

2019

Hypothesis testing interpretations and Renyi differential privacy

B. Balle, G. Barthe, M. Gaboardi, J. Hsu, T. Sato. "Hypothesis Testing Interpretations and Renyi Differential Privacy." Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, PMLR 108:2496-2506, 2020.. International Conference on Artificial Intelligence and Statistics <https://arxiv.org/abs/1905.09982>

<https://hdl.handle.net/2144/43237>

"Downloaded from OpenBU. Boston University's institutional repository."

Hypothesis Testing Interpretations and Rényi Differential Privacy (Long version)

Borja Balle Gilles Barthe Marco Gaboardi Justin Hsu
Tetsuya Sato

October 9, 2019

Abstract

Differential privacy is a de facto standard in data privacy, with applications in the public and private sectors. A way to explain differential privacy, which is particularly appealing to statistician and social scientists, is by means of its statistical hypothesis testing interpretation. Informally, one cannot effectively test whether a specific individual has contributed her data by observing the output of a private mechanism—any test cannot have both high significance and high power.

In this paper, we identify some conditions under which a privacy definition given in terms of a statistical divergence satisfies a similar interpretation. These conditions are useful to analyze the distinguishability power of divergences and we use them to study the hypothesis testing interpretation of some relaxations of differential privacy based on *Rényi divergence*. This analysis also results in an improved conversion rule between these definitions and differential privacy.

1 Introduction

Differential privacy [Dwork et al., 2006] is a formal notion of data privacy that enables accurate statistical analyses on populations while preserving privacy for individuals contributing their data. Differential privacy is supported by a rich theory, with sophisticated algorithms for common data analysis tasks and composition theorems to simplify the design and formal analysis of new private algorithms. This theory has helped make differential privacy a de facto standard for privacy-preserving data analysis. Over the last years, differential privacy has found use in the private sector [Kenthapadi et al., 2019] by companies such as Google [Erlingsson et al., 2014, Papernot et al., 2018], Apple [team at Apple, 2017], and Uber [Johnson et al., 2018], and in the public sector by agencies such as the U.S. Census Bureau [Abowd, 2018, Garfinkel et al., 2018].

A common challenge faced across all uses of differential privacy is to explain its guarantees to users and policy makers. Indeed, differential privacy first emerged in the theoretical computer science community, and was only subsequently considered in other research areas interested in data privacy. For this reason, several works have attempted to provide different interpretations of the *semantics* of differential privacy in an effort to make it more accessible.

One approach that has been particularly successful, especially when introducing differential privacy to people versed in statistical data analysis, is the *hypothesis testing* interpretation of differential privacy [Wasserman and Zhou, 2010, Kairouz et al., 2015]. One can imagine an experiment where one wants to test, based on the output of a differentially private mechanism, the *null hypothesis* that an individual I has contributed her data to a particular dataset x_0 . One can also imagine that an *alternative hypothesis* is that the individual I has not contributed her data. Then, the definition of differential privacy guarantees—and is in fact equivalent to requiring—that every hypothesis test has either low *significance* (it has a high rate of Type I errors), or low *power* (it has a high rate of Type II errors). Under this interpretation, the privacy parameters (ϵ, δ) control the tradeoff between significance and power.

Recently, several variants of differential privacy have been proposed [Dwork and Rothblum, 2016, Bun and Steinke, 2016, Mironov, 2017, Bun et al., 2018, Dong et al., 2019]. Most of these new privacy definitions have been proposed as privacy notions with better composition properties than differential privacy. Having better composition can become a key advantage when a high number of data accesses is needed for a single analysis (e.g., in private deep learning [Abadi et al., 2016]). Technically, many of these variants are formulated as bounds on the *Rényi divergence* between the distribution obtained when running a private mechanism over a dataset where an individual I has contributed her data versus the case when the private mechanism is run over the dataset where I 's data is removed.

In this work we develop some analytical tools to study the hypothesis testing interpretation of privacy definitions based on statistical divergences. The first notion we introduce is the concept of *k-cut* of a divergence. Intuitively this notion corresponds to restricting the distributions that are input to a divergence to a finite domain of cardinality k . We can think about the functions implementing the restrictions as (probabilistic) decision rules with k possible outcomes. The second notion we introduce is the concept of *k-generatedness* for a divergence. Intuitively, a divergence is *k-generated* if it is equal to its *k-cut*. This notion expresses the number of decisions that are needed in decision rules to fully characterize the divergence.

We use these two analytical tools to show that a privacy definition based on a divergence has an hypothesis testing interpretation if and only if it is 2-generated. We show that the divergence characterizing differential privacy is indeed 2-generated and that the notion of 2-generatedness corresponds to the notion of privacy regions introduced in [Kairouz et al., 2015]. On the negative side we show formally that variants of differential privacy based on the Rényi divergence do not admit directly a hypothesis testing interpretation because the Rényi divergence is exactly ∞ -generated (where by ∞ we mean that it is infinitely, but countably generated). Nevertheless, we show that one can achieve the hypothesis testing interpretation by considering the 2-cut of the Rényi divergence. Intuitively, this says that to characterize relaxations of differential privacy based on the Rényi divergence through an experiment similar to the one used in the hypothesis testing interpretation, one needs either to restrict the distinguishability power of the divergence or to consider an infinite number of possible hypothesis. This shows a semantics separation between standard differential privacy and relaxations based on Rényi divergence.

In addition, we use the analytical tools we develop to study the relations between different privacy definitions. Specifically, we use the 2-cut of Rényi divergence to give better conversion rules from Rényi differential privacy to (ϵ, δ) -differential privacy, and to study the relations with *Gaussian Differential Privacy* [Dong et al., 2019] another formal definition of privacy inspired by the hypothesis testing interpretation which was

recently proposed.

Finally, we study a sufficient condition to guarantee that a divergence is k -generated: divergences defined as a supremum of a quasi-convex function F over probabilities of k -partitions are k -generated. This allows one to construct divergences supporting the hypothesis testing interpretation by requiring them to be defined through an function F giving a 2-generated divergence. The condition is also necessary for quasi-convex divergences, characterizing k -generation for all quasi-convex divergences.

Summarizing, our contributions are:

- (1) We first introduce the notions of k -cut and k -generatedness for divergences. These notions allow one to measure the power of divergences in terms of the number of possible decisions that are needed in a test to fully characterize the divergence.
- (2) We show that the divergence used to characterize differential privacy is 2-generated, supporting the usual hypothesis testing interpretation of differential privacy
- (3) We show that Rényi divergence is ∞ -generated, ruling out a direct hypothesis testing interpretation for privacy notions based on it. Nevertheless, we show that one can achieve the hypothesis testing interpretation by considering the 2-cut of Rényi divergence.
- (4) We use our analytic tools to study other notions of privacy and to give better conversion rules between Rényi differential privacy and (ϵ, δ) -differential privacy.
- (5) We give sufficient and necessary conditions for a quasi-convex divergence to be k -generated.

2 Background: hypothesis testing, privacy, and Rényi divergences

2.1 Hypothesis testing interpretation for (ϵ, δ) -differential privacy

We view *randomized algorithms* as functions $\mathcal{M}: X \rightarrow \text{Prob}(Y)$ from a set X of inputs to the set $\text{Prob}(Y)$ of discrete *probability distributions* over a set Y of outputs. We assume that X is equipped with a symmetric *adjacency relation*—informally, inputs are datasets and two inputs x_0 and x_1 are adjacent iff they differ in the data of a single individual.

Definition 1 (Differential Privacy (DP) [Dwork et al., 2006]). *Let $\epsilon > 0$ and $0 \leq \delta \leq 1$. A randomized algorithm $\mathcal{M}: X \rightarrow \text{Prob}(Y)$ is (ϵ, δ) -differentially private if for every pairs of adjacent inputs x_0 and x_1 , and every subset $S \subseteq Y$, we have:*

$$\Pr[\mathcal{M}(x_0) \in S] \leq e^\epsilon \Pr[\mathcal{M}(x_1) \in S] + \delta.$$

[Wasserman and Zhou, 2010, Kairouz et al., 2015] proposed a useful interpretation of this guarantee in terms of *hypothesis testing*. Suppose that x_0 and x_1 are adjacent inputs. The observer sees the output y of running a private mechanism \mathcal{M} on one of these inputs—but does not see the particular input—and wants to guess whether the input was x_0 or x_1 .

In the terminology of hypothesis testing, let $y \in Y$ be an output of a randomized mechanism \mathcal{M} , and take the following *null* and *alternative* hypotheses:

$$\mathbf{H0} : y \text{ came from } \mathcal{M}(x_0), \quad \mathbf{H1} : y \text{ came from } \mathcal{M}(x_1).$$

One simple way of deciding between the two hypotheses is to fix a *rejection region* $S \subseteq Y$; if the observation y is in S then the null hypothesis is rejected, and if the observation y is not in S then the null hypothesis is not rejected. This is an example of a *deterministic decision rule*.

Each decision rule can err in two possible ways. A *false alarm* (i.e. Type I error) is when the null hypothesis is true but rejected. This error rate is defined as $\text{PFA}(x_0, x_1, \mathcal{M}, S) \stackrel{\text{def}}{=} \Pr[\mathcal{M}(x_0) \in S]$. On the other hand, the decision rule may incorrectly fail to reject the null hypothesis, a *false negative* (i.e. Type II error). The probability of missed detection is defined as $\text{PMD}(x_0, x_1, \mathcal{M}, S) \stackrel{\text{def}}{=} \Pr[\mathcal{M}(x_1) \notin S]$. There is a natural tradeoff between these two errors—a rule with a larger rejection region will be less likely to incorrectly fail to reject but more likely to incorrectly reject, while a rule with a smaller rejection region will be less likely to incorrectly reject but more likely to incorrectly fail to reject.

Differential privacy can now be reformulated in terms of these error rates.

Theorem 2 ([Wasserman and Zhou, 2010, Kairouz et al., 2015]). *A randomized algorithm $\mathcal{M}: X \rightarrow \text{Prob}(Y)$ is (ϵ, δ) -differentially private if and only if for every pair of adjacent inputs x_0 and x_1 , and any rejection region $S \subseteq Y$, we have: $\text{PFA}(x_0, x_1, \mathcal{M}, S) + e^\epsilon \text{PMD}(x_0, x_1, \mathcal{M}, S) \geq 1 - \delta$ and $e^\epsilon \text{PFA}(x_0, x_1, \mathcal{M}, S) + \text{PMD}(x_0, x_1, \mathcal{M}, S) \geq 1 - \delta$.*

Intuitively, the lower bound on the sum of the two error rates means that *no* decision rule is capable of achieving low Type I error and low Type II error simultaneously. Thus, the output distributions from any two adjacent inputs are statistically hard to distinguish.

Following [Kairouz et al., 2015], we can also reformulate the definition of differential privacy in terms of a *privacy region* describing the attainable pairs of Type I and Type II errors.

Theorem 3 ([Kairouz et al., 2015]). *A randomized algorithm $\mathcal{M}: X \rightarrow \text{Prob}(Y)$ is (ϵ, δ) -differentially private if and only if for every pair of adjacent inputs x_0 and x_1 , and every rejection region $S \subseteq Y$, we have*

$$(\text{PFA}(x_0, x_1, \mathcal{M}, S), \text{PMD}(x_0, x_1, \mathcal{M}, S)) \in R(\epsilon, \delta),$$

where the privacy region $R(\epsilon, \delta)$ is defined as:

$$R(\epsilon, \delta) = \{ (x, y) \in [0, 1] \times [0, 1] \mid (1 - x) \leq e^\epsilon y + \delta \}.$$

Figure 1 shows an example of a privacy region (the area between the dashed lines) $R(0.67, 0.05)$ and its mirror image and of all the points (PFA, PMD) that can be generated by a randomized response mechanism $\mathcal{M}_{\text{RR}}: \{0, 1\}^3 \rightarrow \text{Prob}(\{0, 1\}^3)$ working on vectors of three bits and flipping each bit with probability 0.34.

Since the original introduction of differential privacy, researchers have proposed several other variants based on Rényi divergence. The central question of this paper is: can we give similar hypothesis testing interpretations to these (and other) variants of differential privacy?

2.2 Variants of differential privacy based on Rényi divergence

We recall here notions of differential privacy based on Rényi divergence.




Figure 1: Pairs (PFA, PMD) of \mathcal{M}_{RR} and $R(0.67, 0.05)$

Definition 4 (Rényi divergence [Renyi, 1961]). *Let $\alpha > 1$. The Rényi divergence of order α between two probability distributions μ_1 and μ_2 on a space X is defined by:*

$$D_X^\alpha(\mu_1 \parallel \mu_2) \stackrel{\text{def}}{=} \frac{1}{\alpha - 1} \log \sum_{x \in X} \mu_2(x) \left(\frac{\mu_1(x)}{\mu_2(x)} \right)^\alpha. \quad (1)$$

The above definition does not consider the cases $\alpha = 1$ and $\alpha = +\infty$. However we can see D_X^α as a function of α for fixed distributions and consider the limits to get:

$$\begin{aligned} D_X^1(\mu_1 \parallel \mu_2) &\stackrel{\text{def}}{=} \mathbf{KL}_X(\mu_1 \parallel \mu_2), \\ D_X^\infty(\mu_1 \parallel \mu_2) &\stackrel{\text{def}}{=} \log \sup_x \frac{\mu_1(x)}{\mu_2(x)}. \end{aligned}$$

The first limit is the well-known KL divergence, while the second limit is the *max divergence* that bounds the pointwise ratio of probabilities; standard $(\epsilon, 0)$ -differential privacy bounds this divergence on distributions from adjacent inputs.

There are several notions of differential privacy based on Rényi divergence, differing in whether the bound holds for all orders α or just some orders. The first notion we consider is Rényi Differential Privacy (RDP) [Mironov, 2017].

Definition 5 (Rényi Differential Privacy (RDP) [Mironov, 2017]). *Let $\alpha \in [1, \infty)$. A randomized algorithm $\mathcal{M} : X \rightarrow \text{Prob}(Y)$ is (α, ρ) -Rényi differentially private if for every pair x_0 and x_1 of adjacent inputs, we have*

$$D_X^\alpha(\mathcal{M}(x_0) \parallel \mathcal{M}(x_1)) \leq \rho.$$

Rényi Differential privacy considers a fixed value of α . In contrast, zero-Concentrated Differential Privacy (zCDP) [Bun and Steinke, 2016], a simplification of Concentrated Differential Privacy (CDP) [Dwork and Rothblum, 2016], quantifies over all possible $\alpha > 1$.

