

Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences

Borja Balle^{*1}, Gilles Barthe², and Marco Gaboardi³

¹*Amazon Research, Cambridge, UK*

²*IMDEA Software Institute, Madrid, Spain*

³*University at Buffalo (SUNY), Buffalo, USA*

Abstract

Differential privacy comes equipped with multiple analytical tools for the design of private data analyses. One important tool is the so-called “privacy amplification by subsampling” principle, which ensures that a differentially private mechanism run on a random subsample of a population provides higher privacy guarantees than when run on the entire population. Several instances of this principle have been studied for different random subsampling methods, each with an ad-hoc analysis. In this paper we present a general method that recovers and improves prior analyses, yields lower bounds and derives new instances of privacy amplification by subsampling. Our method leverages a characterization of differential privacy as a divergence which emerged in the program verification community. Furthermore, it introduces new tools, including advanced joint convexity and privacy profiles, which might be of independent interest.

1 Introduction

Subsampling is a fundamental tool in the design and analysis of differentially private mechanisms. Broadly speaking, the intuition behind the “privacy amplification by subsampling” principle is that the privacy guarantees of a differentially private mechanism can be amplified by applying it to a small random subsample of records from a given dataset. In machine learning, many classes of algorithms involve sampling operations, e.g. stochastic optimization methods and Bayesian inference algorithms, and it is not surprising that results quantifying the privacy amplification obtained via subsampling play a key role in designing differentially private versions of these learning algorithms [Bassily et al., 2014, Wang et al., 2015, Abadi et al., 2016, Jälkö et al., 2017, Park et al., 2016b,a]. Additionally, from a practical standpoint subsampling provides a straightforward method to obtain privacy amplification when the final mechanism is only available as a black-box. For example, in Apple’s iOS and Google’s Chrome deployments of differential privacy for data collection the privacy parameters are hard-coded into the implementation and cannot be modified by the user. In this type of settings, if the default privacy parameters are not satisfactory one could achieve a stronger privacy guarantee by devising a strategy that only submits to the mechanism a random sample of the data.

Despite the practical importance of subsampling, existing tools to bound privacy amplification only work for specific forms of subsampling and typically come with cumbersome proofs providing

*Corresponding e-mail: pigem@amazon.co.uk

no information about the tightness of the resulting bounds. In this paper we remedy this situation by providing a general framework for deriving tight privacy amplification results that can be applied to any of the subsampling strategies considered in the literature. Our framework builds on a characterization of differential privacy in terms of α -divergences [Barthe and Olmedo, 2013]. This characterization has been used before for program verification [Barthe et al., 2012, 2016], while we use it here for the first time in the context of algorithm analysis. In order to do this, we develop several novel analytical tools, including *advanced joint convexity* – a property of α -divergence with respect to mixture distributions – and *privacy profiles* – a general tool describing the privacy guarantees that private algorithms provide.

One of our motivations to initiate a systematic study of privacy amplification by subsampling is that this is an important primitive for the design of differentially private algorithms which has received less attention than other building blocks like composition theorems [Dwork et al., 2010, Kairouz et al., 2017, Murtagh and Vadhan, 2016]. Given the relevance of sampling operations in machine learning, it is important to understand what are the limitations of privacy amplification and develop a fine-grained understanding of its theoretical properties. Our results provide a first step in this direction by showing how privacy amplification resulting from different sampling techniques can be analyzed by means of single set of tools, and by showing how these tools can be used for proving lower bounds. Our analyses also highlight the importance of choosing a sampling technique that is well-adapted to the notion of neighbouring datasets under consideration. A second motivation is that subsampling provides a natural example of mechanisms where the output distribution is a mixture. Because mixtures have an additive structure and differential privacy is defined in terms of a multiplicative guarantee, analyzing the privacy guarantees of mechanisms whose output distribution is a mixture is in general a challenging task. Although our analyses are specialized to mixtures arising from subsampling, we believe the tools we develop in terms of couplings and divergences will also be useful to analyze other types of mechanisms involving mixture distributions. Finally, we want to remark that privacy amplification results also play a role in analyzing the generalization and sample complexity properties of private learning algorithms [Kasiviswanathan et al., 2011, Beimel et al., 2013, Bun et al., 2015, Wang et al., 2016]; an in-depth understanding of the interplay between sampling and differential privacy might also have applications in this direction.

2 Problem Statement and Methodology Overview

A *mechanism* $\mathcal{M} : X \rightarrow \mathbb{P}(Z)$ with input space X and output space Z is a randomized algorithm that on input x outputs a sample from the distribution $\mathcal{M}(x)$ over Z . Here $\mathbb{P}(Z)$ denotes the set of probability measures on the output space Z . We implicitly assume Z is equipped with a sigma-algebra of measurable subsets and a base measure, in which case $\mathbb{P}(Z)$ is restricted to probability measures that are absolutely continuous with respect to the base measure. In most cases of interest Z is either a discrete space equipped with the counting measure or an Euclidean space equipped with the Lebesgue measure. We also assume X is equipped with a binary symmetric relation \simeq_X defining the notion of neighbouring inputs.

Let $\varepsilon \geq 0$ and $\delta \in [0, 1]$. A mechanism \mathcal{M} is said to be (ε, δ) -*differentially private* w.r.t. \simeq_X if for every pair of inputs $x \simeq_X x'$ and every measurable subset $E \subseteq Z$ we have

$$\Pr[\mathcal{M}(x) \in E] \leq e^\varepsilon \Pr[\mathcal{M}(x') \in E] + \delta . \tag{1}$$

For our purposes, it will be more convenient to express differential privacy in terms of α -divergences¹. Concretely, the α -divergence ($\alpha \geq 1$) between two probability measures $\mu, \mu' \in \mathbb{P}(Z)$ is defined as²

$$D_\alpha(\mu \parallel \mu') = \sup_E (\mu(E) - \alpha \mu'(E)) = \int_Z \left[\frac{d\mu}{d\mu'}(z) - \alpha \right]_+ d\mu'(z) = \sum_{z \in Z} [\mu(z) - \alpha \mu'(z)]_+ , \quad (2)$$

where E ranges over all measurable subsets of Z , $[\bullet]_+ = \max\{\bullet, 0\}$, and the last equality is a specialization for discrete Z . It is easy to see [Barthe and Olmedo, 2013] that \mathcal{M} is (ε, δ) -differentially private if and only if $D_{e^\varepsilon}(\mathcal{M}(x) \parallel \mathcal{M}(x')) \leq \delta$ for every x and x' such that $x \simeq_X x'$.

In order to emphasize the relevant properties of \mathcal{M} from a privacy amplification point of view, we introduce the concepts of *privacy profile* and *group-privacy profiles*. The privacy profile $\delta_{\mathcal{M}}$ of a mechanism \mathcal{M} is a function associating to each privacy parameter $\alpha = e^\varepsilon$ a bound on the α -divergence between the results of running the mechanism on two adjacent datasets, i.e. $\delta_{\mathcal{M}}(\varepsilon) = \sup_{x \simeq_X x'} D_{e^\varepsilon}(\mathcal{M}(x) \parallel \mathcal{M}(x'))$ (we will discuss the properties of this tool in more details in the next section). Informally speaking, the privacy profile represents the set of all of privacy parameters under which a mechanism provides differential privacy. In particular, recall that an (ε, δ) -DP mechanism \mathcal{M} is also (ε', δ') -DP for any $\varepsilon' \geq \varepsilon$ and any $\delta' \geq \delta$. The privacy profile $\delta_{\mathcal{M}}$ defines a curve in $[0, \infty) \times [0, 1]$ that separates the space of privacy parameters into two regions: the ones for which \mathcal{M} satisfies differential privacy and the ones for which it does not. This curve exists for every mechanism \mathcal{M} , even for mechanisms that satisfy pure DP for some value of ε . When the mechanism is clear from the context we might slightly abuse our notation and write $\delta(\varepsilon)$ or δ for the corresponding privacy profile. To define group-privacy profiles $\delta_{\mathcal{M},k}$ ($k \geq 1$) we use the path-distance d induced by \simeq_X :

$$d(x, x') = \min\{k : \exists x_1, \dots, x_{k-1}, x \simeq_X x_1, x_1 \simeq_X x_2, \dots, x_{k-1} \simeq_X x'\} .$$

With this notation, we define $\delta_{\mathcal{M},k}(\varepsilon) = \sup_{d(x,x') \leq k} D_{e^\varepsilon}(\mathcal{M}(x) \parallel \mathcal{M}(x'))$. Note that $\delta_{\mathcal{M}} = \delta_{\mathcal{M},1}$.

Problem Statement A well-known method for increasing privacy of a mechanism is to apply the mechanism to a random subsample of the input database, rather than on the database itself. Intuitively, the method decreases the chances of leaking information about a particular individual because nothing about that individual can be leaked in the cases where the individual is not included in the subsample. The question addressed in this paper is to devise methods for quantifying amplification and for proving optimality of the bounds. This turns out to be a surprisingly subtle problem.

Formally, let X and Y be two sets equipped with neighbouring relations \simeq_X and \simeq_Y respectively. We assume that both X and Y contain databases (modelled as sets, multisets, or tuples) over a universe \mathcal{U} that represents all possible records contained in a database. A subsampling mechanism is a randomized algorithm $\mathcal{S} : X \rightarrow \mathbb{P}(Y)$ that takes as input a database x and outputs a finitely supported distribution over datasets. Note that we find it convenient to distinguish between X and Y because x and y might not always have the same type. For example, sampling with replacement from a set x yields a multiset y .

¹Also known in the literature as *elementary divergences* [Österreicher, 2002] and *hockey-stick divergences* [Sason and Verdú, 2016].

²Here $d\mu/d\mu'$ denotes the Radon-Nikodym derivative between μ and μ' . In particular, if μ and μ' have densities $p = d\mu/d\nu$ and $p' = d\mu'/d\nu$ with respect to some base measure ν , then $d\mu/d\mu' = p/p'$.

Subsampling	\simeq_Y	\simeq_X	η	δ'	Theorem
Poisson(γ)	R	R	γ	$\gamma\delta$	8
WOR(n, m)	S	S	$\frac{m}{n}$	$\frac{m}{n}\delta$	9
WR(n, m)	S	S	$1 - (1 - \frac{1}{n})^m$	$\sum_{k=1}^m \binom{m}{k} (\frac{1}{n})^k (1 - \frac{1}{n})^{m-k} \delta_k$	10
WR(n, m)	S	R	$1 - (1 - \frac{1}{n})^m$	$\sum_{k=1}^m \binom{m}{k} (\frac{1}{n})^k (1 - \frac{1}{n})^{m-k} \delta_k$	11

Table 1: Summary of privacy amplification bounds. Amplification parameter η : $e^{\varepsilon'} = 1 + \eta(e^\varepsilon - 1)$. Types of subsampling: without replacement (WOR) and with replacement (WR). Neighbouring relations: remove/add-one (R) and substitute one (S).

