighboring datasets $d \sim d'$,

$$
\Pr[\mathcal{M}(d) \in \mathcal{S}] = \sum_{v \in \mathcal{S}} \Pr[\mathcal{M}(d) = v] \leq \sum_{v \in \mathcal{S}} (e^\epsilon \Pr[\mathcal{M}(d') = v] + \delta)
$$
$$
= e^\epsilon \Pr[\mathcal{M}(d') \in \mathcal{S}] + |\mathcal{S}|\delta.
$$

Since the differential privacy condition trivially holds for $\mathcal{S} = \mathcal{V}$, we can upper bound $|\mathcal{S}|$ by $|\mathcal{V}| - 1$, proving the result. $\square$

In Definition 2, we present the notion of low influence mechanisms. This notion is an adaptation of the following definition, widely used in social choice theory and in the analysis of Boolean functions.

Definition 6. The influence of coordinate i on a function $f : \{-1, 1\}^n \to \{-1, 1\}$ is defined as the probability $\text{Inf}_i[f] = \Pr[f(x) \neq f(x^{\oplus i})]$, where x is a uniformly distributed vector from $\{-1, 1\}^n$, and $x^{\oplus i}$ is the vector such that $x_i^{\oplus i} = -x_i$ and $x_k^{\oplus i} = x_k$ for $k \in [n] - \{i\}$.

If one thinks of f as a voting rule, i.e., a rule for determining how to interpret n votes for two candidates $\{-1, 1\}$ as an election of one of them, then the influence $\text{Inf}_i[f]$ measures the probability that the i-th voter will affect the outcome of the election. In this context, we say a function f is ι-low influence if $\text{Inf}_i[f] \leq \iota$ for every $i \in [n]$.

We note that the probability in the above definition is taken over the argument of the function, and that the function itself is deterministic. The opposite is true in the differential privacy setup. The definition of low influence in Definition 2 in the introduction accounts for this in measuring the influence through $\Pr[\mathcal{M}(d) \neq \mathcal{M}(d')] \leq \iota$, where the randomness is in the randomized mapping $\mathcal{M}$, not the inputs d and d'.

Finally, we define nontrivial mechanisms as a randomized generalization of non-constant mechanisms. More precisely, we say a mechanism is nontrivial if its most likely output is not constant across datasets.

Definition 7. A mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is nontrivial if there exists neighboring datasets $d \sim d'$ such that $\arg\max_v \Pr[\mathcal{M}(d) = v] \neq \arg\max_v \Pr[\mathcal{M}(d') = v]$.

Note this definition of nontriviality is very mild. In practice, the utility of a mechanism is application dependent, and the mechanism being nontrivial is not sufficient, but rather necessary. Recall the voting example, where to nontrivially approximate the majority function, the most likely outcome must be different across two marginal datasets. We saw this can result in high ι influence in independent mechanisms, but can be controlled via joint mechanisms. Across non-marginal datasets, the most likely outcome may be identical without being detrimental to approximating the true majority function, DP or LI constraints.

III. Nontrivial Independent Mechanisms are Not Low Influence

In this section, we show that if an independent mechanism is nontrivial, then it cannot be low influence. We begin by showing that although differential privacy is equivalent to a set of linear constraints on $\mathfrak{I}$, low influence is equivalent to a set of quadratic constraints.

Proposition 4. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be an independent mechanism with matrix representation $M_{ij} = \Pr[\mathcal{M}(d_i) = v_j]$. Then, $\mathcal{M}$ is (ϵ, δ)-DP if and only if, for every $d_i \sim d_k$,*

$$\sum_{j : v_j \in \mathcal{S}} M_{ij} \leq e^{\epsilon} \sum_{j : v_j \in \mathcal{S}} M_{kj} + \delta \quad \text{for every } \mathcal{S} \subseteq \mathcal{V}.$$

Similarly, M is (ϵ, δ)-VDP if and only if, for every $d_i \sim d_k$, it holds that $M_{ij} \leq e^{\epsilon} M_{kj} + \delta$, for every $j \in [|\mathcal{V}|]$. In other words, both (ϵ, δ)-DP and (ϵ, δ)-VDP are linear constraints in the space $\mathfrak{I}$.

Proof. This follows directly from substituting $M_{ij} = \Pr[\mathcal{M}(d_i) = v_j]$ into the definitions of (ϵ, δ)-DP and (ϵ, δ)-VDP. □

Proposition 5. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be an independent mechanism with matrix representation $M_{ij} = \Pr[\mathcal{M}(d_i) = v_j]$. Then $\mathcal{M}$ is ι-LI if and only if, for every $d_i \sim d_k$,*

$$\sum_{j=1}^{|\mathcal{V}|} M_{ij} M_{kj} \geq 1 - \iota.$$

In other words, ι-LI conditions are quadratic constraints over the space $\mathfrak{I}$.

Proof. By the definition of ι-LI, for every $d_i \sim d_k$,

$$\iota \geq \Pr[\mathcal{M}(d_i) \neq \mathcal{M}(d_k)] = 1 - \Pr[\mathcal{M}(d_i) = \mathcal{M}(d_k)]$$

$$= 1 - \sum_{j=1}^{|\mathcal{V}|} \Pr[\mathcal{M}(d_i) = j, \mathcal{M}(d_k) = j]$$

$$= 1 - \sum_{j=1}^{|\mathcal{V}|} \Pr[\mathcal{M}(d_i) = j] \Pr[\mathcal{M}(d_k) = j]$$

$$= 1 - \sum_{j=1}^{|\mathcal{V}|} M_{ij} M_{kj}.$$

□

Before we prove the main result of this section, let us revisit the toy example from the Introduction. We show that for our toy example, a nontrivial mechanism is not ι-LI for $\iota < \frac{1}{2}$.

Example 3. Consider the case where $\mathcal{D}$ consists of two neighboring datasets $d_1 \sim d_2$, and the output space is binary, i.e. $\mathcal{V} = \{1, 2\}$. To be consistent with the simpler notation in the toy example in the Introduction, let $x \doteq M_{11} = \Pr[\mathcal{M}(d_1) = 1]$ and $y \doteq M_{21} = \Pr[\mathcal{M}(d_2) = 1]$. Suppose M is nontrivial. Then, either (i) $\frac{1}{2} \leq x \leq 1$ and $0 \leq y \leq \frac{1}{2}$ or (ii) $0 \leq x \leq \frac{1}{2}$ and $\frac{1}{2} \leq y \leq 1$. Recall that we additionally assume x and y cannot simultaneously be equal to $\frac{1}{2}$. Without loss of generality, we consider case (i) and denote

$$R^* = \left\{ (x, y) \in \mathbb{R}^2 : \frac{1}{2} \leq x \leq 1 \quad \text{and} \quad 0 \leq y \leq \frac{1}{2} \right\}. \tag{3}$$

Let $I : \mathbb{R}^2 \to \mathbb{R}$ be given by $I(x, y) = x + y - 2xy$. Then, a mechanism is ι-LI if and only if $I(x, y) \leq \iota$. Consider the following constrained quadratic optimization problem.

$$
\begin{aligned}
\text{minimize} \quad & I(x, y) = x + y - 2xy \\
\text{subject to} \quad & \frac{1}{2} \leq x \leq 1, \\
& 0 \leq y \leq \frac{1}{2}.
\end{aligned} \tag{4}
$$

The solution to this problem is the minimum influence that a nontrivial mechanism can have.

