Optimal Accounting of Differential Privacy via Characteristic Function

Yuqing Zhu
UC Santa Barbara

Jinshuo Dong
IDEAL Institute, Northwestern University

Yu-Xiang Wang
UC Santa Barbara

Abstract

Characterizing the privacy degradation over compositions, i.e., privacy accounting, is a fundamental topic in differential privacy (DP) with many applications to differentially private machine learning and federated learning. We propose a unification of recent advances (Renyi DP, privacy profiles, f-DP and the PLD formalism) via the *characteristic function* (ϕ-function) of a certain *dominating* privacy loss random variable. We show that our approach allows *natural* adaptive composition like Renyi DP, provides *exactly tight* privacy accounting like PLD, and can be (often *losslessly*) converted to privacy profile and f-DP, thus providing (ϵ, δ)-DP guarantees and interpretable tradeoff functions. Algorithmically, we propose an *analytical Fourier accountant* that represents the *complex* logarithm of ϕ-functions symbolically and uses Gaussian quadrature for numerical computation. On several popular DP mechanisms and their subsampled counterparts, we demonstrate the flexibility and tightness of our approach in theory and experiments.

1 Introduction

Differential privacy (DP) [Dwork et al., 2006] is one of the most promising approaches towards addressing the privacy challenges in the era of artificial intelligence and big data. Recently, DP is going through an exciting transformation from a theoretical construct into a practical technology [see, e.g., Apple, Differential Privacy Team, 2017, Erlingsson et al., 2014, Dajani et al., 2017], which demands constant-tight privacy accounting tools that use the privacy budget with optimal efficiency.

Much of the progress in the recent theory and practice of DP has been driven by Renyi Differential Privacy (RDP) [Mironov, 2017], e.g., it is the major technical component behind the *first practical method* for *deep learning with differential privacy* [Abadi et al., 2016]. More broadly, RDP is among several recent work in differential privacy that conducts fine-grained mechanism specific analysis [Bun and Steinke, 2016, Abadi et al., 2016, Mironov, 2017, Balle and Wang, 2018, Wang et al., 2019, Dong et al., 2021, Sommer et al., 2019, Koskela et al., 2020]. At the heart of these breakthroughs is the idea of using a *function* to describe the privacy guarantee of a randomized procedure, thus produces significantly more favorable privacy-utility tradeoff and tighter bounds under composition. (See Table 1 for a summary their pros and cons).

Note that no single approach dominates others in all dimensions. Renyi DP could be undefined for certain privacy loss distributions, and cannot be used to provide the optimal (ϵ, δ)-DP computation in general (discussed in Section 3). Privacy profiles and f-DP are unwieldy under composition; and the method of [Koskela et al., 2020] is limited to mechanisms with univariate output where $\log(p/q)$ admits a density; or those with discrete outputs. Usually, for a new mechanism, we would be lucky to have any one of these functional descriptions. The need to derive these manually for each new mechanism is clearly limiting the creativity of researchers and practitioners in DP.

In addition, there are some unresolved foundational issues related to the PLD formalism. As is repeatedly articulated by the authors, the PLD formalism is defined for *each pair* of neighboring datasets separately, thus, strictly speaking, does not imply DP unless we can certify that the pair of neighboring datasets is the worst-case. This is challenging because such a

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151.

	Functional view	Pros	Cons
Renyi DP [Mironov, 2017]	$D_\alpha(P\|\|Q) \leq \epsilon(\alpha), \forall \alpha \geq 1$	Natural composition	lossy conversion to (ϵ, δ)-DP.
Privacy profile [Balle and Wang, 2018]	$\mathbb{E}_q[(\frac{p}{q} - e^\epsilon)_+] \leq \delta(e^\epsilon), \forall \epsilon \geq 0$	Interpretable.	messy composition.
f-DP[Dong et al., 2021]	Trade-off function f	Interpretable, CLT	messy composition.
PLD [Sommer et al., 2019, Koskela et al., 2020]	Probability density of $\log(p/q)$	Natural composition via FFT	Limited applicability.

Table 1: Modern functional views of DP guarantees and their pros and cons.

pair of datasets might not exist and it is unclear how we can define a partial ordering of two privacy loss distributions.

In this paper, we provide a unified treatment to these functional representations and resolve the aforementioned subtle issues related to the PLD formalism. Our contributions are summarized below.

1. We formalize and generalize the notion of *"worst-case" pair distributions* discussed in [Sommer et al., 2019] to a *"dominating pair"* and prove several basic properties of the *dominating pairs* including finding such pairs from any privacy-profiles, *adaptive* composition and *amplification by sampling*. These results substantially broaden the applicability of PLD formalism [Sommer et al., 2019] in deriving *worst-case* DP guarantees.

2. We propose a lossless representation of the privacy loss RV by its characteristic function (ϕ-function) and derive optimal conversion formula to (and from) privacy-profile, tradeoff-function (f-DP) and the distribution function of the privacy loss RV. Many of these conversion rules correspond naturally to the classical Fourier / Laplace transforms (and their inverses) from the signal processing literature.

3. We design an Analytical Fourier Accountant (AFA, extending the Fourier accountant of [Koskela et al., 2020, 2021]) which represents the complex logarithm of the ϕ function *symbolically*. AFA can be viewed as an extension of the (analytical) moments-accountant [Abadi et al., 2016, Wang et al., 2019] to complex α, thus allowing straightforward composition. Computing δ as a function of ϵ for (ϵ, δ)-DP boils down to a numerical integral which we use a Gaussian quadrature-based method to solve efficiently and accurately.

4. Experimentally, we demonstrate that our approach provides substantially tighter privacy guarantees over compositions than RDP on both basic mechanisms and their subsampled counterparts. Our results essentially match the results from [Dong et al., 2021] and [Koskela et al., 2021] but neither rely on central-limit-theorem type asymptotic approximation nor require choosing appropriate discretization a priori as in the FFT-based Fourier Accountant.

Related work: The paper builds upon the existing work on RDP-based privacy accounting [Abadi et al., 2016, Mironov, 2017, Wang et al., 2019] as well as f-DP [Dong et al., 2021]. Our main theoretical contribution is to substantially broaden the applicability of the PLD formalism [Sommer et al., 2019] by proposing the notion of *dominating pairs* and providing *general recipes* for constructing these dominating pairs. The closest to algorithmic contribution is the work of Koskela et al. [2020, 2021], who propose Fourier accountant and an FFT-based approximation scheme, the characteristic function view can be seen as an *analytical* version of their Fourier accountant (hence the name AFA). AFA is more generally applicable, and allows more flexible use of existing methods for numerical integral. The recent work of Gopi et al. [2021] improves the FFT accountant substantially. It is complementary to us in that it does not address the foundational issues of the PLD formalism, nor do they propose an analytical representation that allows a more modular design of the privacy accountant. Notably, we can use any blackbox numerical integration tool, e.g., Gaussian quadrature, and set the desired error bound on-the-fly, while an FFT-accountant requires setting the parameters at initialization. Finally, Canonne et al. [2020] considered ϕ function and its numerical / computational properties but the discussion is restricted to the discrete Gaussian mechanism.

Privacy accounting is closely related to the classical advanced composition of (ϵ, δ)-DP [Dwork et al., 2010]; Kairouz et al. [2015] provides the *optimal* k-fold composition of an (ϵ, δ)-DP mechanism and Murtagh and Vadhan [2016] shows that computing the tightest possible bound for the composition of k heterogeneous mechanisms is $\#P$-hard. The recent line of work (that we are building upon) challenges the basic primitive of composing (ϵ_i, δ_i)-DP by composing certain functional descriptions of the mechanisms themselves, which sometimes avoids the computational hardness (but not always) and results in even stronger composition than the best (ϵ, δ)-DP type composition would allow [Bun and Steinke, 2016].

2 Notations and preliminary

In this section, we review the standard definition of differential privacy, its RDP relaxation, introduce the characteristic function and draw connections with RDP.

Symbols and notations. Throughout the paper, we will use standard notations for probability unless otherwise stated, e.g., $\Pr[\cdot]$ for probability, $p[\cdot]$ for density, $\mathbb{E}[\cdot]$ for expectation, $F[\cdot]$ for CDF. ϵ, δ are reserved for privacy budget/loss parameters as in (ϵ, δ)-DP, except in the cases when we write $\epsilon(\cdot)$ or $\delta(\cdot)$, where they become functions of certain arguments. We will use $D, D' \in \mathcal{D}^* := \cup_{n\in\mathbb{N}} \mathcal{Z}^n$ to denote two datasets with an unspecified size. D, D' are neighboring (denoted by $D \simeq D'$) if we can construct D' by adding or removing one data point from Z. $\mathcal{M} : \mathcal{D}^* \to P_{\mathcal{O}}$ is a randomized mechanism which returns an output $o \in \mathcal{O}$ by sampling from distribution $\mathcal{M}(D)$. Sometimes for convenience and clarity we define P, Q and p, q to be the distribution and density functions of $\mathcal{M}(D)$ and $\mathcal{M}(D')$ respectively.

Differential privacy and its equivalent definitions. With these notations clarified, we can now formally define differential privacy.

Definition 1 (Differential Privacy)**.** *A randomized algorithm $\mathcal{M}$ is (ϵ, δ)-DP if for every pair of neighboring datasets D, D', and every possible output set $S \subseteq \mathcal{O}$ the following inequality holds:*

$$\Pr[\mathcal{M}(D) \in S] \le e^{\epsilon}\Pr[\mathcal{M}(D') \in S)] + \delta.$$

We can alternatively interpret DP from the views of a divergence metric of two probability distributions, a hypothesis testing view of a binary-classifier, as well as the distribution of the log-odds ratio. Let us first define these quantities formally.

Definition 2 (Hockey-stick divergence)**.** *For $\alpha > 0$, the Hockey-stick divergence is defined as $H_\alpha(P\|Q) := \mathbb{E}_{o\sim Q}[(\frac{\mathrm{d}P}{\mathrm{d}Q}(o) - \alpha)_+]$, where $(x)_+ := x\mathbf{1}(x \ge 0)$ and $\frac{\mathrm{d}P}{\mathrm{d}Q}$ is the Radon-Nikodym-derivative (or simply the density ratio when density exists for P and Q).*

Definition 3 (Trade-off function)**.** *Let ϕ be a classifier to distinguish two distributions P and Q using a sample. α_ϕ be its Type I error (false positive rate) and β_ϕ be its Type II error (false negative rate). The tradeoff function $T_{P,Q}(\alpha) : [0, 1] \to [0, 1]$ is defined to be $T_{P,Q}(\alpha) := \inf_\phi\{\beta_\phi \mid \alpha_\phi \le \alpha\}$.*

Definition 4 (Privacy loss R.V.)**.** *The privacy loss random variable of for a pair of neighboring dataset D, D' under mechanism $\mathcal{M}$ is defined as $L_{P,Q} := \log\frac{\mathcal{M}(D)(o)}{\mathcal{M}(D')(o)}$ where $o \sim \mathcal{M}(D)$; similarly, we have $L_{Q,P} := \log\frac{\mathcal{M}(D')(o)}{\mathcal{M}(D)(o)}$ where $o \sim \mathcal{M}(D')$.*

These quantities can be used to equivalently define differential privacy [Wasserman and Zhou, 2010, Barthe and Olmedo, 2013, Kairouz et al., 2015, Balle and Wang, 2018, Balle et al., 2018, Dong et al., 2021].

Lemma 5. *The following statements about a randomized algorithm $\mathcal{M}$ are equivalent to Definition 1*

1. $\sup_{D\simeq D'} H_{e^\epsilon}(\mathcal{M}(D)\|\mathcal{M}(D')) \le \delta$.
2. $\sup_{D\simeq D'} T_{\mathcal{M}(D),\mathcal{M}(D')}(\alpha) \ge \max\{0, 1-\delta-e^\epsilon\alpha, e^{-\epsilon}(1-\delta-\alpha)$.
3. $\Pr_{o\sim\mathcal{M}(D)}[L_{P,Q} > \epsilon] - e^\epsilon \Pr_{o\sim\mathcal{M}(D')}[L_{Q,P} < -\epsilon] \le \delta$ *for all neighboring D, D'.*

We highlight that in all these definitions, it is required for the bound to cover all pairs of neighboring datasets D, D'.

Mechanism-specific analysis / Functional representation of DP guarantee. Each of these equivalent interpretations could be used to provide more-fine-grained description of a differential privacy mechanism $\mathcal{M}$. For instance, the privacy profile $\delta_{\mathcal{M}}(\epsilon)$ upper bounds the HS-divergence for all ϵ and the f-DP lowerbounds the tradeoff function for all Type I error α (see Table 1). In addition, Sommer et al. [2019] proposes the PLD formalism, which represents the privacy loss RV by its density function. The PLD formalism can be viewed as another functional representation, but it is qualitatively different from privacy profile and f-DP. We will expand further on PLD in Section 3.

Renyi Differential Privacy and Moments Accountant. Renyi differenital privacy (RDP) [Mironov, 2017] is another generalization of pure-DP via Renyi divergence (denoted by $\mathcal{D}_\alpha(P||Q)$).

Definition 6 (Renyi Differential Privacy [Mironov, 2017])**.** *We say a randomized algorithm $\mathcal{M}$ is $(\alpha, \epsilon(\alpha))$-RDP with order $\alpha \ge 1$ if for neighboring datasets D, D',*

$$\mathcal{D}_\alpha(\mathcal{M}(D)||\mathcal{M}(D')) \le \epsilon(\alpha).$$

(ϵ, α)-RDP implies $(\epsilon(\alpha) + \frac{\log(1/\delta)}{\alpha-1}, \delta)$-DP, thus by viewing RDP as a function $\epsilon_{\mathcal{M}}(\cdot)$, we can find the best ϵ parameter by optimizing over α. Tighter conversion formula had been proposed recently [Balle et al., 2020, Asoodeh et al., 2021], which we discuss in Appendix.

The main advantage of RDP is that it composes *naturally* over multiple adaptively chosen mechanisms

via a straightforward rule $\epsilon_{\mathcal{M}_1 \times \mathcal{M}_2} \leq \epsilon_{\mathcal{M}_1} + \epsilon_{\mathcal{M}_2}$. It recovers the advanced composition when converting to (ϵ, δ)-DP and yields substantial additional savings. These properties, together with the privacy-amplification by sampling, makes RDP the natural choice for privacy accounting in various algorithms of differentially private deep learning. The related algorithm that keeps track of the moment generating function of $L_{P,Q}(o)$ is called "moments accountant" [Abadi et al., 2016, Wang et al., 2019].

3 Motivation of our research

In this section, we discuss a number of limitations of Renyi DP and PLD formalism that, in part, motivated our research.

The limits of RDP. Let us first ask "is the RDP function a *lossless* description?" In particular, does it capture all information in the privacy-profile? Because if it is the case, then we could use RDP for composition, and then find the exact optimal (ϵ, δ)-DP by converting from RDP.

The answer is unfortunately "no". The reasons are twofolds. First, there are mechanisms with non-trivial (ϵ, δ)-DP where RDP parameters partially or entirely do not exist. We give two concrete examples in Appendix A.

The second, and a more troubling issue is that even in the cases when RDP parameters exist everywhere and hence *appears to be* characterizing, it does not lead to a tight conversion to (ϵ, δ)-DP. Gaussian mechanism is such a candidate where its PLD is completely captured by its Renyi divergences. However, in Figure 1 we demonstrate that we cannot, in general, convert the RDP of Gaussian mechanism into an (ϵ, δ)-DP that matches the optimal accounting one can achieve through either the privacy profile or f-DP directly. Specifically, by an example due to [Dong et al., 2021, Proposition B.7], we know that a randomized response mechanism (RR) satisfies 1-zCDP, thus the same RDP as that of a Gaussian mechanism (GM) with $\sigma = 1$. If the RDP conversion is tight, then it will have to apply to RR too, but that will lead to a contradiction with the tradeoff function of RR. More explicitly, when we further convert the f-DP in Figure 1 to (ϵ, δ)-DP, this example shows that while both RR and GM satisfy an RDP with $\epsilon(\alpha) = \frac{\alpha}{2}$, GM obeys $(0.277, 0.3)$-DP but RR does not satisfy $(\epsilon, 0.3)$-DP with $\epsilon < 0.471$.

This example certifies that the conversion rule we used (based on an extension of [Balle et al., 2020]) *cannot be improved* and that *RDP is a lossy representation* even for the Gaussian mechanism.

Trouble with Worst-Cases in the PLD formalism. Recent developments in the PLD formalism show great promises in computing tight (ϵ, δ)-DP with stable numerical algorithms and provable error bounds [Koskela et al., 2020, 2021]. However, as we discussed earlier, PLD is specified for each pair of input datasets separately. To use PLD, the original authors (quoting verbatim) *"require the privacy analyst interested in applying our results (PLD formalism) to provide worst-case distributions."* [Sommer et al., 2019, Section 2]. In a subsequent work [Meiser and Mohammadi, 2018], a subset of the authors further derive the worst-case pair of distributions for basic mechanisms such as Gaussian mechanism and Laplace mechanism [Meiser and Mohammadi, 2018].