The problem of privacy amplification can now be stated as follows: let $\mathcal{M} : Y \rightarrow \mathbb{P}(Z)$ be a mechanism with privacy profile $\delta_{\mathcal{M}}$ with respect to \simeq_Y , and let \mathcal{S} be a subsampling mechanism. Consider the subsampled mechanism $\mathcal{M}^{\mathcal{S}} : X \rightarrow \mathbb{P}(Z)$ given by $\mathcal{M}^{\mathcal{S}}(x) = \mathcal{M}(\mathcal{S}(x))$, where the composition notation means we feed a sample from $\mathcal{S}(x)$ into \mathcal{M} . The goal is to relate the privacy profiles of \mathcal{M} and $\mathcal{M}^{\mathcal{S}}$, via an inequality of the form: for every $\varepsilon \geq 0$, there exists $0 \leq \varepsilon' \leq \varepsilon$ such that $\delta_{\mathcal{M}^{\mathcal{S}}}(\varepsilon') \leq h(\delta_{\mathcal{M}}(\varepsilon))$, where h is some function to be determined. In terms of differential privacy, one can be read as saying that if \mathcal{M} is (ε, δ) -DP, then the subsampled mechanism $\mathcal{M}^{\mathcal{S}}$ is $(\varepsilon', h(\delta))$ -DP for some $\varepsilon' \leq \varepsilon$. This is a privacy amplification statement because the new mechanism has better privacy parameters than the original one.

A full specification of this problem requires formalizing the following three ingredients: (i) *dataset representation* specifying whether the inputs to the mechanism are sets, multisets, or tuples; (ii) *neighbouring relations* in X and Y , including the usual remove/add-one \simeq_r and substitute-one \simeq_s relations; (iii) *subsampling method* and its parameters, with the most commonly used being subsample without replacement, subsampling with replacement, and Poisson subsampling.

Regardless of the specific setting being considered, the main challenge in the analysis of privacy amplification by subsampling resides in the fact that the output distribution of the mechanism $\mu = \mathcal{M}^{\mathcal{S}}(x) \in \mathbb{P}(Z)$ is a *mixture distribution*. In particular, writing $\mu_y = \mathcal{M}(y) \in \mathbb{P}(Z)$ for any $y \in Y$ and taking $\omega = \mathcal{S}(x) \in \mathbb{P}(Y)$ to be the (finitely supported) distribution over subsamples from x produced by the subsampling mechanism, we can write $\mu = \sum_y \omega(y)\mu_y = \omega M$, where M denotes the Markov kernel operating on measures defined by \mathcal{M} . Consequently, proving privacy amplification results requires reasoning about the mixtures obtained when sampling from two neighbouring datasets $x \simeq_X x'$, and how the privacy parameters are affected by the mixture.

Our Contribution We provide a unified method for deriving privacy amplification by subsampling bounds (Section 3). Our method recovers all existing results in the literature and allow us to derive novel amplification bounds (Section 4). In most cases our method also provides optimal constants which are shown to be tight by a generic lower bound (Section 5). Our analysis relies on properties of divergences and privacy profiles, together with two additional ingredients.

The first ingredient is a novel *advanced joint convexity* property providing upper bounds on the α -divergence between overlapping mixture distributions. In the specific context of differential

privacy this result yields for every $x \simeq_X x'$:

$$D_{e^{\varepsilon'}}(\mathcal{M}^{\mathcal{S}}(x) \parallel \mathcal{M}^{\mathcal{S}}(x')) \leq \eta \cdot ((1 - \beta)D_{e^{\varepsilon}}(\mu_1 \parallel \mu_0) + \beta D_{e^{\varepsilon}}(\mu_1 \parallel \mu'_1)) \quad , \quad (3)$$

for $e^{\varepsilon'} = 1 + \eta(e^{\varepsilon} - 1)$, some $\beta \in [0, 1]$, and $\eta = \text{TV}(\mathcal{S}(x), \mathcal{S}(x'))$ being the total variation distance between the distributions over subsamples. Here $\mu_0, \mu_1, \mu'_1 \in \mathbb{P}(Z)$ are suitable measures obtained from $\mathcal{M}^{\mathcal{S}}(x)$ and $\mathcal{M}^{\mathcal{S}}(x')$ through a coupling and projection operation. In particular, the proof of advanced joint convexity uses ideas from probabilistic couplings, and more specifically the maximal coupling construction (see Theorem 2 and its proof for more details). It is also interesting to note that the non-linear relation $\varepsilon' = \log(1 + \eta(e^{\varepsilon} - 1))$ already appears in some existing privacy amplification results (e.g. Li et al. [2012]). Although for small ε and η this relation yields $\varepsilon' = O(\eta\varepsilon)$, our results show that the more complicated non-linear relation is in fact a fundamental aspect of privacy amplification by subsampling.

The second ingredient in our analysis establishes an upper bound for the divergences occurring in the right hand side of (3) in terms of group-privacy profiles. It states that under suitable conditions, we have $D_{e^{\varepsilon}}(\nu M \parallel \nu' M) \leq \sum_{k \geq 1} \lambda_k(\nu) \delta_{\mathcal{M}, k}(e^{\varepsilon})$ for suitable choices of λ_k . Again, the proof of the inequality uses tools from probabilistic couplings.

The combination of these results yields a bound of the privacy profile of $\mathcal{M}^{\mathcal{S}}$ as a function of the group-privacy profiles of \mathcal{M} . Based on this inequality, we will establish several privacy amplification result and prove tightness results. This methodology can be applied to any of the settings discussed above in terms of dataset representation, neighbouring relation, and type of subsampling. Table 1 summarizes several results that can be obtained with our method (see Section 4 for details). The supplementary material also contains plots illustrating our bounds (Figure 1) and proofs of all the results presented in the paper.

3 Tools: Couplings, Divergences and Privacy Profiles

We next introduce several tools that will be used to support our analyses. The first and second tools are known, whereas the remaining tools are new and of independent interest.

Divergences The following characterization follows immediately from the definition of α -divergence in terms of the supremum over E .

Theorem 1 ([Barthe and Olmedo, 2013]). *A mechanism \mathcal{M} is (ε, δ) -differentially private with respect to \simeq_X if and only if $\sup_{x \simeq_X x'} D_{e^{\varepsilon}}(\mathcal{M}(x) \parallel \mathcal{M}(x')) \leq \delta$.*

Note that in the statement of the theorem we take $\alpha = e^{\varepsilon}$. Throughout the paper we sometimes use these two notations interchangeably to make expressions more compact.

We now state consequences of the definition of α -divergence: (i) $0 \leq D_{\alpha}(\mu \parallel \mu') \leq 1$; (ii) the function $\alpha \mapsto D_{\alpha}(\mu \parallel \mu')$ is monotonically decreasing; (iii) the function $(\mu, \mu') \mapsto D_{\alpha}(\mu \parallel \mu')$ is jointly convex. Furthermore, one can show that $\lim_{\alpha \rightarrow \infty} D_{\alpha}(\mu \parallel \mu') = 0$ if and only if $\text{supp}(\mu) \subseteq \text{supp}(\mu')$.

Couplings Couplings are a standard tool for deriving upper bounds for the statistical distance between distributions. Concretely, it is well-known that the total variation distance between two distributions $\nu, \nu' \in \mathbb{P}(Y)$ satisfies $\text{TV}(\nu, \nu') \leq \Pr_{\pi}[y \neq y']$ for any coupling π , where equality is

attained by taking the so-called maximal coupling. We recall the definition of coupling and provide a construction of the maximal coupling, which we shall use in later sections.

A coupling between two distributions $\nu, \nu' \in \mathbb{P}(Y)$ is a distribution $\pi \in \mathbb{P}(Y \times Y)$ whose marginals along the projections $(y, y') \mapsto y$ and $(y, y') \mapsto y'$ are ν and ν' respectively. Couplings always exist, and furthermore, there exists a maximal coupling, which exactly characterizes the total variation distance between ν and ν' . Let $\nu_0(y) = \min\{\nu(y), \nu'(y)\}$ and let $\eta = \text{TV}(\nu, \nu') = 1 - \sum_{y \in Y} \nu_0(y)$, where TV denotes the total variation distance. The maximal coupling between ν and ν' is defined as the mixture $\pi = (1 - \eta)\pi_0 + \eta\pi_1$, where $\pi_0(y, y') = \nu_0(y)\mathbb{1}[y = y']/(1 - \eta)$, and $\nu_1(y, y') = (\nu(y) - \nu_0(y))(\nu'(y') - \nu_0(y'))/\eta$. Projecting the maximal coupling along the marginals yields the overlapping mixture decompositions $\nu = (1 - \eta)\nu_0 + \eta\nu_1$ and $\nu' = (1 - \eta)\nu_0 + \eta\nu'_1$.

Advanced Joint Convexity The privacy amplification phenomenon is tightly connected to an interesting new form of joint convexity for α -divergences, which we call advanced joint convexity.

Theorem 2 (Advanced Joint Convexity of D_α^3). *Let $\mu, \mu' \in \mathbb{P}(Z)$ be measures satisfying $\mu = (1 - \eta)\mu_0 + \eta\mu_1$ and $\mu' = (1 - \eta)\mu_0 + \eta\mu'_1$ for some η, μ_0, μ_1 , and μ'_1 . Given $\alpha \geq 1$, let $\alpha' = 1 + \eta(\alpha - 1)$ and $\beta = \alpha'/\alpha$. Then the following holds:*

$$D_{\alpha'}(\mu \parallel \mu') = \eta D_\alpha(\mu_1 \parallel (1 - \beta)\mu_0 + \beta\mu'_1) . \quad (4)$$

Note that writing $\alpha = e^\varepsilon$ and $\alpha' = e^{\varepsilon'}$ in the above lemma we get the relation $\varepsilon' = \log(1 + \eta(e^\varepsilon - 1))$. Applying standard joint convexity to the right hand side above we conclude: $D_{\alpha'}(\mu \parallel \mu') \leq (1 - \beta)\eta D_\alpha(\mu_1 \parallel \mu_0) + \beta\eta D_\alpha(\mu_1 \parallel \mu'_1)$. Note that applying joint convexity directly on $D_{\alpha'}(\mu \parallel \mu')$ instead of advanced joint convexity yields a weaker bound which implies amplification for the δ privacy parameter, but not for the ε privacy parameter.