We show that the solution to this problem is the set

$$S = \left\{ \left(\frac{1}{2}, t \right) : t \in \left[0, \frac{1}{2} \right[\right\} \bigcup \left\{ \left(t, \frac{1}{2} \right) : t \in \left] \frac{1}{2}, 1 \right] \right\},$$

with $I(S) = \frac{1}{2}$. To do this, we look at the direction of fastest decrease, $-\nabla I = (2y - 1, 2x - 1)$. Note that if $(x, y) \in R^* - S$ then $-(\nabla I)_x < 0$ and $-(\nabla I)_y > 0$. This means that, for every point $(x, y) \in R^*$ the direction $-\nabla I(x, y)$ points towards S. Thus, the solution space of the optimization problem (4) must be contained in S. A direct calculation shows that $I(S) = \frac{1}{2}$, concluding that nontrivial mechanisms cannot have influence smaller than $\frac{1}{2}$. We illustrate this in Fig. 3b.

The lower bound in the example above holds true in general.

Theorem 1. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be an independent, ι-LI, nontrivial mechanism. Then, $\iota \geq \frac{1}{2}$. Furthermore, this bound is tight.*

Proof. We present the full proof in the Appendix. The main idea behind the proof is to characterize the set of nontrivial schemes R and show that the minimum influence in that set is equivalent to the minimum influence in the set R^* of Example 3. Thus, $\iota \geq \frac{1}{2}$. $\qquad\square$

We note that we can obtain a looser lower-bound on the influence of independent nontrivial mechanisms using a simpler argument. For this, we use the mechanism support size $|\mathcal{V}|$ and the nontrivial mechanism assumption to find a simple lower bound on $\min_{d \in \mathcal{D}} \max_{v \in [\mathcal{V}]} \Pr[\mathcal{M}(d) = v]$. However, this bound becomes loose as $|\mathcal{V}|$ increases.

Proposition 6. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be independent, nontrivial, and ι-LI. Then, $\iota \geq \frac{1}{|\mathcal{V}|^2}$.*

Proof. Since $\mathcal{M}$ is nontrivial, there exists neighboring datasets $d_1 \sim d_2$ and two different values

$v_1, v_2 \in \mathcal{V}$ such that $\Pr[\mathcal{M}(d_1) = v_1] \geq 1/|\mathcal{V}|$ and $\Pr[\mathcal{M}(d_2) = v_2] \geq 1/|\mathcal{V}|$. Then,

$$\begin{aligned}
\Pr[\mathcal{M}(d_1) \neq \mathcal{M}(d_2)] &\geq \Pr[\mathcal{M}(d_1) = v_1, \mathcal{M}(d_2) = v_2] \\
&= \Pr[\mathcal{M}(d_1) = v_1]\Pr[\mathcal{M}(d_2) = v_2] \\
&\geq \frac{1}{|\mathcal{V}|}\frac{1}{|\mathcal{V}|} = \frac{1}{|\mathcal{V}|^2}.
\end{aligned}$$

$\square$

The bound in Theorem 1 is strictly larger than the bound above for $|\mathcal{V}| > 2$. The following bound, however, can sometimes outperform the bound in Theorem 1 and is not restricted only to nontrivial mechanisms.

Proposition 7. *Let the mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be independent and ι-LI. Then,*

$$\iota \geq 1 - \min_{d \in \mathcal{D}} \max_{v \in \mathcal{V}} \Pr[\mathcal{M}(d) = v].$$

First, we define a function which measures the local influence of any two datasets.

Definition 8. Let $\mathcal{V}$ be a finite set. We define $I : \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|} \to \mathbb{R}$, such that $I(u, v) = 1 - \sum_{i=1}^{\mathcal{V}} u_i v_i$.

Proof. Let I be the function in Definition 8. For the independent mechanism $\mathcal{M}$ with matrix representation $M_{ij} = \Pr[\mathcal{M}(d_i) = v_j]$ and rows M_i and M_k, we can write

$$\begin{aligned}
I(M_i, M_k) = 1 - \sum_{j=1}^{|\mathcal{V}|} M_{ij} M_{kj} &\geq 1 - \min\left(\max_{j \in [|\mathcal{V}|]} M_{ij}, \max_{j \in [|\mathcal{V}|]} M_{kj}\right) \\
&\geq 1 - \min_{l \in [|\mathcal{D}|]} \max_{j \in [|\mathcal{V}|]} M_{lj} = 1 - \min_{d \in \mathcal{D}} \max_{v \in [\mathcal{V}]} \Pr[\mathcal{M}(d) = v].
\end{aligned}$$

But by Lemma 1 in the Appendix, $\mathcal{M}$ is ι-LI if and only if $I(M_i, M_k) \leq \iota$ for every $d_i \sim d_k$. Thus,

$$\iota \geq 1 - \min_{d \in \mathcal{D}} \max_{v \in \mathcal{V}} \Pr[\mathcal{M}(d) = v].$$

$\square$

This bound can be better than the one in Theorem 1, if $\min_{d \in \mathcal{D}} \max_{v \in \mathcal{V}} \Pr[\mathcal{M}(d) = v] < \frac{1}{2}$. At best, $\min_{d \in \mathcal{D}} \max_{v \in \mathcal{V}} \Pr[\mathcal{M}(d) = v] = \frac{1}{|\mathcal{V}|}$. At worst, it could be arbitrarily close to the trivial bound of 1. Thus, although it can sometimes outperform the bound in Theorem 1, it cannot guarantee that the mechanism is bounded away from 0, as Theorem 1 does.

Remark 1. In [21], the authors studied statistical stability of (independent) differentially private mechanisms in dealing with adaptive data queries. The main question addressed in [21] is how many adaptive statistical queries can one reliably answer using *the same* n data samples which are drawn i.i.d from an underlying data distribution $\mathcal{P}$.

