# Hypothesis Testing Interpretations and Rényi Differential Privacy

**Borja Balle**[†]
—

**Gilles Barthe**
MPI for Security and Privacy and IMDEA Software Institute

**Marco Gaboardi**
Boston University

**Justin Hsu**
University of Wisconsin–Madison

**Tetsuya Sato**
Seikei University

## Abstract

Differential privacy is a de facto standard in data privacy, with applications in the public and private sectors. One way of explaining differential privacy that is particularly appealing to statistician and social scientists is through 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—no test can 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 distinguishing power of divergences and we use them to study the hypothesis testing interpretation of relaxations of differential privacy based on *Rényi divergence*. Our 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 statistical 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 rigorous, 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 in 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]. This is formalized by imagining 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 $x_0$. The test also considers an *alternative hypothesis* where individual $I$ has not contributed her data to $x_0$. The definition of differential privacy guarantees—and is in fact equivalent to requiring—that any 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

[†] Now at DeepMind.

Proceedings of the 23[rd] International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. 

privacy parameters $(\epsilon, \delta)$ control the tradeoff between significance and power.

Recently, motivated by a desire to improve composition properties of private data analysis algorithms, several variants of differential privacy have been proposed [Dwork and Rothblum, 2016, Bun and Steinke, 2016, Mironov, 2017, Bun et al., 2018]. 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]). 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. While these definitions enjoy better theoretical properties, the use of the Rényi divergence makes the privacy guarantees more difficult to interpret.

In this work, we introduce several concepts to help understand whether these variants of differential privacy can be given a hypothesis testing interpretation. The first notion we introduce is the *k-cut* of a divergence. Intuitively, this corresponds to "projecting" the distributions that are compared by a divergence onto a finite domain of size $k$. From a statistical testing point of view, we can think of these projections as encoding (probabilistic) decision rules with $k$ possible outcomes. The second notion we introduce is the concept of *k-generatedness* for a divergence; 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 a hypothesis testing interpretation in the sense described above if and only if it is 2-generated; we show that the divergence characterizing differential privacy is indeed 2-generated and that 2-generatedness corresponds to the notion of privacy regions introduced in [Kairouz et al., 2015]. On the negative side we show that variants of differential privacy based on the Rényi divergence do not directly admit a hypothesis testing interpretation because the Rényi divergence is $\infty$-generated, meaning that it is infinitely, but countably, generated. Nevertheless, we show that one can obtain a hypothesis testing interpretation by considering the 2-cut of the Rényi divergence. Intuitively, this means that to characterize variants of differential privacy based Rényi divergences through a hypothesis testing experiment, one needs either to restrict the distinguishing power of the divergence or to consider an infinite number of possible actions. This shows a separation between the semantics of 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 $(\epsilon, \delta)$-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 show 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 hypothesis 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 $\infty$-generated, ruling out a direct hypothesis testing interpretation for privacy notions based on it. Nevertheless, we show that one can obtain hypothesis testing interpretations by considering the 2-cut of Rényi divergence.

(4) We use our tools to study other notions of privacy and to give better conversion rules between Rényi differential privacy and $(\epsilon, \delta)$-differential privacy.

(5) We give sufficient and necessary conditions for a quasi-convex divergence to be $k$-generated.

**Related work.** Several works have studied the semantics of formal notions of data privacy and differential privacy [Dwork, 2006, Wasserman and Zhou, 2010, Kifer and Machanavajjhala, 2011, Dwork and Roth, 2013, Hsu et al., 2014, Kifer and Machanavajjhala, 2014, Kasiviswanathan and Smith, 2015, Liu et al., 2019]. The hypothesis testing interpretation of differential privacy was first introduced by [Wasserman and Zhou, 2010] and then used to develop the optimal composition theorem for differential privacy [Kairouz et al., 2015]. [Dong et al., 2019] propose to define new notions of privacy based on the hypothesis testing interpretation; our work supports this direction, showing that other existing variants of privacy do not enjoy a hypothesis testing interpretation. [Liu et al., 2019] use the hypothesis testing interpretation to reason about the privacy parameters. The hypothesis testing interpretation of differential privacy has also inspired techniques in formal verification [Sato, 2016, Sato et al., 2017], including techniques to detect violations in differentially private implementations [Ding et al., 2018].