When using advanced joint convexity to analyze privacy amplification we consider two elements x and x' and fix the following notation. Let $\omega = \mathcal{S}(x)$ and $\omega' = \mathcal{S}(x')$ and $\mu = \omega M$ and $\mu' = \omega' M$, where we use the notation M to denote the Markov kernel associated with mechanism \mathcal{M} operating on measures over Y . We then consider the mixture factorization of ω and ω' obtained by taking the decompositions induced by projecting the *maximal coupling* $\pi = (1 - \eta)\pi_0 + \eta\pi_1$ on the first and second marginals: $\omega = (1 - \eta)\omega_0 + \eta\omega_1$ and $\omega' = (1 - \eta)\omega_0 + \eta\omega'_1$. It is easy to see from the construction of the maximal coupling that ω_1 and ω'_1 have disjoint supports and η is the smallest probability such that this condition holds. In this way we obtain the canonical mixture decompositions $\mu = (1 - \eta)\mu_0 + \eta\mu_1$ and $\mu' = (1 - \eta)\mu_0 + \eta\mu'_1$, where $\mu_0 = \omega_0 M$, $\mu_1 = \omega_1 M$ and $\mu'_1 = \omega'_1 M$.

Privacy Profiles We state some important properties of privacy profiles. Our first result illustrates our claim that the “privacy curve” exists for every mechanism \mathcal{M} in the context of the Laplace output perturbation mechanism.

Theorem 3. *Let $f : X \rightarrow \mathbb{R}$ be a function with global sensitivity $\Delta = \sup_{x \simeq_X x'} |f(x) - f(x')|$. Suppose $\mathcal{M}(x) = f(x) + \text{Lap}(b)$ is a Laplace output perturbation mechanism with noise parameter b . The privacy profile of \mathcal{M} is given by $\delta_{\mathcal{M}}(\varepsilon) = [1 - \exp(\frac{\varepsilon - \theta}{2})]_+$, where $\theta = \Delta/b$.*

The well-known fact that the Laplace mechanism with $b \geq \Delta/\varepsilon$ is $(\varepsilon, 0)$ -DP follows from this result by noting that $\delta_{\mathcal{M}}(\varepsilon) = 0$ for any $\varepsilon \geq \theta$. However, Theorem 3 also provides more information:

³Proofs of all our results are presented in the appendix.

it shows that for $\varepsilon < \Delta/b$ the Laplace mechanism with noise parameter b satisfies (ε, δ) -DP with $\delta = \delta_{\mathcal{M}}(\varepsilon)$.

For mechanisms that only satisfy approximate DP, the privacy profile provides information about the behaviour of $\delta_{\mathcal{M}}(\varepsilon)$ as we increase $\varepsilon \rightarrow \infty$. The classical analysis for the Gaussian output perturbation mechanism provides some information in this respect. Recall that for a function $f : X \rightarrow \mathbb{R}^d$ with L_2 global sensitivity $\Delta = \sup_{x \simeq_X x'} \|f(x) - f(x')\|_2$ the mechanism $\mathcal{M}(x) = f(x) + \mathcal{N}(0, \sigma^2 I)$ satisfies (ε, δ) -DP if $\sigma^2 \geq 2\Delta^2 \log(1.25/\delta)/\varepsilon^2$ and $\varepsilon \in (0, 1)$ (cf. [Dwork and Roth, 2014, Theorem A.1]). This can be rewritten as $\delta_{\mathcal{M}}(\varepsilon) \leq 1.25e^{-\varepsilon^2/2\theta^2}$ for $\varepsilon \in (0, 1)$, where $\theta = \Delta/\sigma$. Recently, Balle and Wang [Balle and Wang, 2018] gave a new analysis of the Gaussian mechanism that is valid for all values of ε . Their analysis can be interpreted as providing an expression for the privacy profile of the Gaussian mechanism in terms of the CDF of a standard normal distribution $\Phi(t) = (2\pi)^{-1/2} \int_{-\infty}^t e^{-r^2/2} dr$.

Theorem 4 ([Balle and Wang, 2018]). *Let $f : X \rightarrow \mathbb{R}^d$ be a function with L_2 global sensitivity Δ . For any $\sigma > 0$ let $\theta = \Delta/\sigma$. The privacy profile of the Gaussian mechanism $\mathcal{M}(x) = f(x) + \mathcal{N}(0, \sigma^2 I)$ is given by $\delta_{\mathcal{M}}(e^\varepsilon) = \Phi(\theta/2 - \varepsilon/\theta) - e^\varepsilon \Phi(-\theta/2 - \varepsilon/\theta)$.*

Interestingly, the proof of Theorem 4 implicitly provides a characterization of privacy profiles in terms of privacy loss random variables that holds for any mechanism. Recall that the *privacy loss random variable* of a mechanism \mathcal{M} on inputs $x \simeq_X x'$ is defined as $L_{\mathcal{M}}^{x, x'} = \log(d\mu/d\mu')(\mathbf{z})$, where $\mu = \mathcal{M}(x)$, $\mu' = \mathcal{M}(x')$, and $\mathbf{z} \sim \mu$.

Theorem 5 ([Balle and Wang, 2018]). *The privacy profile of any mechanism \mathcal{M} satisfies*

$$\delta_{\mathcal{M}}(\varepsilon) = \sup_{x \simeq_X x'} \left(\Pr[L_{\mathcal{M}}^{x, x'} > \varepsilon] - e^\varepsilon \Pr[L_{\mathcal{M}}^{x', x} < -\varepsilon] \right) .$$

The characterization above generalizes the well-known inequality $\delta_{\mathcal{M}}(\varepsilon) \leq \sup_{x \simeq_X x'} \Pr[L_{\mathcal{M}}^{x, x'} > \varepsilon]$ (eg. see [Dwork and Roth, 2014]). This bound is often used to derive (ε, δ) -DP guarantees from other notions of privacy defined in terms of the moment generating function of the privacy loss random variable, including concentrated DP [Dwork and Rothblum, 2016], zero-concentrated DP [Bun and Steinke, 2016], Rényi DP [Mironov, 2017], and truncated concentrated DP [Bun et al., 2018]. We now show a reverse implication also holds. Namely, that privacy profiles can be used to recover all the information provided by the moment generating function of the privacy loss random variable.

Theorem 6. *Given a mechanism \mathcal{M} and inputs $x \simeq_X x'$ let $\mu = \mathcal{M}(x)$ and $\mu' = \mathcal{M}(x')$. For $s \geq 0$, define the moment generating function $\varphi_{\mathcal{M}}^{x, x'}(s) = \mathbb{E}[\exp(sL_{\mathcal{M}}^{x, x'})]$. Then we have*

$$\varphi_{\mathcal{M}}^{x, x'}(s) = 1 + s(s+1) \int_0^\infty \left(e^{s\varepsilon} D_{e^\varepsilon}(\mu \parallel \mu') + e^{-(s+1)\varepsilon} D_{e^\varepsilon}(\mu' \parallel \mu) \right) d\varepsilon .$$

In particular, if $D_{e^\varepsilon}(\mu \parallel \mu') = D_{e^\varepsilon}(\mu' \parallel \mu)$ holds⁴ for every $x \simeq_X x'$, then $\sup_{x \simeq_X x'} \varphi_{\mathcal{M}}^{x, x'}(s) = 1 + s(s+1) \int_0^\infty (e^{s\varepsilon} + e^{-(s+1)\varepsilon}) \delta_{\mathcal{M}}(\varepsilon) d\varepsilon$.

⁴For example, this is satisfied by all output perturbation mechanisms with symmetric noise distributions.

Group-privacy Profiles Recall the k th group privacy profile of a mechanism \mathcal{M} is defined as $\delta_{\mathcal{M},k}(\varepsilon) = \sup_{d(x,x') \leq k} D_{e^\varepsilon}(\mathcal{M}(x) \parallel \mathcal{M}(x'))$. A standard group privacy analysis⁵ immediately yields $\delta_{\mathcal{M},k}(\varepsilon) \leq (e^\varepsilon - 1)\delta_{\mathcal{M}}(\varepsilon/k)/(e^{\varepsilon/k} - 1)$. However, “white-box” approaches based on full knowledge of the privacy profile of \mathcal{M} can be used to improve this result for specific mechanisms. For example, it is not hard to see that, combining the expressions from Theorems 3 and 4 with the triangle inequality on the global sensitivity of changing k records in a dataset, one obtains bounds that improve on the “black-box” approach for all ranges of parameters for the Laplace and Gaussian mechanisms. This is one of the reasons why we state our bounds directly in terms of (group-)privacy profiles (a numerical comparison can be found in the supplementary material).

Distance-compatible Coupling The last tool we need to prove general privacy amplification bounds based on α -divergences is the existence of a certain type of couplings between two distributions like the ones occurring in the right hand side of (4). Recall that any coupling π between two distributions $\nu, \nu' \in \mathbb{P}(Y)$ can be used to rewrite the mixture distributions $\tilde{\mu} = \nu M$ and $\tilde{\mu}' = \nu' M$ as $\tilde{\mu} = \sum_{y,y'} \pi_{y,y'} \mathcal{M}(y)$ and $\tilde{\mu}' = \sum_{y,y'} \pi_{y,y'} \mathcal{M}(y')$. Using the joint convexity of D_α and the definition of group-privacy profiles to get the bound

$$D_{e^\varepsilon}(\tilde{\mu} \parallel \tilde{\mu}') \leq \sum_{y,y'} \pi_{y,y'} D_{e^\varepsilon}(\mathcal{M}(y) \parallel \mathcal{M}(y')) \leq \sum_{y,y'} \pi_{y,y'} \delta_{\mathcal{M},d_Y(y,y')}(\varepsilon) . \quad (5)$$

Since this bound holds for any coupling π , one can set out to optimize it by finding a coupling that minimizes the right hand side of (5). We show that the existence of couplings whose support is contained inside a certain subset of $Y \times Y$ is enough to obtain an optimal bound. Furthermore, we show that when this condition is satisfied the resulting bound depends only on ν and the group-privacy profiles of \mathcal{M} .

We say that two distributions $\nu, \nu' \in \mathbb{P}(Y)$ are d_Y -compatible if there exists a coupling π between ν and ν' such for any $(y, y') \in \text{supp}(\pi)$ we have $d_Y(y, y') = d_Y(y, \text{supp}(\nu'))$, where the distance between a point y and the set $\text{supp}(\nu')$ is defined as the distance between y and the closest point in $\text{supp}(\nu')$.