Recall that for a differentially private mechanism, $\Pr[\mathcal{M}(d) \in \mathcal{S}]$ does not change fast between neighboring datasets as a function of d. This fact is used in [21] to prove that the empirical average of a randomly chosen hypothesis ϕ that is adaptively selected *after* observing $\mathcal{M}(d)$ will be close to the true expectation of ϕ with *high probability*. In other words, the chance of observing "unlucky nontypical" datasets is small, where by unlucky we mean the empirical average of a function dependent on d is far away from its true expectation. This can be thought of as a notion of *typical* invariance or stability to particular realizations of the dataset.

We remark that our results in this paper are not inconsistent with those in [21]. The notion of influence in this paper is different from [21] and is a worst-case measure: in the sense that the bound ι on the influence must be universally valid for *all* neighboring datasets. Essentially, the fact that ι for independent mechanisms is bounded away from 0, as shown in Theorem 1 means that in order to preserve utility *there exists at least two (and possibly many more) neighboring datasets which are likely to give very different results to the same query*. That is, for independent mechanisms there exists "unlucky nontypical datasets" which if hit by chance, can lead to unreliable statistical measures.

In Section V, we show that this issue can be overcome through designing mechanisms on the joint space of datasets. That is, nontrivial mechanisms can be universally low influence if the mechanisms are non-independent.

IV. Low Influence Mechanisms are Differentially Private

In this section, we show that low influence functions are differentially private. We will not restrict ourselves to only independent mechanism, considering joint mechanisms too. We begin by showing that any linear constraint on $\mathfrak{I}$ is also linear on $\mathfrak{M}$.

Proposition 8. *If a constraint is linear on $\mathfrak{I}$, it is linear on $\mathfrak{M}$. In particular, both (ϵ, δ)-DP and (ϵ, δ)-VDP are linear on $\mathfrak{M}$.*

Proof. Let $P \in \mathfrak{M}$ with $P_{v_1,\ldots,v_{|\mathcal{D}|}} = \Pr[\mathcal{M}(d_1) = v_1, \ldots \mathcal{M}(d_{|\mathcal{D}|}) = v_{|\mathcal{D}|}]$ and $M \in \mathfrak{I}$ with $M_{kj} = \Pr[\mathcal{M}(d_k) = v_j]$. Then,

$$M_{kj} = \Pr[\mathcal{M}(d_k) = v_j] = \sum_{(w_1,\ldots,w_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}: w_k = v_j} P_{w_1,\ldots,w_{|\mathcal{D}|}}.$$

Thus, each M_{kj} is a linear combination of $P_{v_1,\ldots,v_{|\mathcal{D}|}}$. $\qquad\square$

Thus, looking at $\mathfrak{M}$ does not complicate the differential privacy constraints. Indeed, it will also simplify the low influence constraints into linear constraints in the space $\mathfrak{M}$ (which were non-linear in the space $\mathfrak{I}$). Proposition 1 explains why the low influence constraints are linear on $\mathfrak{M}$, but non-linear (quadratic) on $\mathfrak{I}$. This is illustrated in Figure 6 for $|\mathcal{V}| = 2$.

Proposition 9. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ such that $P_{v_1,\ldots,v_{|\mathcal{D}|}} = \Pr[\mathcal{M}(d_1) = v_1, \ldots, \mathcal{M}(d_{|\mathcal{D}|}) = v_{|\mathcal{D}|}]$. Then $\mathcal{M}$ is ι-LI if and only if, for every $d_i \sim d_k$,*

$$\sum_{(v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}: v_i \neq v_k} P_{v_1,\ldots,v_{|\mathcal{D}|}} \leq \iota.$$

*In other words, ι-LI is a linear constraint in the space $\mathfrak{M}$.*

Proof. The mechanism $\mathcal{M}$ is ι-LI if and only if, for any two neighboring datasets $d_i \sim d_k$,

$$\iota \geq \Pr[\mathcal{M}(d_i) \neq \mathcal{M}(d_k)] = \sum_{(v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}: v_i \neq v_k} P_{v_1,\ldots,v_{|\mathcal{D}|}}.$$

$\square$

We revisit the toy example in the introduction and show that ι-LI implies in $(0, \iota)$-DP.

Example 4. Consider the case where $\mathcal{D}$ consists of two neighboring datasets $d_1 \sim d_2$, and the output space is binary, i.e. $\mathcal{V} = \{1, 2\}$. A general mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ is then determined by
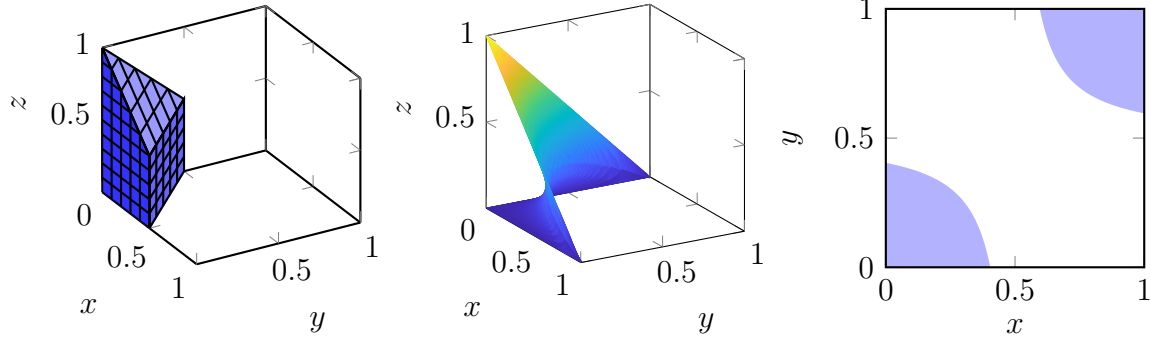
(a) Low influence in the space of general mechanisms.

(b) Independence surface in the space of general mechanisms.

(c) Low influence in the space of independent mechanisms.

Fig. 6: For $|\mathcal{V}| = 2$, Figure (a) shows the region for low influence in the space of general (joint) mechanisms. Figure (b) shows the surface of independent mechanisms. Figure (c) can be interpreted as an intersection of (a) and (b), and shows the region for low influence in the space of independent mechanisms.

the matrix $P = \begin{pmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{pmatrix}$, where $P_{ij} = \Pr[\mathcal{M}(d_1) = i, \mathcal{M}(d_2) = j]$. We note that the marginals are given by $\Pr[\mathcal{M}(d_1) = k] = P_{k1} + P_{k2}$ and $\Pr[\mathcal{M}(d_2) = k] = P_{1k} + P_{2k}$. Thus, a simple calculation shows that, $\mathcal{M}$ is $(0, \delta)$-DP if and only if $|P_{12} - P_{21}| \leq \delta$, i.e. if the entries in the anti-diagonal of P are close. By Proposition 9, the mechanism $\mathcal{M}$ is ι-LI if and only if $P_{12} + P_{21} \leq \iota$, i.e. if the sum of the entries in the anti-diagonal of P is close to zero. Thus, since $|P_{12} - P_{21}| \leq P_{12} + P_{21} \leq \iota$, it follows that if the mechanism $\mathcal{M}$ is ι-LI, it is also $(0, \iota)$-DP.