## 2 Background: hypothesis testing, privacy, and Rényi divergences

### 2.1 Hypothesis testing interpretation for $(\varepsilon, \delta)$-differential privacy

We view *randomized algorithms* as functions $\mathcal{M}\colon X \to \mathrm{Prob}(Y)$ from a set $X$ of inputs to the set $\mathrm{Prob}(Y)$ 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 $\varepsilon > 0$ and $0 \le \delta \le 1$. A randomized algorithm $\mathcal{M}\colon X \to \mathrm{Prob}(Y)$ is $(\varepsilon, \delta)$-*differentially private *if for every pairs of adjacent inputs $x_0$ and $x_1$, and every subset $S \subseteq Y$, we have:*

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

[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:

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

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 $\mathtt{PFA}(x_0, x_1, \mathcal{M}, S) \stackrel{\mathsf{def}}{=} \mathsf{Pr}[\mathcal{M}(x_0) \in S]$. 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 $\mathtt{PMD}(x_0, x_1, \mathcal{M}, S) \stackrel{\mathsf{def}}{=} \mathsf{Pr}[\mathcal{M}(x_1) \notin S]$. 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 $\mathcal{M}\colon X \to \mathrm{Prob}(Y)$ is $(\varepsilon, \delta)$-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: $\mathtt{PFA}(x_0, x_1, \mathcal{M}, S) + e^{\varepsilon}\mathtt{PMD}(x_0, x_1, \mathcal{M}, S) \ge 1 - \delta$ and $e^{\varepsilon}\mathtt{PFA}(x_0, x_1, \mathcal{M}, S) + \mathtt{PMD}(x_0, x_1, \mathcal{M}, S) \ge 1 - \delta$.*

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 $\mathcal{M}\colon X \to \mathrm{Prob}(Y)$ is $(\varepsilon, \delta)$-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*

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

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

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

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 $(\mathtt{PFA}, \mathtt{PMD})$ that can be generated by a randomized response mechanism $\mathcal{M}_{\mathtt{RR}}\colon \{0,1\}^3 \to \mathrm{Prob}(\{0,1\}^3)$ working on vectors of three bits and flipping each bit with probability 0.34.

After the introduction of differential privacy, researchers have proposed alternative definitions 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?

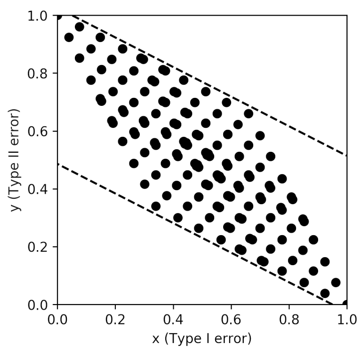


Figure 1: Pairs (PFA, PMD) of $\mathcal{M}_{\mathtt{RR}}$ and $R(0.67, 0.05)$

### 2.2 Variants of differential privacy based on Rényi divergence

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

**Definition 4** (Rényi divergence [Renyi, 1961]). *Let $\alpha > 1$. The* Rényi divergence *of order $\alpha$ between two probability distributions $\mu_1$ and $\mu_2$ on a space $X$ is defined by:*

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

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

$$D_X^1(\mu_1||\mu_2) \stackrel{\mathsf{def}}{=} \lim_{\alpha \to 1+} D_X^\alpha(\mu_1||\mu_2) = \mathbf{KL}_X(\mu_1||\mu_2),$$

$$D_X^\infty(\mu_1||\mu_2) \stackrel{\mathsf{def}}{=} \lim_{\alpha \to \infty} D_X^\alpha(\mu_1||\mu_2) = \log \sup_x \frac{\mu_1(x)}{\mu_2(x)}.$$