Theorem 7. *Let $C(\nu, \nu')$ be the set of all couplings between ν and ν' and for $k \geq 1$ let $Y_k = \{y \in \text{supp}(\nu) : d_Y(y, \text{supp}(\nu')) = k\}$. If ν and ν' are d_Y -compatible, then the following holds:*

$$\min_{\pi \in C(\nu, \nu')} \sum_{y,y'} \pi_{y,y'} \delta_{\mathcal{M},d_Y(y,y')}(\varepsilon) = \sum_{k \geq 1} \nu(Y_k) \delta_{\mathcal{M},k}(\varepsilon) . \quad (6)$$

Applying this result to the bound resulting from the right hand side of (4) yields most of the concrete privacy amplification results presented in the next section.

4 Privacy Amplification Bounds

In this section we provide explicit privacy amplification bounds for the most common subsampling methods and neighbouring relations found in the literature on differential privacy, and provide pointers to existing bounds and other related work. For our analysis we work with order-independent

⁵If \mathcal{M} is (ε, δ) -DP with respect to \simeq_Y , then it is $(k\varepsilon, ((e^{k\varepsilon} - 1)/(e^\varepsilon - 1))\delta)$ -DP with respect to \simeq_Y^k , cf. [Vadhan, 2017, Lemma 2.2]

representations of datasets without repetitions, i.e. sets. This is mostly for technical convenience, since all our results also hold if one considers datasets represented as tuples or multisets. Note however that subsampling with replacement for a set can yield a multiset; hence we introduce suitable notations for sets and multisets.

Fix a universe of records \mathcal{U} and let $\mathcal{2} = \{0, 1\}$. We write $\mathcal{2}^{\mathcal{U}}$ and $\mathbb{N}^{\mathcal{U}}$ for the spaces of all sets and multisets with records from \mathcal{U} . Note every set is also a multiset. For $n \geq 0$ we also write $\mathcal{2}_n^{\mathcal{U}}$ and $\mathbb{N}_n^{\mathcal{U}}$ for the spaces of all sets and multisets containing exactly n records⁶ from \mathcal{U} . Given $x \in \mathbb{N}^{\mathcal{U}}$ we write x_u for the number of occurrences of $u \in \mathcal{U}$ in x . The support of a multiset x is defined as the set $\text{supp}(x) = \{u \in \mathcal{U} : x_u > 0\}$ of elements that occur at least once in x . Given multisets $x, x' \in \mathbb{N}^{\mathcal{U}}$ we write $x' \subseteq x$ to denote that $x'_u \leq x_u$ for all $u \in \mathcal{U}$.

For order-independent datasets represented as multisets it is natural to consider the two following neighbouring relations. The *remove/add-one* relation is obtained by letting $x \simeq_r x'$ hold whenever $x \subseteq x'$ with $|x| = |x'| - 1$ or $x' \subseteq x$ with $|x| = |x'| + 1$; i.e. x' is obtained by removing or adding a single element to x . The *substitute-one* relation is obtained by letting $x \simeq_s x'$ hold whenever $\|x - x'\|_1 = 2$ and $|x| = |x'|$; i.e. x' is obtained by replacing an element in x with a different element from \mathcal{U} . Note how \simeq_r relates pairs of datasets with different sizes, while \simeq_s only relates pairs of datasets with the same size.

Poisson Subsampling Perhaps the most well-known privacy amplification result refers to the analysis of Poisson subsampling with respect to the remove/add-one relation. In this case the subsampling mechanism $\mathcal{S}_\gamma^{\text{po}} : \mathcal{2}^{\mathcal{U}} \rightarrow \mathbb{P}(\mathcal{2}^{\mathcal{U}})$ takes a set x and outputs a sample y from the distribution $\omega = \mathcal{S}_\gamma^{\text{po}}(x)$ supported on all set $y \subseteq x$ given by $\omega(y) = \gamma^{|y|}(1 - \gamma)^{|x| - |y|}$. This corresponds to independently adding to y with probability γ each element from x . Now, given a mechanism $\mathcal{M} : \mathcal{2}^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ with privacy profile $\delta_{\mathcal{M}}$ with respect to \simeq_r , we are interested in bounding the privacy profile of the subsampled mechanism $\mathcal{M}^{\mathcal{S}_\gamma^{\text{po}}}$ with respect to \simeq_r .

Theorem 8. *Let $\mathcal{M}' = \mathcal{M}^{\mathcal{S}_\gamma^{\text{po}}}$. For any $\varepsilon \geq 0$ we have $\delta_{\mathcal{M}'}(\varepsilon') \leq \gamma \delta_{\mathcal{M}}(\varepsilon)$, where $\varepsilon' = \log(1 + \gamma(e^\varepsilon - 1))$.*

Privacy amplification with Poisson sampling was used in [Chaudhuri and Mishra, 2006, Beimel et al., 2010, Kasiviswanathan et al., 2011, Beimel et al., 2014], which considered loose bounds. A proof of this tight result in terms of (ε, δ) -DP was first given in [Li et al., 2012]. In the context of the moments accountant technique based on the moment generating function of the privacy loss random variable, [Abadi et al., 2016] provide an amplification result for Gaussian output perturbation mechanisms under Poisson subsampling.

Sampling Without Replacement Another known results on privacy amplification corresponds to the analysis of sampling without replacement with respect to the substitution relation. In this case one considers the subsampling mechanism $\mathcal{S}_m^{\text{wo}} : \mathcal{2}_n^{\mathcal{U}} \rightarrow \mathbb{P}(\mathcal{2}_m^{\mathcal{U}})$ that given a set $x \in \mathcal{2}_n^{\mathcal{U}}$ of size n outputs a sample from the uniform distribution $\omega = \mathcal{S}_m^{\text{wo}}(x)$ over all subsets $y \subseteq x$ of size $m \leq n$. Then, for a given a mechanism $\mathcal{M} : \mathcal{2}_m^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ with privacy profile $\delta_{\mathcal{M}}$ with respect to the substitution relation \simeq_s on sets of size m , we are interested in bounding the privacy profile of the mechanism $\mathcal{M}^{\mathcal{S}_m^{\text{wo}}}$ with respect to the substitution relation on sets of size n .

Theorem 9. *Let $\mathcal{M}' = \mathcal{M}^{\mathcal{S}_m^{\text{wo}}}$. For any $\varepsilon \geq 0$ we have $\delta_{\mathcal{M}'}(\varepsilon') \leq (m/n)\delta_{\mathcal{M}}(\varepsilon)$, where $\varepsilon' = \log(1 + (m/n)(e^\varepsilon - 1))$.*

⁶In the case of multisets records are counted with multiplicity.

This setting has been used in [Beimel et al., 2013, Bassily et al., 2014, Wang et al., 2016] with non-tight bounds. A proof of this tight bound formulated in terms of (ε, δ) -DP can be directly recovered from Ullman’s class notes [Ullman, 2017], although the stated bound is weaker. Rényi DP amplification bounds for subsampling without replacement were developed in [Wang et al., 2018].

Sampling With Replacement Next we consider the case of sampling with replacement with respect to the substitution relation \simeq_s . The subsampling with replacement mechanism $\mathcal{S}_m^{\text{wr}} : \mathcal{Z}_n^{\mathcal{U}} \rightarrow \mathbb{P}(\mathbb{N}_m^{\mathcal{U}})$ takes a set x of size n and outputs a sample from the multinomial distribution $\omega = \mathcal{S}_m^{\text{wr}}(x)$ over all multisets y of size $m \leq n$ with $\text{supp}(y) \subseteq x$, given by $\omega(y) = (m!/n^m) \prod_{u \in \mathcal{U}} x_u / (y_u!)$. In this case we suppose the base mechanism $\mathcal{M} : \mathbb{N}_m^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ is defined on multisets and has privacy profile $\delta_{\mathcal{M}}$ with respect to \simeq_s . We are interested in bounding the privacy profile of the subsampled mechanism $\mathcal{M}^{\mathcal{S}_m^{\text{wr}}} : \mathcal{Z}_n^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ with respect to \simeq_s .

Theorem 10. *Let $\mathcal{M}' = \mathcal{M}^{\mathcal{S}_m^{\text{wr}}}$. Given $\varepsilon \geq 0$ and $\varepsilon' = \log(1 + (1 - (1 - 1/n)^m)(e^\varepsilon - 1))$ we have*

$$\delta_{\mathcal{M}'}(\varepsilon') \leq \sum_{k=1}^m \binom{m}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{m-k} \delta_{\mathcal{M},k}(\varepsilon) .$$

Note that if $m = \gamma n$, then $1 - (1 - 1/n)^m \approx \gamma$. A version of this bound in terms of (ε, δ) -DP that implicitly uses the group privacy property can be found in [Bun et al., 2015]. Our bound matches the asymptotics of [Bun et al., 2015] while providing optimal constants and allowing for white-box group privacy bounds.

Hybrid Neighbouring Relations Using our method it is also possible to analyze new settings which have not been considered before. One interesting example occurs when there is a mismatch between the two neighbouring relations arising in the analysis. For example, suppose one knows the group-privacy profiles $\delta_{\mathcal{M},k}$ of a base mechanism $\mathcal{M} : \mathbb{N}_m^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ with respect to the substitution relation \simeq_s . In this case one could ask whether it makes sense to study the privacy profile of the subsampled mechanism $\mathcal{M}^{\mathcal{S}_m^{\text{wr}}} : \mathcal{Z}_n^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ with respect to the remove/add relation \simeq_r . In principle, this makes sense in settings where the size of the inputs to \mathcal{M} is restricted due to implementation constraints (eg. limited by the memory available in a GPU used to run a private mechanism that computes a gradient on a mini-batch of size m). In this case one might still be interested in analyzing the privacy loss incurred from releasing such stochastic gradients under the remove/add relation. Note that this setting cannot be implemented using sampling without replacement since under the remove/add relation we cannot a priori guarantee that the input dataset will have at least size m because the size of the dataset must be kept private [Vadhan, 2017]. Furthermore, one cannot hope to get a meaningful result about the privacy profile of the subsampled mechanism across all inputs sets in $\mathcal{Z}_n^{\mathcal{U}}$; instead the privacy guarantee will depend on the size of the input dataset as shown in the following result.

Theorem 11. *Let $\mathcal{M}' = \mathcal{M}^{\mathcal{S}_m^{\text{wr}}}$. For any $\varepsilon \geq 0$ and $n \geq 0$ we have*

$$\sup_{x \in \mathcal{Z}_n^{\mathcal{U}}, x \simeq_r x'} D_{e^{\varepsilon'}}(\mathcal{M}'(x) \| \mathcal{M}'(x')) \leq \sum_{k=1}^m \binom{m}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{m-k} \delta_{\mathcal{M},k}(\varepsilon) ,$$

where $\varepsilon' = \log(1 + (1 - (1 - 1/n)^m)(e^\varepsilon - 1))$.