The statement in the example above holds true in general. We first prove it for value differential privacy and then combine it with Proposition 3 to prove the general result.

Proposition 10. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be an ι-LI mechanism. Then, $\mathcal{M}$ is $(0, \iota)$-VDP.*

Proof. Let $d_1, d_2 \in \mathcal{D}$ be two neighboring databases and P be the joint probability matrix given by $P_{ij} = \Pr[\mathcal{M}(d_1) = v_i, \mathcal{M}(d_2) = v_j]$. Then, for every $k \in [|\mathcal{V}|]$,

$$\iota \geq \Pr[\mathcal{M}(d_1) \neq \mathcal{M}(d_2)] = \sum_{(i,j):i \neq j} P_{ij}$$

$$\geq \sum_{i \neq k}(P_{ki} + P_{ik}) \geq \sum_{i \neq k}|P_{ki} - P_{ik}|$$

$$\geq \left| \sum_{i \neq k}(P_{ki} - P_{ik}) \right| = \left| \sum_{j} P_{kj} - \sum_{i} P_{ik} \right|.$$

But $\sum_j P_{kj} = \Pr[\mathcal{M}(d_1) = v_k]$ and $\sum_i P_{ik} = \Pr[\mathcal{M}(d_2) = v_k]$. Thus,

$$\Pr[\mathcal{M}(d_1) = v_k] \leq \Pr[\mathcal{M}(d_2) = v_k] + \iota \quad \text{and} \quad \Pr[\mathcal{M}(d_2) = v_k] \leq \Pr[\mathcal{M}(d_1) = v_k] + \iota,$$

for every $k \in \mathcal{V}$, i.e., $\mathcal{M}$ is $(0, \iota)$-value differentially private. $\qquad\square$

Theorem 2. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be an ι-LI mechanism. Then, $\mathcal{M}$ is $(0, \iota(|\mathcal{V}| - 1))$-DP.*

$$
\begin{array}{cccc}
P_{11} & P_{12} & P_{13} & P_{14} \\
P_{21} & P_{22} & P_{23} & P_{24} \\
P_{31} & P_{32} & P_{33} & P_{34} \\
P_{41} & P_{42} & P_{43} & P_{44}
\end{array}
\qquad\qquad
\begin{array}{cccc}
P_{11} & P_{12} & P_{13} & P_{14} \\
P_{21} & P_{22} & P_{23} & P_{24} \\
P_{31} & P_{32} & P_{33} & P_{34} \\
P_{41} & P_{42} & P_{43} & P_{44}
\end{array}
$$

(a) $(0,\delta)$-VDP (b) ι-Low Influence

Fig. 7: The main idea behind the proof of Proposition 10. The elements of the matrix signify $P_{ij} = \Pr[\mathcal{M}(d_1) = v_i, \mathcal{M}(d_2) = v_j]$. In Figure (a), $(0,\delta)$-VDP requires the difference between the sum of row k and column k to be small. In Figure (b), ι-LI requires the sum of all non-diagonal elements to be small. The latter implies the former. However, one can easily construct counterexample to show the reverse is generally not true.

Proof. Follows directly from Propositions 3 and 10. □

The idea behind Proposition 10 is shown in Fig. 7. The ι-LI conditions require that for any two $d_1 \sim d_2$, the sum of non-diagonal elements of the joint probability matrix $P_{ij} = \Pr[\mathcal{M}(d_1) = v_i, \mathcal{M}(d_2) = v_j]$ be smaller than ι. On the other hand, $(0,\iota)$-VDP conditions require that for all $k \in [|\mathcal{V}|]$, the difference between the sum of row k and the sum of column k of the same matrix be small. Clearly, if the sum of all non-diagonal elements is small, the difference between the sum of any two subsets of non-diagonal elements is also small. For value differential privacy, it suffices to consider the difference between the sum of row k and column k, noting that the diagonal element P_{kk} (which could be large) will cancel out in the difference.

In most practical (ϵ,δ)-DP scenarios, the value of δ should be rather small, e.g., in the order of 10^{-3}. According to Theorem 10, setting $\iota = \frac{\delta}{|\mathcal{V}|-1}$ when designing a low influence mechanism will ensure the desired δ. However, this method is only valid for (ϵ,δ)-DP with $\delta > 0$. Indeed, if one desires $\delta = 0$, i.e. the classical ϵ-DP condition, a mechanism being low-influence is not sufficient to guarantee privacy unless $\iota = 0$.

V. USEFUL LOW INFLUENCE DIFFERENTIALLY PRIVATE SCHEMES

Most problems in differential privacy exhibit a tradeoff between privacy and some notion of utility. This is typically captured by a utility function $U : \mathfrak{M} \to \mathbb{R}$, which is a function from the space of randomized mechanisms to the real-values which usually captures how well a randomized mechanism approximates a desired output if no privacy mechanism was to be put in place. Problems in differential privacy can then be formulated as: $\underset{\mathcal{M} \in \mathfrak{M}}{\text{maximize }} U(\mathcal{M})$, subject to differential privacy constraint, and potentially other practical constraints. Obviously, the structure of the optimization problem, its computational complexity, and the overall efficacy of the solution are three important considerations when designing differentially private mechanisms.

As shown in Proposition 8, every linear constraint on the space of independent mechanisms $\mathfrak{I}$ is also linear on the space of joint mechanisms $\mathfrak{M}$. In particular, (ϵ,δ)-DP and (ϵ,δ)-VDP are linear constraints over both $\mathfrak{I}$ and $\mathfrak{M}$. However, the opposite is not true. As we have seen, low influence constraints are linear over $\mathfrak{M}$, but non-linear over $\mathfrak{I}$. Linear conditions are desired as efficient linear program solvers can be used to find the solution.

This simplicity in the constraints comes at the price of dimensionality. The degrees of freedom in independent mechanisms is $\mathcal{O}(|\mathcal{D}| \times |\mathcal{V}|)$, whereas the degrees of freedom in general joint mechanisms is $\mathcal{O}(|\mathcal{V}|^{|\mathcal{D}|})$. Clearly, the desired linearity and increase in the degrees of freedom come at the cost of computational complexity. Solving the tradeoff between performance and

complexity in designing joint mechanisms is an interesting topic of research, but is beyond the scope of this paper. However, in the remainder of this section we motivate the potential benefits of considering general joint mechanisms. We start by a general result on the existence of nontrivial mechanisms with arbitrarily low influence (which according to Theorem 1 must necessarily be non-independent).