Definition 6 (zero-Concentrated Differential Privacy (zCDP) [Bun and Steinke, 2016]). A randomized algorithm $\mathcal{M} : X \rightarrow \text{Prob}(Y)$ is (ξ, ρ) -zero concentrated differentially private if for every pairs of adjacent inputs x_0 and x_1 , we have

$$\forall \alpha > 1. D_Y^\alpha(\mathcal{M}(x_0) \parallel \mathcal{M}(x_1)) \leq \xi + \alpha\rho. \quad (2)$$

Truncated Concentrated Differential Privacy (tCDP) [Bun et al., 2018] quantifies over all α below a given threshold.

Definition 7 (Truncated Concentrated Differential Privacy (tCDP) [Bun et al., 2018]). A randomized algorithm $\mathcal{M} : X \rightarrow \text{Prob}(Y)$ is (ρ, ω) -truncated concentrated differentially private if for every pairs of adjacent inputs x_0 and x_1 , we have

$$\forall 1 < \alpha < \omega. D_Y^\alpha(\mathcal{M}(x_0) \parallel \mathcal{M}(x_1)) \leq \alpha\rho. \quad (3)$$

These notions are all motivated by bounds on the *privacy loss* of a randomized algorithm. This quantity is defined by

$$\mathcal{L}^{x_0 \rightarrow x_1}(y) \stackrel{\text{def}}{=} \frac{\Pr[\mathcal{M}(x_0) = y]}{\Pr[\mathcal{M}(x_1) = y]},$$

where x_0 and x_1 are two adjacent inputs. Intuitively, the privacy loss measures how much information is revealed by an output y . While output values with a large privacy loss are highly revealing—they are far more likely to result from a private input x_0 rather than a different private input x_1 —if these outputs are only seen with small probability then it may be reasonable to discount their influence. Each of the privacy definitions above bounds different moments of this privacy loss, treated as a random variable when y is drawn from the output of the algorithm on input x_0 . The following table summarizes these bounds.

Privacy	Bound on privacy loss $\mathcal{L} = \mathcal{L}^{x_0 \rightarrow x_1}$
(ϵ, δ) -DP	$\Pr_{y \sim \mathcal{M}(x_0)}[\mathcal{L}(y) \leq e^\epsilon] \geq 1 - \delta$
(α, ρ) -RDP	$\mathbb{E}_{y \sim \mathcal{M}(x_1)}[\mathcal{L}(y)^\alpha] \leq e^{(\alpha-1)\rho}$
(ξ, ρ) -zCDP	$\forall \alpha \in (1, \infty). \mathbb{E}_{y \sim \mathcal{M}(x_1)}[\mathcal{L}(y)^\alpha] \leq e^{(\alpha-1)(\xi+\alpha\rho)}$
(ω, ρ) -tCDP	$\forall \alpha \in (1, \omega). \mathbb{E}_{y \sim \mathcal{M}(x_1)}[\mathcal{L}(y)^\alpha] \leq e^{(\alpha-1)\alpha\rho}$

In particular, DP bounds the maximum value of the privacy loss,¹ (α, \cdot) -RDP bounds the α -moment, zCDP bounds all moments, and (\cdot, ω) -tCDP bounds the moments up to some cutoff ω . Many conversions are known between these definitions; for instance, RDP, zCDP, and tCDP are known to sit between $(\epsilon, 0)$ and (ϵ, δ) -differential privacy in terms of expressivity, up to some modification in the parameters. While this means that RDP, zCDP, and tCDP can sometimes be analyzed by reduction to standard differential privacy, converting between the different notions requires weakening the parameters and often the privacy analysis is simpler or more precise when working with RDP, zCDP, or tCDP directly. The interested reader can refer to the original papers [Bun and Steinke, 2016, Mironov, 2017, Bun et al., 2018].

3 k -generated divergences

3.1 Background and notation

We use standard notation and terminology from discrete probability. For every $x \in X$, we denote by \mathbf{d}_x the Dirac distribution centered at x defined by $\mathbf{d}_x(x') = 1$ if $x = x'$ and

¹Technically speaking, this is true only for sufficiently well-behaved distributions [Meiser, 2018].

$\mathbf{d}_x(x') = 0$ otherwise. For any probability distribution $\mu \in \text{Prob}(X)$ and $\gamma: X \rightarrow \text{Prob}(Y)$, we define $\gamma(\mu) \in \text{Prob}(Y)$ to be $(\gamma(\mu))(y) = \sum_{x \in X} (\gamma(x))(y) \cdot \mu(x)$ for every $y \in Y$. For any function $\gamma: X \rightarrow Y$, as an abuse of notation, we define $\gamma(\mu) \in \text{Prob}(Y)$ to be $\{x \mapsto \mathbf{d}_{\gamma(x)}(\mu)\}$, equivalently, $(\gamma(\mu))(y) = \sum_{x \in \gamma^{-1}(y)} \mu(x)$ for every $y \in Y$. Every function $\gamma: X \rightarrow Y$ can be regarded as $\{x \mapsto \mathbf{d}_{\gamma(x)}\}: X \rightarrow \text{Prob}(Y)$.

3.2 Divergences between distributions

We start from a very general definition of divergences. Our notation includes the domain of definition of the divergence; this distinction will be important when introducing the concept of k -generatedness.

Definition 8. A divergence is a family $\Delta = \{\Delta_X\}_X$ of functions

$$\Delta_X: \text{Prob}(X) \times \text{Prob}(X) \rightarrow [0, \infty].$$

We use the notation $\Delta_X(\mu_1 \parallel \mu_2)$ to denote the divergence between distributions μ_1 and μ_2 over X .

Our notion of divergence subsumes the general notion of f -divergence from the literature [Csiszár, 1963, Csiszár and Shields, 2004]. In particular, this includes the ε -divergence [Barthe and Olmedo, 2013] used to formulate (ε, δ) -differential privacy:

$$\Delta_X^\varepsilon(\mu_1 \parallel \mu_2) \stackrel{\text{def}}{=} \sup_{S \subseteq X} (\Pr[\mu_1 \in S] - e^\varepsilon \Pr[\mu_2 \in S]).$$

Specifically, a randomized algorithm $\mathcal{M}: X \rightarrow \text{Prob}(Y)$ is (ε, δ) -differentially private if and only if for every pair of adjacent inputs x_0 and x_1 , we have

$$\Delta_Y^\varepsilon(\mathcal{M}(x_0) \parallel \mathcal{M}(x_1)) \leq \delta.$$

Many useful properties of divergences have been explored in the literature. Our technical development will involve the following two properties.

- A divergence Δ satisfies the *data-processing inequality* iff for every $\gamma: X \rightarrow \text{Prob}(Y)$, $\Delta_Y(\gamma(\mu_1) \parallel \gamma(\mu_2)) \leq \Delta_X(\mu_1 \parallel \mu_2)$.
- A divergence Δ is *quasi-convex* iff for every $\alpha_1, \dots, \alpha_m \in [0, 1]$ such that $\sum_{m=1}^N \alpha_m = 1$ and every discrete set X ,

$$\Delta_X(\sum_{m=1}^N \alpha_m d_{1,m} \parallel \sum_{m=1}^N \alpha_m d_{2,m}) \leq \max_m \Delta_X(d_{1,m} \parallel d_{2,m}).$$

These properties are satisfied by many common divergences. Besides Rényi divergences, they also hold for all f -divergences [Csiszár, 1963, Csiszár and Shields, 2004]. We will consider only divergences satisfying them in the following.

3.3 k -cuts of divergences

We now introduce a technical construction that will be useful in the rest of the paper.

Definition 9. Let $k \in \mathbb{N} \cup \{\infty\}$. For any divergence $\Delta = \{\Delta_X\}_{X: \text{set}}$, we define a k -cut $\overline{\Delta}^k = \{\overline{\Delta}_X^k\}_{X: \text{set}}$ as follows: we fix a set Y with cardinality k , i.e. $|Y| = k$, and define

$$\overline{\Delta}_X^k(\mu_1 \parallel \mu_2) \stackrel{\text{def}}{=} \sup_{\gamma: X \rightarrow \text{Prob}(Y)} \Delta_Y(\gamma(\mu_1) \parallel \gamma(\mu_2)).$$

For divergences Δ that satisfy the data-processing inequality, then the k -cut is well-defined: it does not depend on the choice of Y .

Lemma 10. *If a divergence Δ satisfies the data-processing inequality, we have the inequality $\overline{\Delta}^k \leq \Delta$ and the equality $\overline{\Delta}_Y^k = \Delta_Y$ for any set Y with $|Y| = k$.*

So, without loss of generality in the sequel we will refer to this as “the” k -cut.

Another interesting property of k -cuts is that a k -cut $\overline{\Delta}^k$ of a divergence Δ satisfies the data-processing inequality, even if the original divergence Δ does not satisfy it.

Without loss of generality, we can assume the function γ in the definition of a k -cut to be deterministic. This can be proved by a weak version of Birkhoff-von Neumann theorem, which decomposes every probabilistic decision rule into a convex combination of deterministic ones.

Theorem 11 (Weak Birkhoff-von Neumann). *Let $k, l \in \mathbb{N}$ and $k > l$. Let X and Y such that $|X| = k$ and $|Y| = l$. Then for any $\gamma: X \rightarrow \text{Prob}(Y)$, there exist $N \in \mathbb{N}$, $\gamma_1, \dots, \gamma_N: X \rightarrow Y$ and $a_1, \dots, a_N \in [0, 1]$ such that $\sum_{m=1}^N a_m = 1$ and $\gamma(x) = \sum_{m=1}^N a_m \mathbf{d}_{\gamma_m(x)}$ ($x \in X$).*

This fact, allow us to consider simplified formulations of a k -cut of a divergence. Examples of this fact that will be useful in the sequel are the 2-cut and 3-cut of the Rényi divergence of order α . These can be reformulated as follows:

$$\begin{aligned} \overline{D}_X^{\alpha 2}(\mu_1 \| \mu_2) &= \sup_{S \subseteq X} \frac{1}{\alpha - 1} \log \left\{ \begin{array}{l} \Pr[\mu_1 \in S]^\alpha \Pr[\mu_2 \in S]^{1-\alpha} \\ + \Pr[\mu_1 \notin S]^\alpha \Pr[\mu_2 \notin S]^{1-\alpha} \end{array} \right\}, \\ \overline{D}_X^{\alpha 3}(\mu_1 \| \mu_2) &= \sup_{\substack{S_1, S_2 \subseteq X, \\ S_1 \cap S_2 = \emptyset}} \frac{1}{\alpha - 1} \log \left\{ \begin{array}{l} \Pr[\mu_1 \in S_1]^\alpha \Pr[\mu_2 \in S_1]^{1-\alpha} \\ + \Pr[\mu_1 \in S_2]^\alpha \Pr[\mu_2 \in S_2]^{1-\alpha} \\ + \Pr[\mu_1 \notin S_1 \cup S_2]^\alpha \Pr[\mu_2 \notin S_1 \cup S_2]^{1-\alpha} \end{array} \right\}. \end{aligned}$$

3.4 k -generatedness of divergences

We now introduce the notion of k -generatedness. Informally, k -generatedness is a measure of the number of decisions that are needed in an hypothesis test to characterize a divergence.

Definition 12. *Let $k \in \mathbb{N} \cup \{\infty\}$. A divergence Δ is k -generated if a k -cut $\overline{\Delta}^k$ of Δ is equal to Δ itself.*

k -generatedness can also be reformulated as follows:

Lemma 13. *If $\Delta = \{\Delta_X\}_{X: \text{set}}$ is k -generated, for any set Y with $|Y| = k$, we have*

$$\Delta_X(\mu_1 \| \mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(Y)} \Delta_Y(\gamma(\mu_1) \| \gamma(\mu_2)).$$

Lemma 14. *The following basic properties hold for all k -generated divergences.*

- *If Δ is 1-generated, then Δ is constant, i.e. there exists $c \in [0, \infty]$ such that for every X and every $\mu_1, \mu_2 \in \text{Prob}(X)$, we have $\Delta_X(\mu_1 \| \mu_2) = c$.*
- *If Δ is k -generated, then it is also $k + 1$ -generated.*
- *If Δ has the data-processing inequality, then it is at least ∞ -generated.*

- Every k -cut of a divergence Δ is k -generated.

To compare a k -generated divergence and a divergence, we have the following lemma where all the inequalities are defined pointwise.

Lemma 15. *Consider a divergence Δ and a k -generated divergence Δ' . For any k -cut $\overline{\Delta}^k$ of Δ ,*

$$\Delta' \leq \Delta \implies \Delta' \leq \overline{\Delta}^k.$$

Also, if Δ has the data-processing inequality, the k -cut is the greatest k -generated divergence below Δ :

$$\Delta' \leq \Delta \iff \Delta' \leq \overline{\Delta}^k \leq \Delta.$$

3.4.1 A 2-generated divergence for DP

The divergence Δ^ε that can be used to characterize (ε, δ) -DP is 2-generated. This implies that DP can be characterized completely by its hypothesis testing interpretation.

Theorem 16. *The ε -divergence Δ^ε is 2-generated.*

Since Δ^ε is quasi-convex and satisfies data-processing inequality, the 2-cut can be reformulated as:

$$\overline{\Delta}^2_X(\mu_1 \parallel \mu_2) = \sup_{S \subseteq X} (\Pr[\mu_1 \in S] - e^\varepsilon \Pr[\mu_2 \in S])$$

It is easy to show that this is exactly the same as the original definition of Δ^ε , from which follows that it is 2-generated.

3.4.2 Rényi is ∞ -generated

In contrast to the divergence Δ^ε , the 2-cut of the Rényi divergence is not complete with respect to the Rényi divergence.

To see this let $X = \{a, b, c\}$ and let $\mu_1, \mu_2 \in \text{Prob}(X)$ be defined by $\mu_1(a) = \mu_1(b) = \mu_1(c) = \frac{1}{3}$ and $\mu_2(a) = \frac{p^2}{p^2+p+1}$, $\mu_2(b) = \frac{p}{p^2+p+1}$ and $\mu_2(c) = \frac{1}{p^2+p+1}$.