While these are valid arguments, the line of work on PLD formalism does not formally define the worst-case pair of distributions, nor do they provide general recipes for "privacy analysts" to determine which pair of inputs is the worst-case. The issue is more prominent when we consider mechanism-specific analysis, because the pairs of datasets that attain the argmax might be different in different regions of the privacy profile (see an example in Appendix A).

Moreover, in most typical use cases of the privacy accounting tools, the mechanism under consideration is constructed through the composition of a sequence of simpler mechanisms. Even if for each mechanism, we know the worst-case pair distributions, the composition of the individual PLDs may not correspond to the worst-case PLD of the composed mechanism [1]. For this reason, it is unclear how to use PLD for deriving worst-case DP bound under composition except in highly specialized cases (e.g., Gaussian mechanisms and their compositions).

Summary. To reiterate, RDP is lossy when converting to (ϵ, δ)-DP and the PLD formalism cannot be used to handle the composition generically due to issues regarding worst-case distributions. The remainder of the paper will be dedicated to addressing this dilemma.

4 Main results

In this section, we develop a comprehensive solution towards tighter and more flexible *mechanism-specific*

[1] This is an issue we will address later, which shows that it is OK even if it does not.

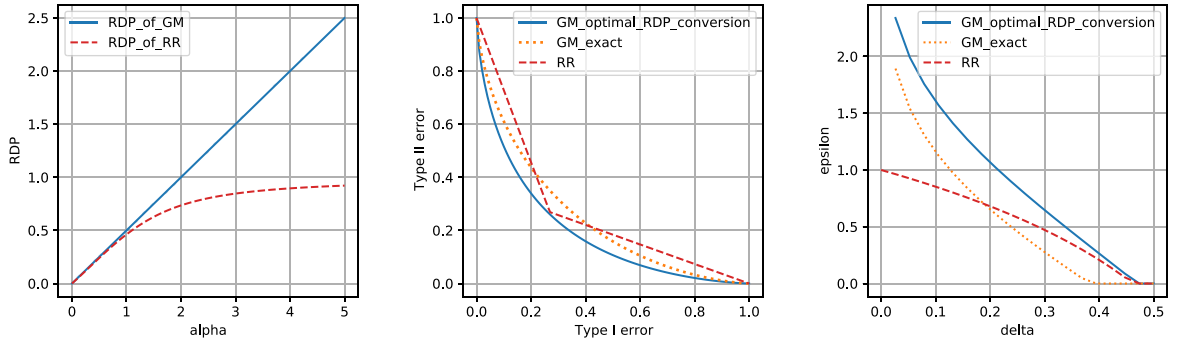

(a) RDP of RR and GM (b) f-DP of RR and GM (c) (ϵ, δ)-DP of RR and GM

Figure 1: The figure illustrates the RDP and f-DP of a Gaussian mechanism with (normalized) $\sigma = 1$, and a randomized response mechanism with $p = \frac{e}{1+e}$. Pane (a) shows the RDP function of RR and GM, clearly, RR also satisfies the same RDP of the Gaussian mechanism for all α. Pane (b) in the middle compares the f-DP of the two mechanisms, as well as the f-DP implied by the optimal conversion from RDP. Pane (c) shows the privacy profile of the two mechanisms, together with Pane (a), it demonstrates that the optimal f-DP and (ϵ, δ)-DP of GM cannot be achieved by a conversion from RDP.

privacy accounting for (ϵ, δ)-DP with a data-structure that allows natural composition.

4.1 Dominating pair of distributions, composition and subsampling

We first patch the PLD formalism by generalizing the idea of worst-case pair (which may not exist) to a *dominating* pair of distributions and prove a number of useful properties.

Definition 7 (Dominating pair of distributions)**.** *We say that (P, Q) is a* dominating *pair of distributions for $\mathcal{M}$ (under neighboring relation $\simeq$) if for all $\alpha \geq 0$*[2]

$$\sup_{D \simeq D'} H_\alpha(\mathcal{M}(D) \| \mathcal{M}(D')) \leq H_\alpha(P \| Q). \quad (1)$$

When P, Q is chosen such that (1) takes "=" for all α, we say that (P, Q) is a tight *dominating pair of distributions or simply,* tightly dominating*. If in addition, there exists a neighboring $(\tilde{D}, \tilde{D}')$ such that $(\mathcal{M}(\tilde{D}), \mathcal{M}(\tilde{D}'))$ is* tightly dominating*, and then we say $(\tilde{D}, \tilde{D}')$ is the worst-case pair of datasets for mechanism $\mathcal{M}$.*

Unless otherwise specified, all subsequent results we present hold for any definitions of neighbors (including asymmetric ones such as *add-only* and *remove-only*, which will be useful later).

A dominating pair of distributions always exists: one can trivially take P and Q that have disjoint supports. What is somewhat surprising is the following

Proposition 8. *Any mechanism has a* tightly *dominating pair of distributions.*

For example, the domintating pair for discrete Gaussian mechanism (DGM) [Canonne et al., 2020] will be two discrete Gaussian, e.g., $P = \mathcal{NZ}(0, \sigma^2), Q = \mathcal{NZ}(\Delta, \sigma^2), \Delta \in \mathbb{Z}_+$ is the sensitivity of the integer-valued query. This follows because the probability mass of the discrete Gaussian is a log-concave sequence. The proof would look very similar to Proposition A.3 of Dong et al. [2021]. On the other hand, worst-case pair of datasets do not always exist, as is shown by Example 16.

Proposition 8 is the direct consequence of the following result which fully characterizes what hockey-stick divergences and privacy profiles look like.

Lemma 9. *For a given $H : \mathbb{R}_{\geqslant 0} \to \mathbb{R}$, there exists P, Q such that $H(\alpha) = H_\alpha(P \| Q)$ if and only if $H \in \mathcal{H}$ where*

$$\mathcal{H} := \left\{ H : \mathbb{R}_{\geqslant 0} \to \mathbb{R} \middle| \begin{array}{l} H \text{ is convex, decreasing,} \\ H(0) = 1 \text{ and } H(x) \geqslant (1-x)_+ \end{array} \right\}.$$

Moreover, one can explicitly construct such P and Q: P has CDF $1 + H^*(x-1)$ in $[0, 1)$ and $Q =$ Uniform$([0, 1])$.

The proof, presented in Appendix C, makes use of the Fenchel duality of the privacy profile with respect to a tradeoff function and a characterization of the tradeoff function due to Dong et al. [2021, Proposition 2.2].

[2] Note that $\alpha \geq 1$ corresponds to the typical range of (ϵ, δ)-DP, but the region for $\alpha < 1$ is important for composition and lossless conversions to other representations.

What makes the specific construction in Lemma 9 (hence Proposition 8) appealing is that even if the output space is complex, the resulting dominating pair of distributions are of univariate random variables defined on $[0, 1]$. This resolves a limitation of Koskela et al. [2020] that requires the mechanism to have either univariate or discrete outputs.

So far, we have shown the existence of a tightly dominating pairs for all mechanisms (Proposition 8), and provided a recipe for constructing such a dominating pair for any valid upper bounds of the privacy profile (Lemma 9 and Corollary 26 in Appendix C). Next we will provide two general primitives on how to construct dominating pairs for more complex mechanisms created by composition and privacy amplification by sampling.

Theorem 10 (Adaptive composition of dominating pairs)**.** *If (P, Q) dominates $\mathcal{M}$ and (P', Q') dominates $\mathcal{M}'$[3], then $(P \times P', Q \times Q')$ dominates the composed mechanism $(\mathcal{M}, \mathcal{M}')$.*

By induction, this theorem implies that if we construct the PLD using a dominating pair of distributions for each individual mechanism, then the composed PLD can be used to obtain a *valid* worst-case DP of the composed mechanism.

Next we present how we can construct a dominating pair of distributions (and datasets) for mechanisms under "privacy-amplification by sampling". This is a powerful primitive that is used widely in differentially private ERM [Bassily et al., 2014], Bayesian learning [Wang et al., 2015] and deep learning [Abadi et al., 2016]. We consider the following two schemes.

Poisson Sampling Denoted by $S^{\gamma}_{\mathbf{Poisson}}$. $S^{\gamma}_{\mathbf{Poisson}}$ takes a dataset of arbitrary size and return a dataset by including *each* data point with probability $0 \leq \gamma \leq 1$ i.i.d. at random.

Subset Sampling Denoted by $S^{\gamma}_{\mathbf{Subset}}$. $S^{\gamma}_{\mathbf{Subset}}$ takes a dataset with size n or $n-1$ and return a subset of size $m < n$ uniformly at random. We define $\gamma := m/n$ as a short-hand. [4]

Somewhat unconventionally, the following theorem not only considers add/remove neighboring relation but also treat them *separately*, which turns out to be crucial in retaining a tight dominating pair with a closed-form expression[5]. Our choice of choosing $\alpha \geq 0$ in Definition 7 ensures that for any mechanism (P, Q) dominates for add neighbors *iff* (Q, P) dominates for removal neighbors.

Theorem 11. *Let $\mathcal{M}$ be a randomized algorithm.*

(1) *If (P, Q) dominates $\mathcal{M}$ for add neighbors then $(P, (1-\gamma)P+\gamma Q)$ dominates $\mathcal{M} \circ S_{\mathbf{Poisson}}$ for add neighbors and $((1-\gamma)Q+\gamma P, Q)$ dominates $\mathcal{M} \circ S_{\mathbf{Poisson}}$ for removal neighbors.*

(2) *If (P, Q) dominates $\mathcal{M}$ for replacing neighbors, then $(P, (1-\gamma)P+\gamma Q)$ dominates $\mathcal{M} \circ S_{\mathbf{Subset}}$ for add neighbors and $((1-\gamma)P+\gamma Q, P)$ dominates $\mathcal{M} \circ S_{\mathbf{Subset}}$ for removal neighbors.*

We can obtain the results for the standard "add/remove" for a k-fold composition of subsampled mechanism by a pointwise maximum of the two:

$$\max\{H_{e^\epsilon}(P_1^k || Q_1^k), H_{e^\epsilon}(P_2^k || Q_2^k))\}$$

where (P_1, Q_1) is the "remove only" version of dominating pair and (P_2, Q_2) is the "add only" version of dominating pair. Existing literature that uses PLD for Poisson-sampled mechanisms while taking $(\gamma P + (1-\gamma)Q, Q)$ as an input are essentially providing privacy guarantees only for the "remove only" neighboring relationship. To the best of our knowledge, this is the first time a dominating pair of distributions under privacy-amplification by sampling is proven generically with an arbitrary base-mechanism $\mathcal{M}$ under the privacy-profile. The result, together with Theorem 10, allows PLD formalism to be applied to a broader family of mechanisms as well as their subsampled versions under adaptive composition.

4.2 Characteristic function representation

Having strengthened the foundation of the PLD formalism with "dominating distribution pairs" and two of its basic primitives, we can now put away RDP and its lossy (ϵ, δ)-DP conversion, then conduct mechanism-specific accounting under (ϵ, δ)-DP directly. Existing computational tools however, either require asymptotic approximation [Dong et al., 2021, Sommer et al., 2019], repeated convolution [Dong et al., 2021] or an *a priori* discretization of the output space [Koskela et al., 2021]. This prompts us to ask:

[3] $\mathcal{M}'$ can be adaptively chosen in that it could depend on the output of $\mathcal{M}$, which requires $\sup_{o \in \mathrm{Range}(\mathcal{M})} H_\alpha(\mathcal{M}'(D, o) \| \mathcal{M}'(D', o)) \leq H_\alpha(P' \| Q')$ for any value of o.

[4] Note that here n, m are public and $\gamma := m/n$ even if $(n-1)$ is the sample size.

[5] See the appendix for a construction of dominating pairs of subsampled mechanisms under "Add/Remove" or "Replace" neighbors and more detailed discussion on the advantage of treating "Add" and "Remove" separately.

"Can we compose mechanisms (with known dominating pairs) naturally just like in RDP? "

To achieve this goal, we propose using the *characteristic function* of the privacy loss RV.

Definition 12 (characteristic function of the privacy loss RV)**.** *Let (P, Q) be a dominating pair of $\mathcal{M}$, and p, q be the probability density (or mass) function of P, Q. The two characteristic functions that describes the PLD are*

$$\phi_{\mathcal{M}}(\alpha) := \mathbb{E}_P[e^{i\alpha \log(p/q)}],\ \phi'_{\mathcal{M}}(\alpha) := \mathbb{E}_Q[e^{i\alpha \log(q/p)}],$$

where i denotes the imaginary unit satisfying $i^2 = -1$ and $\alpha \in \mathbb{R}$.

PLDs are probability measures on the real line, and these ϕ-functions are Fourier transforms of these measures. We provide ϕ-functions for basic mechanisms (see Table 2) and the discrete mechanisms with closed-form expression. For other intricate and continuous mechanisms (e.g., subsample variants), we provide efficient discretization methods with error analysis in Section E.

Advantages over MGF Comparing to the moment generating function used by the RDP, the characteristic function differs only in that we are taking the expectation of the *complex* exponential. At the price of bringing in complex arithmetics, it is now a complex-valued function supported on $\alpha \in \mathbb{R}$ rather than the real-valued Renyi Divergence with order $\alpha > 1$ as was defined in RDP. Unlike MGF, the characteristic function always exists and it characterizes the distribution of the privacy loss R.V., therefore it is always a *lossless* representation. MGF is also characteristic when it exists, but the conversion of MGF to the distribution function is numerically problematic [Epstein and Schotland, 2008].

Moreover, the adaptive composition over multiple heterogeneous mechanisms remains as straightforward as that of the RDP.

Proposition 13. *Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be two randomized algorithms. We have the ϕ-function of the composition $(\mathcal{M}_1, \mathcal{M}_2)$ with order $\alpha \in \mathbb{R}$ satisfies: $\phi_{(\mathcal{M}_1,\mathcal{M}_2)}(\alpha) = \phi_{\mathcal{M}_1}(\alpha) \cdot \phi_{\mathcal{M}_2}(\alpha)$*

Lossless conversion rules. The ϕ-function can be losslessly converted back and forth with other representation such as the privacy-profile, tradeoff function, moment-generating function as well as the distribution function of the privacy loss RV. The conversion rule with prominent interest is the conversion to (ϵ, δ)-DP. Specifically, for finding δ as a function of ϵ (i.e., privacy profile), we invoke the fourth equivalent definition of (ϵ, δ)-DP in Lemma 5, which depends on the cumulative distribution function (CDF) of the privacy loss random variables $L_{P,Q}$ and $L_{Q,P}$. In Appendix B, we establish that these CDFs can be evaluated through an integration of ϕ-functions via Levy's formula. The lossless conversions to other quantities are summarized in Figure 2 and we provide more details in Appendix B. Moreover, most of the conversion formula correspond to well-known transforms such as the Fourier transform, Laplace transform and its double-sided variant. Except for those involve RDP and hence Laplace transform, numerical algorithms for implementing these transforms are often available.

4.3 Analytical Fourier Accountant and numerical algorithms

We now propose our analytical Fourier Accoutant (AFA) in Algorithm 1, which is a combination of the *lossless conversion rules* and the *analytical composition rule* (Proposition 13). Given a sequence of mechanisms (can be varied) applied to the same dataset, the data structure tracks the log characteristic function of each mechanism in a symbolic form. When there is a $\delta(\epsilon)$ query, the accountant first constructs two analytical CDFs (with respect to the privacy loss RV $L_{P,Q}$ and $L_{Q,P}$) using Theorem 18 in Appendix B. Then the conversion to (ϵ, δ)-DP is obtained using Lemma 5. For computing ϵ given δ, we use bisection to solve $\delta_{\mathcal{M}}(\epsilon) = \delta$.