Theorem 3. *There exists nontrivial mechanisms $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ with arbitrarily low influence.*

Proof. Without loss of generality, let $\mathcal{D} = \{d_1, \ldots, d_{|\mathcal{D}|}\}$ such that $d_1 \sim d_2$, $\mathcal{V} = \{1, \ldots, |\mathcal{V}|\}$, and choose $0 < \alpha < 1$. Consider the mechanism $P_{v_1,\ldots,v_{|\mathcal{D}|}} \in \mathfrak{M}$ such that $P_{1,\ldots,1} = P_{2,\ldots,2} = \frac{1-\alpha}{2}$, $P_{1,2,1,\ldots,1} = \alpha$, and $P_{v_1,\ldots,v_{|\mathcal{D}|}} = 0$ for the other $(v_1, \ldots, v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}$. Then,

$$\Pr[\mathcal{M}(d_1) = 1] = \sum_{(v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|} : v_1 = 1} P_{v_1,\ldots,v_{|\mathcal{D}|}} = P_{1,\ldots,1} + P_{1,2,1,\ldots,1} = \frac{1}{2} + \frac{\alpha}{2}$$

and

$$\Pr[\mathcal{M}(d_2) = 2] = \sum_{(v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|} : v_2 = 2} P_{v_1,\ldots,v_{|\mathcal{D}|}} = P_{2,\ldots,2} + P_{1,2,1,\ldots,1} = \frac{1}{2} + \frac{\alpha}{2}.$$

Since both values are larger than $1/2$, it follows that

$$\arg\max_{v \in \mathcal{V}} \Pr[\mathcal{M}(d_1) = v] = 1 \neq 2 = \arg\max_{v \in \mathcal{V}} \Pr[\mathcal{M}(d_2) = v].$$

And thus, $\mathcal{M}$ is nontrivial.

We now show that $\mathcal{M}$ is α-LI. This follows directly from the fact that

$$\sum_{(v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|} : v_i \neq v_k} P_{v_1,\ldots,v_{|\mathcal{D}|}} = \begin{cases} P_{1,2,1,\ldots,1} = \alpha & \text{if } v_i = 2 \text{ or } v_k = 2, \\ 0 & \text{otherwise.} \end{cases}$$

$\square$

Example 5. Let $\mathcal{D}$ consist of two neighboring datasets $d_1 \sim d_2$ and the output set be the binary set $\mathcal{V} = \{1, 2\}$. We wish to design a mechanism which does not favor one database over another. This can be achieved through imposing $\Pr[\mathcal{M}(d_1) = 1] = \Pr[\mathcal{M}(d_2) = 2]$, which we call the balancing constraint. Consider the utility function given by $U(\mathcal{M}) = \Pr[\mathcal{M}(d_1) = 1]$. Restricting ourselves to independent mechanisms we consider the following optimization problem.

$$\begin{aligned} \underset{\mathcal{M} \in \mathfrak{I}}{\text{maximize}} \quad & U(\mathcal{M}) \\ \text{subject to} \quad & \mathcal{M} \text{ satisfies } (\epsilon, \delta)\text{-DP and} \\ & \Pr[\mathcal{M}(d_1) = 1] = \Pr[\mathcal{M}(d_2) = 2]. \end{aligned}$$

Using the characterization of $\mathfrak{I}$ given in Proposition 3, and setting $x = \Pr[\mathcal{M}(d_1) = 1] = \Pr[\mathcal{M}(d_2) = 2]$, the optimization problem is as follows.

$$\begin{aligned} \underset{x \in \mathbb{R}}{\text{maximize}} \quad & x \\ \text{subject to} \quad & x \leq e^\epsilon (1 - x) + \delta \\ & (1 - x) \leq e^\epsilon x + \delta \\ & 0 \leq x \leq 1. \end{aligned}$$

The solution to this optimization problem is $x^* = \frac{e^\epsilon + \delta}{e^\epsilon + 1} \geq \frac{1}{2}$, confirming that the mechanism is nontrivial. As we previously saw, the influence of this mechanism is lower bounded by $\frac{1}{2}$. Indeed, the influence is given by

$$I(x^*) = 1 - \frac{2(1 - \delta)(e^\epsilon + \delta)}{(e^\epsilon + 1)^2} \geq \frac{1}{2}.$$

We now show that if we consider general mechanisms, we can lower the influence.

Example 6. Let $\mathcal{D}$ consist of two neighboring datasets $d_1 \sim d_2$ and the output set be the binary set $\mathcal{V} = \{1, 2\}$. We consider the utility function given by $U(\mathcal{M}) = \Pr[\mathcal{M}(d_1) = 1]$ with the balancing constraint that $\Pr[\mathcal{M}(d_1) = 1] = \Pr[\mathcal{M}(d_2) = 2]$ as in Example 5. We now consider general mechanisms.

$$\begin{aligned}
\underset{\mathcal{M} \in \mathfrak{M}}{\text{maximize}} \quad & U(\mathcal{M}) \\
\text{subject to} \quad & \mathcal{M} \text{ satisfies } (\epsilon, \delta)\text{-DP and} \\
& \Pr[\mathcal{M}(d_1) = 1] = \Pr[\mathcal{M}(d_2) = 2].
\end{aligned}$$

Using the characterization of $\mathfrak{M}$ given in Proposition 4, and setting $x = \Pr[\mathcal{M}(d_1) = 1, \mathcal{M}(d_2) = 2]$ and $y = \Pr[\mathcal{M}(d_1) = 2, \mathcal{M}(d_2) = 1]$, the optimization problem is as follows.

$$\begin{aligned}
\underset{(x,y) \in \mathbb{R}^2}{\text{maximize}} \quad & \frac{1 + x - y}{2} \\
\text{subject to} \quad & \frac{1 + x - y}{2} \leq e^\epsilon \frac{1 + y - x}{2} + \delta \\
& \frac{1 + y - x}{2} \leq e^\epsilon \frac{1 + x - y}{2} + \delta \\
& 0 \leq x \leq 1 \\
& 0 \leq y \leq 1 \\
& 0 \leq x + y \leq 1.
\end{aligned}$$

If we also minimize the influence, given by $I(x, y) = x + y$, we get the solution $x^* = \frac{e^\epsilon + 2\delta - 1}{e^\epsilon + 1}$ and $y^* = 0$ with influence

$$I(x^*, y^*) = \frac{e^\epsilon + 2\delta - 1}{e^\epsilon + 1} \leq 1 - \frac{2(1 - \delta)(e^\epsilon + \delta)}{(e^\epsilon + 1)^2}.$$

It can be verified that the above x^* and y^* ensure nontrivial mechanisms for all $\epsilon \geq 0$ and $0 \leq \delta \leq 1$. Note that for $\epsilon = 0$ and $\delta = 0$, $I(x, y) = 0$. Thus, in this case, the influence is not lower bounded.