When the Neighbouring Relation is “Incompatible” Now we consider a simple example where distance-compatible couplings are not available: Poisson subsampling with respect to the substitution relation. Suppose $x, x' \in 2^{\mathcal{U}}$ are sets of size n related by the substitution relation \simeq_s . Let $\omega = \mathcal{S}_\eta^{\text{po}}(x)$ and $\omega' = \mathcal{S}_\eta^{\text{po}}(x')$ and note that $\text{TV}(\omega, \omega') = \eta$. Let $x_0 = x \cap x'$ and $v = x \setminus x_0$, $v' = x' \setminus x_0$. In this case the factorization induced by the maximal coupling is obtained by taking $\omega_0 = \mathcal{S}_\eta^{\text{po}}(x_0)$, $\omega_1(y \cup \{v\}) = \omega_0(y)$, and $\omega'_1(y \cup \{v'\}) = \omega_0(y)$. Now the support of ω_0 contains sets of sizes between 0 and $n - 1$, while the supports of ω_1 and ω'_1 contain sets of sizes between 1 and n . From this observation one can deduce that ω_1 and ω_0 are not d_{\simeq_s} -compatible, and ω_1 and ω'_1 are not d_{\simeq_r} -compatible.

This argument shows that the method we used to analyze the previous settings cannot be extended to analyze Poisson subsampling under the substitution relation, regardless of whether the privacy profile of the base mechanism is given in terms of the replacement/addition or the substitution relation. This observation is saying that some pairings between subsampling method and neighbouring relation are more natural than others. Nonetheless, even without distance-compatible couplings it is possible to provide privacy amplification bounds for Poisson subsampling with respect to the substitution relation, although the resulting bound is quite cumbersome. The corresponding statement and analysis can be found in the supplementary material.

5 Lower Bounds

In this section we show that many of the results given in the previous section are tight by constructing a randomized membership mechanism that attains these upper bounds. For the sake of generality, we state the main construction in terms of tuples instead of multisets. In fact, we prove a general lemma that can be used to obtain tightness results for any subsampling mechanism and any neighbouring relation satisfying two natural assumptions.

For $p \in [0, 1]$ let $\mathcal{R}_p : \{0, 1\} \rightarrow \mathbb{P}(\{0, 1\})$ be the randomized response mechanism that given $b \in \{0, 1\}$ returns b with probability p and $1 - b$ with probability $1 - p$. Note that for $p = (e^\varepsilon + \delta)/(e^\varepsilon + 1)$ this mechanism is (ε, δ) -DP. Let $\nu_0 = \mathcal{R}_p(0)$ and $\nu_1 = \mathcal{R}_p(1)$. For any $\varepsilon \geq 0$ and $p \in [0, 1]$ define $\psi_p(\varepsilon) = [p - e^\varepsilon(1 - p)]_+$. It is easy to verify that $D_{e^\varepsilon}(\nu_0 \| \nu_1) = D_{e^\varepsilon}(\nu_1 \| \nu_0) = \psi_p(\varepsilon)$. Now let \mathcal{U} be a universe containing at least two elements. For $v \in \mathcal{U}$ and $p \in [0, 1]$ we define the *randomized membership* mechanism $\mathcal{M}_{v,p}$ that given a tuple $x = (u_1, \dots, u_n) \in \mathcal{U}^*$ returns $\mathcal{M}_{v,p}(x) = \mathcal{R}_p(\mathbb{I}[v \in x])$. We say that a subsampling mechanism $\mathcal{S} : X \rightarrow \mathbb{P}(\mathcal{U}^*)$ defined on some set $X \subseteq \mathcal{U}^*$ is *natural* if the following two conditions are satisfied: (1) for any $x \in X$ and $u \in \mathcal{U}$, if $u \in x$ then there exists $y \in \text{supp}(\mathcal{S}(x))$ such that $u \in y$; (2) for any $x \in X$ and $u \in \mathcal{U}$, if $u \notin x$ then we have $u \notin y$ for every $y \in \text{supp}(\mathcal{S}(x))$.

Lemma 12. *Let $X \subseteq \mathcal{U}^*$ be equipped with a neighbouring relation \simeq_X such that there exist $x \simeq_X x'$ with $v \in x$ and $v \notin x'$. Suppose $\mathcal{S} : X \rightarrow \mathbb{P}(\mathcal{U}^*)$ is a natural subsampling mechanism and let $\eta = \sup_{x \simeq_X x'} \text{TV}(\mathcal{S}(x), \mathcal{S}(x'))$. For any $\varepsilon \geq 0$ and $\varepsilon' = \log(1 + \eta(e^\varepsilon - 1))$ we have*

$$\delta_{\mathcal{M}_{v,p}^{\mathcal{S}}}(\varepsilon') = \sup_{x \simeq_X x'} D_{e^{\varepsilon'}}(\mathcal{M}_{v,p}^{\mathcal{S}}(x) \| \mathcal{M}_{v,p}^{\mathcal{S}}(x')) = \eta \psi_p(\varepsilon) .$$

We can now apply this lemma to show that the first three results from previous section are tight. This requires specializing from tuples to (multi)sets, and plugging in the definitions of neighbouring relation, subsampling mechanism, and η used in each of these theorems.

Theorem 13. *The mechanism $\mathcal{M}_{v,p}$ attains the bounds in Theorems 8, 9, 10 for any p and η .*

6 Conclusions

We have developed a general method for reasoning about privacy amplification by subsampling. Our method is applicable to many different settings, some which have already been studied in the literature, and others which are new. Technically, our method leverages two new tools of independent interest: advanced joint convexity and privacy profiles. In the future, it would be interesting to study whether our tools can be extended to give concrete bounds on privacy amplification for other privacy notions such as concentrated DP [Dwork and Rothblum, 2016], zero-concentrated DP [Bun and Steinke, 2016], Rényi DP [Mironov, 2017], and truncated concentrated DP [Bun et al., 2018]. A good starting point is Theorem 6 establishing relations between privacy profiles and moment generating functions of the privacy loss random variable. An alternative approach is to extend the recent results for Rényi DP amplification by subsampling without replacement given in [Wang et al., 2018] to more general notions of subsampling and neighbouring relations.

Acknowledgments

This research was initiated during the 2017 Probabilistic Programming Languages workshop hosted by McGill University’s Bellairs Research Institute.

References

- Martín Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318. ACM, 2016.
- Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *Proceedings of the 35th International Conference on Machine Learning, ICML, 2018*.
- Gilles Barthe and Federico Olmedo. Beyond differential privacy: Composition theorems and relational logic for f-divergences between probabilistic programs. In *International Colloquium on Automata, Languages, and Programming*, pages 49–60. Springer, 2013.
- Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. Probabilistic relational reasoning for differential privacy. In *Symposium on Principles of Programming Languages (POPL)*, pages 97–110, 2012.
- Gilles Barthe, Marco Gaboardi, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. Proving differential privacy via probabilistic couplings. In *Symposium on Logic in Computer Science (LICS)*, pages 749–758, 2016.
- Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 464–473. IEEE, 2014.
- Amos Beimel, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. In *Theory of Cryptography Conference*, pages 437–454. Springer, 2010.

- Amos Beimel, Kobbi Nissim, and Uri Stemmer. Characterizing the sample complexity of private learners. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 97–110. ACM, 2013.
- Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. *Machine learning*, 94(3):401–437, 2014.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, pages 635–658, 2016.
- Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil Vadhan. Differentially private release and learning of threshold functions. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 634–649. IEEE, 2015.
- Mark Bun, Cynthia Dwork, Guy Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated cdp. In *Symposium on Theory of Computing, STOC*, 2018.
- Kamalika Chaudhuri and Nina Mishra. When random sampling preserves privacy. In *Annual International Cryptology Conference*, pages 198–213. Springer, 2006.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 51–60. IEEE, 2010.
- Joonas Jälkö, Antti Honkela, and Onur Dikmen. Differentially private variational inference for non-conjugate models. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, August 11-15, 2017*, 2017.
- Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017.
- Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- Ninghui Li, Wahbeh Qardaji, and Dong Su. On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security*, pages 32–33. ACM, 2012.
- Ilya Mironov. Rényi differential privacy. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pages 263–275, 2017.
- Jack Murtagh and Salil Vadhan. The complexity of computing the optimal composition of differential privacy. In *Theory of Cryptography Conference*, pages 157–175. Springer, 2016.

- Ferdinand Österreicher. Csiszár’s f -divergences-basic properties. *RGMA Res. Rep. Coll*, 2002.
- Mijung Park, James R. Foulds, Kamalika Chaudhuri, and Max Welling. Private topic modeling. *CoRR*, abs/1609.04120, 2016a.
- Mijung Park, James R. Foulds, Kamalika Chaudhuri, and Max Welling. Variational bayes in private settings (VIPS). *CoRR*, abs/1611.00340, 2016b.
- Igal Sason and Sergio Verdú. f -divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016.
- Jonathan Ullman. Cs7880: Rigorous approaches to data privacy. <http://www.ccs.neu.edu/home/jullman/PrivacyS17/HW1sol.pdf>, 2017.
- Salil P. Vadhan. The complexity of differential privacy. In *Tutorials on the Foundations of Cryptography.*, pages 347–450. 2017.
- Yu-Xiang Wang, Stephen Fienberg, and Alex Smola. Privacy for free: Posterior sampling and stochastic gradient monte carlo. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2493–2502, 2015.
- Yu-Xiang Wang, Jing Lei, and Stephen E. Fienberg. Learning with differential privacy: Stability, learnability and the sufficiency and necessity of erm principle. *Journal of Machine Learning Research*, 17(183):1–40, 2016.
- Yu-Xiang Wang, Borja Balle, and Shiva Kasiviswanathan. Subsampled rényi differential privacy and analytical moments accountant. *ArXiv e-prints*, 2018.