We set $\beta > \alpha + 1$ and $p = (1/2)^{\beta/(\alpha-1)}$, a simple calculation shows:

$$\overline{D}^{\alpha}_X(\mu_1 \parallel \mu_2) + \frac{1}{\alpha-1} \log \frac{2^\beta + 2^{-\beta} + 1}{\max(2^{\alpha+1}, 2^\beta + 1)} \leq D^{\alpha}_X(\mu_1 \parallel \mu_2)$$

The difference is quantitatively small, but it is nevertheless strictly positive. This shows also that the Rényi divergence is not 2-generated.

Similarly, one can show that the 3-cut is not complete, that the 4-cut is not complete, etc. In fact, Rényi divergence is exactly ∞ -generated. Indeed, Rényi divergence satisfies the data-processing inequality, hence it is at most ∞ -generated. Moreover, any f -divergences whose weight function f is strictly convex is not k -generated for any finite k . The formulation of Rényi divergence of order α given by $\exp((\alpha-1)D^{\alpha}_X(\mu_1 \parallel \mu_2))$ is an f -divergence related to the weight function $t \mapsto t^\alpha$, which is strictly convex. Since the logarithm function is continuous on $(0, \infty)$ and strictly monotone, we conclude that the Rényi divergence is ∞ -generated. The formal details can be found in the appendix.

4 Hypothesis Testing Interpretation of Divergences

In this section, we give an hypothesis testing characterization similar to the one that differential privacy satisfies for the 2-cut of an arbitrary divergence.

We first define privacy regions for divergences using their 2-cuts.

Definition 17. For any divergence Δ , we define its privacy region $R^\Delta(\rho) \subseteq [0, 1] \times [0, 1]$ by

$$R^\Delta(\rho) \stackrel{\text{def}}{=} \left\{ (x, y) \mid \bar{\Delta}_{\{\text{Acc}, \text{Rej}\}}^{-2}((1-x)\mathbf{d}_{\text{Acc}} + x\mathbf{d}_{\text{Rej}} \| y\mathbf{d}_{\text{Acc}} + (1-y)\mathbf{d}_{\text{Rej}}) \leq \rho \right\}.$$

Notice that if Δ satisfies the data-processing inequality, or is 2-generated, then $\bar{\Delta}_{\{\text{Acc}, \text{Rej}\}}^{-2}$ in the definition above can be replaced by $\Delta_{\{\text{Acc}, \text{Rej}\}}$.

As an example, we can give the privacy region of DP.

$$R^{\Delta^\varepsilon}(\delta) = \{ (x, y) \mid 1 - x \leq e^\varepsilon y + \delta, \quad x \leq e^\varepsilon(1 - y) + \delta \}$$

Privacy regions are intimately related to the hypothesis testing interpretation of privacy definitions based on divergences.

Theorem 18. Let $\mu_1, \mu_2 \in \text{Prob}(X)$. $\bar{\Delta}_X^{-2}(\mu_1 \| \mu_2) \leq \rho$ holds if and only if for any $\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})$,

$$(\Pr[\gamma(\mu_1) = \text{Rej}], \Pr[\gamma(\mu_2) = \text{Acc}]) \in R^\Delta(\rho).$$

In the theorem above, the functions $\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})$ can be seen as probabilistic decision rules. Moreover, the privacy region can be actually relaxed to $R^\Delta(\rho) \cup \{ (x, y) \mid x + y \geq 1 \}$ since for any decision rule $\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})$, we can take its negation $\neg\gamma$. Hence we do not need to check the cases of $\Pr[\gamma(\mu_1) = \text{Rej}] + \Pr[\gamma(\mu_2) = \text{Acc}] > 1$. This also corresponds to the symmetry that we have in their graphical representations.

Finally, if a divergence Δ is quasi-convex, we also have the equivalent of Theorem 18 under deterministic decision rules. In this case, we have the following reformulation. Let $\mu_1, \mu_2 \in \text{Prob}(X)$. $\bar{\Delta}_X^{-2}(\mu_1 \| \mu_2) \leq \rho$ iff for any $S \subseteq X$,

$$(\Pr[\mu_1 \in S], \Pr[\mu_2 \notin S]) \in R^\Delta(\rho).$$

This give us the hypothesis testing characterization of DP, since the ε -divergence is 2-generated and quasi-convex.

Corollary 19. Let $\mu_1, \mu_2 \in \text{Prob}(X)$. Set $\varepsilon, \delta \geq 0$. $\Delta_X^\varepsilon(\mu_1 \| \mu_2) \leq \delta$ iff for any $S \subseteq X$,

$$(\Pr[\mu_1 \in S], \Pr[\mu_2 \notin S]) \in R^{\Delta^\varepsilon}(\delta).$$

We conclude this section by stressing that Theorem 18 tell us two important things:

- Every privacy definition similar to differential privacy but based on a 2-generated divergence is characterized completely by its hypothesis testing interpretation.
- For every privacy definition similar to differential privacy but based on an arbitrary divergence we can have an hypothesis testing interpretations by considering its 2-cut. However, this characterization will not be necessarily complete.

The second remark applies in particular to relaxations of differential privacy based on the Rényi divergence: if we want to have the hypothesis testing interpretation for one of these relaxations we can use the 2-cut of the Rényi divergence.

5 Applications

In this section we will use the technical tools we developed in the previous sections to better study the relations between different privacy definitions.

5.1 Conversions from Divergences to DP

Privacy regions can be used to give better conversion rules between privacy definitions based on divergences and differential privacy. Let $\Delta' = \{\Delta'_X\}_{X: \text{set}}$ be a divergence satisfying the data-processing inequality. We want to find the minimal parameters $(\varepsilon(\rho), \delta(\rho))$ such that $\Delta'_X(\mu_1 \parallel \mu_2) \leq \rho$ implies $\Delta_X^{\varepsilon(\rho)}(\mu_1 \parallel \mu_2) \leq \delta(\rho)$.

By Theorem 18 and Lemma 15, $R^{\Delta'}(\rho) \subseteq R^{\Delta^{\varepsilon(\rho)}}(\delta(\rho))$ holds if and only if for any pair $\mu_1, \mu_2 \in \text{Prob}(X)$,

$$\Delta'_X(\mu_1 \parallel \mu_2) \leq \rho \implies \Delta_X^{\varepsilon(\rho)}(\mu_1 \parallel \mu_2) \leq \delta(\rho).$$

This means that to find a good conversion law we can just compare the privacy regions.

5.1.1 Better Conversion from RDP to DP




Figure 2: A refined conversion law from RDP to DP. The gray region is $R^{D^\alpha}(\rho)$. The gray and black lines show original and refined DP-bounds for the same δ .

Using privacy regions, we can refine Mironov's conversion law from RDP to DP in a simple way.

Lemma 20 ([Mironov, 2017, Prop. 3]). *If a mechanism M is (α, ρ) -RDP then the mechanism is also $(\rho - \log \delta / (\alpha - 1), \delta)$ -DP for any $0 < \delta < 1$.*

The privacy region of Rényi divergence is given by

$$R^{D^\alpha}(\rho) = \left\{ (x, y) \mid x^\alpha (1-y)^{1-\alpha} + (1-x)^\alpha y^{1-\alpha} \leq e^{\rho(\alpha-1)} \right\}.$$

Here we assume $0^{1-\alpha} = 0$.

By an extension of Lemma 15, to find ε satisfying

$$\forall X. \forall \mu_1, \mu_2 \in \text{Prob}(X). D_X^\alpha(\mu_1 \| \mu_2) \leq \rho \implies \Delta_X^\varepsilon(\mu_1 \| \mu_2) \leq \delta,$$

it is necessary and sufficient to find ε satisfying

$$\forall X. \forall \mu_1, \mu_2 \in \text{Prob}(X). \overline{D}_X^{\alpha, 2}(\mu_1 \| \mu_2) \leq \rho \implies \Delta_X^\varepsilon(\mu_1 \| \mu_2) \leq \delta.$$

By Theorem 18, this is equivalent to find ε satisfying $R^{D^\alpha}(\rho) \subseteq R^{\Delta^\varepsilon}(\delta)$. Inspired from Mironov's proof of conversion law from RDP to DP [Mironov, 2017, Propositon 3]: we obtain,

$$\begin{aligned} x^\alpha(1-y)^{1-\alpha} + (1-x)^\alpha y^{1-\alpha} &\leq e^{\rho(\alpha-1)} \\ \implies (1-x)^\alpha y^{1-\alpha} &\leq e^{\rho(\alpha-1)} \\ \implies (1-x) &\leq (e^\rho y)^{\frac{\alpha-1}{\alpha}} & (\dagger) \\ \implies (e^\rho y > \delta^{\frac{\alpha}{\alpha-1}} \implies (1-x) &\leq e^{\rho - \log d / (\alpha-1)} y) \\ \wedge (e^\rho y \leq \delta^{\frac{\alpha}{\alpha-1}} \implies (1-x) &\leq \delta) \\ \implies (1-x) &\leq e^{\rho - \log d / (\alpha-1)} y + \delta. & (\ddagger) \end{aligned}$$

The equality (\ddagger) derives original Mironov's result [Mironov, 2017, Propositon 3]. Now, starting from (\dagger) , we have a better bound for DP as follows: consider a curve C given by the equation

$$1-x = (e^\rho y)^{\frac{\alpha-1}{\alpha}} \iff x = 1 - (e^\rho y)^{\frac{\alpha-1}{\alpha}}$$

. We have the derivative of x as follows:

$$\frac{dx}{dy} = -\frac{\alpha-1}{\alpha} e^{\frac{\alpha-1}{\alpha}\rho} y^{-\frac{1}{\alpha}}$$

We can take the tangent of the curve C by

$$x = \frac{dx}{dy}(t)(y-t) + (e^\rho(1-t))^{\frac{\alpha-1}{\alpha}}$$

We will find parameters that a tangent of C meets $(1-x) = e^\varepsilon y + \delta$. $x = -e^\varepsilon y - \delta + 1$
We first solve

$$-e^\varepsilon = \frac{dx}{dy}(t) = -\frac{\alpha-1}{\alpha} e^{\frac{\alpha-1}{\alpha}\rho} t^{-\frac{1}{\alpha}} \iff \varepsilon = \log\left(\frac{\alpha-1}{\alpha}\right) + \frac{\alpha-1}{\alpha}\rho - \frac{1}{\alpha} \log t.$$

Next we solve

$$1-\delta = -t \frac{dx}{dy}(t) + 1 - (e^\rho t)^{\frac{\alpha-1}{\alpha}} \iff 1-\delta = \frac{\alpha-1}{\alpha} e^{\frac{\alpha-1}{\alpha}\rho} t^{-\frac{1}{\alpha}} t + 1 - (e^\rho t)^{\frac{\alpha-1}{\alpha}}$$

We then have

$$\delta = (e^\rho t)^{\frac{\alpha-1}{\alpha}} - \frac{\alpha-1}{\alpha} e^{\frac{\alpha-1}{\alpha}\rho} t^{-\frac{1}{\alpha}} t = \frac{1}{\alpha} (e^\rho t)^{\frac{\alpha-1}{\alpha}} \iff t = (\delta \alpha e^{-\frac{\alpha-1}{\alpha}\rho})^{\frac{\alpha}{\alpha-1}}$$

Simple computations give the following:

$$\varepsilon = \log\left(\frac{\alpha-1}{\alpha}\right) + \rho - \frac{\log \delta + \log \alpha}{\alpha-1}.$$

By the symmetry of $R^{D^\alpha}(\rho)$ and $R^{\Delta^\epsilon}(\delta)$, we have

$$R^{D^\alpha}(\rho) \subseteq R^{\Delta^\epsilon}(\delta).$$

As we mentioned, it is equivalent to

$$\forall X. \forall \mu_1, \mu_2 \in \text{Prob}(X). D_X^\alpha(\mu_1 \parallel \mu_2) \leq \rho \implies \Delta_X^\epsilon(\mu_1 \parallel \mu_2) \leq \delta.$$

Therefore, we have the following better conversion law:

Theorem 21. *If a mechanism \mathcal{M} is (α, ρ) -RDP then it is $(\rho + \log((\alpha - 1)/\alpha) - (\log \delta + \log \alpha)/(\alpha - 1), \delta)$ -DP for any $0 < \delta < 1$.*

As a conjecture, if we calculate tangents of the boundary of the privacy region $R^{D^\alpha}(\rho)$, we have *optimal* conversion law from (α, ρ) -RDP to DP. The boundary of $R^{D^\alpha}(\rho)$ is given by the equation

$$x^\alpha(1-y)^{1-\alpha} + (1-x)^\alpha y^{1-\alpha} = e^{\rho(\alpha-1)}.$$

5.2 On Gaussian Differential Privacy

Gaussian differential privacy (GDP) [Dong et al., 2019, Def. 2.6] has been recently proposed as a privacy definition trading-off PMD and PFA. This can be characterized by means of privacy regions. We have seen that privacy regions correspond to 2-generated divergence. Thus, a natural question is: can we characterize GDP using a 2-generated divergence. The answer is yes. We can characterize GDP by the following divergence:

$$\begin{aligned} & \Delta_X^{\text{Gauss}}(\mu_1 \parallel \mu_2) \\ &= \sup \left\{ \delta \mid \begin{array}{l} \exists \gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\}). \\ \Pr[\gamma(\mu_2) = \text{Acc}] \\ \geq \Phi(\Phi^{-1}(\Pr[\gamma(\mu_1) = \text{Rej}]) - \delta) \end{array} \right\}. \end{aligned}$$

where Φ is the standard normal CDF. The data-processing inequality of the divergence Δ^{Gauss} is proved from [Dong et al., 2019, Lem. 2.6] Hence, the privacy region is given as follows:

$$R^{\Delta^{\text{Gauss}}}(\delta) = \left\{ (x, y) \mid \begin{array}{l} y \geq \Phi(\Phi^{-1}(1-x) - \delta) \\ 1-y \geq \Phi(\Phi^{-1}(x) - \delta) \end{array} \right\}.$$

By Theorem 18, Δ^{Gauss} is 2-generated.

5.3 Informativeness of k -cuts

The concept of k -cut can be related to the ability that a divergence has of distinguishing two distributions.