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 $(\varepsilon, 0)$-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 $\alpha$ 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 $\mathcal{M} : X \to$* Prob$(Y)$ *is* $(\alpha, \rho)$-Rényi differentially private *if for every pair $x_0$ and $x_1$ of adjacent inputs, we have*

$$D_Y^\alpha(\mathcal{M}(x_0)||\mathcal{M}(x_1)) \leq \rho.$$

Renyi Differential privacy considers a fixed value of $\alpha$. 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 $\mathcal{M} : X \to$* Prob$(Y)$ *is* $(\xi, \rho)$-zero concentrated differentially private *if for every pairs of adjacent inputs $x_0$ and $x_1$, we have*

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

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

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

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

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

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

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 $\mathcal{L} = \mathcal{L}^{x_0 \to x_1}$ |
|---|---|
| $(\varepsilon, \delta)$-DP | $\mathsf{Pr}_{\mathcal{M}(x_0)}[\mathcal{L}(y) \leq e^\varepsilon] \geq 1 - \delta$ |
| $(\alpha, \rho)$-RDP | $\mathbb{E}_{\mathcal{M}(x_1)}[\mathcal{L}(y)^\alpha] \leq e^{(\alpha-1)\rho}$ |
| $(\xi, \rho)$-zCDP | $\forall \alpha \in (1, \infty).\ \mathbb{E}_{\mathcal{M}(x_1)}[\mathcal{L}(y)^\alpha] \leq e^{(\alpha-1)(\xi+\alpha\rho)}$ |
| $(\omega, \rho)$-tCDP | $\forall \alpha \in (1, \omega).\ \mathbb{E}_{\mathcal{M}(x_1)}[\mathcal{L}(y)^\alpha] \leq e^{(\alpha-1)\alpha\rho}$ |

In particular, DP bounds the maximum value of the privacy loss,[1] $(\alpha, \cdot)$-RDP bounds the $\alpha$-moment, zCDP bounds all moments, and $(\cdot, \omega)$-tCDP bounds the moments up to some cutoff $\omega$. Many conversions are known between these definitions; for instance, RDP,

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

zCDP, and tCDP are known to sit between $(\varepsilon, 0)$ and $(\varepsilon, \delta)$-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 [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 $\mathbf{d}_x(x') = 0$ otherwise. For any probability distribution $\mu \in \mathrm{Prob}(X)$ and $\gamma\colon X \to \mathrm{Prob}(Y)$, we define $\gamma(\mu) \in \mathrm{Prob}(Y)$ to be $(\gamma(\mu))(y) = \sum_{x\in X}(\gamma(x))(y) \cdot \mu(x)$ for every $y \in Y$. For any function $\gamma\colon X \to Y$, as an abuse of notation, we define $\gamma(\mu) \in \mathrm{Prob}(Y)$ 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 $\gamma\colon X \to Y$ can be regarded as $\{x \mapsto \mathbf{d}_\gamma(x)\}\colon X \to \mathrm{Prob}(Y)$.

## 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 (indexed by all sets):*

$$\Delta_X\colon \mathrm{Prob}(X) \times \mathrm{Prob}(X) \to [0, \infty].$$

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

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

$$\Delta^{\varepsilon}_X(\mu_1||\mu_2) \stackrel{\mathsf{def}}{=} \sup_{S\subseteq X} (\mathsf{Pr}[\mu_1 \in S] - e^{\varepsilon}\mathsf{Pr}[\mu_2 \in S]).$$