Algorithm 1 Analytical Fourier Accountant

Input Mechanisms $\mathcal{M}_1, ..., \mathcal{M}_K$ and δ .
for $i = 1, ..., K$ **do**
  Maintain the symbolic accountant
  $\log \phi_{(\mathcal{M})}(\alpha) \leftarrow \log \phi_{(\mathcal{M})}(\alpha) + \log \phi_{(\mathcal{M}_i)}(\alpha)$
  $\log \phi'_{(\mathcal{M})}(\alpha) \leftarrow \log \phi'_{(\mathcal{M})}(\alpha) + \log \phi'_{(\mathcal{M}_i)}(\alpha)$
  if query (ϵ, δ)-DP **then**
    Compute the CDF $F_{L_{P,Q}}(\cdot)$ and $F_{L_{Q,P}}(\cdot)$ by integrating $\log \phi_{(\mathcal{M})}(\alpha)$ and $\log \phi'_{(\mathcal{M})}(\alpha)$ using Theorem 18.
    Return δ by Lemma 5.
  end if
end for

AFA vs FFT. Comparing to the FFT-accountant approach [Koskela et al., 2020, 2021, Koskela and Honkela, 2021], our approach decouples representation and numerical computation. We do not make any approximation when tracking the mechanisms, and use numerical computation only when converting to (ϵ, δ)-DP. This avoids the need for setting appropriate discretization parameters of FFT ahead of

Mechanism	Dominating Pair	ϕ function
Randomized Response	$P: \Pr_P[0] = p; Q: \Pr_Q[1] = p$	$\phi_{\mathcal{M}}(\alpha) = \phi'_{\mathcal{M}}(\alpha) = pe^{\alpha i \log(\frac{p}{1-p})} + (1-p)e^{\alpha i \log(\frac{1-p}{p})}$
Laplace Mechanism	$P: p(x) = \frac{1}{2\lambda}e^{-\lvert x\rvert/\lambda}; Q: q(x) = \frac{1}{2\lambda}e^{-\lvert x-1\rvert/\lambda}$	$\phi_{\mathcal{M}}(\alpha) = \phi'_{\mathcal{M}}(\alpha) = \frac{1}{2}\left(e^{\frac{\alpha i}{\lambda}} + e^{\frac{-\alpha i-1}{\lambda}} + \frac{1}{2\alpha i+1}(e^{\frac{\alpha i}{\lambda}} - e^{\frac{-\alpha i-1}{\lambda}})\right)$
Gaussian Mechanism	$P: \mathcal{N}(1,\sigma^2); Q: \mathcal{N}(0,\sigma^2)$	$\phi_{\mathcal{M}}(\alpha) = \phi'_{\mathcal{M}}(\alpha) = e^{\frac{-1}{2\sigma^2}(\alpha^2 - i\alpha)}$

Table 2: ϕ functions and dominating pairs for basic mechanisms.

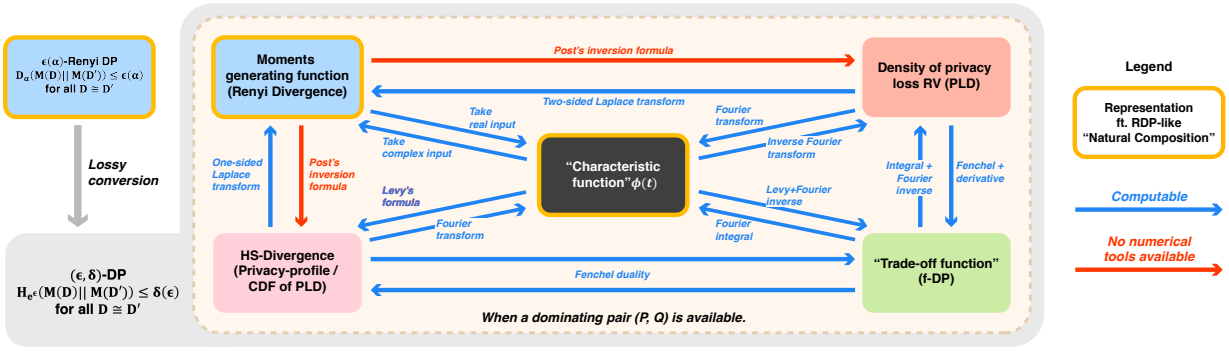

Figure 2: Summary of the various functional descriptions and their conversion rules.

time before knowing which sequence $\mathcal{M}_1, ..., \mathcal{M}_K$ we will receive.

Gaussian quadrature For fast and numerically stable evaluation of the CDF, we propose to use Gaussian quadrature which adaptively selects the intervals between interpolation points, rather than the FFT approach which requires equally spaced discretization. When we apply this approach to efficiently evaluate integral in computing CDFs, where the numerical error is often negligible, i.e., $O(10^{-13})$ for CDFs in our experiments, even if we only sample a few hundreds points. We defer a more detailed error analysis to Section E.

5 Experiments

In this section, we conduct numerical experiments to illustrate the behaviors of our analytical Fourier Accountant. We will have three sets of experiments.

Exp. 1 (Gaussian mechanism) We compare the privacy cost over compositions between RDP accountant and AFA accountant on Gaussian mechanism.

Exp. 2 (Compositions of discrete and continuous mechanisms) We evaluate the Fourier accountant variants and RDP accountant on heterogeneous mechanisms.

Exp. 3 (Compositions over Poisson Subsample mechanisms) Comparison of our AFA with discretization-based ϕ-function to the Fourier accountant (FA) and the RDP accountant.

In Exp1, we compare our AFA method to the RDP-based accoutant[Mironov, 2017] and the exact accountant from the analytical Gaussian mechanism [Balle and Wang, 2018]. In Figure 3(a), we evaluate ϵ with a fixed $\delta = 10^{-4}$ and use $\sigma \in \{50, 100\}$.

Observation: In Figure 3(a), our ϕ function-based AFA exactly matches the result from the analytical Gaussian mechanism and strictly outperforms the RDP accountant in different privacy regimes.

In Exp2, motivated by [Koskela and Honkela, 2021], we consider an adaptive composition of the form $\mathcal{M}(X) = \left(\mathcal{M}_1(X), \tilde{\mathcal{M}}_2(X), ..., \mathcal{M}_{k-1}(X), \tilde{\mathcal{M}}_k(X)\right)$, where each $\mathcal{M}_i$ is a Gaussian mechanism with sensitivity 1, and each $\tilde{\mathcal{M}}_i$ is a randomized mechanism with probability p. We consider $\sigma = 5.0, p = 0.52, \epsilon = 2.0$ and compare $\delta(\epsilon)$ between the RDP accountant, Fourier Accountant [Koskela and Honkela, 2021] and our AFA.

Unlike the FA, our AFA allows an analytical composition over discrete and continuous mechanisms without sampling discretisation points over the privacy loss distribution, therefore achieves an exact privacy accountant. In Figure 3(b), we plot the $\delta(\epsilon)$ over k compositions given by FA and the moments accountant with RDP. We use $n = 10^5$ discretisations points and $L = 10$ for FA. Our numerical result matches FA as $n = 10^5$ is already a very accurate estimation as stated in [Koskela and Honkela, 2021].

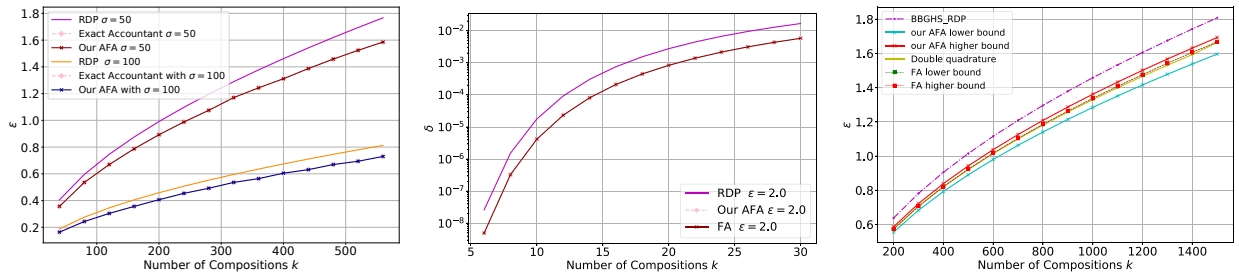

(a) Exp1 Gaussian mechanism (b) Exp2 heterogeneous mechanisms (c) Exp3 Poisson Subsample

Figure 3: Pane 3(a) compares privacy cost over compositions in Exp 1. Pane 3(b) is for the heterogeneous composition in Exp 2. Pane 3(c) is for Poisson subsampled Gaussian mechanism in Exp 3.

There are cases when the closed-form ϕ-functions do not exist. In Exp 3, we consider this problem by analyzing the Poisson Subsample Gaussian mechanism using our discretization-based approach (Algorithm 2) and "Double quadrature" in Appendix E. We discuss the dominating distribution, the construction on ϕ-function, and its discretization in Appendix E. Figure 3(c) shows a comparsion of our AFA to the Fourier accountant method [Koskela et al., 2021] and the moments accountant method [Zhu and Wang, 2019]. The sampling probability is $\gamma = 0.01$, the noise scale is $\sigma = 2.0$ and we evaluate ϵ with $\delta = 10^{-5}$. We use the tighter conversion rule from Balle et al. [2020] to convert the RDP back to (ϵ, δ)-DP. The numerical issues induced by Gaussian quadrature are at most $O(10^{-14})$. Our lower and upper bounds of $\delta(\epsilon)$ shown in Figure 3(c) already incorporate the error induced by discretization and ignoring the tail integral. We emphasize that the lower and upper bounds can match the bounds from FA by increasing sample points n. Moreover, "Double quadrature" is our proposed efficient approximation method. We only unevenly sample 700 points for each ϕ-function and the result of the "Double quadrature" lines between our lower and upper bounds and matches the result from FA. Lastly, all Fourier accountant-based approaches improve over the RDP-based accountant.

Runtime and space analysis of AFA We first compare the time complexity and memory when we have analytical expressions of ϕ-functions. In Exp 2, each mechanism admits an analytical ϕ-function and can be represented in $O(1)$ memory and evaluated in $O(1)$ time. Therefore, the memory cost is $O(\#$ unique mechanisms). We analyze the runtime by decomposing it into the "composition" and "conversion to $\delta(\epsilon)$" separately.

Let k denote the number of compositions. Regarding the runtime in the conversion to $\delta(\epsilon)$ query, we apply Gaussian quadrature to compute the CDF, which requires $O(\frac{1}{\delta_{err}^{1/\alpha}})$ runtime complexity for the αth order differentiable functions. The following composition runtime for Koskela and Honkela [2021] and Gopi et al. [2021] denote the runtime for discretization and convolution via FFT for a homogeneous composition of a mechanism for k rounds. We use n to denote the size of grid discretization in the FFT approximation.

Of course, this is by no means a fair discussion because the FFT approach computes the entire (discretized) PLD of the composed mechanisms together while AFA computes just one point. In terms of the approximation error, our method is the only approach that adapts to the structures of the ϕ functions being integrated and achieves a faster convergence rate.

For the cases when the analytical expressions of ϕ-functions do not exist (see EXP3), we need to approximate the ϕ function too. Thus one single evaluation calls require $O(n)$, and our method is slower than Koskela and Honkela [2021], Gopi et al. [2021], because we do not use FFT. The space and time complexity of the adaptive discretization approach via double quadrature is unclear, though very fast in practice.

6 Conclusion

In this paper, we studied the problem of privacy accounting with mechanism-specific analysis. We introduced the notion of *dominating pair distributions*, showed that each mechanism's privacy profile is *characterized* by a tight dominating pair, and derived a number of useful algebra of dominating pairs including *adaptive composition* and *amplification by sampling*. These results strengthen the foundation of the PLD formalism and make it more widely ap-

Privacy accountant	Composition runtime	$\delta(\epsilon)$ conversion runtime	Memory	Choice of n
Our AFA	$O(1)$	$O(\frac{1}{\delta_{err}^{1/\alpha}})$	$O(1)$	Not applicable
Koskela and Honkela [2021]	$O(n \log n)$	$O(n \log n)$	$O(n)$	$n = O(k/\delta_{err})$
Gopi et al. [2021]	$O(n \log n)$	$O(n \log n)$	$O(n)$	$n = O(\sqrt{k} \log(1/\delta_{err})/\epsilon_{err})$

Table 3: The runtime/space complexity comparisons of different algoirthms

plicable. Algorithmically, we proposed an analytical Fourier accountant that represents the characteristic functions of a dominating pair *symbolically*, which features RDP-like natural composition and allows us to leverage off-the-shelf numerical tools. Our experiments demonstrate the merits of AFA and suggest that it can flexibly and efficiently fit into every DP application.

This work also leaves several open questions. Among those

- As Lemma 9 demonstrates, the construction of the domaining pair is severely constrained when trade-off functions are not clear. For example, characterizing high-dimension discrete Gaussian mechanism remains a tricky open problem.
- Moreover, there are cases where our approach requires much more quadrature points: We apply Gaussian quadrature to compute the CDF of the privacy loss RV through integration over ϕ-functions. If the composed ϕ functions have large values at the tail of integral (e.g., near ∞), we need to sample more quadrature points. We hope to solve this issue using numerical tools in the next step.

Acknowledgments

The work was partially supported by NSF CAREER Award # 2048091, Google Research Scholar Award and a gift from NEC Labs. The authors thank the anonymous reviewers for catching a subtle issue in defining dominating pairs for $\alpha \geq 1$ *only* in an earlier version of the paper. In the hindsight, defining $\alpha \geq 0$ is more natural and elegant. We thank Antti Koskela and Thomas Steinke for helpful discussion. We also thank Salil Vadhan for sharing a shorter alternative proof of the composition theorem based on a deep result due to Blackwell.

References

Martin 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, 2016.

Apple, Differential Privacy Team. Learning with privacy at scale. *Apple Machine Learning Journal*, 2017.

Shahab Asoodeh, Jiachun Liao, Flavio P Calmon, Oliver Kosut, and Lalitha Sankar. Three variants of differential privacy: Lossless conversion and applications. *IEEE Journal on Selected Areas in Information Theory*, 2(1):208–222, 2021.

Borja Balle and Yu-Xiang Wang. Improving gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. *International Conference in Machine Learning (ICML)*, 2018.

Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *Advances in Neural Information Processing Systems (NIPS-18)*, 2018.

Borja Balle, Gilles Barthe, Marco Gaboardi, Justin Hsu, and Tetsuya Sato. Hypothesis testing interpretations and rényi differential privacy. In *International Conference on Artificial Intelligence and Statistics*, pages 2496–2506. PMLR, 2020.

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.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, pages 464–473, 2014.

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.

Clément L Canonne, Gautam Kamath, and Thomas Steinke. The discrete gaussian for differential privacy. *arXiv preprint arXiv:2004.00010*, 2020.

Aref Dajani, Amy Lauger, Phyllis Singer, Daniel Kifer, Jerome Reiter, Ashwin Machanavajjhala, Simon Garfinkel, Scot Dahl, Matthew Graham, Vishesh Karwa, Hang Kim, Philip Leclerc, Ian Schmutte, William Sexton, Lars Vilhuber, and John Abowd. The modernization of statistical disclosure limitation at the u.s. census bureau. *Census Scientific Advisory Commitee Meetings*, 2017. URL https://www2.census.gov/cac/sac/meetings/2017-09/statistical-disclosure-limitation.pdf.

Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society, Series B*, 2021. to appear.

Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 371–380. ACM, 2009.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.

Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Symposium on Foundations of Computer Science (STOC-10)*, pages 51–60. IEEE, 2010.

Charles L Epstein and John Schotland. The bad truth about laplace's transform. *SIAM review*, 50 (3):504–520, 2008.

Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067. ACM, 2014.

Sivakanth Gopi, Yin Tat Lee, and Lukas Wutschitz. Numerical composition of differential privacy. *arXiv preprint arXiv:2106.02848*, 2021.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *International Conference on Machine Learning (ICML-15)*, 2015.

Antti Koskela and Antti Honkela. Computing differential privacy guarantees for heterogeneous compositions using fft. *arXiv preprint arXiv:2102.12412*, 2021.

Antti Koskela, Joonas Jälkö, and Antti Honkela. Computing tight differential privacy guarantees using fft. In *International Conference on Artificial Intelligence and Statistics*, pages 2560–2569. PMLR, 2020.

Antti Koskela, Joonas Jälkö, Lukas Prediger, and Antti Honkela. Tight differential privacy for discrete-valued mechanisms and for the subsampled gaussian mechanism using fft. In *International Conference on Artificial Intelligence and Statistics*, pages 3358–3366. PMLR, 2021.

Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science (FOCS-07)*, pages 94–103. IEEE, 2007.

Sebastian Meiser and Esfandiar Mohammadi. Tight on budget? tight bounds for r-fold approximate differential privacy. In *ACM SIGSAC Conference on Computer and Communications Security (CCS-18)*, pages 247–264, 2018.

Ilya Mironov. Rényi differential privacy. In *Computer Security Foundations Symposium (CSF), 2017 IEEE 30th*, pages 263–275. IEEE, 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.

Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *ACM symposium on Theory of computing (STOC-07)*, pages 75–84. ACM, 2007.

David M Sommer, Sebastian Meiser, and Esfandiar Mohammadi. Privacy loss classes: The central limit theorem in differential privacy. *Proceedings on privacy enhancing technologies*, 2019(2):245–269, 2019.

Josef Stoer and Roland Bulirsch. Interpolation. In *Introduction to Numerical Analysis*, pages 37–144. Springer, 2002.

Abhradeep Guha Thakurta and Adam Smith. Differentially private feature selection via stability arguments, and the robustness of the lasso. In *Conference on Learning Theory*, pages 819–850. PMLR, 2013.

Tim Van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.

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*, pages 2493–2502, 2015.

Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled rényi differential privacy and analytical moments accountant. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1226–1235. PMLR, 2019.

Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489): 375–389, 2010.

Yuqing Zhu and Yu-Xiang Wang. Poisson subsampled rényi differential privacy. In *International Conference on Machine Learning*, pages 7634–7642. PMLR, 2019.

Yuqing Zhu, Xiang Yu, Manmohan Chandraker, and Yu-Xiang Wang. Private-knn: Practical differential privacy for computer vision. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR-20)*, pages 11854–11862, 2020.