Definition 22. *We say that a divergence Δ is δ -distinguishing a pair $\mu_1, \mu_2 \in \text{Prob}(X)$ of probability distributions if $\Delta_X(\mu_1, \mu_2) > \delta$.*

Now, consider a divergence Δ satisfying the data-processing inequality. Then the k -cuts form a monotone increasing sequence: $\bar{\Delta}^1 \leq \bar{\Delta}^2 \leq \bar{\Delta}^3 \leq \dots \leq \bar{\Delta}^k \leq \bar{\Delta}^{k+1} \leq \dots$. Thus for any divergence with data-processing inequality, $k+1$ -cut of Δ is always more informative than the k -cut of Δ for every $k \in \mathbb{N}$ in the following sense. If the

k -cut of a divergence is δ -distinguishing a pair $\mu_1, \mu_2 \in \text{Prob}(X)$ then the $k + 1$ cut is δ -distinguishing them too:

$$\overline{\Delta}_X^{k+1}(\mu_1 \parallel \mu_2) \geq \overline{\Delta}_X^k(\mu_1 \parallel \mu_2) > \delta.$$

For example, for the pair $\mu_1, \mu_2 \in \text{Prob}(\{a, b, c\})$ in the counterexample given in Section 3.4.2, we can find a value of δ such that \overline{D}^{α^3} is δ -distinguishing μ_1, μ_2 and \overline{D}^{α^2} is *not* δ -distinguishing them.

6 A characterization of k -generated divergences

As we have seen, k -generated divergences satisfy a number of useful properties; known divergences from the literature can be classified according to this parameter k - we have shown some examples here, more examples are in the supplemental material. In the other direction, we give a simple condition to ensure that a divergence is k -generated: the suprema of quasi-convex functions over size k -partitions determine k -generated divergences.

Theorem 23. *Let $F: [0, 1]^{2k} \rightarrow [0, \infty]$ be a quasi-convex function. Then the divergence Δ^F defined below is k -generated and quasi-convex.*

$$\Delta_X^F(\mu_1 \parallel \mu_2) \stackrel{\text{def}}{=} \sup_{\substack{\{A_i\}_{i=1}^k \\ \text{partition of } X}} F(\mu_1(A_1), \dots, \mu_1(A_k), \mu_2(A_1), \dots, \mu_2(A_k)).$$

Proof. The quasi-convexity is obvious from the quasi-convexity of $F: [0, 1]^{2k} \rightarrow [0, \infty]$. We show the k -generatedness. We take the k -cut with respect to the k -element set $\{1, 2, \dots, k\}$. We may assume X is countable. For any $\mu_1, \mu_2 \in \text{Prob}(X)$,

$$\begin{aligned} & \overline{\Delta}_X^k(\mu_1 \parallel \mu_2) \\ &= \sup_{\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\})} \Delta_{\{1, 2, \dots, k\}}^F(\gamma(\mu_1) \parallel \gamma(\mu_2)) \\ &= \sup_{\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\})} \sup_{\substack{\{A_i\}_{i=1}^k \\ \text{partition of} \\ \{1, 2, \dots, k\}}} F \left(\begin{array}{c} (\gamma(\mu_1))(A_1), \dots, (\gamma(\mu_1))(A_k), \\ (\gamma(\mu_2))(A_1), \dots, (\gamma(\mu_2))(A_k) \end{array} \right) \\ &= \sup_{\substack{\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\}) \\ p: \{1, 2, \dots, k\} \rightarrow \{1, 2, \dots, k\}}} F \left(\begin{array}{c} (\gamma(\mu_1))(p^{-1}(1)), \dots, (\gamma(\mu_1))(p^{-1}(k)), \\ (\gamma(\mu_2))(p^{-1}(1)), \dots, (\gamma(\mu_2))(p^{-1}(k)) \end{array} \right) \\ &= \sup_{\substack{\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\}) \\ p: \{1, 2, \dots, k\} \rightarrow \{1, 2, \dots, k\}}} F \left(\begin{array}{c} ((p \bullet \gamma)(\mu_1))(1), \dots, ((p \bullet \gamma)(\mu_1))(k), \\ ((p \bullet \gamma)(\mu_2))(1), \dots, ((p \bullet \gamma)(\mu_2))(k) \end{array} \right) \\ &= \sup_{\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\})} F \left(\begin{array}{c} ((\gamma(\mu_1))(1), \dots, (\gamma(\mu_1))(k)), \\ (\gamma(\mu_2))(1), \dots, (\gamma(\mu_2))(k) \end{array} \right). \end{aligned}$$

Here by weak Birkhoff-von Neumann theorem (countable version), every function $\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\})$ is decomposed into a (countable) convex combination

$\sum_{i \in I} a_i(\eta_{\{1,2,\dots,k\}} \circ \gamma_i)$ of $\gamma_i: X \rightarrow \{1, 2, \dots, k\}$. Hence,

$$\begin{aligned}
& \overline{\Delta_X^F}^k(\mu_1 \parallel \mu_2) \\
&= \sup_{\gamma: X \rightarrow \text{Prob}(\{1,2,\dots,k\})} F \left(\begin{array}{c} ((\gamma(\mu_1))(1), \dots, (\gamma(\mu_1))(k)), \\ (\gamma(\mu_2))(1), \dots, (\gamma(\mu_2))(k) \end{array} \right) \tag{\dagger} \\
&= \sup_{\gamma: X \rightarrow \text{Prob}(\{1,2,\dots,k\})} F \left(\begin{array}{c} ((\sum_{i \in I} a_i(\eta_{\{1,2,\dots,k\}} \circ \gamma_i)(\mu_1))(1), \\ \dots, (\sum_{i \in I} a_i(\eta_{\{1,2,\dots,k\}} \circ \gamma_i)(\mu_1))(k), \\ (\sum_{i \in I} a_i(\eta_{\{1,2,\dots,k\}} \circ \gamma_i)(\mu_2))(1), \\ \dots, (\sum_{i \in I} a_i(\eta_{\{1,2,\dots,k\}} \circ \gamma_i)(\mu_2))(k) \end{array} \right) \\
&= \sup_{\gamma: X \rightarrow \text{Prob}(\{1,2,\dots,k\})} F \left(\begin{array}{c} (\sum_{i \in I} a_i(\gamma_i(\mu_1))(1), \dots, \sum_{i \in I} a_i(\gamma_i(\mu_1))(k)), \\ (\sum_{i \in I} a_i(\gamma_i(\mu_2))(1), \dots, \sum_{i \in I} a_i(\gamma_i(\mu_2))(k) \end{array} \right) \\
&\leq \sup_{\gamma: X \rightarrow \text{Prob}(\{1,2,\dots,k\})} \sup_{i \in I} F \left(\begin{array}{c} ((\gamma_i(\mu_1))(1), \dots, (\gamma_i(\mu_1))(k)), \\ (\gamma_i(\mu_2))(1), \dots, (\gamma_i(\mu_2))(k) \end{array} \right) \\
&\leq \sup_{\gamma: X \rightarrow \{1,2,\dots,k\}} F \left(\begin{array}{c} ((\gamma(\mu_1))(1), \dots, (\gamma(\mu_1))(k)), \\ (\gamma(\mu_2))(1), \dots, (\gamma(\mu_2))(k) \end{array} \right) \\
&= \sup_{\gamma: X \rightarrow \{1,2,\dots,k\}} F \left(\begin{array}{c} (\mu_1(\gamma^{-1}(1)), \dots, \mu_1(\gamma^{-1}(k))), \\ (\mu_2(\gamma^{-1}(1)), \dots, \mu_2(\gamma^{-1}(k))) \end{array} \right) \\
&\leq \sup_{\substack{\{A_i\}_{i=1}^k \\ \text{partition of} \\ \{1,2,\dots,k\}}} F \left(\begin{array}{c} (\gamma(\mu_1))(A_1), \dots, (\gamma(\mu_1))(A_k), \\ (\gamma(\mu_2))(A_1), \dots, (\gamma(\mu_2))(A_k) \end{array} \right) \\
&= \Delta_X^F(\mu_1 \parallel \mu_2).
\end{aligned}$$

We have $\overline{\Delta_X^F}^k(\mu_1 \parallel \mu_2) \leq \Delta_X^F(\mu_1 \parallel \mu_2)$. Conversely, by equality (\dagger), we also have $\overline{\Delta_X^F}^k(\mu_1 \parallel \mu_2) \geq \Delta_X^F(\mu_1 \parallel \mu_2)$. This completes the proof. \square

This result characterizes k -generated quasi-convex divergences. It also serves as a useful tool to construct new divergences with a hypothesis testing interpretation, by varying the quasi-convex function F .

7 Conclusion

In this paper we have developed analytical tools to study the hypothesis testing interpretation of privacy definitions similar to differential privacy but measured with another statistical divergence. We introduced the notions of k -cut and k -generatedness for divergences. These notions quantifies the number of decisions that are needed in an experiment similar to the ones used in hypothesis testing to fully characterize the divergence. We used these notions to study the hypothesis testing interpretation of relaxations of differential privacy based on the Rényi divergence. These notions give a measure of the complexity that tools for formal verification may have. We leave the study of this connection for future work.

A Weak version of Birkhoff-von Neumann Theorem

Theorem 24 (Weak Birkhoff-von Neumann theorem (Theorem 11)). *Let $k, l \in \mathbb{N}$ and $k > l$. For any $\gamma: k \rightarrow \text{Prob}(l)$, there are $\gamma_1, \gamma_2, \dots, \gamma_N: k \rightarrow l$ and $0 \leq a_1, a_2, \dots, a_N \leq 1$ such that $\sum_{m=1}^N a_m = 1$ and $\gamma(i) = \sum_{m=1}^N a_m \mathbf{d}_{\gamma_m(i)}$ for any $1 \leq i \leq k$.*

The cardinal k can be relaxed to countable infinite cardinal ω , and then the families $\{\gamma_j\}_j$ and $\{a_j\}_j$ may be countable infinite.

Proof. Consider the following matrix representation f of γ :

$$f = \begin{pmatrix} f_{1,1} & \cdots & f_{l,1} \\ \vdots & & \vdots \\ f_{1,k} & \cdots & f_{l,k} \end{pmatrix}.$$

where $f_{i,j} = \gamma(i)(j)$ and $\sum_{j=1}^l f_{i,j} = 1$ for any $1 \leq i \leq l$.

For any $h: k \rightarrow l$, the matrix representation g of $(\{x \mapsto \mathbf{d}_x\} \circ h)$ is

$$g = \begin{pmatrix} g_{1,1} & \cdots & g_{l,1} \\ \vdots & & \vdots \\ g_{1,k} & \cdots & g_{l,k} \end{pmatrix}$$

satisfying that for any $1 \leq i \leq l$, there is exactly $1 \leq j \leq k$ such that $g_{i,j} = 1$ and $g_{i,s} = 0$ for $s \neq j$. Conversely, any matrix g satisfying this condition corresponds to some function $h: k \rightarrow l$. Consider the family G of matrix representations of maps of the form $(\{x \mapsto \mathbf{d}_x\} \circ h)$. We give an algorithm decomposing f to a convex sum of g :

1. Let $r_0 = 1$ and $\tilde{f}_0 = f$. We have $\sum_j (\tilde{f}_0)_{i,j} = r_0$ for all $1 \leq i \leq l$.
2. For given $0 \leq r_m \leq 1$ and \tilde{f}_m satisfying $\sum_j (\tilde{f}_m)_{i,j} = r_m$ for all $1 \leq i \leq l$, we define $g_{m+1} \in G$, $\alpha_{m+1} \in [0, 1]$, \tilde{f}_{m+1} and $r_{m+1} \in [0, 1]$ as follows:

$$\alpha_{m+1} = \min_s \max_t (\tilde{f}_m)_{s,t}, \quad r_{m+1} = r_m - \alpha_{m+1},$$

$$(g_{m+1})_{i,j} = \begin{cases} 1 & j = \underset{s}{\operatorname{argmax}} (\tilde{f}_m)_{i,s} \\ 0 & \text{(otherwise)} \end{cases}, \quad \tilde{f}_{m+1} = \tilde{f}_m - \alpha_{m+1} \cdot g_{m+1}.$$

3. If $r_{s+1} = 0$ then we terminate. Otherwise, we repeat the previous step.

In each step, we obtain the following conditions:

- We have $g_{m+1} \in G$ because g_{m+1} can be written as $g_{m+1} = \{x \mapsto \mathbf{d}_x\} \circ (\lambda i. \underset{s}{\operatorname{argmax}} (\tilde{f}_m)_{i,s})$.

- We have $0 < \alpha_{m+1}$ whenever $0 < r_m$ because

$$\alpha_{m+1} = 0 \iff \exists i. \max_j (\tilde{f}_m)_{i,j} = 0 \implies \exists i. r_m = \sum_j (\tilde{f}_m)_{i,j} = 0.$$

- We have $0 \leq (\tilde{f}_{m+1})_{i,j} \leq 1$ for any (i, j) from the following equation:

$$(\tilde{f}_{m+1})_{i,j} = \begin{cases} (\tilde{f}_m)_{i,j} - \min_t \max_s (\tilde{f}_m)_{t,s} & \text{if } j = \underset{s}{\operatorname{argmax}} (\tilde{f}_m)_{i,s} \\ (\tilde{f}_m)_{i,j} & \text{otherwise} \end{cases}.$$

When $i = \underset{s}{\operatorname{argmin}} \max_t (\tilde{f}_m)_{s,t}$ and $j = \underset{s}{\operatorname{argmax}} (\tilde{f}_m)_{i,s}$, we obtain $(\tilde{f}_{m+1})_{i,j} = 0$ while $0 < (\tilde{f}_{m+1})_{i,j}$. This implies that the number of 0 in \tilde{f}_m increases in this operation.

- We also have $\sum_j (\tilde{f}_{m+1})_{i,j} = r_{m+1}$ for all $1 \leq i \leq k$ because

$$\sum_j (\tilde{f}_{m+1})_{i,j} = \sum_j (\tilde{f}_m)_{i,j} - \alpha_{m+1} \cdot \sum_j (\tilde{g}_{m+1})_{i,j} = r_m - \alpha_{m+1} \cdot 1 = r_{m+1}.$$

Therefore the construction of $g_l \in G$, $\alpha_l \in [0, 1]$, \tilde{f}_l and $r_l \in [0, 1]$ terminates within $k \cdot l$ steps. When the construction terminates at the step N ($r_N = 0$ also holds), we have a convex decomposition of f by $f = \sum_{m=1}^N \alpha_m \cdot g_m$ where $\sum_{m=1}^N \alpha_m = 1$. This implies By taking $\gamma_1, \gamma_2, \dots, \gamma_N: k \rightarrow l$ such that g_m is a matrix representation of $(\{x \mapsto \mathbf{d}_x\} \circ \gamma_m)$, we obtain $\gamma(i) = \sum_{m=1}^N a_m \mathbf{d}_{\gamma_m(i)}$ for any $1 \leq i \leq k$ with $0 \leq a_1, a_2, \dots, a_N \leq 1$ and $\sum_{m=1}^N a_m = 1$. \square

B Omitted Proofs

B.1 Compositions of probabilistic processes

For simplicity, we introduce the composition operator of probabilistic processes (inspired from [Giry, 1982]). For any $\gamma_1: X \rightarrow \operatorname{Prob}(Z)$ and $\gamma: Z \rightarrow \operatorname{Prob}(Y)$, we define their composition $(\gamma \bullet \gamma_1): X \rightarrow \operatorname{Prob}(Y)$ by $(\gamma \bullet \gamma_1)(x) \stackrel{\text{def}}{=} \gamma(\gamma_1(x))$. It is easy to check that the composition $(\gamma \bullet \gamma_1)$ satisfies $(\gamma \bullet \gamma_1)(\mu) = \gamma(\gamma_1(\mu))$ for every $\mu \in \operatorname{Prob}(X)$.