Specifically, a randomized algorithm $\mathcal{M}\colon X \to \mathrm{Prob}(Y)$ is $(\varepsilon, \delta)$-differentially private if and only if for every pair of adjacent inputs $x_0$ and $x_1$, we have

$$\Delta^{\varepsilon}_Y(\mathcal{M}(x_0)||\mathcal{M}(x_1)) \le \delta.$$

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

- A divergence $\Delta$ satisfies the *data-processing inequality* iff for every $\gamma\colon X \to \mathrm{Prob}(Y)$, $\Delta_Y(\gamma(\mu_1)||\gamma(\mu_2)) \le \Delta_X(\mu_1||\mu_2)$.
- A divergence $\Delta$ 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 $X$,

$$\Delta_X(\sum_m \alpha_m d_{1,m}|| \sum_m \alpha_m d_{2,m}) \le \max_m \Delta_X(d_{1,m}||d_{2,m}).$$

These properties are satisfied by many common divergences. Besides Rényi divergences, they also hold for all $f$-divergences [Csiszár, 1963, Csiszàr 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 $k \in \mathbb{N} \cup \{\infty\}$. We define* a $k$-cut $\overline{\Delta}^k = \{\overline{\Delta}^k_X\}_{X:\ \mathrm{set}}$ *of a divergence $\Delta$ as follows: fix a set $Y$ with cardinality $k$ (i.e. $|Y| = k$), and define*

$$\overline{\Delta}^k_X(\mu_1||\mu_2) \stackrel{\mathsf{def}}{=} \sup_{\gamma\colon X\to\mathrm{Prob}(Y)} \Delta_Y(\gamma(\mu_1)||\gamma(\mu_2)).$$

For divergences $\Delta$ 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 $\Delta$ satisfies the data-processing inequality, we have the inequality $\overline{\Delta}^k \le \Delta$ and the equality $\overline{\Delta}^k_Y = \Delta_Y$ for any set $Y$ with $|Y| = k$.*

Hence, 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 $\Delta$ satisfies the data-processing inequality, even if the original $\Delta$ does not satisfy it.

Without loss of generality, we can assume the function $\gamma$ 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 $Y$ with $|X| = k$ and $|Y| = l$. For any $\gamma\colon X \to \mathrm{Prob}(Y)$, there exist $N \in \mathbb{N}$, $\gamma_1, \dots, \gamma_N\colon X \to Y$ 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 result allows 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 $\alpha$. These can be reformulated as follows:

$$\overline{D^\alpha}^2_X(\mu_1||\mu_2) = \sup_{S\subseteq X} \frac{1}{\alpha-1}\log\left\{\begin{array}{l}\mathsf{Pr}[\mu_1\in S]^\alpha \mathsf{Pr}[\mu_2\in S]^{1-\alpha}\\ +\mathsf{Pr}[\mu_1\notin S]^\alpha \mathsf{Pr}[\mu_2\notin S]^{1-\alpha}\end{array}\right\},$$

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

### 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 $\Delta$ is* $k$-generated *if a $k$-cut $\overline{\Delta}^k$ of $\Delta$ is equal to $\Delta$ 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*

$$\Delta_X(\mu_1||\mu_2) = \sup_{\gamma\colon X\to\mathrm{Prob}(Y)} \Delta_Y(\gamma(\mu_1)||\gamma(\mu_2)).$$

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

- *If $\Delta$ is 1-generated, then $\Delta$ is constant, i.e. there exists $c\in[0,\infty]$ such that for every $X$ and every $\mu_1,\mu_2\in\mathrm{Prob}(X)$, we have $\Delta_X(\mu_1||\mu_2)=c$.*
- *If $\Delta$ is $k$-generated, then it is also $k+1$-generated.*
- *If $\Delta$ has the data-processing inequality, then it is at least $\infty$-generated.*
- *Every $k$-cut of a divergence $\Delta$ 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 $\Delta$ and a $k$-generated divergence $\Delta'$. For any $k$-cut $\overline{\Delta}^k$ of $\Delta$,*

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

*Also, if $\Delta$ has the data-processing inequality, the $k$-cut is the greatest $k$-generated divergence below $\Delta$:*

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

#### 3.4.1 A 2-generated divergence for DP

The divergence $\Delta^\varepsilon$ that can be used to characterize $(\varepsilon,\delta)$-DP is 2-generated. This implies that DP can be characterized completely by its hypothesis testing interpretation.