A Proofs from Section 3

Proof of Theorem 2. It suffices to check that for any $z \in Z$,

$$[\mu(z) - \alpha' \mu'(z)]_+ = \eta [\mu_1(z) - \alpha ((1 - \beta)\mu_0(z) + \beta\mu'_1(z))]_+ .$$

Plugging this identity in the definition of $D_{\alpha'}$ we get the desired equality

$$D_{\alpha'}(\mu \parallel \mu') = \eta D_{\alpha}(\mu_1 \parallel (1 - \beta)\mu_0 + \beta\mu'_1) .$$

□

Proof of Theorem 3. Suppose $x \simeq_X x'$ and assume without loss of generality that $y = f(x) = 0$ and $y' = f(x) = \Delta > 0$. Plugging the density of the Laplace distribution in the definition of α -divergence we get

$$D_{e^\varepsilon}(\text{Lap}(b) \parallel \Delta + \text{Lap}(b)) = \frac{1}{2b} \int_{\mathbb{R}} \left[e^{-\frac{|z|}{b}} - e^\varepsilon e^{-\frac{|z-\Delta|}{b}} \right]_+ dz .$$

Now we observe that the quantity inside the integral above is positive if and only if $|z - \Delta| - |z| \geq \varepsilon b$. Since $||z + \Delta| - |z|| \leq \Delta$, we see that the divergence is zero for $\varepsilon > \Delta/b$. On the other hand, for $\varepsilon \in [0, \Delta/b]$ we have $\{z : |z - \Delta| - |z| \geq \varepsilon b\} = (-\infty, (\Delta - \varepsilon b)/2]$. Thus, we have

$$\frac{1}{2b} \int_{\mathbb{R}} \left[e^{-\frac{|z|}{b}} - e^\varepsilon e^{-\frac{|z-\Delta|}{b}} \right]_+ dz = \frac{1}{2b} \int_{-\infty}^{(\Delta - \varepsilon b)/2} e^{-\frac{|z|}{b}} dz - \frac{e^\varepsilon}{2b} \int_{-\infty}^{(\Delta - \varepsilon b)/2} e^{-\frac{|z-\Delta|}{b}} dz .$$

Now we can compute both integrals as probabilities under the Laplace distribution:

$$\begin{aligned} \frac{1}{2b} \int_{-\infty}^{(\Delta - \varepsilon b)/2} e^{-\frac{|z|}{b}} dz &= \Pr \left[\text{Lap}(b) \leq \frac{\Delta - \varepsilon b}{2} \right] \\ &= 1 - \frac{1}{2} \exp \left(\frac{\varepsilon b - \Delta}{2b} \right) , \\ \frac{e^\varepsilon}{2b} \int_{-\infty}^{(\Delta - \varepsilon b)/2} e^{-\frac{|z-\Delta|}{b}} dz &= e^\varepsilon \Pr \left[\text{Lap}(b) \leq \frac{-\Delta - \varepsilon b}{2} \right] \\ &= \frac{e^\varepsilon}{2} \exp \left(\frac{-\varepsilon b - \Delta}{2b} \right) . \end{aligned}$$

Putting these two quantities together we finally get, for $\varepsilon \leq \Delta/b$:

$$D_{e^\varepsilon}(\text{Lap}(b) \parallel \Delta + \text{Lap}(b)) = 1 - \exp \left(\frac{\varepsilon}{2} - \frac{\Delta}{2b} \right) .$$

□

Proof of Theorem 6. Let $\varphi = \varphi_{\mathcal{M}}^{x, x'}$, $L = L_{\mathcal{M}}^{x, x'}$, $\tilde{\varphi} = \varphi_{\mathcal{M}}^{x', x}$, and $\tilde{L} = L_{\mathcal{M}}^{x', x}$. Recall that for any non-negative random variable \mathbf{z} one has $\mathbf{E}[\mathbf{z}] = \int_0^\infty \Pr[\mathbf{z} > t] dt$. We use this to write the moment generating function of the corresponding privacy loss random variable for $s \geq 0$ as follows:

$$\begin{aligned} \varphi(s) &= \int_0^\infty \Pr[e^{sL} > t] dt \\ &= \int_0^\infty \Pr \left[\frac{p(\mathbf{z})}{q(\mathbf{z})} > t^{1/s} \right] dt , \end{aligned}$$

where $\mathbf{z} \sim \mu$, and p and q represent the densities of μ and ν with respect to a fixed base measure. Next we observe the probability inside the integral above can be decomposed in terms of a divergence and a second integral with respect to q :

$$\begin{aligned}
\Pr \left[\frac{p(\mathbf{z})}{q(\mathbf{z})} > t^{1/s} \right] &= \Pr[p(\mathbf{z}) > t^{1/s}q(\mathbf{z})] \\
&= \mathbb{E}_\mu \left[\mathbb{I}[p > t^{1/s}q] \right] \\
&= \int \mathbb{I}[p(z) > t^{1/s}q(z)]p(z)dz \\
&= \int \mathbb{I}[p(z) > t^{1/s}q(z)](p(z) - t^{1/s}q(z))dz + t^{1/s} \int \mathbb{I}[p(z) > t^{1/s}q(z)]q(z)dz \\
&= \int [p(z) - t^{1/s}q(z)]_+ dz + t^{1/s} \int \mathbb{I}[p(z) > t^{1/s}q(z)]q(z)dz \\
&= D_{t^{1/s}}(\mu||\mu') + t^{1/s} \int \mathbb{I}[p(z) > t^{1/s}q(z)]q(z)dz .
\end{aligned}$$

Note the term $D_{t^{1/s}}(\mu||\mu')$ above is not a divergence when $t^{1/s} < 1$. The integral term above can be re-written as a probability in terms of \tilde{L} as follows:

$$\begin{aligned}
\int \mathbb{I}[p(z) > t^{1/s}q(z)]q(z)dz &= \Pr[p(\mathbf{z}') > t^{1/s}q(\mathbf{z}')] \\
&= \Pr \left[\frac{p(\mathbf{z}')}{q(\mathbf{z}')} > t^{1/s} \right] \\
&= \Pr \left[e^{-\tilde{L}} > t^{1/s} \right] ,
\end{aligned}$$

where $\mathbf{z}' \sim \mu'$. Thus, integrating with respect to t we get an expression for $\varphi(s)$ involving two terms that we will need to massage further:

$$\varphi(s) = \int_0^\infty D_{t^{1/s}}(\mu||\mu')dt + \int_0^\infty t^{1/s} \Pr \left[e^{-\tilde{L}} > t^{1/s} \right] dt .$$

To compute the second integral in the RHS above we perform the change of variables $dt' = t^{1/s}dt$, which comes from taking $t' = t^{1+1/s}/(1+1/s)$, or, equivalently, $t = ((1+1/s)t')^{1/(1+1/s)}$. This allows us to introduce the moment generating function of \tilde{L} as follows:

$$\begin{aligned}
\int_0^\infty t^{1/s} \Pr \left[e^{-\tilde{L}} > t^{1/s} \right] dt &= \int_0^\infty \Pr \left[e^{-\tilde{L}} > ((1+1/s)t')^{1/(s+1)} \right] dt' \\
&= \int_0^\infty \Pr \left[\frac{s}{s+1} e^{-(s+1)\tilde{L}} > t' \right] dt' \\
&= \frac{s}{s+1} \mathbb{E} \left[e^{-(s+1)\tilde{L}} \right] \\
&= \frac{s}{s+1} \tilde{\varphi}(-s-1) .
\end{aligned}$$

Putting the derivations above together and substituting $\tilde{\varphi}(-s-1)$ for $\varphi(s)$ we see that

$$\varphi(s) = \frac{s}{s+1} \varphi(s) + \int_0^\infty D_{t^{1/s}}(\mu||\mu')dt ,$$

or equivalently:

$$\varphi(s) = (s+1) \int_0^\infty D_{t^{1/s}}(\mu\|\mu') dt .$$

Now we observe that some terms in the integral above cannot be bounded using an α -divergence between μ and μ' , e.g. for $t \in (0,1)$ the term $D_{t^{1/s}}(\mu\|\mu')$ is not a divergence. Instead, using the definition of $D_{t^{1/s}}(\mu\|\mu')$ we can see that these terms are equal to by $1 - t^{1/s} + t^{1/s} D_{t^{-1/s}}(\mu'\|\mu)$, where the last term is now a divergence. Thus, we split the integral in the expression for $\varphi(s)$ into two parts and obtain

$$\begin{aligned} \varphi(s) &= (s+1) \int_0^1 \left(1 - t^{1/s} + t^{1/s} D_{t^{-1/s}}(\mu'\|\mu)\right) dt' + (s+1) \int_1^\infty D_{t^{1/s}}(\mu\|\mu') dt \\ &= 1 + (s+1) \int_0^1 t^{1/s} D_{t^{-1/s}}(\mu'\|\mu) dt' + (s+1) \int_1^\infty D_{t^{1/s}}(\mu\|\mu') dt . \end{aligned}$$

Finally, we can obtain the desired equation by performing a series of simple changes of variables $t' = 1/t$, $\alpha = t^{1/s}$, and $\varepsilon = e^\varepsilon$:

$$\begin{aligned} \varphi(s) &= 1 + (s+1) \int_1^\infty t^{-2-1/s} D_{t^{1/s}}(\mu'\|\mu) dt + (s+1) \int_1^\infty D_{t^{1/s}}(\mu\|\mu') dt \\ &= 1 + s(s+1) \int_1^\infty \left(\alpha^{s-1} D_\alpha(\mu\|\mu') + \alpha^{-s-2} D_\alpha(\mu'\|\mu)\right) d\alpha \\ &= 1 + s(s+1) \int_0^\infty \left(e^{s\varepsilon} D_{e^\varepsilon}(\mu\|\mu') + e^{-(s+1)\varepsilon} D_{e^\varepsilon}(\mu'\|\mu)\right) d\varepsilon . \end{aligned}$$

□

Proof of Theorem 7. The result follows from a few simple observations. The first observation is that for any coupling $\pi \in C(\nu, \nu')$ and $y \in \text{supp}(\nu')$ we have

$$\begin{aligned} \sum_{y'} \pi_{y,y'} \delta_{\mathcal{M},d(y,y')}(\varepsilon) &\geq \sum_{y'} \pi_{y,y'} \delta_{\mathcal{M},d(y,\text{supp}(\nu'))}(\varepsilon) \\ &= \sum_y \nu_y \delta_{\mathcal{M},d(y,\text{supp}(\nu'))}(\varepsilon) , \end{aligned}$$

where the first inequality follows from $d(y, y') \geq d(y, \text{supp}(\nu'))$ and the fact that $\delta_{\mathcal{M},k}(\varepsilon)$ is monotonically increasing with k . Thus the RHS of (6) is always a lower bound for the LHS. Now let π be a d_Y -compatible coupling. Since the support of π only contains pairs (y, y') such that $d(y, y') = d(y, \text{supp}(\nu'))$, we see that

$$\sum_{y,y'} \pi_{y,y'} \delta_{\mathcal{M},d(y,y')}(\varepsilon) = \sum_{y,y'} \pi_{y,y'} \delta_{\mathcal{M},d(y,\text{supp}(\nu'))}(\varepsilon) = \sum_y \nu_y \delta_{\mathcal{M},d(y,\text{supp}(\nu'))}(\varepsilon) .$$