- The composition operator \bullet is associative: $\gamma \bullet (\gamma_1 \bullet \gamma_2) = (\gamma \bullet \gamma_1) \bullet \gamma_2$ holds for all $\gamma_2: W \rightarrow \operatorname{Prob}(X)$, $\gamma_1: X \rightarrow \operatorname{Prob}(Z)$, and $\gamma: Z \rightarrow \operatorname{Prob}(Y)$.
- The function $\eta_X: X \rightarrow \operatorname{Prob}(X)$ defined by $\eta_X = \{x \mapsto \mathbf{d}_x\}$ is the unit of operator \bullet : we have $\gamma \bullet \eta_X = \gamma$ and $\eta_Y \bullet \gamma = \gamma$ for all $\gamma: X \rightarrow \operatorname{Prob}(Y)$

Thanks to the unit law and associativity of \bullet as an abuse of notations, we define

- $(\gamma \bullet \gamma_1): X \rightarrow \operatorname{Prob}(Z)$ for $\gamma_1: X \rightarrow Z$ and $\gamma: Z \rightarrow \operatorname{Prob}(Y)$ by $\gamma \bullet (\eta_Z \circ \gamma_1)$.
- $(\gamma \bullet \gamma_1): X \rightarrow \operatorname{Prob}(Z)$ for $\gamma_1: X \rightarrow \operatorname{Prob}(Z)$ and $\gamma: Z \rightarrow Y$ by $(\eta_Y \circ \gamma) \bullet \gamma_1$.
- $(\gamma \bullet \gamma_1): X \rightarrow \operatorname{Prob}(Z)$ for $\gamma_1: X \rightarrow Z$ and $\gamma: Z \rightarrow Y$ by $(\eta_Y \circ \gamma) \bullet (\eta_Z \circ \gamma_1)$, which is equal to $\eta_Y \circ (\gamma \circ \gamma_1)$.

Notice that, $\gamma(\mu) \in \operatorname{Prob}(Y)$ defined under $\gamma: X \rightarrow Y$ and $\mu \in \operatorname{Prob}(X)$ is exactly $(\eta_Y \circ \gamma)(\mu)$.

B.2 Proof of the data-processing inequality of k -cuts

Lemma 25. *For any divergence Δ , every k -cut $\overline{\Delta}^k$ satisfies data-processing inequality.*

Proof. We consider the k -cut of Δ with respect to a set Y satisfying $|Y| = k$

$$\overline{\Delta}_X^k(\mu_1 \parallel \mu_2) \stackrel{\text{def}}{=} \sup_{\gamma: X \rightarrow \operatorname{Prob}(Y)} \Delta_Y(\gamma(\mu_1) \parallel \gamma(\mu_2)).$$

For every pair $\mu_1, \mu_2 \in \text{Prob}(X)$, and any function $\gamma_1: X \rightarrow \text{Prob}(Z)$, we obtain the data-processing inequality

$$\begin{aligned} \overline{\Delta}_Z^k(\gamma_1(\mu_1) \parallel \gamma_1(\mu_2)) &= \sup_{\gamma: Z \rightarrow \text{Prob}(Y)} \Delta_Y(\gamma(\gamma_1(\mu_1)) \parallel \gamma(\gamma_1(\mu_2))) \\ &= \sup_{\gamma: Z \rightarrow \text{Prob}(Y)} \Delta_Y((\gamma \bullet \gamma_1)(\mu_1) \parallel (\gamma \bullet \gamma_1)(\mu_2)) \\ &\leq \sup_{\gamma': X \rightarrow \text{Prob}(Y)} \Delta_Y(\gamma'(\mu_1) \parallel \gamma'(\mu_2)) = \overline{\Delta}_X^k(\mu_1 \parallel \mu_2). \end{aligned}$$

The inequality is obtained by the inclusion

$$\{(\gamma \bullet \gamma_1): X \rightarrow \text{Prob}Y \mid \gamma: Z \rightarrow \text{Prob}(Y)\} \subseteq \{\gamma': X \rightarrow \text{Prob}(Y)\}.$$

□

B.3 Proof of Lemma 10

Lemma 26 (Lemma 10). *If a divergence Δ has the data-processing inequality, we have the inequality $\overline{\Delta}^k \leq \Delta$ and the equality $\overline{\Delta}_Y^k = \Delta_Y$ for any set Y with $|Y| = k$.*

Proof. We consider the k -cut of Δ with respect to a set W satisfying $|W| = k$

$$\overline{\Delta}_X^k(\mu_1 \parallel \mu_2) \stackrel{\text{def}}{=} \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_W(\gamma(\mu_1) \parallel \gamma(\mu_2)).$$

Thanks to the data-processing inequality of Δ , we have $\overline{\Delta}^k \leq \Delta$: for every pair $\mu_1, \mu_2 \in \text{Prob}(X)$, we obtain

$$\overline{\Delta}_X^k(\mu_1 \parallel \mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_W(\gamma(\mu_1) \parallel \gamma(\mu_2)) \leq \Delta_X(\mu_1, \mu_2).$$

Now, we consider a set Y with $|Y| = k$. We already have $\overline{\Delta}_Y^k \leq \Delta_Y$. We want to prove $\Delta_Y \leq \overline{\Delta}_Y^k$. Since $|Y| = |W| = k$, there is a bijection $f: Y \rightarrow W$. We then obtain for every pair $\nu_1, \nu_2 \in \text{Prob}(Y)$,

$$\Delta_Y(\nu_1 \parallel \nu_2) = \Delta_Y(f^{-1}(f(\nu_1)) \parallel f^{-1}(f(\nu_2))) \leq \Delta_W(f(\nu_1) \parallel f(\nu_2)) \leq \overline{\Delta}_Y^k(\nu_1 \parallel \nu_2)$$

The first and second inequalities are obtained by the dataprocessing inequality and the definition of k -cut respectively. □

B.4 Proof of Lemma 13

Lemma 27 (Lemma 13). *If $\Delta = \{\Delta_X\}_{X: \text{set}}$ is k -generated, for any set Y with $|Y| = k$, we have*

$$\Delta_X(\mu_1 \parallel \mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(Y)} \Delta_Y(\gamma(\mu_1) \parallel \gamma(\mu_2)).$$

Proof. Suppose that Δ is equal to the k -cut of Δ with respect to a set W satisfying $|W| = k$.

$$\Delta_X(\mu_1 \parallel \mu_2) = \overline{\Delta}_X^k(\mu_1 \parallel \mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_W(\gamma(\mu_1) \parallel \gamma(\mu_2)).$$

Since $\overline{\Delta}^k$ always satisfies data-processing inequality, the divergence Δ itself do so. We fix an arbitrary set $|Y|$ with $|Y| = k$. Since $|Y| = |W| = k$, there is a bijection $f: Y \rightarrow W$. For every pair $\mu_1, \mu_2 \in \text{Prob}(X)$, we obtain

$$\begin{aligned}
\Delta_X(\mu_1 \|\mu_2) &= \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_W(\gamma(\mu_1) \|\gamma(\mu_2)) \\
&= \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_W(f(f^{-1}(\gamma(\mu_1))) \|\ f(f^{-1}(\gamma(\mu_2)))) \\
&\leq \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_Y(f^{-1}(\gamma(\mu_1)) \|\ f^{-1}(\gamma(\mu_2))) \\
&= \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_Y((f^{-1} \bullet \gamma)(\mu_1) \|\ (f^{-1} \bullet \gamma)(\mu_2)) \\
&\leq \sup_{\gamma: X \rightarrow \text{Prob}(Y)} \Delta_Y(\gamma(\mu_1) \|\ \gamma(\mu_2)) \leq \Delta_X(\mu_1 \|\mu_2).
\end{aligned}$$

Here, the first and last inequalities are obtained from the data-processing inequality of Δ . The second inequality is proved from the inclusion

$$\left\{ (f^{-1} \bullet \gamma): X \rightarrow \text{Prob}(Y) \mid \gamma: X \rightarrow \text{Prob}(W) \right\} \subseteq \{ \gamma: X \rightarrow \text{Prob}(Y) \}.$$

□

B.5 Proof of Basic Properties of k -generatedness (Lemma 14)

Lemma 28 (Lemma 14 (1)). *If Δ is 1-generated, then Δ is constant, i.e. there exists $c \in [0, \infty]$ such that for every X and every $\mu_1, \mu_2 \in \text{Prob}(X)$, we have $\Delta_X(\mu_1 \|\mu_2) = c$.*

Proof. When Δ is 1-generated, there is a singleton set $\{a\}$ such that, for every pair $\mu_1, \mu_2 \in \text{Prob}(X)$,

$$\Delta_X(\mu_1 \|\mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(\{a\})} \Delta_{\{a\}}(\gamma(\mu_1) \|\gamma(\mu_2)).$$

Now, the set $\text{Prob}(\{a\})$ is a singleton set $\{\mathbf{d}_a\}$, and therefore both $\gamma(\mu_1)$ and $\gamma(\mu_2)$ are equal to \mathbf{d}_a for every $\gamma: X \rightarrow \text{Prob}(\{a\})$ and every pair $\mu_1, \mu_2 \in \text{Prob}(X)$. Hence, $\Delta_X(\mu_1 \|\mu_2) = c$ where $c = \Delta_{\{a\}}(\mathbf{d}_a \|\mathbf{d}_a)$. □

Lemma 29 (Lemma 14 (2)). *If Δ is k -generated, then it is also $k + 1$ -generated.*

Proof. Suppose that Δ is equal to the k -cut of Δ with respect to a set W satisfying $|W| = k$.

$$\Delta_X(\mu_1 \|\mu_2) = \overline{\Delta}_X^k(\mu_1 \|\mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_W(\gamma(\mu_1) \|\gamma(\mu_2)).$$

Let V be an arbitrary set with $|V| = k + 1$. We define the $k + 1$ -cut of $\overline{\Delta}^k$ with respect to the set V .

$$\overline{\overline{\Delta}}_X^{k+1}(\mu_1 \|\mu_2) \stackrel{\text{def}}{=} \sup_{\gamma: X \rightarrow \text{Prob}(V)} \overline{\Delta}_V^k(\gamma(\mu_1) \|\gamma(\mu_2)).$$

We then have

$$\begin{aligned}
\overline{\Delta}_X^{k+1}(\mu_1 \parallel \mu_2) &= \sup_{\gamma: X \rightarrow \text{Prob}(V)} \overline{\Delta}_V^k(\gamma(\mu_1) \parallel \gamma(\mu_2)) \\
&= \sup_{\gamma: X \rightarrow \text{Prob}(V)} \sup_{\gamma_1: V \rightarrow \text{Prob}(W)} \Delta_W(\gamma_1(\gamma(\mu_1)) \parallel \gamma_1(\gamma(\mu_2))) \\
&= \sup_{\gamma: X \rightarrow \text{Prob}(V)} \sup_{\gamma_1: V \rightarrow \text{Prob}(W)} \Delta_W((\gamma_1 \bullet \gamma)(\mu_1) \parallel (\gamma_1 \bullet \gamma)(\mu_2)) \\
&\stackrel{(*)}{=} \sup_{\gamma': X \rightarrow \text{Prob}(W)} \Delta_W(\gamma'(\mu_1) \parallel \gamma'(\mu_2)) = \overline{\Delta}_X^k(\mu_1 \parallel \mu_2)
\end{aligned}$$

The equality is obtained by the equality

$$\left\{ (\gamma_1 \bullet \gamma): X \rightarrow \text{Prob}(W) \mid \begin{array}{l} \gamma: X \rightarrow \text{Prob}(V), \\ \gamma_1: V \rightarrow \text{Prob}(W) \end{array} \right\} = \{ \gamma': X \rightarrow \text{Prob}(W) \}$$

The inclusion \subseteq is obvious. We show the reverse inclusion \supseteq . Since $|V| \geq |W|$, there is a pair of function $f: W \rightarrow V$ and $g: V \rightarrow W$ such that $g \circ f = \text{id}_W$. Then, every $\gamma': X \rightarrow \text{Prob}(W)$ can be decomposed into $\gamma' = \gamma_1 \bullet \gamma$ where $\gamma = (f \bullet \gamma')$ and $\gamma_1 = g$ (strictly, $\gamma = ((\eta_V \circ f) \bullet \gamma')$ and $\gamma_1 = \eta_W \circ g$). \square

Lemma 30 (Lemma 14 (3)). *If Δ has the data-processing inequality, then it is at least ∞ -generated.*

Proof. We fix a pair $\mu_1, \mu_2 \in \text{Prob}(X)$. The set $Y = \text{supp}(\mu_1) \cup \text{supp}(\mu_2)$ is at most countable. Hence there are two functions $f: X \rightarrow \mathbb{N}$ and $g: \mathbb{N} \rightarrow X$ such that $(g \circ f)(x) = x$ for every $x \in Y$. We then have $\mu_1 = (g \circ f)(\mu_1)$ and $\mu_2 = (g \circ f)(\mu_2)$. Thus,

$$\begin{aligned}
\Delta_X(\mu_1 \parallel \mu_2) &= \Delta_X((g \circ f)(\mu_1) \parallel (g \circ f)(\mu_2)) \\
&\leq \Delta_{\mathbb{N}}(f(\mu_1) \parallel f(\mu_2)) \\
&\leq \sup_{\gamma: X \rightarrow \text{Prob}(\mathbb{N})} \Delta_{\mathbb{N}}(\gamma(\mu_1) \parallel \gamma(\mu_2)) \\
&\leq \Delta_X(\mu_1 \parallel \mu_2).
\end{aligned}$$