**Theorem 16.** *The $\varepsilon$-divergence $\Delta^\varepsilon$ is 2-generated.*

Since $\Delta^\varepsilon$ is quasi-convex and satisfies data-processing inequality, the 2-cut can be reformulated as:

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

This is exactly the same as the original definition of $\Delta^\varepsilon$, from which follows that it is 2-generated.

#### 3.4.2 Rényi divergences are $\infty$-generated

In contrast to the divergence $\Delta^\varepsilon$, 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 $\mu_1,\mu_2\in\mathrm{Prob}(X)$ 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 $\mu_2(b)=\frac{1}{p^2+p+1}$.

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

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

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 $\infty$-generated. Indeed, Rényi divergence satisfies the data-processing inequality, hence it is at most $\infty$-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 $\alpha$ given by $\exp((\alpha-1)D^\alpha_X(\mu_1||\mu_2))$ 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 $\infty$-generated. The formal details can be found in the supplementary material.

## 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* $\Delta$*, we define its* privacy region $R^{\Delta}(\rho) \subseteq [0,1] \times [0,1]$ *by*

$$R^{\Delta}(\rho) \stackrel{\text{def}}{=} \left\{ (x,y) \;\middle|\; \begin{array}{l} \overline{\Delta}^{2}_{\{\mathtt{Acc},\mathtt{Rej}\}}(\mu_1 || \mu_2) \leq \rho, \\ \mu_1 = (1-x)\mathbf{d}_{\mathtt{Acc}} + x\mathbf{d}_{\mathtt{Rej}} \\ \mu_2 = y\mathbf{d}_{\mathtt{Acc}} + (1-y)\mathbf{d}_{\mathtt{Rej}} \end{array} \right\}.$$

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

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

$$R^{\Delta^{\varepsilon}}(\delta) = \left\{ (x,y) \;\middle|\; \begin{array}{l} 1-x \leq e^{\varepsilon}y + \delta \\ x \leq e^{\varepsilon}(1-y) + \delta \end{array} \right\}$$

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

**Theorem 18.** *Let* $\mu_1, \mu_2 \in \mathrm{Prob}(X)$*.* $\overline{\Delta}^{2}_{X}(\mu_1 || \mu_2) \leq \rho$ *holds if and only if for any* $\gamma\colon X \to \mathrm{Prob}(\{\mathtt{Acc},\mathtt{Rej}\})$*,*

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

In the theorem above, the functions $\gamma\colon X \to \mathrm{Prob}(\{\mathtt{Acc},\mathtt{Rej}\})$ 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 $\gamma\colon X \to \mathrm{Prob}(\{\mathtt{Acc},\mathtt{Rej}\})$, we can take its negation $\neg\gamma$. Hence we do not need to check the cases of $\Pr[\gamma(\mu_1) = \mathtt{Rej}] + \Pr[\gamma(\mu_2) = \mathtt{Acc}] > 1$. This also corresponds to the symmetry that we have in their graphical representations.

Finally, if a divergence $\Delta$ is quasi-convex, we also have the equivalent of Theorem 18 under deterministic decision rules. In this case, we have the following reformulation. Let $\mu_1, \mu_2 \in \mathrm{Prob}(X)$. $\overline{\Delta}^{2}_{X}(\mu_1 || \mu_2) \leq \rho$ 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 $\varepsilon$-divergence is 2-generated and quasi-convex.