Contents

1 Introduction **1**

2 Notations and preliminary **2**

3 Motivation of our research **3**

4 Main results **5**
4.1 Dominating pair of distributions, composition and subsampling 5
4.2 Characteristic function representation 6
4.3 Analytical Fourier Accountant and numerical algorithms 7

5 Experiments **7**

6 Conclusion **9**

A Limits of RDP and the PLD formalism **12**

B Conversion rules between functional representations **13**

C Omitted proofs in the main body **16**
C.1 Characterization of privacy profiles 16
C.2 Composition theorem of dominating pairs 18
C.3 Privacy-amplification for dominating pairs 19
C.4 Other schemes in privacy-amplification for dominating pairs 22

D The characteristic function of basic mechanisms **25**
D.1 ϕ-function of basic mechanisms 25
D.2 Handling probability mass at ∞ 26

E ϕ-function with discretization and experimental details **27**
E.1 Error analysis 29
E.2 Experimental details in Exp 3 30

F An "optimal" Renyi DP to DP conversion? **31**
F.1 RDP to f-DP 31
F.2 f-DP to (ϵ, δ)-DP 32
F.3 Optimality of this conversion rule? 33

G Omitted proofs in Appendix B **33**

A Limits of RDP and the PLD formalism

In Section 3 we omitted a few examples when we talk about the limitation of Renyi Differential Privacy (RDP) in describing common mechanisms. Specifically, we commented that there are mechanisms where RDP either does not exist or does not exist for most order α that implies stronger privacy guarantees.

We give two concrete examples below.

Example 14 (Distance-to-Instability)**.** *The stability-based argument of query release first add noise to a special integer-valued function* $\mathrm{dist}_q(D)$ *which measures the number of data points to add / remove before the local sensitivity of query* $q(D)$ *becomes non-zero. No matter that* q *is,* dist_q *always has a global sensitivity of at most* 1*. The stability-based query release outputs* $\perp$ *(nothing) if* $\mathrm{dist}_q(D) + \mathrm{Lap}(1/\epsilon) \leq {}^{\log(1/\delta)}/_{\epsilon}$ *otherwise outputs the answer* $q(D)$ *without adding noise. This algorithm is satisfies* (ϵ, δ)*-DP [Thakurta and Smith, 2013], but since there is a probability mass at the* $+\infty$ *for the case when* $q(D) \neq q(D')$*, RDP is* $+\infty$ *for all* α*.*

Example 15 (Gaussian-noise adding with data-dependent variance). *In smooth sensitivity-based query release [Nissim et al., 2007], one perturbs the output with a noise with a data-dependent variance. Consider, for example, $P = \mathcal{N}(0, \sigma_1^2), Q = \mathcal{N}(0, \sigma_2^2)$, then the Renyi-divergence $D_\alpha(P\|Q)$ is undefined for all α such that $\alpha\sigma_2^2 + (1-\alpha)\sigma_1^2 < 0$. Specifically, if $\sigma_1^2 = 2, \sigma_2^2 = 1$, then $D_\alpha(P\|Q) = +\infty$ for all $\alpha \geq 2$.*

These examples demonstrate the deficiency of RDP in analyzing flexible algorithm design tools such as the proposed-test-release [Dwork and Lei, 2009]), which typically introduces a heavier-tailed privacy-loss distributions for which the moment generating function is not defined.

On the contrary, the privacy-profile is well-defined in both examples and imply nontrivial (ϵ, δ)-DP. The characteristic function exists no matter how heavy-tailed the distribution of the privacy loss random variable is so it naturally handles the second example. In Section D.2, we describe how we can handle a probability mass at $+\inf$ in our approach.

We also omitted an example for which there are no single pair of neighboring datasets that attain the argmax might be different in different regions of the privacy profile.

Example 16 (Distance to Instability). *Distance to instability $\mathrm{dist}_q(D)$ is a special function that measures the number of data points to add / remove before the local sensitivity of query $q(D)$ becomes non-zero. The stability-based query release outputs $\perp$ (nothing) if $\mathrm{dist}_q(D) + \mathrm{Lap}(1/\epsilon) \leq {}^{\log(1/\delta)}/_{\epsilon}$ otherwise outputs the answer $q(D)$ without adding noise. In this algorithm, the privacy loss distribution has exactly two modes.*

Mode 1 *When $\mathrm{dist}_q(D) > 0$, then for all D' neighboring to D, $q(D) = q(D')$, which implies that the PLD is from the post-processing of a Laplace mechanism (for releasing the perturbed $\mathrm{dist}_q(D)$), i.e., $(\epsilon, 0)$-DP.*

Mode 2 *When $\mathrm{dist}_q(D) = 0$, then for those neighboring D' such that $q(D) \neq q(D')$, it must hold that $\mathrm{dist}_q(D') = 0$, thus the privacy loss distribution is a point mass of $1-\delta$ at 0 (for outputting $\perp$) and a point mass of δ at $+\infty$, i.e., $(0, \delta)$-DP.*

Clearly, there is no single pair of datasets that attains the privacy-profile of this mechanism for all input parameter $\tilde{\epsilon}$. When $\tilde{\epsilon} > \epsilon$, $\delta_{\mathcal{M}}(\tilde{\epsilon}) = \delta$ and is attained by the second mode. On the other hand, if we choose $\tilde{\epsilon}$ such that $\delta_{Lap.\ Mech.(1/\epsilon)}(\tilde{\epsilon}) > \delta$, then $\delta_{\mathcal{M}}(\tilde{\epsilon}) = \delta_{Lap.\ Mech.(1/\epsilon)}(\tilde{\epsilon})$ and the equal sign is attained by a pair of distributions in the first mode.

B Conversion rules between functional representations

In this section we give the details of conversions between various functional representations of the privacy loss distribution (of a dominating pair of distributions). These conversions are summarized in Figure 2 and repeated here.

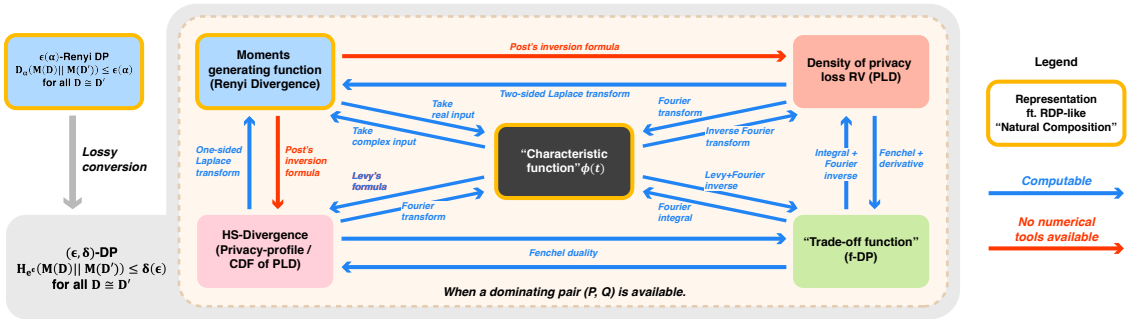

Before we proceed to the details of all these arrows, we would like to emphasize a important distinction:

> These conversions rules are *not* about converting between different DP definitions, but rather converting between different representations of the privacy loss r.v. under the same DP definition — in our case, (ϵ, δ)-DP.

More precisely, we mean that the conversion from RDP to DP (leftmost grey arrow in the figure, which we will talk about in details in Section F) is *qualitatively different* from the conversion from Renyi divergence to Hockey-Stick divergence (red arrow labeled "Post's inversion formula").

Modulo some details[6], a conversion from RDP to DP is about finding function $\delta(\cdot)$ that upper bounds the Hockey-stick divergence for all pairs of neighboring datasets using an RDP function $\epsilon(\cdot)$.

$$\text{If } \sup_{D\simeq D'} \mathcal{D}_\alpha(\mathcal{M}(D)\|\mathcal{M}(D')) \leqslant \epsilon(\alpha), \text{ then } \sup_{D\simeq D'} H_{\mathrm{e}^\varepsilon}(\mathcal{M}(D)\|\mathcal{M}(D')) \leqslant \delta(\varepsilon).$$

In contrast, a conversion from Renyi divergence to hockey-stick divergence is about a given pair of P, Q, and the input function $\epsilon(\alpha)$ is expected to be the exact Renyi-divergence of order α. The goal of the divergence-to-divergence conversion rule is to find a different divergence of the same pair of distribuiton P, Q, i.e.

$$\text{If } \mathcal{D}_\alpha(P\|Q) = \epsilon(\alpha), \text{ then } H_{\mathrm{e}^\varepsilon}(P\|Q) = \delta(\varepsilon).$$

Both conversions aim to compute a function δ from a function ϵ. The seemingly harmless distinction of inequalities and identities is actually the devil in the details. It has two major consequences

1. When applied to privacy, divergence conversion requires a dominating pair of distributions as a prerequisite, which may or may not be a tight dominating pair. In the figure, results that require a dominating pair are enclosed in the light yellow region labeled "When a dominating pair (P, Q) is available".
2. DP conversion is lossy even when converting the statement "standard randomized response is 1-zCDP" to (ε, δ)-DP, as demonstrated by Figure 1. On the other hand, divergence conversion is generically lossless (under some regularity condition), though numerical issues often arise since the inverse Laplace transform is involved Epstein and Schotland [2008].

In alignment with the focus of this paper, in this section we focus on the light yellow region assuming (P, Q) is a dominating pair. DP conversion is discussed in more detail in Appendix F.

From	To	Result
ϕ, ϕ'	F, G	Lemma 18, a direct consequence of Levy's formula
F, G	ϕ, ϕ'	Fourier transform, by definition
F, G	H_α	Lemma 19
H_α	F, G	Lemma 23
F, G	f	Lemma 22
f	F, G	Lemma 23
H_α	f	Proposition 2.12 of Dong et al. [2021], restated as Lemma 21
f	H_α	Proposition 2.12 of Dong et al. [2021], restated as Lemma 20
H_α	$\mathcal{D}_\alpha$	Theorem 8 of Balle et al. [2018], restated as Lemma 24
$\mathcal{D}_\alpha$	H_α	Post's formula. In fact, any inverse Laplace transform works.
$\mathcal{D}_\alpha$	ϕ	take pure imaginary input. Need analytic extension in general and not always possible.
ϕ, ϕ'	$\mathcal{D}_\alpha$	take pure imaginary input. Need analytic extension in general and not always possible.
f	ϕ, ϕ'	Lemma 25
ϕ, ϕ'	f	first use Levy's formula to compute F and G, then use Lemma 22

Table 4: References for conversion results declared in Figure 2. Notations: ch.f. of PLD are denoted by ϕ, ϕ'. CDFs of PLD are denoted by F, G. The trade-off function is denoted by f (note that f is *not* the derivative of F). Hockey-stick divergences are denoted by H_α. Renyi divergences are denoted by $\mathcal{D}_\alpha$.

We recall some definitions. Let P, Q be two probability distributions on the same measurable space. For

[6] such as the symmetry of $P = M(D)$ and $Q = M(D')$ and the domains of α and ε.

$\alpha > 0$, their hockey-stick divergence is defined as

$$H_\alpha(P\|Q) = \mathbb{E}_{\omega\sim Q}[(\tfrac{\mathrm{d}P}{\mathrm{d}Q}(\omega) - \alpha)_+].$$

For $\alpha > 1$, their Renyi divergence is defined as

$$\mathcal{D}_\alpha(P\|Q) := \tfrac{1}{\alpha-1}\log \mathbb{E}_P\left(\tfrac{\mathrm{d}P}{\mathrm{d}Q}\right)^\alpha.$$

Let F and G be the CDFs of the privacy loss random variables. Namely,

$$F(x) := P[\log \tfrac{\mathrm{d}Q}{\mathrm{d}P} \leqslant x]$$
$$G(x) := Q[\log \tfrac{\mathrm{d}Q}{\mathrm{d}P} \leqslant x]$$

The corresponding densities (if exist) will be F' and G'. The corresponding characteristic functions (ch.f.) are the Fourier transforms of the two measures, i.e.

$$\phi(t) = \int \mathrm{e}^{itx}\,\mathrm{d}F(x)$$
$$\phi'(t) = \int \mathrm{e}^{itx}\,\mathrm{d}G(x)$$

Trade-off functions are $T[P,Q]$ and $T[Q,P]$, which map the type I error to the corresponding minimal type II error in testing problems P vs Q and Q vs P respectively.

From these definitions we see that all five functional representations actually require **two** functions for each pair of distributions. Below we summarize how one determines the other.

- $H_\alpha(Q\|P) = \alpha H_{\alpha^{-1}}(P\|Q) - \alpha + 1$, which is stated as Lemma 46 in Appendix G.
- For $\alpha \in (0,1)$, $\mathcal{D}_\alpha(Q\|P) = \frac{\alpha}{1-\alpha}\mathcal{D}_{1-\alpha}(P\|Q)$. See Proposition 2 of Van Erven and Harremos [2014].
- $G'(x) = \mathrm{e}^x F'(x)$, which is stated as Lemma 47 in Appendix G.
- Using the above formula, ϕ' can be obtained by the following process: $\phi \xrightarrow{\text{Levy's formula}} F \to G \to \phi'$.
- If $T[P,Q] = f$ then $T[Q,P] = f^{-1}$. See Lemma A.2 of Dong et al. [2021]

We now consider the conversion from the ϕ-function to CDFs using the following Levy's theorem.

Theorem 17 (Levy). *Let ϕ be the ch.f. of the distribution function F and $a < b$, then*

$$F(b) - F(a) = \lim_{T\to\infty}\frac{1}{2\pi}\int_{-T}^{T}\frac{\mathrm{e}^{-ita} - \mathrm{e}^{-itb}}{it}\cdot\phi(t)\,\mathrm{d}t.$$

Note that $\lim_{T\to\infty,\alpha\to\infty}\int_{-\infty}^{\infty}\frac{e^{-i\alpha a}}{i\alpha}\phi(\alpha)d\alpha = \pi$. To compute the CDF of the privacy loss RV $L_{P,Q}$ at b, we can substitude a with $-\infty$ and obtain the following result.

Lemma 18.

$$F(x) = \frac{1}{2} + \lim_{T\to\infty}\frac{1}{2\pi}\int_{-T}^{T}\frac{i\mathrm{e}^{-itx}}{t}\phi(t)\,\mathrm{d}t$$
$$G(x) = \frac{1}{2} + \lim_{T\to\infty}\frac{1}{2\pi}\int_{-T}^{T}\frac{i\mathrm{e}^{-itx}}{t}\phi'(t)\,\mathrm{d}t$$

Lemma 19.

$$H_\alpha(P\|Q) = F(-\log\alpha) - \alpha G(-\log\alpha)$$
$$H_\alpha(Q\|P) = 1 - G(\log\alpha) - \alpha(1 - F(\log\alpha))$$

Lemma 20. $H_{e^\varepsilon}(Q\|P) = 1 + f^*(-e^\varepsilon)$

Lemma 21.

$$f(x) = \sup_{\varepsilon \geqslant 0} \max\{0, 1 - H_{e^\varepsilon}(P\|Q) - e^\varepsilon x, e^{-\epsilon}(1 - H_{e^\varepsilon}(P\|Q) - x)\}.$$

Lemma 22. $f(\alpha) = G(F^{-1}(1-\alpha))$.

Lemma 23.

$$\begin{aligned} F(x) &= 1 - (f^*)'(-e^x) \\ &= 1 + e^{-x} \cdot \frac{\mathrm{d}}{\mathrm{d}x} H_{e^x}(Q\|P) \\ G(x) &= f\big((f^*)'(-e^x)\big) \\ &= 1 - H_{e^x}(Q\|P) + \frac{\mathrm{d}}{\mathrm{d}x} H_{e^x}(Q\|P) \end{aligned}$$

Lemma 24 (Theorem 6 of Balle et al. [2018])**.**

$$\mathcal{D}_\alpha(P\|Q) = \tfrac{1}{\alpha-1} \log \left(1 + \alpha(\alpha+1) \int_0^\infty \left(e^{\alpha\varepsilon} H_{e^\varepsilon}(P\|Q) + e^{-(\alpha+1)\varepsilon} H_{e^\varepsilon}(Q\|P)\right) \mathrm{d}\varepsilon\right).$$

Lemma 25. *The characteristic functions are determined by the trade-off function f via the following formula:*

$$\begin{aligned} \phi(t) &= \int_0^1 e^{-it \log|f'(x)|} \,\mathrm{d}x \\ \phi'(t) &= \int_0^1 e^{it \log|f'(x)|} \cdot |f'(x)| \,\mathrm{d}x \end{aligned}$$

C Omitted proofs in the main body

C.1 Characterization of privacy profiles

Proof of Proposition 9. Let

$$\begin{aligned} \mathcal{H} &:= \{h : \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0} \mid \exists P, Q \text{ s.t. } h(\alpha) = H_\alpha(P\|Q)\}, \\ \mathcal{F} &:= \{f : [0,1] \to [0,1] \mid \exists P, Q \text{ s.t. } f = T[P,Q]\}. \end{aligned}$$

By Lemma 20, $H_\alpha(P\|Q)$ can be related to $f = T[Q,P]$ as follows:

$$H_{e^\varepsilon}(P\|Q) = 1 + f^*(-e^\varepsilon)$$

where ε ranges over the whole real line. By a simple change of variable, we see that $h \in \mathcal{H}$ iff there exists $f \in \mathcal{F}$ such that $h(\alpha) = 1 + f^*(-\alpha)$, or equivalently,

$$\mathcal{H} = \{h : \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0} \mid \exists f \in \mathcal{F}, h(\alpha) = 1 + f^*(-\alpha)\}.$$

By Proposition 2.2 of Dong et al. [2021], we know

$$\mathcal{F} = \{f : [0,1] \to [0,1] \mid f \text{ is convex, decreasing, continuous and } f(x) \leqslant 1 - x\}.$$

Let $\mathcal{G} := \{g : (-\infty, 0] \to \mathbb{R} \mid g(0) = 0, g \text{ is convex, increasing, continuous and } g(x) \geqslant \max\{x, -1\}\}$.