The last part is an ∞ -cut. The first and last inequality is obtained by the data-processing inequality. The second one is obvious ($f: X \rightarrow \mathbb{N}$ is regarded as $\{x \mapsto \mathbf{d}_{f(x)}\}: X \rightarrow \text{Prob}(\mathbb{N})$). \square

Lemma 31 (Lemma 14 (4)). *Every k -cut of a divergence Δ is always k -generated.*

Proof. We can prove $\overline{\Delta}^k = \overline{\Delta}^k$ in a almost the same way as Lemma 14 (2). \square

Continuity of divergence (Lemma 14(3) in general setting) We can extend the results on divergences in the discrete setting to general measurable setting using the continuity of divergences. We say that a divergence Δ is continuous if for any pair $\mu_1, \mu_2 \in \text{Prob}(X)$,

$$\Delta_X(\mu_1 \parallel \mu_2) = \sup_{n \in \mathbb{N}} \sup_{\gamma: X \rightarrow \{0,1,2,\dots,n\}} \Delta_{\{0,1,2,\dots,n\}}(\gamma(\mu_1) \parallel \gamma(\mu_2)).$$

If Δ is continuous and satisfies data-processing inequality we have ∞ -generatedness (moreover we show the “countable”-generatedness) as follows:

$$\begin{aligned}
& \Delta_X(\mu_1 \parallel \mu_2) \\
&= \sup_{n \in \mathbb{N}} \sup_{\gamma: X \rightarrow \{0,1,2,\dots,n-1\}} \Delta_{\{0,1,2,\dots,n-1\}}(\gamma(\mu_1) \parallel \gamma(\mu_2)) \\
&= \sup_{n \in \mathbb{N}} \sup_{\gamma: X \rightarrow \{0,1,2,\dots,n-1\}} \Delta_{\{0,1,2,\dots,n-1\}}((g_n \circ f_n)(\gamma(\mu_1)) \parallel (g_n \circ f_n)(\gamma(\mu_2))) \\
&= \sup_{n \in \mathbb{N}} \sup_{\gamma: X \rightarrow \{0,1,2,\dots,n\}} \Delta_{\{0,1,2,\dots,n-1\}}(g_n((f_n \bullet \gamma)(\mu_1)) \parallel g_n((f_n \bullet \gamma)(\mu_1))) \\
&\leq \sup_{n \in \mathbb{N}} \sup_{\gamma: X \rightarrow \{0,1,2,\dots,n-1\}} \Delta_{\mathbb{N}}((f_n \bullet \gamma)(\mu_1) \parallel (f_n \bullet \gamma)(\mu_1)) \\
&\leq \sup_{\gamma: X \rightarrow \mathbb{N}} \Delta_{\mathbb{N}}(\gamma(\mu_1) \parallel \gamma(\mu_2)) \leq \overline{\Delta}_X^\infty(\mu_1, \mu_2) \\
&\leq \Delta_X(\mu_1, \mu_2).
\end{aligned}$$

Here $f_n: \{0, 1, 2, \dots, n-1\} \rightarrow \mathbb{N}$ is the inclusion mapping, and $g_n: \mathbb{N} \rightarrow \{0, 1, 2, \dots, n-1\}$ is defined by $g_n(k) = k$ if $(k < n)$ and $g_n(k) = n-1$ otherwise. We have $(g_n \circ f_n) = \text{id}_{\{0,1,2,\dots,n-1\}}$.

The first and last inequalities are obtained from data-processing inequality. The second inequality is obvious.

B.6 Proof of Lemma 15

Lemma 32 (Lemma 15). *Consider a divergence Δ and a k -generated divergence Δ' . For any k -cut $\overline{\Delta}^k$ of Δ ,*

$$\Delta' \leq \Delta \implies \Delta' \leq \overline{\Delta}^k.$$

Also, if Δ has the data-processing inequality, the k -cut is the greatest k -generated divergence below Δ :

$$\Delta' \leq \Delta \iff \Delta' \leq \overline{\Delta}^k \leq \Delta.$$

Proof. Since Δ' is k -generated, for any choice of Y with $|Y| = k$, we have

$$\Delta' \leq \Delta \implies \Delta'_Y \leq \Delta_Y \implies \overline{\Delta'}^k \leq \overline{\Delta}^k \iff \Delta' \leq \overline{\Delta}^k.$$

The second statement is proved as follows: From the first statement of this lemma and Lemma 3 (Lemma 10 in the paper), We have

$$\Delta' \leq \Delta \implies \Delta' \leq \overline{\Delta}^k \leq \Delta$$

The converse direction is obvious. \square

An extended version. We can extend this theorem to more suitable for conversion laws of differential privacy.

Lemma 33 (Lemma 15, extended). *Consider a divergence Δ satisfying data-processing inequality and a k -generated divergence Δ' .*

$$\begin{aligned}
& \forall X. \forall \mu_1, \mu_2 \in \text{Prob}(X). (\Delta_X(\mu_1 \parallel \mu_2) \leq \delta \implies \Delta'_X(\mu_1 \parallel \mu_2) \leq \rho) \\
& \iff \forall X. \forall \mu_1, \mu_2 \in \text{Prob}(X). (\overline{\Delta}_X^k(\mu_1 \parallel \mu_2) \leq \delta \implies \Delta'_X(\mu_1 \parallel \mu_2) \leq \rho)
\end{aligned}$$

Proof. (\Leftarrow) Obvious from Lemma 3 (Lemma 10 in the paper). (\Rightarrow) From the assumption, we obtain

$$\begin{aligned} \forall X. \forall \mu_1, \mu_2 \in \text{Prob}(X). \forall \gamma: X \rightarrow \text{Prob}(Y). \\ \Delta_Y(\gamma(\mu_1) \parallel \gamma(\mu_2)) \leq \delta \implies \Delta'_Y(\gamma(\mu_1) \parallel \gamma(\mu_2)) \leq \rho. \end{aligned}$$

This implies

$$\forall X. \forall \mu_1, \mu_2 \in \text{Prob}(X). (\overline{\Delta}_X^k(\mu_1 \parallel \mu_2) \leq \delta \implies \overline{\Delta}'_X^k(\mu_1 \parallel \mu_2) \leq \rho)$$

Thanks to the k -generatedness of Δ' , we conclude the statement of this lemma. \square

B.7 Proof of 2-generatedness of ε -divergence

Theorem 34. *The ε -divergence Δ^ε is 2-generated for all ε .*

Proof. We recall that the ε -divergence Δ^ε is quasi-convex (moreover, jointly convex) and satisfies data-processing inequality. We choose a set $Y = \{\text{Acc}, \text{Rej}\}$, and take the 2-cut of Δ^ε by

$$\overline{\Delta}_X^2(\mu_1 \parallel \mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})} \Delta_Y^\varepsilon(\gamma(\mu_1) \parallel \gamma(\mu_2))$$

We show this is equal to the original $\Delta_X^\varepsilon(\mu_1 \parallel \mu_2)$. Without loss of generality we may assume X is at most countable. If X is an arbitrary set, we can restrict it to countable set in a similar way as the proof of Lemma 7 (Lemma 14(3) in the paper).

By the weak Birkhoff-von Neumann Theorem (Theorem 1 in the appendix), each $\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})$ can be decomposed into a convex combination $\gamma(x) = \sum_{i \in I} \alpha_i \mathbf{d}_{\gamma_i(x)}$ of functions $\gamma_i: X \rightarrow \{\text{Acc}, \text{Rej}\}$ ($i \in I$) where I is a countable set and $\sum_{i \in I} \alpha_i = 1$. By combining this and quasi-convexity and data-processing inequality of Δ^ε , we obtain

$$\begin{aligned} \Delta^\varepsilon(\gamma(\mu_1) \parallel \gamma(\mu_2)) &= \Delta^\varepsilon(\sum_{i \in I} \alpha_i \gamma_i(\mu_1) \parallel \sum_{i \in I} \alpha_i \gamma_i(\mu_2)) \\ &= \sup_{i \in I} \Delta^\varepsilon(\gamma_i(\mu_1) \parallel \gamma_i(\mu_2)) \\ &\leq \sup_{\gamma: X \rightarrow \{\text{Acc}, \text{Rej}\}} \Delta_X^\varepsilon(\gamma(\mu_1) \parallel \gamma(\mu_2)). \end{aligned}$$

This implies

$$\begin{aligned} \overline{\Delta}_X^2(\mu_1 \parallel \mu_2) &= \sup_{\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})} \Delta^\varepsilon(\gamma(\mu_1) \parallel \gamma(\mu_2)) \\ &= \sup_{\gamma: X \rightarrow \{\text{Acc}, \text{Rej}\}} \Delta_{\{\text{Acc}, \text{Rej}\}}^\varepsilon(\gamma(\mu_1) \parallel \gamma(\mu_2)) \\ &= \sup_{\gamma: X \rightarrow \{\text{Acc}, \text{Rej}\}} \sup_{A \subseteq \{\text{Acc}, \text{Rej}\}} (\Pr[\gamma(\mu_1) \in A] - e^\varepsilon \Pr[\gamma(\mu_2) \in A]) \\ &= \sup_{\gamma: X \rightarrow \{\text{Acc}, \text{Rej}\}} \sup_{A \subseteq \{\text{Acc}, \text{Rej}\}} (\Pr[\mu_1 \in \gamma^{-1}(A)] - e^\varepsilon \Pr[\mu_2 \in \gamma^{-1}(A)]) \\ &\stackrel{(*)}{=} \sup_{S \subseteq X} (\Pr[\mu_1 \in S] - e^\varepsilon \Pr[\mu_2 \in S]) \\ &= \Delta_X^\varepsilon(\mu_1 \parallel \mu_2) \end{aligned}$$

We have the 2-generatedness: $\overline{\Delta}^\varepsilon = \Delta^\varepsilon$. The equality (*) is proved as follows: for given γ and A , we take $S = \gamma^{-1}(A)$. Conversely, for any $S \subseteq X$ we take $A = \{\text{Acc}\}$ and $\gamma = \chi_S$, which is the indicator function of S defined by $\chi_S(x) = 1$ if $x \in S$ and $\chi_S(x) = 0$ otherwise. \square

General version We can extend this result to general measurable setting by using the continuity of Δ^ε (see also [Liese and Vajda, 2006]), which is obtained by f -divergence characterization of Δ^ε [Barthe and Olmedo, 2013]. For general measurable space X and every pair $\mu_1, \mu_2 \in \text{Prob}(X)$ we have

$$\begin{aligned}\Delta_X^\varepsilon(\mu_1 \parallel \mu_2) &= \sup_{\gamma: X \rightarrow \mathbb{N}} \Delta_{\mathbb{N}}^\varepsilon(\gamma(\mu_1) \parallel \gamma(\mu_2)) \\ &= \sup_{\gamma: X \rightarrow \mathbb{N}} \sup_{\gamma': \mathbb{N} \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})} \Delta_{\{\text{Acc}, \text{Rej}\}}^\varepsilon((\gamma' \bullet \gamma)(\mu_1) \parallel (\gamma' \bullet \gamma)(\mu_2)) \\ &= \sup_{\gamma: X \rightarrow \{\text{Acc}, \text{Rej}\}} \Delta_{\{\text{Acc}, \text{Rej}\}}^\varepsilon(\gamma(\mu_1) \parallel \gamma(\mu_2))\end{aligned}$$

Functions are assumed to be measurable.

B.8 Counterexample: Rényi-divergence is not 2-generated

Theorem 35. *There are $\mu_1, \mu_2 \in \text{Prob}(\{a, b, c\})$ such that*

$$\overline{D}_{\{a,b,c\}}^{\alpha^2}(\mu_1 \parallel \mu_2) < D_{\{a,b,c\}}^\alpha(\mu_1 \parallel \mu_2)$$

Proof. Let $p = (1/2)^{\beta/(\alpha-1)}$ and $\alpha + 1 < \beta$ and define

$$\begin{aligned}\mu_1 &= \frac{1}{3} \mathbf{d}_a + \frac{1}{3} \mathbf{d}_b + \frac{1}{3} \mathbf{d}_c, \\ \mu_2 &= \frac{p^2}{p^2 + p + 1} \mathbf{d}_a + \frac{p}{p^2 + p + 1} \mathbf{d}_b + \frac{1}{p^2 + p + 1} \mathbf{d}_c\end{aligned}$$

Since Rényi divergence is quasi-convex and satisfies data-processing inequality, it suffices to show the proper inequality $D_{\{\text{Acc}, \text{Rej}\}}^\alpha(\gamma(\mu_1) \parallel \gamma(\mu_2)) < D_{\{a,b,c\}}^\alpha(\mu_1 \parallel \mu_2)$ holds for any *deterministic decision rule* $\gamma: \{a, b, c\} \rightarrow \{\text{Acc}, \text{Rej}\}$. There are 8 cases of $\gamma: \{a, b, c\} \rightarrow \{\text{Acc}, \text{Rej}\}$, but thanks to the data-processing inequality and reflexivity of Rényi divergence, it suffices to consider 3 cases: $(\gamma(a), \gamma(b), \gamma(c)) = (\text{Acc}, \text{Acc}, \text{Rej}), (\text{Acc}, \text{Rej}, \text{Acc}), (\text{Rej}, \text{Acc}, \text{Acc})$. Hence,

$$\begin{aligned}& \frac{\exp((\alpha-1)D_{\{a,b,c\}}^\alpha(\mu_1 \parallel \mu_2))}{\exp((\alpha-1)D_{\{\text{Acc}, \text{Rej}\}}^\alpha(\gamma(\mu_1) \parallel \gamma(\mu_2)))} \\ & \geq \min \left(\frac{p^{2(1-\alpha)} + p^{1-\alpha} + 1}{2^\alpha(p^2 + p)^{1-\alpha} + 1}, \frac{p^{2(1-\alpha)} + p^{1-\alpha} + 1}{2^\alpha(p^2 + 1)^{1-\alpha} + p^{1-\alpha}}, \frac{p^{2(1-\alpha)} + p^{1-\alpha} + 1}{2^\alpha(p+1)^{1-\alpha} + p^{2(1-\alpha)}} \right) \\ & \geq \min \left(\frac{2^\beta + 2^{-\beta} + 1}{2^\alpha(p+1)^{1-\alpha} + 2^{-\beta}}, \frac{2^\beta + 2^{-\beta} + 1}{2^{\alpha-\beta}(p^2 + 1)^{1-\alpha} + 1}, \frac{2^\beta + 2^{-\beta} + 1}{2^\beta + 2^{\alpha-\beta}(p+1)^{1-\alpha}} \right) \\ & \geq \min \left(\frac{2^\beta + 2^{-\beta} + 1}{2^{\alpha+1}}, \frac{2^\beta + 2^{-\beta} + 1}{2^\beta + 1} \right) > 1.\end{aligned}$$