**Corollary 19.** *Let* $\mu_1, \mu_2 \in \mathrm{Prob}(X)$*. Set* $\varepsilon, \delta \geq 0$*.* $\Delta^{\varepsilon}_{X}(\mu_1 || \mu_2) \leq \delta$ *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 $\Delta' = \{\Delta'_X\}_{X\colon \mathrm{set}}$ be a divergence satisfying the data-processing inequality. We want to find the minimal parameters $(\varepsilon(\rho), \delta(\rho))$ such that $\Delta'_X(\mu_1 || \mu_2) \leq \rho$ implies $\Delta^{\varepsilon(\rho)}_X(\mu_1 || \mu_2) \leq \delta(\rho)$.

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

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

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

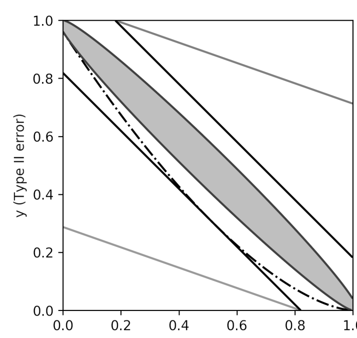


Figure 2: A refined conversion law from RDP to DP. The gray region is $R^{D^{\alpha}}(\rho)$. The gray and black lines show original and refined DP-bounds for the same $\delta$.

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

**Theorem 20.** *If a mechanism $\mathcal{M}$ is $(\alpha, \rho)$-RDP then it is $(\rho + \log((\alpha-1)/\alpha) - (\log\delta + \log\alpha)/(\alpha-1), \delta)$-DP for any $0 < \delta < 1$.*

The privacy region $R^{D^\alpha}(\rho)$ of Rényi divergence is:

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

Inspired from Mironov's original proof, we compute

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

The inequality (‡) gives Mironov's original conversion law. However, starting from (†), we have a better bound for DP as follows: consider a curve $C$ defined by the equation $1-x = (e^\rho y)^{\frac{\alpha-1}{\alpha}}$. We can find parameters that the line $(1-x) = e^\varepsilon y + \delta$ meets a tangent of $C$. Simple computations give $\varepsilon^* = \log\left(\frac{\alpha-1}{\alpha}\right) + \rho - \frac{\log\delta + \log\alpha}{\alpha-1}$. We then have $e^{\varepsilon^*} y + \delta \le (e^\rho y)^{\frac{\alpha-1}{\alpha}}$. From this, by the symmetries of geometrical presentations of $R^{D^\alpha}(\rho)$ and $R^{\Delta^\varepsilon}(\delta)$, we obtain $R^{D^\alpha}(\rho) \subseteq R^{\Delta^\varepsilon}(\delta)$. This is equivalent to: $D^\alpha_X(\mu_1||\mu_2) \le \rho \implies \Delta^{\varepsilon^*}_X(\mu_1||\mu_2) \le \delta$.

### 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? Indeed, we can characterize GDP by the following divergence:

$$\Delta^{\text{Gauss}}_X(\mu_1||\mu_2) = \sup\left\{ \delta \;\middle|\; \begin{array}{l} \exists\gamma\colon X \to \text{Prob}(\{\texttt{Acc}, \texttt{Rej}\}). \\ \quad \Pr[\gamma(\mu_2) = \texttt{Acc}] \\ \quad\quad \ge \Phi(\Phi^{-1}(\Pr[\gamma(\mu_1) = \texttt{Rej}]) - \delta) \end{array} \right\}.$$

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

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

By Theorem 18, $\Delta^{\text{Gauss}}$ 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 21.** *We say that a divergence $\Delta$ is $\delta$-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 $\Delta$ satisfying the data-processing inequality. Then the $k$-cuts form a monotone increasing sequence: $\overline{\Delta}^1 \le \overline{\Delta}^2 \le \overline{\Delta}^3 \le \cdots \le \overline{\Delta}^k \le \overline{\Delta}^{k+1} \le \cdots$. Thus for any divergence with data-processing inequality, $k+1$-cut of $\Delta$ is always more informative than the $k$-cut of $\Delta$ for every $k \in \mathbb{N}$ in the following sense. If the $k$-cut of a divergence is $\delta$-distinguishing a pair $\mu_1, \mu_2 \in \text{Prob}(X)$ then the $k+1$ cut is $\delta$-distinguishing them too:

$$\overline{\Delta}^{k+1}_X(\mu_1||\mu_2) \ge \overline{\Delta}^k_X(\mu_1||\mu_2) > \delta.$$

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 $\delta$ such that $\overline{D^\alpha}^3$ is $\delta$-distinguishing $\mu_1, \mu_2$ and $\overline{D^\alpha}^2$ is *not* $\delta$-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 show 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.