In Figure 8 we show the optimal tradeoffs between $(\epsilon, 0)$-DP, the utility $U(\mathcal{M})$, and the influence ι of joint, and independent randomized mechanisms. Note that this tradeoff can be obtained by letting $\delta = 0$ in the solutions in Examples 5 and 6, and by then varying the value of ϵ. As can be seen in the figure, both independent and joint mechanisms exhibit the same privacy-utility tradeoff, i.e. for any ϵ, there is an $(\epsilon, 0)$-DP independent mechanism that achieves the same utility as the best $(\epsilon, 0)$-DP joint mechanism. On the other hand, joint mechanisms achieve a strictly better optimal tradeoff in terms of their influence ι. That is, while independent mechanisms are sufficient to achieve the optimal privacy (captured by the parameter ϵ in the $(\epsilon, 0)$-DP guarantee) and utility tradeoff, they fall short at providing the best influence level ι.
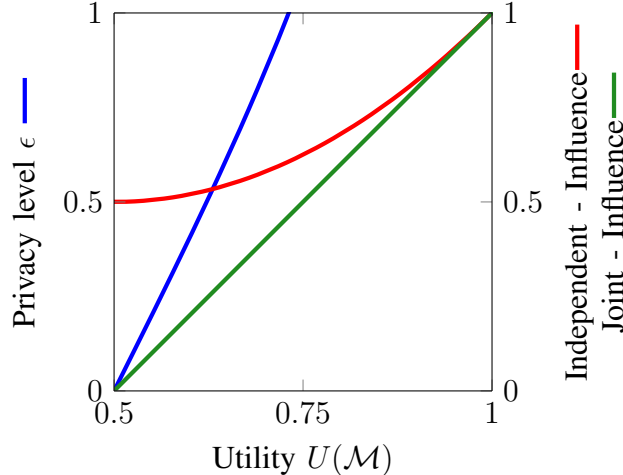
Fig. 8: The privacy-utility-influence tradeoff for joint and independent mechanisms. While independent mechanisms are sufficient to achieve the best privacy-utility tradeoff, they fall short in terms of influence – for any utility level, there is a joint scheme which achieves the same privacy and a lower influence.

VI. CONCLUDING REMARKS

The study in this paper reveals the interesting interactions between differential privacy and utility, on the one hand, and mechanism influence and independence on the other. In the independent differentially private mechanisms prevalent in the literature, the additive noise is independent across input databases. Such independence enables relatively easy design of the mechanism in terms of the privacy-utility tradeoff. While such mechanisms are casually understood to offer "low influence" of an individual on the query output, we show that they are not necessarily statistically low influence in a formal sense. Even further, we show that if one were to enforce an arbitrarily low influence on an independent differentially private mechanism, a heavy price would have to be paid: the mechanism becomes trivial (statistically-speaking constant) and hence loses any practical utility. The more general class of joint mechanisms not only allows differential privacy, low influence, and utility to be jointly achieved, they also make the optimization of influence linear over the mechanism space. However, this increases the dimension of the mechanism space over which such optimization takes place.

There are several interesting directions for future research. In [22], the authors studied the relation between individual *identifiability* and differential privacy. Roughly speaking, identifiability aims to measures the guessing capability of an adversary who knows all elements of the dataset d that are common with its neighboring datasets and wants to guess the remaining unknown element of d after observing the mechanism output $\mathcal{M}(d)$. Identifiability level, denoted by ϵ_d, of a randomized mechanism depends on the input distribution of elements in d and the ϵ level in differential privacy. In [23], the authors studied the *worst-case* mutual information (MI) leakage between the i-th entry of dataset d and the randomized mechanism output $\mathcal{M}(d)$, given the knowledge of dataset entries excluding i and when maximized over all possible distributions of dataset entries. They proved that such defined ϵ_m-MI privacy level is weaker than $(\epsilon, 0)$-DP, but does imply $(\epsilon, \sqrt{2\epsilon_m})$-DP for any $\epsilon, \epsilon_m \geq 0$. Exploring the relation between low influence, ϵ_d-identifiability and ϵ_m-MI privacy is an interesting research direction.

Another interesting direction for future research will be systematic low-complexity design of non-independent, high-utility, low-influence, and differentially-private mechanisms. In addition, it is worthwhile to formalize the counterpart of our results for continuous-valued queries and mechanisms. Finally, we remark that the general framework of input-dependent joint mechanisms that we studied in this paper may be of value in time-varying systems where data changes over time according to an inherent dynamical model or is gradually acquired and released (e.g., in social networks or health systems), see for example [24], [25]. We propose further investigation of a formal connection as future research.

APPENDIX

In this Appendix we prove Propositions 1, 2, and Theorem 1 of the main text.

Proposition 1. *It holds that,*

$$
\mathfrak{I} \cong \left\{ P_{v_1,\ldots,v_{|\mathcal{D}|}} \in \mathfrak{M} : P_{v_1,\ldots,v_{|\mathcal{D}|}} = \prod_{i=1}^{|\mathcal{D}|} \sum_{(w_1,\ldots,w_{|\mathcal{D}|})\in\mathcal{V}^{|\mathcal{D}|}:w_i=v_i} P_{w_1,\ldots,w_{|\mathcal{D}|}}, \right.
$$
$$
\left. \textit{for every } (v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|} \right\}.
$$

In other words, $\mathfrak{I}$ is a polynomial of degree $|\mathcal{D}|$ in $\mathbb{R}^{|\mathcal{V}|^{|\mathcal{D}|}}$.

Proof. Let $P \in \mathfrak{M}$ where $P_{v_1,\ldots,v_{|\mathcal{D}|}} = \Pr[\mathcal{M}(d_1) = v_1, \ldots \mathcal{M}(d_{|\mathcal{D}|}) = v_{|\mathcal{D}|}]$. Then,

$$
\Pr[\mathcal{M}(d_k) = v] = \sum_{(w_1,\ldots,w_{|\mathcal{D}|})\in\mathcal{V}^{|\mathcal{D}|}:w_k=v} P_{w_1,\ldots,w_{|\mathcal{D}|}}.
$$

Thus, the mechanism is independent if and only if

$$
P_{v_1,\ldots,v_{|\mathcal{D}|}} = \prod_{i=1}^{|\mathcal{D}|} \sum_{(w_1,\ldots,w_{|\mathcal{D}|})\in\mathcal{V}^{|\mathcal{D}|}:w_i=v_i} P_{w_1,\ldots,w_{|\mathcal{D}|}}, \text{for every } (v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}.
$$