Claim: Convex conjugacy is a bijection between $\mathcal{F}$ and $\mathcal{G}$.

Proof of the claim. Since both $\mathcal{F}$ and $\mathcal{G}$ consist of convex functions, double convex conjugacy brings back the function, it suffices to show that $f \in \mathcal{F} \implies f^* \in \mathcal{G}$ and $g \in \mathcal{G} \implies g^* \in \mathcal{F}$. Now suppose $f \in \mathcal{F}$. f is extended to be $+\infty$ in $(-\infty, 0)$ and 0 in $(1, +\infty)$. Thus f is a convex function on $\mathbb{R}$. By definition f^* is convex, and we can calculate

$$f^*(y) = \sup_{x \in \mathbb{R}} yx - f(x) = \sup_{x \geqslant 0} yx - f(x) = \begin{cases} +\infty, & \text{if } y > 0 \\ 0, & \text{if } y = 0 \end{cases}$$

With $y_1 < y_2$, we have $y_1 x - f(x) \leqslant y_2 x - f(x)$. Taking supremum over $x \geqslant 0$, we have $f^*(y_1) \leqslant f^*(y_2)$. This shows f^* is monotone and finite on $(-\infty, 0]$. Let

$$I(x) = \begin{cases} +\infty, & \text{if } x < 0 \\ \max\{1 - x, 0\}, & \text{if } x \geqslant 0 \end{cases}$$

It is straightforward to compute that

$$I^*(y) = \begin{cases} \max\{y, -1\}, & \text{if } y \leqslant 0 \\ +\infty, & \text{if } y > 0 \end{cases}$$

Since $f \leqslant I$, we conclude that $f^*(x) \geqslant I^*(x) = \max\{x, -1\}$.

Now suppose $g \in \mathcal{G}$. Similarly, g is extended to be $+\infty$ in $(0, +\infty)$. $g^*(y) = \sup_{x \leqslant 0} yx - g(x)$ and $g^*(y) = +\infty$ if $y < 0$. By a similar argument, g^* is increasing. Since $g \geqslant I^*$, we have $g^* \leqslant I^{**} = I$. That is, $g^*(x) \leqslant 1 - x$. Let J be zero on $(-\infty, 0]$ and infinity otherwise. We have J^* is zero on $[0, +\infty)$ and infinity otherwise. We know that $g(0) = 0$ and g is increasing so $g \leqslant J$. Hence $g^* \geqslant J^*$, i.e. $g^*(y) \geqslant 0$ if $y \geqslant 0$. This justifies that $g^*(y) \in [0, 1]$ if $y \in [0, 1]$ and $g^*(y) = 0$ if $y \geqslant 1$. □

Now with the help of this claim, $\mathcal{H}$ and $\mathcal{G}$ are simply related: $h \in \mathcal{H}$ iff $\alpha \mapsto h(-\alpha) - 1$ is in $\mathcal{G}$. Therefore we can get the description of $\mathcal{H}$. The proof of the first statement is complete.

Explicit construction. Next we derive the specific choice of P, Q as stated works using the result from Dong et al. [2021].

Continuing with the notations in the proof above, when H satisfies the conditions, i.e. $H \in \mathcal{H}$, we know there is a $f \in \mathcal{F}$ such that $H(\alpha) = 1 + f^*(-\alpha)$. Let $g(\alpha) = H(-\alpha) - 1$ and we will have $g = f^*$ and hence $f = g^*$ as f is convex. Therefore,

$$f(x) = g^*(x) = \sup_y yx - H(-y) + 1 = \sup_z -zx - H(z) + 1 = 1 + H^*(-x).$$

From Dong et al. [2021, Proposition 2.2], we know that $f = T[Q, P]$ where $Q = U[0, 1]$ is the uniform distribution over $[0, 1]$ and P has CDF

$$F_P(x) = \begin{cases} 0, & \text{if } x < 0, \\ f(1 - x), & \text{if } x \in [0, 1), \\ 1, & \text{if } x \geqslant 1. \end{cases}$$

Plugging in $f(x) = 1 + H^*(-x)$, we have the CDF of P being

$$F_P(x) = \begin{cases} 0, & \text{if } x < 0, \\ 1 + H^*(1 - x), & \text{if } x \in [0, 1), \\ 1, & \text{if } x \geqslant 1. \end{cases}$$

Note that when the infimum of H is positive, $H^*(1 - x) < 0$ and P has an atom at 1. This completes the proof. □

Another interesting consequence of Lemma 9 is one can often get a stronger bound on the hockey-stick divergence or privacy profile *for free*. Recall that for a function g, its convex hull $\mathrm{conv}(g)$ (a.k.a., the lower convex envelope) is defined as the greatest convex lower bound of g and satisfies $\mathrm{conv}(g) = g^{**}$ where the double star means taking Fenchel conjugate twice.

For a function $h : \mathbb{R}_+ \to \mathbb{R}$, let $g(x) = \inf_{y\in[0,x]} h(y)$ and $\mathrm{HS}(h) = (\min\{1, g\})^{**}$. It turns out that $\mathrm{HS}(h)$ is the greatest lower bound of h that lies in $\mathcal{H}$, and we have

Corollary 26 (Dominating pairs from any privacy profile upper bounds)**.** *If the privacy profile of a mechanism $\mathcal{M}$ is bounded by $h : \mathbb{R}_+ \to \mathbb{R}$, i.e. $\delta_{\mathcal{M}}(\alpha) \leqslant h(\alpha), \forall \alpha \geqslant 0$, then $\delta_{\mathcal{M}}$ is also bounded by* $\mathrm{HS}(h)$.

Note that $\mathrm{HS}(h)$ can be significantly smaller than the original bound h, and it admits a dominating pair by Proposition 9, even if h does not.

Proof. We know that $\delta_{\mathcal{M}} \in \mathcal{H}$. It suffices to show that

$$f \in \mathcal{H}, f \leqslant h \implies f \leqslant \mathrm{HS}(h).$$

Recall that we let $g(x) = \inf_{y\in[0,x]} h(y)$ and $\mathrm{HS}(h) = (\min\{1, g\})^{**}$. Since $f \in \mathcal{H}$ is decreasing, $f(x) = \inf_{y\in[0,x]} f(y) \leqslant \inf_{y\in[0,x]} h(y) = g(x)$. Furthermore, $f(x) \leqslant f(0) = 1$, so $f \leqslant \min\{1, g\}$. Since f is convex, it also holds that $f \leqslant \min\{1, g\}^{**} = \mathrm{HS}(h)$. □

C.2 Composition theorem of dominating pairs

Theorem 27 (Restatement of Theorem 10 Adaptive composition of dominating pairs)**.** *Let P, Q be a dominating pair distributions for $\mathcal{M}$ and P', Q' be a dominating pair distributions for $\mathcal{M}'$[7], then $(P\times P', Q\times Q')$ is a dominating pair distributions for the composed mechanism $(\mathcal{M}, \mathcal{M}')$.*

Proof.

$$H_\alpha(P\|P') = \int_\Omega [p(\omega) - \alpha p'(\omega)]_+ \,\mathrm{d}\omega.$$

Integration with respect to a dominating measure of both P and Q and p, q are the densities (Radon-Nikodym derivatives) for the probability measures P, Q respectively.

Our goal is to show $H_\alpha\big(M(D), M(D')\big) \leqslant H_\alpha\big(P \times R, Q \times S\big)$. We break it into the following two parts.

$$H_\alpha\big(M(D), M(D')\big) \leqslant H_\alpha\big(M_1(D) \times R, M_1(D') \times S\big) \leqslant H_\alpha\big(P \times R, Q \times S\big).$$

Starting from the first part, we have

$$\begin{aligned}
H_\alpha\big(M(D), M(D')\big) &= \iint_{X\times Y} [p_1(x)p_2(x,y) - \alpha p_1'(x)p_2'(x,y)]_+ \,\mathrm{d}x\,\mathrm{d}y \\
&= \int_X p_1(x) \cdot \left(\int_Y \left[p_2(x,y) - \alpha \cdot \frac{p_1'(x)}{p_1(x)} \cdot p_2'(x,y) \right]_+ \mathrm{d}y \right) \mathrm{d}x \\
&= \int_X p_1(x) \cdot \left(H_{\alpha\cdot\frac{p_1'(x)}{p_1(x)}} \big(M_2(D,x)\|M_2(D',x)\big) \right) \mathrm{d}x \\
&\leqslant \int_X p_1(x) \cdot \left(H_{\alpha\cdot\frac{p_1'(x)}{p_1(x)}} \big(R\|S\big) \right) \mathrm{d}x \\
&= \int_X p_1(x) \cdot \left(\int_{\Omega_2} \left[r(\omega_2) - \alpha \cdot \frac{p_1'(x)}{p_1(x)} \cdot s(\omega_2) \right]_+ \mathrm{d}\omega_2 \right) \mathrm{d}x \\
&= \iint_{X\times\Omega_2} [p_1(x)r(\omega_2) - \alpha p_1'(x)s(\omega_2)]_+ \,\mathrm{d}x\,\mathrm{d}\omega_2 \\
&= H_\alpha\big(M_1(D) \times R, M_1(D') \times S\big).
\end{aligned}$$

[7] $\mathcal{M}'$ can be adaptively chosen in that it could depend on the output of $\mathcal{M}$, which requires $\sup_{o\in\mathrm{Range}(\mathcal{M})} H_{e^\epsilon}(\mathcal{M}'(D,o)\|\mathcal{M}'(D',o)) \leq H_{e^\epsilon}(P'\|Q')$ for any value of o.

Continuing this argument, we have

$$\begin{aligned}
H_\alpha\big(M_1(D)\times R, M_1(D')\times S\big) &= \iint_{X\times\Omega_2} [p_1(x)r(\omega_2) - \alpha p_1'(x)s(\omega_2)]_+ \,\mathrm{d}x\,\mathrm{d}\omega_2 \\
&= \int_{\Omega_2} r(\omega_2)\cdot\left(\int_X \left[p_1(x) - \alpha\cdot\frac{s(\omega_2)}{r(\omega_2)}\cdot p_1'(x)\right]_+ \mathrm{d}x\right)\mathrm{d}\omega_2 \\
&= \int_{\Omega_2} r(\omega_2)\cdot\left(H_{\alpha\cdot\frac{s(\omega_2)}{r(\omega_2)}}\big(M_1(D)\|M_1(D')\big)\right)\mathrm{d}\omega_2 \\
&\leqslant \int_{\Omega_2} r(\omega_2)\cdot\left(H_{\alpha\cdot\frac{s(\omega_2)}{r(\omega_2)}}\big(P\|Q\big)\right)\mathrm{d}\omega_2 \\
&= \int_{\Omega_2} r(\omega_2)\cdot\left(\int_X \left[p(\omega_1) - \alpha\cdot\frac{s(\omega_2)}{r(\omega_2)}\cdot q(\omega_1)\right]_+ \mathrm{d}\omega_1\right)\mathrm{d}\omega_2 \\
&= \iint_{\Omega_1\times\Omega_2} [p(\omega_1)r(\omega_2) - \alpha q(\omega_1)s(\omega_2)]_+ \,\mathrm{d}\omega_1\,\mathrm{d}\omega_2 \\
&= H_\alpha\big(P\times R, Q\times S\big).
\end{aligned}$$

The proof is complete. □

C.3 Privacy-amplification for dominating pairs

Recall we stated the following theorem in the main body:

Theorem 11. *Let $\mathcal{M}$ be a randomized algorithm.*

(1) *If (P,Q) dominates $\mathcal{M}$ for add neighbors then $(P,(1-\gamma)P+\gamma Q)$ dominates $\mathcal{M}\circ S_{\textbf{Poisson}}$ for add neighbors and $((1-\gamma)Q+\gamma P,Q)$ dominates $\mathcal{M}\circ S_{\textbf{Poisson}}$ for removal neighbors.*

(2) *If (P,Q) dominates $\mathcal{M}$ for replacing neighbors, then $(P,(1-\gamma)P+\gamma Q)$ dominates $\mathcal{M}\circ S_{\textbf{Subset}}$ for add neighbors and $((1-\gamma)P+\gamma Q,P)$ dominates $\mathcal{M}\circ S_{\textbf{Subset}}$ for removal neighbors.*

The proof we present here is written in the language of trade-off functions [Dong et al., 2021]. However, with Lemma 20, everything can be conveniently translated to the language of (ε,δ). We made the choice because some parameters have slightly easier forms in the language of trade-off functions.

We begin with a lemma that cut our workload in half — dominance for removal neighbors is actually equivalent to the dominance of add neighbors, so it suffices to show either one of them.

Lemma 28. *The followings are equivalent*

1. *(P,Q) dominates mechanism $\mathcal{M}$ for add neighbors.*
2. *(Q,P) dominates mechanism $\mathcal{M}$ for removal neighbors.*
3. *$T[M(S),M(S\cup\{x\})]\geqslant T[P,Q]$ for any dataset S and data entry x.*

Proof of Lemma 28. Recall that Lemma 20 says for any $\alpha>0$,

$$H_\alpha(P\|Q) = 1 + T[P,Q]^*(-\alpha).$$

$$\begin{aligned}
\text{condition 1} &\Leftrightarrow H_\alpha(M(S)\|M(S\cup\{x\}))\leqslant H_\alpha(P\|Q) \\
&\Leftrightarrow T[M(S\cup\{x\}),M(S)]\geqslant T[Q,P] && \text{(Lemma 20)} \\
&\Leftrightarrow T[M(S),M(S\cup\{x\})]\geqslant T[P,Q]\Leftrightarrow\text{condition 3} && (*) \\
&\Leftrightarrow H_\alpha(M(S\cup\{x\}),M(S))\leqslant H_\alpha(Q\|P)\Leftrightarrow\text{condition 2}
\end{aligned}$$

Here $(*)$ uses the fact that $T[P,Q]$ and $T[Q,P]$ are inverse functions of each other. □

The next lemma plays the central role in both parts of Theorem 11. Suppose we have probability distributions $P_1, \dots, P_n$ and $Q_1, \dots, Q_n$, all on the same domain and let $\bar{P} = \sum_{i=1}^n p_i P_i$ and $\bar{Q} = \sum_{i=1}^n p_i Q_i$ be the corresponding mixture distributions with the same coefficients $p_1, \dots, p_n$ where $\sum_{i=1}^n p_i = 1$ and all $p_i \geqslant 0$. Then we have

Lemma 29. *If $T[P_i, Q_i] \geqslant T[P, Q]$ for all $i \in [n]$, then for any $\gamma \in [0,1]$, we have*

$$T\big[\bar{P}, (1-\gamma)\bar{P} + \gamma\bar{Q}\big] \geqslant T[P, (1-\gamma)P + \gamma Q].$$

Proof of Lemma 29. Let $f = T[P, Q]$. First we claim that

$$T[P, (1-\gamma)P + \gamma Q](x) = (1-\gamma)(1-x) + \gamma f(x). \tag{2}$$

To see this, consider any testing rule ϕ such that $\mathbb{E}_P[\phi] = x$. We need to show that

$$\inf_{\phi : \mathbb{E}_P[\phi] = x} \mathbb{E}_{(1-\gamma)P+\gamma Q}[1-\phi] = (1-\gamma)(1-x) + \gamma f(x).$$

By definition of $f = T[P, Q]$, we have $\inf_{\phi:\mathbb{E}_P[\phi]=x} \mathbb{E}_Q[1-\phi] = f(x)$. Therefore,

$$\begin{aligned}
\mathbb{E}_{(1-\gamma)P+\gamma Q}[1-\phi] &= (1-\gamma)\mathbb{E}_P[1-\phi] + \gamma\mathbb{E}_Q[1-\phi] \\
&= (1-\gamma)(1-x) + \gamma\mathbb{E}_Q[1-\phi] \\
\inf_{\phi:\mathbb{E}_P[\phi]=x} \mathbb{E}_{(1-\gamma)P+\gamma Q}[1-\phi] &= (1-\gamma)(1-x) + \gamma \inf_{\phi:\mathbb{E}_P[\phi]=x} \mathbb{E}_Q[1-\phi] \\
&= (1-\gamma)(1-x) + \gamma f(x).
\end{aligned}$$