Hence,

$$\begin{aligned}D_{\{\text{Acc}, \text{Rej}\}}^\alpha(\gamma(\mu_1) \parallel \gamma(\mu_2)) &+ \frac{1}{\alpha-1} \log \min \left(\frac{2^\beta + 2^{-\beta} + 1}{2^{\alpha+1}}, \frac{2^\beta + 2^{-\beta} + 1}{2^\beta + 1} \right) \\ &\leq D_{\{a,b,c\}}^\alpha(\mu_1 \parallel \mu_2).\end{aligned}$$

holds for any $\gamma: \{a, b, c\} \rightarrow \{\text{Acc}, \text{Rej}\}$. By the data-processing inequality of Rényi divergence, this discussion does not depend on the choice of $\{\text{Acc}, \text{Rej}\}$. By weak

Birkhoff-von Neumann theorem, and the quasi-convexity Rényi divergence, we conclude

$$\overline{D}_X^{\alpha^2}(\mu_1\|\mu_2) + \frac{1}{\alpha-1} \log \min\left(\frac{2^\beta + 2^{-\beta} + 1}{2^{\alpha+1}}, \frac{2^\beta + 2^{-\beta} + 1}{2^\beta + 1}\right) \leq D_X^\alpha(\mu_1\|\mu_2).$$

□

B.9 Proof of ∞ -generatedness of Rényi-divergence

f -divergences is a class of divergences that are characterized by convex functions. For a given convex function $f: [0, \infty) \rightarrow \mathbb{R}$ satisfying $\lim_{t \rightarrow 0^+} tf(0/t) = 0$ (this function is called weight function), we define an f -divergence Δ^f corresponding the function f ,

$$\Delta_X^f(\mu_1\|\mu_2) \stackrel{\text{def}}{=} \sum_{x \in X} \mu_2(x) f\left(\frac{\mu_1(x)}{\mu_2(x)}\right).$$

The α -Rényi divergence D^α can also be characterized using f -divergence as follows:

$$D^\alpha(\mu_1\|\mu_2) = \frac{1}{\alpha-1} \log \sum_{x \in X} \mu_2(x) \left(\frac{\mu_1(x)}{\mu_2(x)}\right)^\alpha = \frac{1}{\alpha-1} \log \Delta_X^{t \mapsto t^\alpha}(\mu_1\|\mu_2).$$

Remark that every f -divergence is quasi-convex (moreover jointly convex) and continuous, and satisfies data-processing inequality (see also [Liese and Vajda, 2006, Theorems 14–16]).

Since the mapping $t \mapsto \frac{1}{\alpha-1} \log t$ is monotone, every α -Rényi divergence D^α is also quasi-convex and satisfies data-processing inequality. Thanks to the data-processing inequality, every α -Rényi divergence D^α is at least ∞ -generated. We need to prove that for every finite k , every α -Rényi divergence D^α is not k -generated. To prove this, we use that the mapping $t \mapsto t^\alpha$ is strictly convex.

Lemma 36. *If a weight function is strictly convex, its f -divergence Δ^f is not k -generated for every finite k .*

Proof. Without loss of generality, we may assume $k > 1$.

Consider a pair $\mu_1, \mu_2 \in \text{Prob}(\{0, 1, 2, \dots, k\})$ satisfying $\text{supp}(\mu_1) = \text{supp}(\mu_2) = \{0, 1, 2, \dots, k\}$ and $\mu_1(i)/\mu_2(i) \neq \mu_1(j)/\mu_2(j)$ where $1 \leq i, j \leq k+1$ and $i \neq j$. We can give such distributions. Then we obtain,

$$\begin{aligned} & \overline{\Delta}_{\{0,1,2,\dots,k\}}^k(\mu_1\|\mu_2) \\ &= \sup_{\gamma: \{0,1,2,\dots,k\} \rightarrow \text{Prob}(\{0,1,2,\dots,k-1\})} \Delta_{\{0,1,2,\dots,k-1\}}^f(\gamma(\mu_1)\|\gamma(\mu_2)) \\ & \quad \{\text{Weak Birkhoff-von Neumann theorem and the joint convexity of } \Delta^f\} \\ &= \max_{\gamma: \{0,1,2,\dots,k\} \rightarrow \{0,1,2,\dots,k-1\}} \Delta^f(\gamma(\mu_1)\|\gamma(\mu_2)) \\ &= \max_{\gamma: \{0,1,2,\dots,k\} \rightarrow \{0,1,2,\dots,k-1\}} \sum_{j=0}^{k-1} f\left(\frac{\sum_{\gamma(i)=j} \mu_1(i)}{\sum_{\gamma(i)=j} \mu_2(i)}\right) (\sum_{\gamma(i)=j} \mu_2(i)) \\ & \quad \{\text{Jensen's inequality with the strict convexity of the weight function } f\} \\ &< \sum_{i=0}^k f\left(\frac{\mu_1(i)}{\mu_2(i)}\right) \mu_2(i) = \Delta^f(\mu_1\|\mu_2). \end{aligned}$$

Since $k + 1 > k$, by Dirichlet's pigeonhole principle, for any $\gamma: \{0, 1, 2, \dots, k\} \rightarrow \{0, 1, 2, \dots, k - 1\}$, for some $j \in \{0, 1, 2, \dots, k\}$, there are at least two different $i_1, i_2 \in \{0, 1, 2, \dots, k - 1\}$ such that $\gamma(i_1) = j$ and $\gamma(i_2) = j$. From the assumption on μ_1 and μ_2 , we have $(\mu_1(i_1)/\mu_2(i_1)) \neq (\mu_1(i_2)/\mu_2(i_2))$. Since the function f is strictly convex, by the condition for equality of Jensen's inequality, we have the strict inequality

$$f\left(\frac{\mu_1(i_1) + \mu_1(i_2)}{\mu_2(i_1) + \mu_2(i_2)}\right)(\mu_2(i_1) + \mu_2(i_2)) < f\left(\frac{\mu_1(i_1)}{\mu_2(i_1)}\right)\mu_2(i_1) + f\left(\frac{\mu_1(i_2)}{\mu_2(i_2)}\right)\mu_2(i_2).$$

Therefore, for any $\gamma: \{0, 1, 2, \dots, k\} \rightarrow \{0, 1, 2, \dots, k - 1\}$, we have

$$\sum_{j=1}^k \left(\frac{\sum_{\gamma(i)=j} \mu_1(i)}{\sum_{\gamma(i)=j} \mu_2(i)} \right)^\alpha (\sum_{\gamma(i)=j} \mu_2(i)) < \sum_{i=1}^{k+1} \left(\frac{\mu_1(i)}{\mu_2(i)} \right)^\alpha \mu_2(i).$$

Since there only finite case of $\gamma: \{0, 1, 2, \dots, k\} \rightarrow \{0, 1, 2, \dots, k - 1\}$, we conclude $\overline{\Delta}_X^k(\mu_1 \parallel \mu_2) < \Delta_{\{0,1,2,\dots,k\}}^f(\mu_1 \parallel \mu_2)$. Since every f -divergence satisfies data-processing inequality, this discussion does not depend on the choice of set Y with $|Y| = k$ in the construction of the k -cut $\overline{\Delta}^k$. Thus, Δ^f is not k -generated for any finite k . \square

Since the mapping $t \mapsto \frac{1}{\alpha-1} \log t$ is strict, we conclude,

Corollary 37. *For any alpha > 1 , the α -Rényi divergence D^α is not k -generated for every finite k .*

B.10 Proof of Theorem 18

Theorem 38 (Theorem 18). *Let $\mu_1, \mu_2 \in \text{Prob}(X)$. $\overline{\Delta}_X^2(\mu_1 \parallel \mu_2) \leq \rho$ holds if and only if for any $\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})$,*

$$(\Pr[\gamma(\mu_1) = \text{Rej}], \Pr[\gamma(\mu_2) = \text{Acc}]) \in R^\Delta(\rho).$$

Proof. We fix a 2-cut $\overline{\Delta}^2$ of a divergence Δ . Suppose that it is defined with a set W satisfying $|W| = 2$.

$$\overline{\Delta}_X^2(\mu_1 \parallel \mu_2) = \sup_{\gamma: X \rightarrow \text{Prob}(W)} \Delta_W(\gamma(\mu_1) \parallel \gamma(\mu_2)).$$

We recall the definition of privacy region

$$R^\Delta(\rho) = \left\{ (x, y) \mid \overline{\Delta}_{\{\text{Acc}, \text{Rej}\}}^2((1-x)\mathbf{d}_{\text{Acc}} + x\mathbf{d}_{\text{Rej}} \parallel y\mathbf{d}_{\text{Acc}} + (1-y)\mathbf{d}_{\text{Rej}}) \leq \rho \right\}.$$

Since every probability distribution $\nu \in \text{Prob}(\{\text{Acc}, \text{Rej}\})$ can be rewritten as $\nu = \Pr[\nu = \text{Acc}]\mathbf{d}_{\text{Acc}} + \Pr[\nu = \text{Rej}]\mathbf{d}_{\text{Rej}}$, we obtain

$$\begin{aligned} \overline{\Delta}_{\{\text{Acc}, \text{Rej}\}}^2(\gamma(\mu_1) \parallel \gamma(\mu_2)) &\leq \rho \\ \iff (\Pr[\gamma(\mu_1) = \text{Rej}], \Pr[\gamma(\mu_2) = \text{Acc}]) &\in R^\Delta(\rho). \end{aligned}$$

Hence, it suffices to show

$$\begin{aligned} (\overline{\Delta}_X^2(\mu_1 \parallel \mu_2) \leq \rho) \\ \iff \forall \gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\}). (\overline{\Delta}_{\{\text{Acc}, \text{Rej}\}}^2(\gamma(\mu_1) \parallel \gamma(\mu_2)) \leq \rho) \end{aligned}$$

- (\implies) Obvious by the data-processing inequality of the 2-cut $\overline{\Delta}^2$.
 (\impliedby) The assumption is equivalent to

$$\forall \gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\}). \forall \gamma': \{\text{Acc}, \text{Rej}\} \rightarrow \text{Prob}(W). \\
 (\Delta_W(\gamma'(\gamma(\mu_1)) \| \gamma'(\gamma(\mu_2))) \leq \rho)$$

Since $|W| = |\{\text{Acc}, \text{Rej}\}| = 2$, this is equivalent to

$$\gamma'': X \rightarrow \text{Prob}(W). \Delta_W(\gamma''(\mu_1) \| \gamma''(\mu_2)) \leq \rho.$$

For any $\gamma: X \rightarrow \text{Prob}(\{\text{Acc}, \text{Rej}\})$. and $\gamma': \{\text{Acc}, \text{Rej}\} \rightarrow \text{Prob}(W)$. we take $\gamma'' = \gamma' \bullet \gamma$.
 Conversely for any $\gamma'': X \rightarrow \text{Prob}(W)$ we take $\gamma = f \bullet \gamma''$ and $\gamma' = f^{-1}$ where
 $f: \{\text{Acc}, \text{Rej}\} \rightarrow W$ is a bijection. \square

B.11 Proof of Theorem 23 in general setting

If the quasi-convex function $F: [0, 1]^{2k} \rightarrow [0, \infty]$ is also continuous, we can extend Theorem 23 to general measurable setting.

Theorem 39 (*k-generatedness in general setting*). *Assume that $F: [0, 1]^{2k} \rightarrow [0, \infty]$ is quasi-convex and continuous. For any measurable space X , we have*

$$\Delta_X(\mu_1, \mu_2) = \sup_{\substack{\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\}) \\ \text{measurable function}}} \Delta_X(\gamma(\mu_1), \gamma(\mu_2)).$$

Proof. We easily calculate as follows (functions are assumed to be measurable):

$$\begin{aligned} & \Delta_X(\mu_1, \mu_2) \\ &= \sup \left\{ F(\mu_1(A_1), \dots, \mu_1(A_k), \mu_2(A_1), \dots, \mu_2(A_k)) \mid \{A_i\}_{i=1}^k: \text{m'ble partition of } X \right\} \\ &= \sup \left\{ F(\mu_1(f^{-1}(1)), \dots, \mu_1(f^{-1}(k)), \mu_2(f^{-1}(1)), \dots, \mu_2(f^{-1}(k))) \mid f: X \rightarrow \{1, 2, \dots, k\} \right\} \\ &= \sup \left\{ F((f(\mu_1))(1), \dots, (f(\mu_1))(k), (f(\mu_2))(1), \dots, (f(\mu_2))(k)) \mid f: X \rightarrow \{1, 2, \dots, k\} \right\} \\ &\leq \sup \left\{ F((\gamma(\mu_1))(1), \dots, (\gamma(\mu_1))(k), (\gamma(\mu_2))(1), \dots, (\gamma(\mu_2))(k)) \mid \gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\}) \right\} \\ &\leq \sup \left\{ F \left(\begin{array}{c} (\gamma(\mu_1))(A_1), \dots, (\gamma(\mu_1))(A_k), \\ (\gamma(\mu_2))(A_1), \dots, (\gamma(\mu_2))(A_k) \end{array} \right) \mid \begin{array}{c} \gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\}), \\ \{A_i\}_{i=1}^k: \text{m'ble partition of } X \end{array} \right\} \\ &= \sup_{\gamma: X \rightarrow \text{Prob}(\{1, 2, \dots, k\})} \Delta_{\{1, 2, \dots, k\}}(\gamma(\mu_1), \gamma(\mu_2)) \end{aligned}$$

Note that we treat $\{1, 2, \dots, k\}$ as a finite discrete space. Consider the family $\{J_n\}_{n=1}^\infty$ of finite sets (discrete spaces) defined as follows:

$$J_n = \left\{ (j_1, \dots, j_k) \mid j_1, \dots, j_k \in \{0, 1, \dots, 2^n - 1\}, C_{j_1 \dots j_k}^n \neq \emptyset \right\}.$$