$\square$

Proposition 2. *A randomized mechanism is equivalent to a probability distribution on the space of all functions $h : \mathcal{D} \to \mathcal{V}$.*

Proof. The space of functions from $h : \mathcal{D} \to \mathcal{V}$ is isomorphic to the space $\mathcal{V}^{|\mathcal{D}|}$. Indeed, if we order $\mathcal{D} = \{d_1,\ldots,d_{|\mathcal{D}|}\}$, then each vector $(v_1,\ldots,v_{|\mathcal{D}|}) \in \mathcal{V}^{|\mathcal{D}|}$ corresponds to the function $h : \mathcal{D} \to \mathcal{V}$ such that $h(d_1) = v_1,\ldots,h(d_{|\mathcal{D}|}) = v_{|\mathcal{D}|}$. By Equation 1 in the main text, every $P \in \mathfrak{M}$ is a probability distribution on $\mathcal{V}^{|\mathcal{D}|}$. Thus, using the above equivalence, every mechanism $P \in \mathfrak{M}$ is a probability distribution on the space of functions from $h : \mathcal{D} \to \mathcal{V}$. $\square$

We now prove Theorem 1. To do this, we present some preliminary definitions and lemmas. First, recall Definition 9 of the main text $I(u,v) = 1 - \sum_{i=1}^{\mathcal{V}} u_i v_i$, which can be used to measure the local influence of any two datasets. We relate this function to ι-LI for independent mechanisms in the following lemma.

Lemma 1. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be an independent mechanism and denote $M_{ij} = \Pr[\mathcal{M}(d_i) = v_j]$. Then, $\mathcal{M}$ is ι-LI if and only if $I(M_i, M_k) \leq \iota$ for every $d_i \sim d_k$.*

Proof. This is a direct restatement of Proposition 5 in the main text. $\qquad \square$

It follows from the lemma that a lower bound on I for two neighboring datasets gives us a lower bound on ι.

We now define generalized versions of the sets R^* and S in Example 1 of the main text.

Definition 9. Let $\mathcal{V}$ be a finite set. We define

$$R = \left\{ (u, v) \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|} : \max_{i \in [|\mathcal{V}|]} u_i = u_1, \max_{i \in [|\mathcal{V}|]} v_i = v_2, \sum_{i=1}^{|\mathcal{V}|} u_i = \sum_{i=1}^{|\mathcal{V}|} v_i = 1, \right.$$
$$\left. 0 \le u_i \le 1, 0 \le v_i \le 1, \quad \text{for every } i \in [|\mathcal{V}|] \right\}.$$

This set corresponds to nontrivial independent mechanisms in the following way.

Lemma 2. *Let $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ be a nontrivial independent mechanism and $M_{ij} = \Pr[M(d_i) = v_j]$. Then, there exist two neighboring datasets $d_i \sim d_j$ such that $(M_i, M_j) \in R$.*

Proof. Since $\mathcal{M}$ is nontrivial, there exist two neighboring datasets $d_i \sim d_j$ such that

$$\arg \max_{v \in \mathcal{V}} \Pr[M(d_i) = v] \ne \arg \max_{v \in \mathcal{V}} \Pr[M(d_j) = v].$$

We order $\mathcal{V} = \{v_1, \ldots, v_{|V|}\}$ such that

$$v_1 = \arg \max_{v \in \mathcal{V}} \Pr[M(d_i) = v], \quad v_2 = \arg \max_{v \in \mathcal{V}} \Pr[M(d_j) = v].$$

Then, by construction, $(M_i, M_j) \in R$. $\qquad \square$

Thus, a lower bound on I subject to R would give us a lower bound on ι for all nontrivial independent mechanisms. To do this, we will consider the following set.

Definition 10. Let R be as in Definition 9. We define

$$S = \{(u, v) \in R : u_i = v_i = 0 \quad \text{for every } i \ge 3\}.$$

We show in Lemma 4 that the minimum of I over R is in S. To do this we need the following construction.

Definition 11. Let $(u, v) \in R$, $S_u = \sum_{i=3}^{|\mathcal{V}|} u_i$, $S_v = \sum_{i=3}^{|\mathcal{V}|} v_i$. We define $u^* \in \mathbb{R}^{|\mathcal{V}|}$ such that $u_1^* = u_1 + \frac{S_u}{2}$, $u_2^* = u_2 + \frac{S_u}{2}$, and $u_i^* = 0$ for $i \in |\mathcal{V}| - \{1, 2\}$. We also define $v^* \in \mathbb{R}^{|\mathcal{V}|}$ such that $v_1^* = v_1 + \frac{S_v}{2}$, $v_2^* = v_2 + \frac{S_v}{2}$, and $v_i^* = 0$ for $i \in |\mathcal{V}| - \{1, 2\}$.

Thus, for any $(u, v) \in R$ we can construct the point (u^*, v^*).

Lemma 3. *Let $(u, v) \in R$. Then, $(u^*, v^*) \in S \subseteq R$.*

Proof. Since $u_1 \ge u_2$, it follows that

$$u_1^* = u_1 + \frac{S_u}{2} \ge u_2 + \frac{S_u}{2} = u_2^* \ge u_i = 0 \quad \text{for every } i \in |\mathcal{V}| - \{1, 2\}.$$

Also, $\sum_{i=1}^{|\mathcal{V}|} u_i = u_1 + \frac{S_u}{2} + u_2 + \frac{S_u}{2} = 1$. The same arguments hold for v^*. Thus, $(u^*, v^*) \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$. $\qquad \square$

We now show that the minimum of I over R is in S.

Lemma 4. *Let* $(u, v) \in R$. *Then,* $I(u^*, v^*) \leq I(u, v)$. *And therefore,*

$$\min_{(u,v) \in R} I(u, v) = \min_{(u,v) \in S} I(u, v).$$

Proof. Since,

$$\sum_{i=3}^{|\mathcal{V}|} u_1 v_i \geq \sum_{i=3}^{|\mathcal{V}|} u_i v_i \quad \text{and} \quad \sum_{i=3}^{|\mathcal{V}|} u_i v_2 \geq \sum_{i=3}^{|\mathcal{V}|} u_i v_i,$$

it follows that,

$$I(u^*, v^*) = 1 - \sum_{i=1}^{|\mathcal{V}|} u_i^* v_i^* = 1 - \left(u_1 + \frac{S_u}{2}\right)\left(v_1 + \frac{S_v}{2}\right) - \left(u_2 + \frac{S_u}{2}\right)\left(v_2 + \frac{S_v}{2}\right)$$

$$= 1 - u_1 v_1 - u_1 \frac{S_v}{2} - v_1 \frac{S_u}{2} - \frac{S_u}{2}\frac{S_v}{2} - u_2 v_2 - u_2 \frac{S_v}{2} - v_2 \frac{S_u}{2} - \frac{S_u}{2}\frac{S_v}{2}$$

$$\leq 1 - u_1 v_1 - u_2 v_2 - u_1 \frac{S_v}{2} - v_2 \frac{S_u}{2} = 1 - u_1 v_1 - u_2 v_2 - \frac{1}{2}\sum_{i=3}^{|\mathcal{V}|} u_1 v_i - \frac{1}{2}\sum_{i=3}^{|\mathcal{V}|} u_i v_2$$

$$\leq 1 - \sum_{i=1}^{|\mathcal{V}|} u_i v_i = I(u, v).$$

$\square$

To conclude our proof we must minimize I over S.