The result follows. □

B Proofs from Section 4

Proof of Theorem 8. Using the tools from Section 3, the analysis is quite straightforward. Given $x, x' \in \mathcal{Z}^{\mathcal{U}}$ with $x \simeq_r x'$, we write $\omega = \mathcal{S}_\eta^{\text{wo}}(x)$ and $\omega' = \mathcal{S}_\eta^{\text{wo}}(x')$ and note that $\text{TV}(\omega, \omega') = \eta$. Next we define $x_0 = x \cap x'$ and observe that either $x_0 = x$ or $x_0 = x'$ by the definition of \simeq_r . Let $\omega_0 = \mathcal{S}_\eta^{\text{po}}(x_0)$. Then the decompositions of ω and ω' induced by their maximal coupling have either $\omega_1 = \omega_0$ when $x = x_0$ or $\omega'_1 = \omega_0$ when $x' = x_0$. Noting that applying advanced joined convexity in the former case leads to an additional cancellation we see that the maximum will be attained when $x' = x_0$. In this case the distribution ω_1 is given by $\omega_1(y \cup \{v\}) = \omega_0(y)$. This observation yields an obvious d_{\simeq_r} -compatible coupling between ω_1 and $\omega_0 = \omega'_1$: first sample y' from ω_0 and then build y by adding v to y' . Since every pair of datasets generated by this coupling has distance one with respect to d_{\simeq_r} , Theorem 7 yields the bound $\delta_{\mathcal{M}'}(\varepsilon') \leq \eta \delta_{\mathcal{M}}(\varepsilon)$. \square

Proof of Theorem 9. The analysis proceeds along the lines of the previous proof. First we note that for any $x, x' \in \mathcal{Z}_n^{\mathcal{U}}$ with $x \simeq_s x'$, the total variation distance between $\omega = \mathcal{S}_m^{\text{wo}}(x)$ and $\omega' = \mathcal{S}_m^{\text{wo}}(x')$ is given by $\eta = \text{TV}(\omega, \omega') = m/n$. Applying advanced joint convexity (Theorem 2) with the decompositions $\omega = (1 - \eta)\omega_0 + \eta\omega_1$ and $\omega' = (1 - \eta)\omega_0 + \eta\omega'_1$ given by the maximal coupling, the analysis of $D_{e^{\varepsilon'}}(\omega M \parallel \omega' M)$ reduces to bounding the divergences $D_{e^{\varepsilon}}(\omega_1 M \parallel \omega_0 M)$ and $D_{e^{\varepsilon}}(\omega_1 M \parallel \omega'_1 M)$. In this case both quantities can be bounded by $\delta_{\mathcal{M}}(\varepsilon)$ by constructing appropriate d_{\simeq_s} -compatible couplings and combining (5) with Theorem 7.

We construct the couplings as follows. Suppose $v, v' \in \mathcal{U}$ are the elements where x and x' differ: $x_v = x'_v + 1$ and $x'_{v'} = x_v + 1$. Let $x_0 = x \cap x'$. Then we have $\omega_0 = \mathcal{S}_m^{\text{wo}}(x_0)$. Furthermore, writing $\tilde{\omega}_1 = \mathcal{S}_{m-1}^{\text{wo}}(x_0)$ we have $\omega_1(y) = \tilde{\omega}_1(y \cap x_0)$ and $\omega'_1(y) = \tilde{\omega}_1(y \cap x_0)$. Using these definitions we build a coupling $\pi_{1,1}$ between ω_1 and ω'_1 through the following generative process: sample y_0 from $\tilde{\omega}_1$ and then let $y = y_0 \cup \{v\}$ and $y' = y_0 \cup \{v'\}$. Similarly, we build a coupling $\pi_{1,0}$ between ω_1 and ω_0 as follows: sample y_0 from $\tilde{\omega}_1$, sample u uniformly from $x_0 \setminus y_0$, and then let $y = y_0 \cup \{v\}$ and $y' = y_0 \cup \{u\}$. It is obvious from these constructions that $\pi_{1,1}$ and $\pi_{0,1}$ are both d_{\simeq_s} -compatible. Plugging these observations together, we get $\delta_{\mathcal{M}'}(\varepsilon') \leq (m/n)\delta_{\mathcal{M}}(\varepsilon)$. \square

Proof of Theorem 10. To bound the privacy profile of the subsampled mechanism $\mathcal{M}^{\text{S}^{\text{wr}}}$ on $\mathcal{Z}_n^{\mathcal{U}}$ with respect to \simeq_s we start by noting that taking $x, x' \in \mathcal{Z}_n^{\mathcal{U}}$, $x \simeq_s x'$, the total variation distance between $\omega = \mathcal{S}_m^{\text{wr}}(x)$ and $\omega' = \mathcal{S}_m^{\text{wr}}(x')$ is given by $\eta = \text{TV}(\omega, \omega') = 1 - (1 - 1/n)^m$. To define appropriate mixture components for applying the advanced joint composition property we write v and v' for the elements where x and x' differ and $x_0 = x \cap x'$ for the common part between both datasets. Then we have $\omega_0 = \mathcal{S}_m^{\text{wr}}(x_0)$. Furthermore, ω_1 is the distribution obtained from sampling \tilde{y} from $\tilde{\omega}_1 = \mathcal{S}_{m-1}^{\text{wr}}(x_0)$ and building y by adding one occurrence of v to \tilde{y} . Similarly, sampling y' from ω'_1 corresponds to adding v' to a multiset sampled from $\mathcal{S}_{m-1}^{\text{wr}}(x')$.

Now we construct appropriate distance-compatible couplings. First we let $\pi_{1,1} \in \mathbb{P}(\mathbb{N}_m^{\mathcal{U}} \times \mathbb{N}_m^{\mathcal{U}})$ be the distribution given by sampling y from ω_1 as above and outputting the pair (y, y') obtained by replacing each v in y by v' . It is immediate from this construction that $\pi_{1,1}$ is a d_{\simeq_s} -compatible coupling between ω_1 and ω'_1 . Furthermore, using the notation from Theorem 7 and the construction of the maximal coupling, we see that for $k \geq 1$:

$$\omega_1(Y_k) = \frac{\omega(Y_k) - (1 - \eta)\omega_0(Y_k)}{\eta} = \frac{\Pr_{y \sim \omega}[y_v = k]}{\eta} = \frac{1}{\eta} \binom{m}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{m-k},$$

where we used $\omega_0(Y_k) = 0$ since ω_0 is supported on multisets that do not include v . Therefore, the distributions $\mu_1 = \omega_1 M$ and $\mu'_1 = \omega'_1 M$ satisfy

$$\eta D_{e^\varepsilon}(\mu_1 \| \mu'_1) \leq \sum_{k=1}^m \binom{m}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{m-k} \delta_{\mathcal{M},k}(\varepsilon) . \quad (7)$$

On the other hand, we can build a d_{\simeq_s} -compatible coupling between ω_1 and ω_0 by first sampling y from ω_1 and then replacing each occurrence of v by an element picked uniformly at random from x_0 . Again, this shows that $D_{e^\varepsilon}(\mu_1 \| \mu_0)$ is upper bounded by the right hand side of (7).

Therefore, we conclude that

$$\delta_{\mathcal{M}'}(\varepsilon') \leq \sum_{k=1}^m \binom{m}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{m-k} \delta_{\mathcal{M},k}(\varepsilon) .$$

□

Proof of Theorem 11. Suppose $x \simeq_r x'$ with $|x| = n$ and $|x'| = n - 1$. This is the worst-case direction for the neighbouring relation like in the proof of Theorem 8. Let $\omega = \mathcal{S}_m^{\text{wr}}(x)$ and $\omega' = \mathcal{S}_m^{\text{wr}}(x')$. We have $\eta = \text{TV}(\omega, \omega') = 1 - (1 - 1/n)^m$, and the factorization induced by the maximal coupling has $\omega_0 = \omega'_1 = \omega'$ and ω_1 is given by first sampling \tilde{y} from $\mathcal{S}_{m-1}^{\text{wr}}(x)$ and then producing y by adding to \tilde{y} a copy of the element v where x and x' differ. This definition of ω_1 suggests the following coupling between ω_1 and ω_0 : first sample y from ω_1 , then produce y' by replacing each copy of v with a element from x' sampled independently and uniformly. By construction we see that this coupling is d_{\simeq_s} -compatible, so we can apply Theorem 7. Using the same argument as in the proof of Theorem 10 we see that $\eta \omega_1(Y_k) = \binom{m}{k} (1/n)^k (1 - 1/n)^{m-k}$. Thus, we finally get

$$\begin{aligned} D_{e^{\varepsilon'}}(\mathcal{M}^{\mathcal{S}_m^{\text{wr}}}(x) \| \mathcal{M}^{\mathcal{S}_m^{\text{wr}}}(x')) &= \eta D_{e^\varepsilon}(\omega_1 M \| \omega_0 M) \\ &\leq \eta \sum_{k=1}^m \omega_1(Y_k) \delta_{\mathcal{M},k}(\varepsilon) \\ &= \sum_{k=1}^m \binom{m}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{m-k} \delta_{\mathcal{M},k}(\varepsilon) . \end{aligned}$$

□

Theorem 14. Let $\mathcal{M} : \mathcal{Z}^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ be a mechanism with privacy profile $\delta_{\mathcal{M}}$ with respect to \simeq_s . Then the privacy profile with respect of \simeq_s of the subsampled mechanism $\mathcal{M}' = \mathcal{M}^{\mathcal{S}_n^{\text{po}}} : \mathcal{Z}_n^{\mathcal{U}} \rightarrow \mathbb{P}(Z)$ on datasets of size n satisfies the following:

$$\delta_{\mathcal{M}'}(\varepsilon') \leq \gamma \beta \delta_{\mathcal{M}}(\varepsilon) + \gamma(1 - \beta) \left(\sum_{k=1}^{n-1} \tilde{\gamma}_k \delta_{\mathcal{M}}(\varepsilon_k) + \tilde{\gamma}_n \right) ,$$

where $\varepsilon' = \log(1 + \gamma(e^\varepsilon - 1))$, $\beta = e^{\varepsilon'}/e^\varepsilon$, $\varepsilon_k = \varepsilon + \log(\frac{\gamma}{1-\gamma}(\frac{n}{k} - 1))$, and $\tilde{\gamma}_k = \binom{n-1}{k-1} \gamma^{k-1} (1 - \gamma)^{n-k}$.