This verifies (2). Next we proceed to the proof of the lemma. Similarly, it suffices to consider arbitrary testing rules ϕ with $\mathbb{E}_{\bar{P}}[\phi] \leqslant x$ and show

$$\mathbb{E}_{(1-\gamma)\bar{P}+\gamma\bar{Q}}[1-\phi] \geqslant T[P, (1-\gamma)P + \gamma Q](x) = (1-\gamma)(1-x) + \gamma f(x). \tag{3}$$

Expanding the convex combination, we have

$$\begin{aligned}
\mathbb{E}_{(1-\gamma)\bar{P}+\gamma\bar{Q}}[1-\phi] &= (1-\gamma)\mathbb{E}_{\bar{P}}[1-\phi] + \gamma\sum p_i \mathbb{E}_{Q_i}[1-\phi] \\
&\geqslant (1-\gamma)(1-x) + \gamma\sum p_i \mathbb{E}_{Q_i}[1-\phi]
\end{aligned}$$

Comparing to (3), it suffices to show

$$\sum p_i \mathbb{E}_{Q_i}[1-\phi] \geqslant f(x)$$

We know that $T[P_i, Q_i] \geqslant f$. Hence $\mathbb{E}_{Q_i}[1-\phi] \geqslant f(\mathbb{E}_{P_i}[\phi])$. By convexity of f,

$$\sum p_i \mathbb{E}_{Q_i}[1-\phi] \geqslant \sum p_i f(\mathbb{E}_{P_i}[\phi]) \geqslant f\left(\sum p_i \mathbb{E}_{P_i}[\phi]\right) = f(\mathbb{E}_{\bar{P}}[\phi]) \geqslant f(x).$$

The last inequality follows from the monotonicity of trade-off functions. Hence (3) is verified and the proof is complete. □

Proof of Theorem 11 (1). By Lemma 28, it suffices to prove

$$T\big[M(S_{\mathbf{Poisson}}), M(S'_{\mathbf{Poisson}})\big] \geqslant T[P, (1-\gamma)P + \gamma Q] \tag{4}$$

where $S = \{x_1, \dots, x_{n-1}\}$ and $S' = S \cup \{x_n\}$. The outcome of repeated coin flips can be labeled as $\vec{b} \in \{0,1\}^{n-1}$. We use $S_{\vec{b}}$ to denote the corresponding subset of S and $S'_{\vec{b}0}$ or $S'_{\vec{b}1}$ for that of S', depending on whether x_n is included. Note that $S'_{\vec{b}0} = S_{\vec{b}}$. Furthermore, let $p_{\vec{b}}$ be the probability of the outcome $\vec{b}$ (recall that each coin is a Bernoulli γ random variable). In fact, $p_{\vec{b}} = \prod_{i=1}^{n-1} \gamma^{b_i}(1-\gamma)^{1-b_i}$ but we will not use it.

Both $M(S_{\mathbf{Poisson}})$ and $M(S'_{\mathbf{Poisson}})$ are mixtures. We have the following decompositions

$$\begin{aligned} M(S_{\mathbf{Poisson}}) &= \sum_{\vec{b}\in\{0,1\}^{n-1}} p_{\vec{b}} M(S_{\vec{b}}) \\ M(S'_{\mathbf{Poisson}}) &= \sum_{\vec{b}\in\{0,1\}^{n-1}} (1-\gamma) p_{\vec{b}} M(S'_{\vec{b}0}) + \gamma p_{\vec{b}} M(S'_{\vec{b}1}) \\ &= (1-\gamma) M(S_{\mathbf{Poisson}}) + \gamma \sum_{\vec{b}\in\{0,1\}^{n-1}} p_{\vec{b}} M(S'_{\vec{b}1}) \end{aligned}$$

Now we are ready to use Lemma 29, with the family of P_i being $M(S_{\vec{b}})$ and the family of Q_i being $M(S'_{\vec{b}1})$. By the calculation above, $M(S_{\mathbf{Poisson}})$ will be the $\bar{P}$ in Lemma 29 and $M(S'_{\mathbf{Poisson}})$ is exactly the $(1-\gamma)\bar{P}+\gamma\bar{Q}$ in Lemma 29. We still need to verify the condition in Lemma 29: since (P, Q) dominates M for add neighbors, for each $\vec{b} \in \{0,1\}^{n-1}$ we have

$$T\big[M(S_{\vec{b}}), M(S'_{\vec{b}1})\big] \geqslant T[P, Q].$$

Therefore, Lemma 29 gives us (4), which is exactly what we want. □

Proof of Theorem 11 (2). By Lemma 28, it suffices to prove

$$T\big[M(S_{\mathbf{Subset}}), M(S'_{\mathbf{Subset}})\big] \geqslant T[P, (1-\gamma)P + \gamma Q] \tag{5}$$

where $S = \{x_1, \ldots, x_{n-1}\}$ and $S' = S \cup \{x_n\}$.

Below we use notations such as $S_I, S_{J\cup\{n\}}$ to denote the (obvious) subsets of S where $I, J \subseteq [n-1]$ and $|I| = m, |J| = m-1$ consistently. Both $M(S_{\mathbf{Subset}})$ and $M(S'_{\mathbf{Subset}})$ are mixtures. We have the following decompositions, where the latter is further decomposed into two parts depending on whether n is selected.

$$M(S_{\mathbf{Subset}}) = \sum_{|I|=m} \frac{1}{\binom{n-1}{m}} M(S_I) \tag{6}$$

$$M(S'_{\mathbf{Subset}}) = \sum_{|J|=m-1} \frac{1}{\binom{n}{m}} M(S_{J\cup\{n\}}) + \sum_{|I|=m} \frac{1}{\binom{n}{m}} M(S_I) \tag{7}$$

It's not in a ready shape to use Lemma 29. We need to further break the summands by carefully creating copies of the components. Let

$$\begin{aligned} \mathcal{I} &:= \{I^{(k)} : I \subseteq [n-1], |I| = m, 1 \leqslant k \leqslant m\} \\ \mathcal{J} &:= \{J^{(l)} \cup \{n\} : J \subseteq [n-1], |J| = m-1, 1 \leqslant l \leqslant n-m\}. \end{aligned}$$

That is, we create m copies of each subset of $[n-1]$ of cardinality m and collect as $\mathcal{I}$; create $n-m$ copies of each subset of $[n]$ of cardinality m that includes n and collect as $\mathcal{J}$. Now we claim two things

1. There is a bijection $F : \mathcal{I} \to \mathcal{J}$ such that $F(I)$ differs from I by one element for all $I \in \mathcal{I}$.
2. Let $p_I = \frac{1}{\binom{n-1}{m} m}$ for all $I \in \mathcal{I}$ and $\gamma = \frac{m}{n}$. Then

$$M(S_{\mathbf{Subset}}) = \sum_{I\in\mathcal{I}} p_I M(S_I) \tag{8}$$

$$M(S'_{\mathbf{Subset}}) = \gamma \sum_{I\in\mathcal{I}} p_I M(S_{F(I)}) + (1-\gamma) \sum_{I\in\mathcal{I}} p_I M(S_I) \tag{9}$$

Now we are ready to use Lemma 29: the collection $\{P_i\}$ is $\{M(S_I)\}$ and $\{Q_i\}$ is $\{M(S_{F(I)})$. The conclusion is exactly (5).

Next we turn our attention to the proofs of the two claims.

For claim 1, let's first construct a $(n-m)$-to-one surjective map $\tilde{F}: \mathcal{I} \to \tilde{\mathcal{J}}$ where $\tilde{\mathcal{J}} = \{J \cup \{n\} : J \subseteq [n-1], |J| = m-1\}$. $\tilde{F}(I^{(k)})$ is obtained by replacing the k-th element in I by n. We see that $\tilde{F}(I^{(k)})$ is indeed in $\tilde{\mathcal{J}}$. For any $J \cup \{n\} \in \tilde{\mathcal{J}}$, it is hit by $\tilde{F}$ exactly $n-m$ times since the n could have been any of the $(n-1)-(m-1) = n-m$ indices not already in J.

Since $\mathcal{J}$ contains $n-m$ copies of $\tilde{\mathcal{J}}$, the $(n-m)$-to-one mapping $\tilde{F}: \mathcal{I} \to \tilde{\mathcal{J}}$ can be "redirected" in an obvious way and become a bijection between $\mathcal{I}$ and $\mathcal{J}$. By construction, $F(I)$ and I differ in exactly one element.

For claim 2, (8) is easier and we only show (9). Since we have created copies, by splitting the summands of (7), we have

$$\begin{aligned} M(S'_{\textbf{Subset}}) &= \sum_{|J|=m-1} \frac{1}{\binom{n}{m}} M(S_{J\cup\{n\}}) + \sum_{|I|=m} \frac{1}{\binom{n}{m}} M(S_I) \\ &= \sum_{J\in\mathcal{J}} \frac{1}{\binom{n}{m}\cdot(n-m)} M(S_J) + \sum_{I\in\mathcal{I}} \frac{1}{\binom{n}{m} m} M(S_I) \\ &= \sum_{I\in\mathcal{I}} \frac{1}{\binom{n}{m}\cdot(n-m)} M(S_{F(I)}) + \sum_{I\in\mathcal{I}} \frac{1}{\binom{n}{m} m} M(S_I) \end{aligned}$$

Comparing to (9), it suffices to show

$$\gamma p_I = \frac{1}{\binom{n}{m}\cdot(n-m)} \text{ and } (1-\gamma)p_I = \frac{1}{\binom{n}{m} m}$$

which are elementary combination identities. □

Remark (Exact optimality of the bounds). *If (P, Q) is a* tightly dominating pair *for $\mathcal{M}$, for both "Removal"-neighboring relation or "Add"-neighboring relation, then under some mild regularity conditions on $\mathcal{M}$ and the space of the input datasets, Theorem 11 can be strengthened to show that that $((1-\gamma)Q + \gamma P, Q)$ and $(P, (1-\gamma)P + \gamma Q)$ are* tight dominating pairs *for the "Removal"-neighboring relation and "Add"-neighboring relation respectively — i.e., the dominating pair is realized by some concrete datasets. For example, consider $\mathcal{M}$ to be Gaussian mechanism or Laplace mechanism that releases the total number of 1s in a dataset. Then two neighboring datasets $D = [0, ..., 0, 0, 1] \in \mathbb{R}^{n+1}$, $D' = [0, 0, ..., 0] \in \mathbb{R}^n$ for "removal" and $D = [0, ..., 0] \in \mathbb{R}^n$, $D' = [0, 0, ..., 0, 1] \in \mathbb{R}^{n+1}$ for "addition" attains the upper bound for all $\alpha > 0$ in each category.*

Remark (Renyi DP and Optimal Moments Accountant for subsampled mechanisms). *Renyi-DP and moments accountant are closely related concepts that are often considered identical. However, our results suggest that there is a distinction. The above pair of P, Q we constructed are not necessarily attaining the Renyi-DP bounds (see a concrete example from Zhu and Wang [2019], but as moments accountant focuses only on computing (ϵ, δ)-DP, it suffices use the Renyi-divergence functions $R_\alpha(P\|Q)$. Specifically, this closes the constant gap between the moments accountant for subsampled mechanisms and Poisson sampled mechanisms.*

C.4 Other schemes in privacy-amplification for dominating pairs

Theorem 11 provides new results and a novel composition algorithm for the popular Poisson sampling under "add/remove" neighboring relations by treating "add" and "remove" separately. It also shows that the same result hold in a not-so-typical but practically relevant scheme of the *random-subset* sampled mechanism under the "add / remove" neighboring relation for the case when the base mechanism's privacy is defined by the "replace" neighboring relation.

What we left unspecified is whether there is a clean dominating pair under the two alternative schemes: (a) Poisson-sampling + "Add/remove" *without* treating "Add" or "Remove" separately; (b) Subset sampling + "Replace" neighboring relation for both $\mathcal{M} \circ \text{Sample}_\gamma$ and $\mathcal{M}$.

It turns out that while we can construct these dominating pairs of $\mathcal{M} \circ \text{Sample}_\gamma$ explicitly based on the dominating pair (P, Q) of $\mathcal{M}$, but the expression is not simple. We present these results in this section.

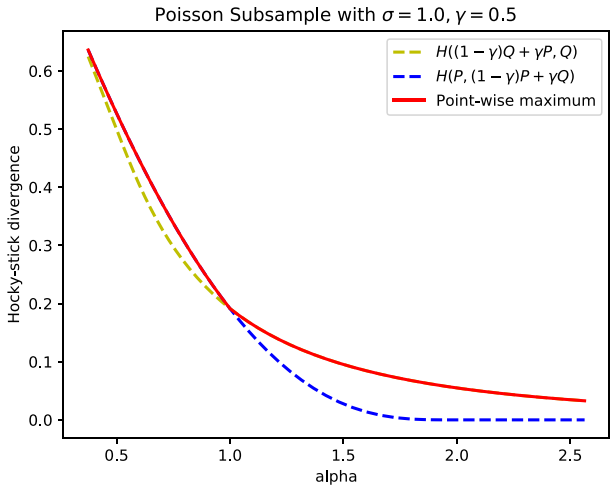

(a) HS divergence

Figure 4: We consider Poisson subsample Gaussian mechanism with $\sigma = 1.0$ and $\gamma = 0.5$. The red line denotes the pointwise maximum over two HS curves. Figure shows that $H_\alpha((1-\gamma)Q + \gamma P, Q)$ dominates the region $\alpha \geq 1$ and $H_\alpha(P, (1-\gamma)P + \gamma Q)$ dominates the region $\alpha < 1$.

To avoid any confusion, for all practical purposes, the result in Theorem 11 suffices because we can always compose "Add" or "Remove" separately and only take the pointwise maximum in the end, while only incurring twice as much computation, but the results in this section are interesting from a purely scientific perspective and they are included for the completeness in our understanding of the problem.

Proposition 30. *If (P, Q) is a dominating pair of $\mathcal{M}$ under "Add/remove" Relation, then*

$$\delta_{\mathcal{M} \circ S_{\textbf{Poisson}}}(\alpha) \leq \begin{cases} H_\alpha((1-\gamma)Q + \gamma P, Q) & \text{for } \alpha \geq 1; \\ H_\alpha(P, (1-\gamma)P + \gamma Q) & \text{for } 0 < \alpha < 1. \end{cases}$$

under the "Add/Remove" relation. Similarly, if (P, Q) is a dominating pair of $\mathcal{M}$ under "Replace" relation for dataset of size γn, then

$$\delta_{\mathcal{M} \circ S_{\textbf{Subset}}}(\alpha) \leq \begin{cases} H_\alpha((1-\gamma)Q + \gamma P, Q) & \text{for } \alpha \geq 1; \\ H_\alpha(P, (1-\gamma)P + \gamma Q) & \text{for } 0 < \alpha < 1. \end{cases}$$

under "Replace" relation for dataset of size n.

We plot $H_\alpha((1-\gamma)Q + \gamma P, Q)$ and $H_\alpha(P, (1-\gamma)P + \gamma Q)$ for $\delta_{\mathcal{M} \circ S_{\textbf{Poisson}}}(\alpha)$ in Figure 4(a).

The proof of the above result requires the use of the following general result that establishes the relationship between pairs that dominate only one half of the range for α and those that dominate the other half.

Lemma 31 (Properties for "symmetric neighbors")**.** *Let $\mathcal{M}$ be a mechanism and $\simeq$ be a* symmetric *neighboring relationships , i.e., $D \simeq D' \Leftrightarrow D \simeq D'$. Then*

1. *If (P, Q) is a dominating pair of $\mathcal{M}$, then (Q, P) is also a dominating pair of $\mathcal{M}$,*
2. *If $\sup_{D \simeq D'} H_\alpha(\mathcal{M}(D) \| \mathcal{M}(D')) \leq \min\{H_\alpha(P \| Q), H_\alpha(Q \| P)\}$ for all $\alpha \geq 1$, then (P, Q) and (Q, P) are both dominating pairs of $\mathcal{M}$.*
3. *The following two statements are equivalent.*
   (a) $\sup_{D \simeq D'} H_\alpha(\mathcal{M}(D) \| \mathcal{M}(D')) \leq H_\alpha(P \| Q)$ *for all* $\alpha \geq 1$.
   (b) $\sup_{D \simeq D'} H_\alpha(\mathcal{M}(D) \| \mathcal{M}(D')) \leq H_\alpha(Q \| P)$ *for all* $0 < \alpha \leq 1$.

Proof. First notice that the first and second statements are both implied by the third. For the first statement, notice that by the definition of the dominating pair, the upper bound applies for all $\alpha > 0$. Thus by applying $(a) \Rightarrow (b)$ and $(b) \Rightarrow (a)$ for (P, Q), we get that (Q, P) is also a dominating pair. For the second statement, we apply $(a) \Rightarrow (b)$ for both (P, Q) and (Q, P), then the result extends the bound to the full range.