Lemma 5. *Let* $(u, v) \in S$. *Then,* $\min_{(u,v) \in S} I(u, v) = \frac{1}{2}$

Proof. Since $u_2^* = 1 - u_1^*$, $v_2^* = 1 - v_1^*$, minimizing I over R is equivalent to minimizing over the same two-dimensional space defined in (2) of Example 1 in the main text,

$$R^* = \left\{(u, v) \in \mathbb{R}^2 : \frac{1}{2} \leq u \leq 1 \quad \text{and} \quad 0 \leq v \leq \frac{1}{2}\right\}.$$

Thus, as shown in the example, $\min_{(u,v) \in R^*} I(u, v) = \frac{1}{2}$.

$\square$

We now prove Theorem 1 of the main text.

Theorem 1. *Let* $\mathcal{M} : \mathcal{D} \to \mathcal{V}$ *be an independent,* ι*-LI, nontrivial mechanism. Then,* $\iota \geq \frac{1}{2}$. *Furthermore, this bound is tight.*

Proof. By Lemma 2, there exists two neighboring datasets $d_i \sim d_k$ such that $(M_i, M_j) \in R$. Lemmas 4 and 5 imply that $I(M_i, M_j) \geq \frac{1}{2}$. But then, by Lemma 1, $\iota \geq \frac{1}{2}$.

To see that the bound is tight, assume $d_1 \sim d_2$ and set $M_{11} = M_{12} = \frac{1}{2}$, $M_{1j} = 0$ for $j > 2$, $M_{i2} = 1$, and $M_{ij} = 0$ for all the other i and j. Then,

$$I(M_i, M_k) = \begin{cases} \frac{1}{2} & \text{if } i = 1 \text{ or } k = 1, \\ 0 & \text{otherwise.} \end{cases}$$

$\square$

ACKNOWLEDGEMENT

REFERENCES

[1] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography*, S. Halevi and T. Rabin, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 265–284.

[2] C. Dwork, "Differential privacy: A survey of results," in *International Conference on Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.

[3] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[4] D. P. Team, *Learning with Privacy at Scale*, 2017 (last accessed May 2020). [Online]. Available: https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html

[5] *Disclosure Avoidance and the 2020 Census*, 2020 (last accessed May 2020). [Online]. Available: https://www.census.gov/about/policies/privacy/statistical_safeguards/disclosure-avoidance-2020-census.html

[6] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *Annual Symposium on Foundations of Computer Science*, 2010, pp. 51–60.

[7] Z. Ding, Y. Wang, G. Wang, D. Zhang, and D. Kifer, "Detecting violations of differential privacy," in *ACM SIGSAC Conference on Computer and Communications*, 2018, pp. 475–489.

[8] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and N. Moni, "Our data, ourselves: Privacy via distributed noise generation," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2006, pp. 486–503.

[9] L. S. Penrose, "The elementary statistics of majority voting," *Journal of the Royal Statistical Society*, vol. 109, no. 1, pp. 53–57, 1946.

[10] R. O'Donnell, *Analysis of Boolean functions*. Cambridge University Press, 2014.

[11] K. Nissim, S. Raskhodnikova, and A. Smith, "Smooth sensitivity and sampling in private data analysis," in *ACM Symp. Theory Comput. (STOC)*, 2007, p. 75–84.

[12] Q. Geng and P. Viswanath, "The optimal noise-adding mechanism in differential privacy," *IEEE Trans. on Inf. Theory*, vol. 62, no. 2, pp. 925–951, 2016.

[13] J. Bailie and C.-H. Chien, "ABS perturbation methodology through the lens of differential privacy," 2019. [Online]. Available: http://www.unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2019/mtg1/SDC2019_S2_ABS_Bailie_D.pdf

[14] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility- maximizing privacy mechanisms," in *Proceedings of the forty-first annual ACM symposium on Theory of computing (STOC)*, 2009, p. 351–360. [Online]. Available: http://doi.acm.org/10.1145/1536414.1536464

[15] Q. Geng and P. Viswanath, "Differential privacy on finite computers," *Journal of Privacy and Confidentiality*, vol. 9, no. 2, pp. 1–46, 2019.

[16] M. Gupte and M. Sundararajan, "Universally optimal privacy mechanisms for minimax agents," in *ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, NY, USA, 2010, p. 135–146.

[17] C. Canonne, G. Kamath, and T. Steinke. (2020) The discrete gaussian for differential privacy. [Online]. Available: https://arxiv.org/abs/2004.00010v2

[18] G. Wezerek and D. Van Riper, *Changes to the Census Could Make Small Towns Disappear*, 2020 (last accessed May 2020). [Online]. Available: https://www.nytimes.com/interactive/2020/02/06/opinion/census-algorithm-privacy.html

[19] G. Cormode, T. Kulkarni, and D. Srivastava. (2017) Constrained differential privacy for count data. [Online]. Available: https://arxiv.org/abs/1710.00608

[20] P. Sadeghi, S. Asoodeh, and F. du Pin Calmon. (2020) Differentially private mechanisms for count queries. [Online]. Available: https://arXiv:2007.09374v1

[21] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. L. Roth, "Preserving statistical validity in adaptive data analysis," in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 2015, pp. 117–126.

[22] W. Wang, L. Ying, and J. Zhang, "On the relation between identifiability, differential privacy, and mutual-information privacy," *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5018–5029, 2016.

[23] P. Cuff and L. Yu, "Differential privacy as a mutual information constraint," in *CSS*, 2016, pp. 43–54.

[24] F. Koufogiannis and G. J. Pappas, "Differential privacy for dynamical sensitive data," in *Conference on Decision and Control (CDC)*, 2017, pp. 1118–1125.

[25] F. Koufogiannis, S. Han, and G. J. Pappas, "Gradual release of sensitive data under differential privacy," *Journal of Privacy and Confidentiality*, vol. 7, no. 2, Jan. 2017. [Online]. Available: https://journalprivacyconfidentiality.org/index.php/jpc/article/view/649