Proof of Theorem 14. Suppose $x, x' \in \mathcal{Z}_n^{\mathcal{U}}$ are sets of size n related by the substitution relation \simeq_s . Let $\omega = \mathcal{S}_n^{\text{po}}(x)$ and $\omega' = \mathcal{S}_n^{\text{po}}(x')$ and note that $\text{TV}(\omega, \omega') = \eta$. Let $x_0 = x \cap x'$ and $v = x \setminus x_0$, $v' = x' \setminus x_0$. In this case the factorization induced by the maximal coupling is obtained by taking

$\omega_0 = \mathcal{S}_\eta^{\text{po}}(x_0)$, $\omega_1(y \cup \{v\}) = \omega_0(y)$, and $\omega'_1(y \cup \{v'\}) = \omega_0(y)$. From this factorization we see it is easy to construct a coupling $\pi_{1,1}$ between ω_1 and ω'_1 that is d_{\sim_s} -compatible. Therefore we have $D_{e^\varepsilon}(\omega_1 M \|\omega'_1 M) \leq \delta_{\mathcal{M}}(\varepsilon)$.

Since we have already identified that no d_{\sim_s} -compatible coupling between ω_1 and ω_0 can exist, we shall further decompose these distributions “by hand”. Let $\nu_k = \mathcal{S}_k^{\text{wo}}(x_0)$ and note that ν_k corresponds to the distribution ω_0 conditioned on $|y| = k$. Similarly, we define $\tilde{\nu}_k$ as the distribution corresponding to sampling \tilde{y} from $\mathcal{S}_{k-1}^{\text{wo}}(x_0)$ and outputting the set y obtained by adding v to \tilde{y} . Then $\tilde{\nu}_k$ equals the distribution of ω_1 conditioned on $|y| = k$. Now we write $\gamma_k = \Pr_{y \sim \omega_0}[|y| = k] = \binom{n-1}{k} \gamma^k (1-\gamma)^{n-1-k}$ and $\tilde{\gamma}_k = \Pr_{y \sim \omega_1}[|y| = k] = \binom{n-1}{k-1} \gamma^{k-1} (1-\gamma)^{n-k}$. With these notations we can write the decompositions $\omega_0 = \sum_{k=0}^{n-1} \gamma_k \nu_k$ and $\omega_1 = \sum_{k=1}^n \tilde{\gamma}_k \tilde{\nu}_k$. Further, we observe that the construction of $\tilde{\nu}_k$ and ν_k shows there exist d_{\sim_s} -compatible couplings between these pairs of distributions when $1 \leq k \leq n-1$, leading to $D_{e^\varepsilon}(\tilde{\nu}_k M \|\nu_k M) \leq \delta_{\mathcal{M}}(\varepsilon)$. To exploit this fact we first write

$$D_{e^\varepsilon}(\omega_1 M \|\omega_0 M) = D_{e^\varepsilon} \left(\sum_{k=1}^{n-1} \tilde{\gamma}_k \tilde{\nu}_k M + \tilde{\gamma}_n \tilde{\nu}_n M \left\| \gamma_0 \nu_0 M + \sum_{k=1}^{n-1} \gamma_k \nu_k M \right. \right) .$$

Now we use that α -divergences can be applied to arbitrary non-negative measures, which are not necessarily probability measures, using the same definition we have used so far. Under this relaxation, given non-negative measures ν_i, ν'_i , $i = 1, 2$, on a measure space Z we have $D_\alpha(\nu_1 + \nu_2 \|\nu'_1 + \nu'_2) \leq D_\alpha(\nu_1 \|\nu'_1) + D_\alpha(\nu_2 \|\nu'_2)$, $D_\alpha(a\nu_1 \|\ b\nu_2) = aD_{\alpha b/a}(\nu_1 \|\nu_2)$ for $a \geq 0$ and $b > 0$, and $D_\alpha(\nu_1 \|\ 0) = \nu_1(Z)$. Using these properties on the decomposition above we see that

$$\begin{aligned} D_{e^\varepsilon}(\omega_1 M \|\omega_0 M) &\leq \sum_{k=1}^{n-1} \tilde{\gamma}_k D_{e^{\varepsilon_k}}(\tilde{\nu}_k M \|\nu_k M) + \tilde{\gamma}_n \\ &\leq \sum_{k=1}^{n-1} \tilde{\gamma}_k \delta_{\mathcal{M}}(\varepsilon_k) + \tilde{\gamma}_n , \end{aligned}$$

where $e^{\varepsilon_k} = (\gamma_k/\tilde{\gamma}_k)e^\varepsilon = (\gamma/(1-\gamma))(n/k-1)e^\varepsilon$. □

C Proofs from Section 5

Proof of Lemma 12. We start by observing that for any $x \in X$ the distribution $\mu = \mathcal{M}_{v,p}^S(x)$ must be a mixture $\mu = (1-\theta)\nu_0 + \theta\nu_1$ for some $\theta \in [0, 1]$. This follows from the fact that there are only two possibilities ν_0 and ν_1 for $\mathcal{M}_{v,p}(y)$ depending on whether $v \notin y$ or $v \in y$. Similarly, taking $x \simeq_X x'$ we get $\mu' = \mathcal{M}_{v,p}^S(x')$ with $\mu' = (1-\theta')\nu_0 + \theta'\nu_1$ for some $\theta' \in [0, 1]$. Assuming (without loss of generality) $\theta \geq \theta'$, we use the advanced joint convexity property of D_α to get

$$\begin{aligned} D_{e^{\varepsilon'}}(\mu \|\mu') &= \theta D_{e^\varepsilon}(\nu_1 \|(1-\theta'/\theta)\nu_0 + (\theta'/\theta)\nu_1) \\ &\leq \theta(1-\theta'/\theta) D_{e^\varepsilon}(\nu_1 \|\nu_0) = (\theta-\theta')\psi_p(\varepsilon) \leq \theta\psi_p(\varepsilon) , \end{aligned}$$

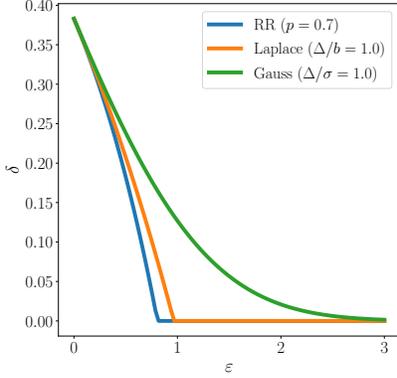
where $\varepsilon' = \log(1 + \theta(e^\varepsilon - 1))$ and $\beta = e^{\varepsilon'}/e^\varepsilon$, and the inequality follows from joint convexity. Now note the inequalities above are in fact equalities when $\theta' = 0$, which is equivalent to the fact $v \notin x'$ because \mathcal{S} is a natural subsampling mechanism. Thus, observing that the function

$\theta \mapsto \theta\psi_p(\log(1 + (e^{\varepsilon'} - 1)/\theta))$ is monotonically increasing, we get

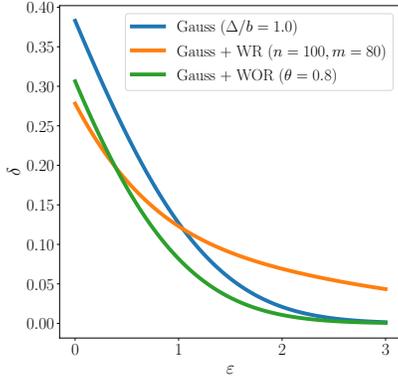
$$\begin{aligned} \sup_{x \simeq_X x'} D_{e^{\varepsilon'}}(\mathcal{M}_{v,p}^{\mathcal{S}}(x) \parallel \mathcal{M}_{v,p}^{\mathcal{S}}(x')) &= \sup_{x \simeq_X x', v \notin x'} \theta\psi_p(\log(1 + (e^{\varepsilon'} - 1)/\theta)) \\ &= \eta\psi_p(\log(1 + (e^{\varepsilon'} - 1)/\eta)) = \eta\psi_p(\varepsilon) . \end{aligned}$$

□

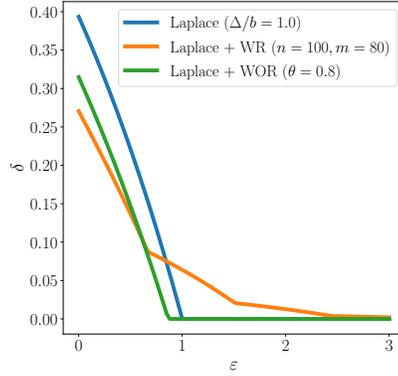
D Plots of Privacy Profiles



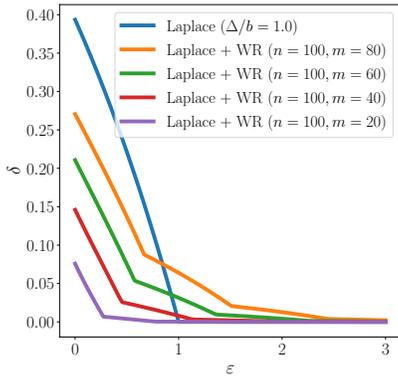
(a) Privacy profiles with mechanisms calibrated to provide the same δ at $\varepsilon = 0$. Profile expressions are given in Section 5 (RR), Theorem 3 (Laplace), and Theorem 4 (Gauss).



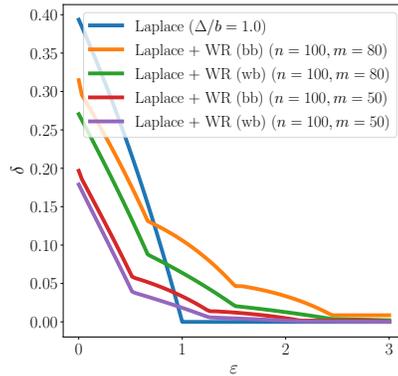
(b) Subsampled Gaussian mechanism. Comparison between sampling without replacement (Theorem 9) and with replacement (Theorem 10, with white-box group privacy), both with the same subsampled dataset sizes.



(c) Subsampled Laplace mechanism. Comparison between sampling without replacement (Theorem 9) and with replacement (Theorem 10, with white-box group privacy), both with the same subsampled dataset sizes.



(d) Subsampled Laplace mechanism. Impact of group-privacy effect in sampling with replacement (white-box group privacy).



(e) Subsampled Laplace mechanism. Impact of white-box vs. black-box group-privacy in sampling with replacement.

Figure 1: Plots of privacy profiles. Results illustrate the notion of privacy profile and the different subsampling bounds derived in the paper.