It remains to prove the third statement. For any pair of neighboring D, D' and $0 < \alpha \leq 1$, by Lemma 46.

$$\begin{aligned} H_\alpha(\mathcal{M}(D)\|\mathcal{M}(D')) &= \alpha H_{\alpha-1}(\mathcal{M}(D')\|\mathcal{M}(D)) - \alpha + 1 \\ &\leq \alpha H_{\alpha^{-1}}(P\|Q) - \alpha + 1 \\ &= H_\alpha(Q\|P), \end{aligned}$$

where the inequality uses the fact that $\alpha^{-1} \geq 1$, $\simeq$ is symmetric, and that (P, Q) dominates for order ≥ 1. The converse follows the same argument but starts with $\alpha > 1$. □

Proof of Proposition 30. We first prove the case for $\alpha \geq 1$. By Theorem 8 of [Balle et al., 2018] (Standard amplification by sampling bound in DP), we know that for any $D \simeq_{\textbf{add/remove}} D'$ and all $\epsilon \geq 0$ (thus $\alpha \geq 1$!)

$$\begin{aligned} &H_{e^\epsilon}(\mathcal{M} \circ S^\gamma_{\textbf{Poisson}}(D)\|\mathcal{M} \circ S^\gamma_{\textbf{Poisson}}(D')) \leq \gamma \sup_{\tilde{D} \simeq_{\text{add/remove}} \tilde{D}'} H_{1+\frac{e^\epsilon-1}{\gamma}}(\mathcal{M}(\tilde{D})\|\mathcal{M}(\tilde{D}')) \\ \leq &\gamma H_{1+\frac{e^\epsilon-1}{\gamma}}(P\|Q) = \gamma \int \Big(p(x) - (1+\tfrac{e^\epsilon-1}{\gamma})q(x)\Big)_+ \,\mathrm{d}x \\ = &\int ((1-\gamma)q(x) + \gamma p(x) - e^\epsilon q(x))_+ \,\mathrm{d}x = H_{e^\epsilon}((1-\gamma)Q + \gamma P\|Q). \end{aligned}$$

where the inequality in the second line is due to that (P, Q) is a dominating pair of $\mathcal{M}$.

The same proof works line by line for the random subset sampling when we use Theorem 9 of [Balle et al., 2018] instead and adopt the $D \simeq_{\textbf{replace}} D'$ neighboring relationship.

Next we prove the statement for $0 < \alpha < 1$. Check that we can apply Lemma 31 because both $\simeq_{\textbf{replace}}$ and $\simeq_{\textbf{add/remove}}$ are symmetric. By the third statement of Lemma 31, $(Q, (1-\gamma)Q + \gamma P)$ dominates the subsampled mechanism for $0 < \alpha < 1$. Also by the first statement of Lemma 31 we know that (Q, P) is also a dominating pair for $\mathcal{M}$, thus by repeating the same argument, we get that $(P, (1-\gamma)P + \gamma Q)$ also dominates the subsampled mechanism for $0 < \alpha < 1$, which completes the proof. □

The above discussion characterizes the tight[8] upper bound of the privacy profile of subsampled mechanisms using two (ordered) pairs of distributions, rather than just one pair. This is insufficient for us to apply the composition theorem because the region between $\alpha > 1$ and $\alpha < 1$ in some sense "mixes with each other" during composition, as we have clearly seen from the proof of Theorem 10. What we do know is that neither $((1-\gamma)Q + \gamma P, Q)$ nor $(Q, (1-\gamma)Q + \gamma P)$ is a dominating pair for the sampled mechanism under "Add/Remove" or "Replace" neighboring relation. This is the reason why we proposed the more elegant approach for handling "add" and "remove" separately in the first place.

For completeness, and to also handle the case when we want the "replace one" neighboring relation for $\mathcal{M} \circ S_{\textbf{Subset}}$, we state the following result which constructs an explicit but not-so-clean dominating pair.

Corollary 32 (Dominating pair of sampled mechanisms under symmetric neighbor relations)**.** *Let (P, Q) be a dominating pair of $\mathcal{M}$ and $\simeq$ is "Add/Remove" (or "Replace"). Then $(\tilde{P}, \tilde{Q})$ is a dominating pair of $\mathcal{M} \circ S_{\textbf{Poisson}}$ (or $\mathcal{M} \circ S_{\textbf{Subset}}$) where $\tilde{Q} = \mathit{Uniform}[0, 1]$ and $\tilde{P}$ has a CDF of*

$$F_{\tilde{P}}(x) = \begin{cases} 0, & \text{if } x < 0; \\ 1 + \Big[\max\{H_\alpha((1-\gamma)Q + \gamma P, Q), H_\alpha(P, (1-\gamma)P + \gamma Q)\}\Big]^*(-x), & \text{if } x \in [0, 1), \\ 1, & \text{if } x \geq 1; \end{cases}$$

where $\big[\,\cdot\,\big]^*$ denotes the Fenchel conjugate *of a function.*

[8] They are tight for the reason we described in the Remark on the "Exact optimality of the bounds" above.

The proof, which we omit, is a direct application of Proposition 30 with the generic approach of Proposition 9 for constructing the dominating pair.

D The characteristic function of basic mechanisms

D.1 ϕ-function of basic mechanisms

We now derive ϕ-function for three basic mechanisms: randomized response, Laplace and Gaussian mechanism. The results and their dominating distributions are summarized in Table 2.

Let f be a predicate, i.e., $f : \mathcal{D}^* \to \{0, 1\}$. The Randomized Response mechanism for f is defined as

$$\mathcal{M}(D) = \begin{cases} f(D) & \text{with probabilty } p \\ 1 - f(D) & \text{with probability} 1 - p \end{cases}$$

Lemma 33 (Randomized response)**.** *The ϕ function of Randomized Response mechanism $\mathcal{M}$ with the parameter p satisfies $\phi_{\mathcal{M}}(\alpha) = \phi'_{\mathcal{M}}(\alpha) = pe^{\alpha i \log(\frac{p}{1-p})} + (1-p)e^{\alpha i \log(\frac{1-p}{p})}$.*

Proof. First of all, the dominating pair will be:

$$P : \Pr_P[0] = p; Q : \Pr_Q[1] = p$$

Then, follow the definition of ϕ-function, we have

$$\phi_{\mathcal{M}}(\alpha) = \mathbb{E}_{x\sim P} e^{i\alpha \log(p(x)/q(x))} = pe^{\alpha i \log(\frac{p}{1-p})} + (1-p)e^{\alpha i \log(\frac{1-p}{p})}$$

□

For Laplace and Gaussian mechanisms, we assume that $f : \mathcal{D}^* \to \mathcal{R}$ is a function of sensitivity 1.

Lemma 34 (Laplace Mechanism)**.** *Let Laplace Mechanism for f is defined as $\mathcal{M}(D) = f(D) + \mathrm{Lap}(\lambda)$ where $\mathrm{Lap}(\lambda)$ is Laplace dstribution with scale λ, i.e., its density function is $\frac{1}{2\lambda}\exp(-|x|/\lambda)$. For any $\alpha \in \mathcal{R}$ and $\lambda > 0$, we have*

$$\phi_{\mathcal{M}}(\alpha) = \phi'_{\mathcal{M}}(\alpha) = \frac{1}{2}\left(e^{\frac{i\alpha}{\lambda}} + e^{\frac{-i\alpha-1}{\lambda}} + \frac{1}{2\alpha i + 1}(e^{\frac{i\alpha}{\lambda}} - e^{\frac{-i\alpha-1}{\lambda}})\right).$$

Proof. we consider the dominating distribution $P(o) = \frac{1}{2\lambda}\exp(-|o|/\lambda)$ and $Q(o) = \frac{1}{2\lambda}\exp(-|o-1|/\lambda)$. We can show that the dominating pair for $\alpha \geq 1$ also dominates $0 < \alpha < 1$ using the third statement of Lemma 31 and the symmetry of Laplace mechanism. To calcuate the characteristic function $\phi_{\mathcal{M}}$, we define the privacy loss RV $L_{P,Q}$ as follows

$$L_{P,Q}(o) = \log\left(\frac{p(o)}{q(o)}\right) = \begin{cases} \frac{1}{\lambda}, & o \leq 0 \\ \frac{1-2o}{\lambda}, & 0 < o \leq 1 \\ \frac{-1}{\lambda}, & o > 1 \end{cases}$$

The characteristic function $\phi_{\mathcal{M}}(\alpha)$ is calculated as follows

$$\begin{aligned}
&\phi_{\mathcal{M}}(\alpha) = \mathbb{E}_P[e^{i\alpha L_{P,Q}}] \\
&= \int_{-\infty}^{0} e^{\frac{i\alpha}{\lambda}} \frac{1}{2\lambda} e^{\frac{o}{\lambda}} do + \int_{0}^{1} e^{\frac{i\alpha(1-2o)}{\lambda}} \frac{1}{2\lambda} e^{\frac{-o}{\lambda}} do + \int_{1}^{\infty} e^{\frac{-i\alpha}{\lambda}} \frac{1}{\lambda} e^{\frac{-o}{\lambda}} do \\
&= \frac{1}{2\lambda} \cdot \left(e^{\frac{i\alpha}{\lambda}}\lambda + \frac{\lambda}{2it+1}(e^{\frac{i\alpha}{\lambda}} - e^{\frac{-i\alpha-1}{\lambda}}) + \lambda e^{-1/\lambda} e^{\frac{-i\alpha}{\lambda}}\right) \\
&= \frac{1}{2}\left(e^{\frac{i\alpha}{\lambda}} + \frac{1}{2i\alpha+1}(e^{\frac{i\alpha}{\lambda}} - e^{\frac{-i\alpha-1}{\lambda}}) + e^{\frac{i\alpha+1}{-\lambda}}\right)
\end{aligned}$$

Similarly, for $\phi'_{\mathcal{M}}$, we define privacy loss RV $L_{Q,P}$ as follows

$$L_{Q,P}(o) = \log\left(\frac{q(o)}{p(o)}\right) = \begin{cases} \frac{-1}{\lambda}, & o \le 0 \\ \frac{2o-1}{\lambda}, & 0 < o \le 1 \\ \frac{1}{\lambda}, & o > 1 \end{cases}$$

Therefore, we have

$$\begin{aligned}\phi'_{\mathcal{M}}(\alpha) &= \int_{-\infty}^{0} e^{\frac{-|o-1|}{\lambda}} \cdot e^{\frac{-\alpha i}{\lambda}} do + \frac{1}{2\lambda}\int_{0}^{1} e^{\frac{-|o-1|}{\lambda}} e^{\frac{2o-1}{\lambda} i\alpha} do + \int_{1}^{\infty} e^{\frac{-|o-1|}{\lambda}} e^{\frac{i\alpha}{\lambda}} do \\ &= \frac{1}{2}\left(e^{\frac{i\alpha}{\lambda}} + \frac{1}{2i\alpha+1}(e^{\frac{i\alpha}{\lambda}} - e^{\frac{-i\alpha-1}{\lambda}}) + e^{\frac{i\alpha+1}{-\lambda}}\right)\end{aligned}$$

□

Lemma 35 (Gaussian Mechanism). *Let Gaussian mechanism is defined as $\mathcal{M}(D) = f(D) + \mathcal{N}(0, \sigma^2)$. For any $\alpha \in \mathcal{R}$ and $\sigma > 0$, we have $\phi_{\mathcal{M}}(\alpha) = \phi'_{\mathcal{M}}(\alpha) = e^{\frac{-1}{2\sigma^2}(\alpha^2 - i\alpha)}$.*

Proof. For Gaussian mechanism, no matter what the dimensionality of the output sapce is, the dominating distributions will always be 1D, and that extends to subsampled-gaussian as well. In the proof, we consider the worst-case pair $p(o) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(o-1)^2}{2\sigma^2}}, q(o) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{o^2}{2\sigma^2}}$.

Follow the definion of privacy loss RV, we have $L_{P,Q}(o) = \frac{2o-1}{2\sigma^2}$ and $L_{Q,P}(o) = \frac{1-2o}{2\sigma^2}$. Then we have

$$\phi_{\mathcal{M}}(\alpha)[e^{i\alpha \log(p/q)}] = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{(o-1)^2}{-2\sigma^2}} e^{\frac{(2o-1)i\alpha}{2\sigma^2}} do = e^{\frac{\alpha^2 - i\alpha}{-2\sigma^2}}$$

Similarly,

$$\phi'_{\mathcal{M}}(\alpha)[e^{i\alpha \log(p/q)}] = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{(o)^2}{-2\sigma^2}} e^{\frac{(1-2o)i\alpha}{2\sigma^2}} do = e^{\frac{\alpha^2 - i\alpha}{-2\sigma^2}}$$

□

Besides basic mechanisms, all mechanisms with discrete outputs admit an analytical ϕ function by definition.

Definition 36 (ϕ-function of mechanisms with discrete outputs). *Let p, q be the probability mass function induced by $\mathcal{M}$,*

$$\phi_{\mathcal{M}}(\alpha) = \sum_{o \in \mathcal{O}} e^{\log p(o) + i\alpha(\log p(o) - \log q(o))} \text{ and } \phi'_{\mathcal{M}}(\alpha) = \sum_{o \in \mathcal{O}} e^{\log q(o) + i\alpha(-\log p(o) + \log q(o))}.$$

This function can be represented by two vectors that lists the probability masses at o from p and q. When evaluating $\log \phi_{\mathcal{M}}(\alpha)$ at a given α, we could use the log-sum-exp trick to improve the numerical stability. Overall the space and time in representing these functions are linear in the size of the output space.

Provided that the worst-case pair of distributions are known, this procedure allows us to compose over exponential mechanisms [McSherry and Talwar, 2007], Report-noisy-max, as well as other complex mechanisms that arise out of post-processing of continuous output mechanisms, e.g., NoisyScreening [Zhu et al., 2020] that had been used as a practical alternative to sparse vector techniques.

D.2 Handling probability mass at ∞

One of the motivations of the work is for us to handle the situations where there is a non-zero probability mass where the privacy loss r.v. is at infinity. This naturally happens in propose-test-rease-style algorithm [Dwork and Lei, 2009] where we first construct a differentially private upper bound of the local sensitivity (or other data-dependent quantities) then calibrate noise according to the local sensitivity. The issue is that there is always a non-zero probability where the upper bound is not valid. Standard (ϵ, δ)-DP handles this case at ease, but modern techniques such as RDP struggles as such a mechanism does not satisfy RDP for any $\alpha > 0$.

In this case, Lemma 5 can be more explicitly rewritten into.

Lemma 37. *Let D, D' be the worst-case pair of datasets for $\mathcal{M}$, and $P = \mathcal{M}(D), Q = \mathcal{M}(D')$ then*

$$\delta_{\mathcal{M}}(\epsilon) = \max\{\delta_P(\epsilon), \delta_Q(\epsilon)\}$$

where

$$\delta_P(\epsilon) = \Pr[L_{P,Q} = \infty] + \left(\Pr[\epsilon < L_{P,Q} < \infty] - e^{\epsilon}\Pr[L_{Q,P} < -\epsilon]\right)$$
$$\delta_Q(\epsilon) = \Pr[L_{Q,P} = \infty] + \left(\Pr[\epsilon < L_{Q,P} < \infty] - e^{\epsilon}\Pr[L_{P,Q} < -\epsilon]\right)$$

Proposition 38 (Composition)**.** *Let $L^{(1)}_{P,Q'}, ..., L^{(k)}_{P,Q'}$ be the privacy loss R.V. of mechanism $\mathcal{M}_1, ..., \mathcal{M}_k$ ($L^{(i)}_{Q',P}$ is defined analogously). Then, the $\delta_P(\epsilon)$ and $\delta_Q(\epsilon)$ of the composed mechansims are as follow:*

$$\delta_P(\epsilon) = 1 - \prod_{i=1}^{k}(1 - \Pr[L^{(i)}_{P,Q'} = \infty]) + \left(\Pr[\epsilon < \tilde{L}_{P,Q'} < \infty] - e^{\epsilon}\Pr[\tilde{L}_{Q',P} < -\epsilon]\right)$$
$$\delta_Q(\epsilon) = 1 - \prod_{i=1}^{k}(1 - \Pr[L^{(i)}_{Q',P} = \infty]) + \left(\Pr[\epsilon < \tilde{L}_{Q',P} < \infty] - e^{\epsilon}\Pr[\tilde{L}_{P,Q'} < -\epsilon]\right)$$

where $\tilde{L}_{P,Q'}$ is defined as the privacy loss R.V. of the composed mechanisms that excludes ∞.

The above essentially says that we can handle the cases with ∞ separately, and compose the characteristic function of the sub-probability measure that excludes ∞.

E ϕ-function with discretization and experimental details

There are cases when the closed-form ϕ-functions do not exist. For example, in the subsample mechanisms, the privacy loss distribution is complicated and continuous, suggesting that we cannot derive an exact closed-form expression naively. In this section, we first provide a discretization-based solution and analyze its error bound. Later, we develop an efficient approximation method, "Double quadrature".

Algorithm 2 ϕ-function approximation

Input: The output interval $[-S, S]$, pdf $p(\cdot), q(\cdot)$ and N.