**Theorem 22.** *Let $F\colon [0,1]^{2k} \to [0,\infty]$ be a quasi-convex function. Then the divergence $\Delta^F$ defined below is $k$-generated and quasi-convex.*

$$\Delta^F_X(\mu_1||\mu_2) \stackrel{\text{def}}{=} \sup_{\substack{\{A_i\}_{i=1}^k \\ \textit{partition of } X}} F\begin{pmatrix} \mu_1(A_1), \cdots, \mu_1(A_k), \\ \mu_2(A_1), \cdots, \mu_2(A_k) \end{pmatrix}.$$

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 other statistical divergences. We introduced the notions of $k$-cut and $k$-generatedness for divergences. These notions quantify 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.

## Acknowledgements

This work was partially funded by the University of Wisconsin–Madison, a Facebook TAV Award, and by the National Science Foundation under CNS Grant No. 1565365. Sato was supported by JST ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603).

### 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.

[Ding et al., 2018] Ding, Z., Wang, Y., Wang, G., Zhang, D., and Kifer, D. (2018). Detecting violations of differential privacy. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, pages 475–489.

[Dong et al., 2019] Dong, J., Roth, A., and Su, W. J. (2019). Gaussian Differential Privacy. *arXiv e-prints*.

[Dwork, 2006] Dwork, C. (2006). Differential privacy. In *International Colloquium on Automata, Languages and Programming (ICALP), Venice, Italy*, pages 1–12.

[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 Roth, 2013] Dwork, C. and Roth, A. (2013). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4).

[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.

[Hsu et al., 2014] Hsu, J., Gaboardi, M., Haeberlen, A., Khanna, S., Narayan, A., Pierce, B. C., and Roth, A. (2014). Differential privacy: An economic method for choosing epsilon. In *IEEE 27th Computer Security Foundations Symposium, CSF 2014, Vienna, Austria, 19-22 July, 2014*, pages 398–410.

[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.

[Kasiviswanathan and Smith, 2015] Kasiviswanathan, S. P. and Smith, A. D. (2015). On the 'semantics' of differential privacy: A bayesian formulation. *CoRR*, abs/0803.3946.

[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.

[Kifer and Machanavajjhala, 2011] Kifer, D. and Machanavajjhala, A. (2011). No free lunch in data privacy. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, pages 193–204.

[Kifer and Machanavajjhala, 2014] Kifer, D. and Machanavajjhala, A. (2014). Pufferfish: A framework for mathematical privacy definitions. *ACM Trans. Database Syst.*, 39(1):3:1–3:36.

[Liu et al., 2019] Liu, C., Chanyaswad, T., Wang, S., and Mittal, P. (2019). Investigating statistical privacy frameworks from the perspective of hypothesis testing. In *Proceedings on Privacy Enhancing Technologies*, volume 2019, pages 233–254.

[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.

[Sato, 2016] Sato, T. (2016). Approximate relational Hoare logic for continuous random samplings. *Electronic Notes in Theoretical Computer Science*, 325:277–298. Conference on the Mathematical Foundations of Programming Semantics (MFPS), Pittsburgh, Pennsylvania.

[Sato et al., 2017] Sato, T., Barthe, G., Gaboardi, M., Hsu, J., and Katsumata, S. (2017). Approximate span liftings: Compositional semantics for relaxations of differential privacy. *CoRR*, abs/1710.09010. (The extended version of this paper).

[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.