We fix a measurable function $\gamma: X \rightarrow \text{Prob}(k)$ and treat $\text{Prob}(k)$ as a subset of $[0, 1]^k$. For each $n \in \mathbb{N}$, we define a measurable partition $\{C_{j_1 \dots j_k}^n\}_{j_1, \dots, j_k \in \{0, 1, \dots, 2^n - 1\}}$ of X by

$$\begin{aligned} C_{j_1 \dots j_k}^n &= \gamma^{-1}(B_{j_1 \dots j_k}^n) \\ \text{where } B_{j_1 \dots j_k}^n &= D_{j_1} \times \dots \times D_{j_k} \quad ((j_1 \dots j_k) \in J_n), \\ D_0^n &= \{0\} \text{ and } D_{l+1}^n = (l/2^n, (l+1)/2^n] \quad (l = 0, 1, 2, \dots, 2^n - 1). \end{aligned}$$

We next define $m_n^*: X \rightarrow J_n$ and $m_n: J_n \rightarrow X$ as follows: $m_n^*(x)$ is the unique element $(j_1, \dots, j_k) \in J_n$ satisfying $x \in C_{j_1, \dots, j_k}^n$, and we choose $m_n(j_1, \dots, j_k)$ is an element of C_{j_1, \dots, j_k}^n . Thanks to the measurability of each C_{j_1, \dots, j_k}^n , the function m_n^* is measurable, and the measurability of m_n follows from the discreteness of J_n . From the construction of $\{C_{j_1, \dots, j_k}^n\}_{j_1, \dots, j_k \in \{0, 1, \dots, 2^n - 1\}}$, for any $n \in \mathbb{N}$, $x \in X$, and $i \in I$, we have,

$$|\gamma(x)(i) - (\gamma \circ m_n \circ m_n^*)(x)(i)| \leq 1/2^n$$

This implies that the sequence $\{\gamma \circ m_n \circ m_n^*\}_{n=1}^\infty$ of measurable function converges uniformly to γ . Hence, for any $n \in \mathbb{N}$ and $D \subseteq k$, we have

$$\left| \int \gamma(x)(D) d\mu_1(x) - \int (\gamma \circ m_n \circ m_n^*)(x)(D) d\mu_1(x) \right| \leq 1/2^n$$

Hence the sequence of probability measures $\{(\gamma \circ m_n \circ m_n^*)(\mu_1)\}_{n=1}^\infty$ converges to the probability measure $\gamma(\mu_1)$. Similarly, $\{(\gamma \circ m_n \circ m_n^*)(\mu_2)\}_{n=1}^\infty$ converges to $\gamma(\mu_2)$.

By the continuity of F , we obtain

$$\begin{aligned} & F((\gamma(\mu_1))(A_1), \dots, (\gamma(\mu_1))(A_k), (\gamma(\mu_2))(A_1), \dots, (\gamma(\mu_2))(A_k)) \\ &= \lim_{n \rightarrow \infty} F \left(\begin{array}{c} ((\gamma \circ m_n \circ m_n^*)(\mu_1))(A_1), \dots, ((\gamma \circ m_n \circ m_n^*)(\mu_1))(A_k), \\ ((\gamma \circ m_n \circ m_n^*)(\mu_2))(A_1), \dots, ((\gamma \circ m_n \circ m_n^*)(\mu_2))(A_k) \end{array} \right) \\ &= \lim_{n \rightarrow \infty} F \left(\begin{array}{c} ((\gamma \circ m_n)(m_n^*(\mu_1)))(A_1), \dots, ((\gamma \circ m_n)(m_n^*(\mu_1)))(A_k), \\ ((\gamma \circ m_n)(m_n^*(\mu_2)))(A_1), \dots, ((\gamma \circ m_n)(m_n^*(\mu_2)))(A_k) \end{array} \right) \\ &\leq \sup_{n \in \mathbb{N}} \Delta_{\{1, 2, \dots, k\}}(((\gamma \circ m_n)(m_n^*(\mu_1))), ((\gamma \circ m_n)(m_n^*(\mu_2)))) \\ &\quad \{\text{Since } J_n \text{ is finite (countable and discrete), we can apply Theorem 23.}\} \\ &\leq \sup_{n \in \mathbb{N}} \Delta_{J_n}(m_n^*(\mu_1), m_n^*(\mu_2)) \\ &= \sup_{n \in \mathbb{N}} \sup \left\{ F \left(\begin{array}{c} f(m_n^*(\mu_1))(1), \dots, f(m_n^*(\mu_1))(k), \\ f(m_n^*(\mu_2))(1), \dots, f(m_n^*(\mu_1))(k) \end{array} \right) \middle| f: J_n \rightarrow \{1, 2, \dots, k\} \right\} \\ &\leq \sup \{ F((g(\mu_1))(1), \dots, (g(\mu_1))(k), (g(\mu_2))(1), \dots, (g(\mu_2))(k)) \mid g: X \rightarrow \{1, 2, \dots, k\} \} \\ &= \Delta_X(\mu_1, \mu_2). \end{aligned}$$

This implies $\sup_{\gamma: X \rightarrow \text{Prob}(k)} \Delta_k(\gamma(\mu_1), \gamma(\mu_2)) \leq \Delta_X(\mu_1, \mu_2)$. \square

C Additional Results

C.1 Total variation distance is 2-generated

We recall the definition of the total variation distance

$$\text{TV}_X(\mu_1 \parallel \mu_2) = \sup_{S \subseteq X} |\Pr[\mu_1 \in S] - \Pr[\mu_2 \in S]|.$$

In a similar way as ε -divergence Δ^ε , we can prove 2-generatedness of the total variation distance TV , but we can prove it easily by applying Theorems 16–17 (Theorem 23 in the paper).

Define $F: [0, 1]^4 \rightarrow [0, \infty]$ by $F(x, x', y, y') = |x - y|$. It is easy to check that the function is obviously quasi-convex, and that we have $\text{TV} = \Delta^F$.

C.2 An optimal conversion law from Hellinger to DP

We recall the definition of the Hellinger distance

$$\text{HD}_X(\mu_1 \parallel \mu_2) = 1 - \sum_{x \in X} \sqrt{\mu_1(x)\mu_2(x)}.$$

Since it is the f -divergence of weight function $w(t) = \sqrt{t} - 1$ (strict convex), the Hellinger distance is exactly ∞ -generated, quasi-convex and continuous.

Here is the essence of an optimal conversion law from the Hellinger distance to DP.

Lemma 40. *We have $R^{\text{HD}}(\rho) \subseteq R^{\Delta^\varepsilon}(\delta(\varepsilon, \rho))$ where*

$$\delta(\varepsilon, \rho) = 1 - t - \frac{f(t)}{g(t)} \quad (4)$$

$$t = \frac{z^2 + 4 - z\sqrt{z^2 + 4}}{2(z^2 + 4)} \quad (5)$$

$$z = \frac{1/e^\varepsilon - 2(1 - \rho) + 1}{(1 - \rho)\sqrt{\rho(2 - \rho)}}$$

$$f(x) = (1 - \rho)^2(1 - 2x) + x - 2(1 - \rho)\sqrt{d(2 - d)x(1 - x)}$$

$$g(x) = \frac{df}{dx}(x) = (1 - \rho)^2(1 - 2x) + xf - 2(1 - \rho)\sqrt{d(2 - d)x(1 - x)}$$

Proof. We may regard

$$R^{\text{HD}}(\rho) = \left\{ (x, y) \in [0, 1]^2 \mid 1 - \sqrt{x(1 - y)} - \sqrt{(1 - x)y} \leq \rho \right\},$$

$$R^{\Delta^\varepsilon}(\delta) = \left\{ (x, y) \in [0, 1]^2 \mid \max((1 - x) - e^\varepsilon y, x - e^\varepsilon(1 - y)) \leq \delta \right\}.$$

We first calculate the boundary of $R^{\text{HD}}(\rho)$. Thus, we solve the following equation for y :

$$1 - \sqrt{x(1 - y)} - \sqrt{(1 - x)y} = \rho.$$

We first have

$$1 - \sqrt{x(1 - y)} - \sqrt{(1 - x)y} = \rho$$

$$\iff (1 - \rho)^2 - x(1 - y) - y(1 - x) = 2\sqrt{x(1 - x)y(1 - y)}$$

$$\iff (1 - \rho)^4 + x^2(1 - y)^2 + y^2(1 - x)^2 - 2x(1 - y)(1 - \rho)^2 - 2y(1 - x)(1 - \rho)^2$$

The degree of this equation is 2, so we can solve it. For given $x \in [0, 1]$, we have

$$y = (1 - \rho)^2(1 - 2x) + x \pm 2(1 - \rho)\sqrt{x(1 - x)\rho(2 - \rho)}.$$

Thanks to the Symmetry of $R^{\text{HD}}(\rho)$ and $R^{\Delta^\varepsilon}(\delta)$, we may consider the curve:

$$y = (1 - \rho)^2(1 - 2x) + x - 2(1 - \rho)\sqrt{x(1 - x)\rho(2 - \rho)} = f(x).$$

The tangent of the curve $y = f(x)$ that passes the point $(t, f(t))$ is given by the equation $x - \frac{y}{g(t)} = t - \frac{f(t)}{g(t)}$ where $g(x) = \frac{df}{dx}(x)$. We next find t and δ that the lower boundary

$$(1 - x) - e^\varepsilon y = \delta \iff x + e^\varepsilon y = 1 - \delta$$

of $R^{\Delta^\varepsilon}(\delta(\varepsilon, \rho))$ is the same as the line $x - \frac{y}{g(t)} = t - \frac{f(t)}{g(t)}$. We solve the equation $e^\varepsilon = \frac{1}{g(t)}$ on t about the slope as (5). Finally, we obtain δ as (4). \square

We conclude an optimal conversion law from the Hellinger distance to DP.

Theorem 41. *We always have $\text{HD}_X(d_1, d_2) \leq \rho \implies \Delta_X^\varepsilon(d_1, d_2) \leq \delta(\varepsilon, \rho)$ where $\delta(\varepsilon, \rho)$ is given by (4).*




Figure 3: Comparison of the privacy region for DP and the one for 2-cut of Hellinger distance.

References

- [Abadi et al., 2016] Abadi, M., Chu, A., Goodfellow, I. J., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. (2016). Deep learning with differential privacy. In *ACM SIGSAC Conference on Computer and Communications Security (CCS), Vienna, Austria*, pages 308–318.
- [Abowd, 2018] Abowd, J. M. (2018). The U.S. census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, page 2867.
- [Barthe and Olmedo, 2013] Barthe, G. and Olmedo, F. (2013). Beyond differential privacy: Composition theorems and relational logic for f -divergences between probabilistic programs. In *International Colloquium on Automata, Languages and Programming (ICALP), Riga, Latvia*, volume 7966 of *Lecture Notes in Computer Science*, pages 49–60. Springer-Verlag.
- [Bun et al., 2018] Bun, M., Dwork, C., Rothblum, G. N., and Steinke, T. (2018). Composable and versatile privacy via truncated CDP. In *ACM SIGACT Symposium on Theory of Computing (STOC), Los Angeles, California*.
- [Bun and Steinke, 2016] Bun, M. and Steinke, T. (2016). Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *IACR Theory of Cryptography Conference (TCC), Beijing, China*, volume 9985 of *Lecture Notes in Computer Science*, pages 635–658. Springer-Verlag.
- [Csiszár, 1963] Csiszár, I. (1963). Eine informationstheoretische ungleichung und ihre anwendung auf den beweis der ergodizitat von markoffschen ketten. *Magyar. Tud. Akad. Mat. Kutató Int. Közl*, 8:85–108.

- [Csiszàr and Shields, 2004] Csiszàr, I. and Shields, P. (2004). Information theory and statistics: A tutorial. *Foundations and Trends® in Communications and Information Theory*, 1(4):417–528.
- [Dong et al., 2019] Dong, J., Roth, A., and Su, W. J. (2019). Gaussian Differential Privacy. *arXiv e-prints*.
- [Dwork et al., 2006] Dwork, C., McSherry, F., Nissim, K., and Smith, A. D. (2006). Calibrating noise to sensitivity in private data analysis. In *IACR Theory of Cryptography Conference (TCC), New York, New York*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer-Verlag.
- [Dwork and Rothblum, 2016] Dwork, C. and Rothblum, G. N. (2016). Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*.
- [Erlingsson et al., 2014] Erlingsson, Ú., Pihur, V., and Korolova, A. (2014). RAPPOR: randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 1054–1067.
- [Garfinkel et al., 2018] Garfinkel, S. L., Abowd, J. M., and Powazek, S. (2018). Issues encountered deploying differential privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society, WPES@CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, pages 133–137.
- [Giry, 1982] Giry, M. (1982). A categorical approach to probability theory. In Banaschewski, B., editor, *Categorical Aspects of Topology and Analysis*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85. Springer-Verlag.
- [Johnson et al., 2018] Johnson, N. M., Near, J. P., and Song, D. (2018). Towards practical differential privacy for SQL queries. *PVLDB*, 11(5):526–539.
- [Kairouz et al., 2015] Kairouz, P., Oh, S., and Viswanath, P. (2015). The composition theorem for differential privacy. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 1376–1385.
- [Kenthapadi et al., 2019] Kenthapadi, K., Mironov, I., and Thakurta, A. G. (2019). Privacy-preserving data mining in industry. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11-15, 2019*, pages 840–841.
- [Liese and Vajda, 2006] Liese, F. and Vajda, I. (2006). On divergences and informations in statistics and information theory. *IEEE Transactions on Information Theory*, 52(10):4394–4412.
- [Meiser, 2018] Meiser, S. (2018). Approximate and probabilistic differential privacy definitions. *IACR Cryptology ePrint Archive*, 2018:277.
- [Mironov, 2017] Mironov, I. (2017). Rényi differential privacy. In *IEEE Computer Security Foundations Symposium (CSF), Santa Barbara, California*, pages 263–275.
- [Papernot et al., 2018] Papernot, N., Song, S., Mironov, I., Raghunathan, A., Talwar, K., and Erlingsson, Ú. (2018). Scalable private learning with PATE. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.

- [Renyi, 1961] Renyi, A. (1961). On measures of entropy and information. In *Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 547–561, Berkeley, Calif. University of California Press.
- [team at Apple, 2017] team at Apple, D. P. (2017). Learning with privacy at scale. <https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html>.
- [Wasserman and Zhou, 2010] Wasserman, L. and Zhou, S. (2010). A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389.