1: Set N equidistant points: $o_j = -S + j\triangle_o$, where $\triangle_o = 2S/N, j = 0, ..., N-1$.
2: $L^-_{P,Q}, L^-_{Q,P} \leftarrow$ the lower approximations of $L_{P,Q}, L_{Q,P}$ using Definition 39.
3: Construct $\phi^-_{\mathcal{M}_P}(\alpha)$: $\phi^-_{\mathcal{M}_P}(\alpha) \leftarrow \sum_{o_j} e^{\log \tilde{p}(o_j) + i\alpha L^-_{P,Q}(o_j)}$, where $\tilde{p}(o_j) = p(o_j)\triangle_o$.
4: Construct $\phi^-_{\mathcal{M}_Q}(\alpha)$: $\phi^-_{\mathcal{M}_Q}(\alpha) \leftarrow \sum_{o_j} e^{\log \tilde{q}(o_j) + i\alpha L^-_{Q,P}(o_j)}$, where $\tilde{q}(o_j) = q(o_j)\triangle_o$.
5: Return $\phi^-_{\mathcal{M}_P}$ and $\phi^-_{\mathcal{M}_Q}$.

The main challenge in approximating ϕ-function for continuous mechanisms is that an upper/lower bound of ϕ-function does not necessarily attain the upper/lower bound of privacy costs. Motivated by recent

work [Koskela et al., 2021] that truncates the privacy loss R.V. range to discretize subsample Gaussian mechanism, we consider discretizing the output domain. Our choice on the output domain instead of the privacy loss R.V. range is because many recent advances in communication-efficient private learning require the output space to be discrete. Note that the output domain is the output domain of the dominating pair. There always exists a one-dimension tightly dominating pair for any mechanisms as we have shown in the main results. The approximation procedure given in Algorithm 2 takes as input a truncated output domain $[-S, S] \subset \mathcal{O}$, the partition parameter N and the pdf measure of two privacy loss random variables. The algorithm first introduces the grid approximation of privacy loss RVs using the following definition.

Definition 39 (Grid approximation of privacy loss R.V.). *Given N equidistant points $o_1, ..., o_N$ over the interval $[-S, S]$, We define $L^-_{P,Q}(o_j) = \min_{t\in[j-1,j]} L_{P,Q}(o_t)$ and $L^+_{P,Q}(o_j) = \max_{t\in[j-1,j]} L_{P,Q}(o_t)$ as the grid approximation of the original privacy loss R.V., where $L_{P,Q}(o_t)$ denotes the density of $L_{P,Q}$ when evaluated at o_t.*

$L^-_{Q,P}(o_j)$ and $L^+_{Q,P}(o_j)$ are defined analogously. Next, the algorithm constructs ϕ-function approximations $\phi^-_{\mathcal{M}_P}$ and $\phi^-_{\mathcal{M}_Q}$ using similar ideas as that in Definition 36, except that we replace privacy loss R.V. with their approximation alternatives. The idea behinds the approximation is to construct Riemann sum style lower and upper bounds of $\delta_{\mathcal{M}}(\epsilon)$ using sampled points in the output interval. We formalize the idea using the following lemma.

Lemma 40. *Consider the truncation parameter goes to infinity, i.e., $S \to \infty$, and privacy loss random variable $L_{P,Q}$ and $L_{Q,P}$ is a monotonical function of the output random variable o. We have for all $\epsilon > 0$, the privacy profile $\delta(\epsilon)$ of a continuous mechanism $\mathcal{M}$ is bounded by*

$$\delta_{\mathcal{M}_{min}}(\epsilon) \le \delta_{\mathcal{M}}(\epsilon) \le \delta_{\mathcal{M}_{max}}(\epsilon)$$

where $\delta_{\mathcal{M}_{min}}(\epsilon)$ is constructed using Algorithm 1 with$(\phi^-_{\mathcal{M}_P}, \phi^-_{\mathcal{M}_Q})$ pair and $\delta_{\mathcal{M}_{max}}(\epsilon)$ is constructed using $(\phi^+_{\mathcal{M}_P}, \phi^+_{\mathcal{M}_Q})$ pair.

Proof. The proof sketch is first to show that the CDF $F_{L_{P,Q}}(\epsilon)$ is always smaller bounded by $F_{L^-_{P,Q}}(\epsilon)$ when the monotonical condition is satisfied. Then we rewrite the privacy profile into the CDF forms and demonstrate that a larger CDF leads to a smaller δ for any $\epsilon > 0$.

$$F_{L_{P,Q}}(\epsilon) = \int_{-\infty}^{\infty} \mathbb{I}\left[\log\left(\frac{p(o)}{q(o)}\right) \le \epsilon\right] p(o)do \tag{10}$$

$$\le \sum_{j\in\mathbb{Z}} \triangle_o \mathbb{I}\left[\log\left(\min_{t\in[j-1,j]} \frac{p(o_t)}{q(o_t)}\right) \le \epsilon\right] p(o_j) \tag{11}$$

$$= F_{L^-_{P,Q}}(\epsilon) \tag{12}$$

Note that the indicator function preserves the monotonic property, which implies that we can lower bound $F_{L_{P,Q}}(\epsilon)$ using the left Riemann sum. Analogously, the CDF $F_{L_{Q,P}}(-\epsilon)$ is smaller than $F_{L^-_{Q,P}}(-\epsilon)$ for all $\epsilon > 0$. Therefore, we have

$$\begin{aligned}
\delta_{\mathcal{M}}(\epsilon) &= \sup_{D\simeq D'} \left(\Pr[L_{P,Q} > \epsilon] - e^\epsilon \Pr[L_{Q,P} < -\epsilon]\right) \\
&= 1 - \sup_{P\simeq Q} \left(F_{L_{P,Q}}(\epsilon) + e^\epsilon F_{L_{Q,P}}(-\epsilon)\right) \\
&\ge 1 - \sup_{P\simeq Q} \left(F_{L^-_{P,Q}}(\epsilon) + e^\epsilon F_{L^-_{Q,P}}(-\epsilon)\right) \\
&= \delta_{\mathcal{M}_{\min}}(\epsilon)
\end{aligned}$$

□

The following corollary allows us to upper and lower bound privacy cost over composition.

Corollary 41. *Consider a simple composition of $\mathcal{M}_1$ and $\mathcal{M}_2$ on datasets D and D', we have*

$$\delta_{\mathcal{M}_{1_{\min}}+\mathcal{M}_{2_{\min}}}(\epsilon) \le \delta_{\mathcal{M}_1+\mathcal{M}_2}(\epsilon) \le \delta_{\mathcal{M}_{1_{\max}}+\mathcal{M}_{2_{\max}}}(\epsilon)$$

Proof. We overload the $L_{P,Q}$ to denote the privacy R.V. over composition. The lower bound of the CDF $F_{L_{P,Q}}$ is given as follows:

$$\begin{aligned}
&F_{L_{P,Q}}(\epsilon) = \int_{-\infty}^{\infty} \mathbb{I}\left[\log\left(\frac{p_{\mathcal{M}_1\cdot\mathcal{M}_2}(o)}{q_{\mathcal{M}_1\cdot\mathcal{M}_2}(o)}\right) \le \epsilon\right] p_{\mathcal{M}_1,\mathcal{M}_2}(o)do \\
&= \triangle_o \sum_{j\in\mathbb{Z}}\sum_{r\in\mathbb{Z}} \mathbb{I}\Bigg[\log\left(\frac{p_{\mathcal{M}_1}(o_r)}{q_{\mathcal{M}_1}(o_r)}\right) + \log\left(\frac{p_{\mathcal{M}_2}(o_j-o_r)}{q_{\mathcal{M}_2}(o_j-o_r)}\right) \\
&\le \epsilon\Bigg] p_{\mathcal{M}_1\mathcal{M}_2}(o_j) \\
&\le \triangle_o \sum_{j\in\mathbb{Z}}\sum_{r\in\mathbb{Z}} \mathbb{I}\Bigg[\log\left(\min_{t_1\in[r-1,r]}\frac{p_{\mathcal{M}_1}(o_{t_1})}{q_{\mathcal{M}_1}(o_{t_1})}\right) \\
&+ \log\left(\min_{t_2\in[o_j-o_r-1,o_j-o_r]}\frac{p_{\mathcal{M}_2}(t_2)}{q_{\mathcal{M}_2}(t_2)}\right) \le \epsilon\Bigg] p_{\mathcal{M}_1\mathcal{M}_2}(o_j) \\
&= F_{L^-_{P,Q}}(\epsilon)
\end{aligned}$$

□

Gaussian quadrature: We apply Gaussian quadrature to efficiently evaluate integral in computing CDFs. Note that the integral is defined over an infinite integral (see Theorem 18), we will use the following lemma to convert the integral range to $[-1, 1]$ before we apply Gaussian quadrature.

Lemma 42 (Integrals over infinite intervals)**.** *Let $f(x)$ is defined over the infinite interval. We have*

$$\int_{-\infty}^{\infty} f(x)dx = \int_{-1}^{+1} f\left(\frac{t}{1-t^2}\right)\frac{1+t^2}{(1-t^2)^2}dt$$

By a change of variables, we can convert the infinite integral to a finite integral, which can be easily implemented using numerical integration methods.

Double quadrature: To improve the time complexity of Algorithm 2, which is linear with a sufficiently large N, we next apply Gaussian quadrature to approximate ϕ-function when its closed-form expression is not available. We call this algorithm "Double quadrature", as we use it twice, one is in the approximation for ϕ-function, and another is for the CDF computation. In the experiment section, we show that the "Double quadrature" algorithm matches the Fourier Accountant approach [Koskela and Honkela, 2021] while only samples hundreds of points in evaluating Poisson Subsample Mechanisms.

E.1 Error analysis

In this section, we provide the end-to-end error analysis of AFA by looking into the following three scenarios.

1. ϕ-functions have closed-form expressions.
2. ϕ-functions with discretization is used (see Algorithm 2).
3. Double quadrature algorithm: ϕ-function is approximated using Gaussian quarature.

Error analysis closed-form ϕ functions. When we have closed-form ϕ-functions (see Exp1 and Exp2), the numerical error is only caused by the quadrature method, which is used to convert ϕ-function to CDFs. Therefore, we can tap into the classical results from numerical analysis that bounds the error in quadrature methods (see, e.g., Chapter 7 of Conte and de Boor "Elementary Numerical Analysis") with an asymptotic scaling more or less $O(1/n^\alpha)$ for integrating an αth times differentiable functions, and faster for more advanced rules, see the following Lemma.

Lemma 43. *[Stoer and Bulirsch, 2002] Let the function $f(\cdot)$ has $2n$ continuous derivatives over the integral $[a,b]$ and n is the number sample points. We have the error estimate*

$$\frac{(b-a)^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3} f^{(2n)}(\zeta), a < \zeta < b$$

From a practical standpoint, "Scipy.integrate" allows us to specify a desired error tolerance (e.g., $O(10^{-14})$ used in experiments). The error induced in CDFs will be amplified by e^{ϵ} in the final $\delta(\epsilon)$ evaluation according to Lemma 5, which is still negligible in practice. The inverse (ϵ as a function of δ) requires an additional binary search, which calls $\delta(\epsilon)$ a handful of times.

In the cases when the ϕ-function does not have a closed-form expression, we proposed two approaches to approximate the -function: Algorithm 2 and Double Quadrature. In Algorithm 2, we discretize the support of the dominating pairs using an equispaced grid approximation.

Error analysis with approximated ϕ functions. We analyzed the error caused by truncation, approximation of Algorithm 2 as follows.

1. The error caused by truncation (ignore $o \geq S$ or $o \leq -S$).
2. The error arising from Riemann Sum approximation.
3. The error caused by using Gaussian quadrature to compute CDF.

The first two errors arising from the approximation of ϕ-function and the third error is the numerical error when we apply Gaussian quadrature to compute the integral in privacy profile. The third term is often negligible as we discuss earlier. When p and q are determined, we bound the tail integral $\int_S^{\infty} \max(p(o), q(o))do$ and $\int_{-S}^{-\infty} \max(p(o), q(o))do$ using the Chernoff bound and denote it is upper bounded by δ_{tail}. Then a union bound over k compositions will bound all bad events that the output happens to be out of $[-S, S]^k$. Lastly, we estimate the total error by subtracting and adding $k\delta_{tail}$ to the lower and the upper Riemann sums bound, respectively.

Theorem 44. *Denote the privacy profile of machanism $\mathcal{M}$ over k-fold composition as $\delta_{\mathcal{M}}^{(k)}$. Let $\delta_{\mathcal{M}_{min}}^{(k)}(\epsilon)$ and $\delta_{\mathcal{M}_{max}}^{(k)}(\epsilon)$ denote the lower and upper bounds of Riemann sum approximation over k-fold composition. Denote δ_{tail} be the upper bound of tail integral, i.e.,*$\Pr[o > S] + \Pr[o < S] \leq \delta_{tail}$. *Then we have*

$$\delta_{\mathcal{M}_{min}}^{(k)}(\epsilon) - k\delta_{tail} < \delta_{\mathcal{M}}^{(k)} < k\delta_{tail} + \delta_{\mathcal{M}_{max}}^{(k)}(\epsilon)$$

Double Quadrature For the second approach — adaptive approximation via double quadrature, we do not have an asymptotic analysis of the bounds but 'scipy.integrate.dblquad' does provide us valid bounds. The "double quadrature" approach is the best-performing algorithm we recommend in practice.

E.2 Experimental details in Exp 3

Theorem 11 allows us to consider the following one dimension[9] distribution as the worst-case pair neighboring distribution for the Poisson subsampling Gaussian mechanism.

$$p(o) = \gamma \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{(o-1)^2}{-2\sigma^2}} + (1-\gamma)\frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{o^2}{-2\sigma^2}}$$
$$q(o) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{o^2}{-2\sigma^2}}$$

Therefore, the pairing privacy loss RV is given by

$$L_{P,Q} = \log\left(\gamma e^{\frac{2o-1}{2\sigma^2}} + 1 - \gamma\right), L_{Q,P} = -\log\left(\gamma e^{\frac{2o-1}{2\sigma^2}} + 1 - \gamma\right)$$

[9] the error analysis of the $(P, (1-\gamma)P + \gamma Q)$ dominating pair is similar

We cannot derive a closed-form expression of ϕ-function naively for the above privacy loss RVs. Therefore, we consider the discretization method (Algorithm 2) to approximate the Poisson subsampling.

We set $S = 100$ and $N = 10^5$ to discretize the exact integral of o, which is $[-\infty, \infty]$ for the Gaussian mechanism. We use $\gamma = 0.01$ and $\sigma = 2.0$, for number of compositions up to 1500. We now provide the error analysis of our Poisson subsample Gaussian experiments. Recall the error analysis in Theorem 44, we first bound the error caused by truncation (ignore $o \geq S$ or $o \leq -S$).

The error caused by trunction consists of the tail integral $t_1 := \int_S^\infty \max(p(o), q(o))do$ and $t_2 := \int_{-S}^{-\infty} \max(p(o), q(o))do$. We can upper bound t_1 and t_2 using the tail bound of Gaussian distribution.

$$t_1 \leq \Pr[\mathcal{N}(1, \sigma^2) > S] \leq e^{\frac{(S-1)^2}{-2\sigma^2}}$$

The first inequality is because $p(o)$ is a mixture of two Gaussian distribution $\mathcal{N}(0, \sigma^2)$ and $\mathcal{N}(1, \sigma^2)$. The second inequality follows the tail bound of Gaussian distribution $\Pr[\mathcal{N}(0, \sigma) > x] \leq e^{\frac{-x^2}{2\sigma^2}}$. Similarly, we can upper bound t_2 by $e^{\frac{-S^2}{-2\sigma^2}}$.

We use δ_{tail} to denote the failure probability when o happens to be out of the range $[-S, S]$. Here, we have $\delta_{tail} \leq e^{\frac{-S^2}{-2\sigma^2}} + e^{\frac{(S-1)^2}{-2\sigma^2}}$. Substituting $S = 100, k = 1500$ and $\sigma = 2.0$ into Theorem 44, we have $k \cdot \delta_{tail}$ is upper bounded by e^{-1000}, which is neglectable in the δ term. For the error caused by discretization, we plot the valid lower and upper bound using Algorithm 2.

In "Double quadrature", we apply Gaussian quadrature to compute $\phi(\mathcal{M})$ and $\phi'(\mathcal{M})$ directly. That is, we apply Gaussian quadrature to solve the integration

$$\phi(\mathcal{M})(\alpha) = \int_{-\infty}^{\infty} e^{i\alpha \log(p(o)/q(o))} q(o)do$$

and

$$\phi'(\mathcal{M})(\alpha) = \int_{-\infty}^{\infty} e^{i\alpha \log(q(o)/p(o))} p(o)do$$

Though we did not include an error analysis of the Double quadrature algorithm, the algorithm exactly matches the result from Koskela et al. [2021] and its computation time for each $\delta(\epsilon)$ query is only around 0.2 sec.

F An "optimal" Renyi DP to DP conversion?

In this section, we provide the detailed description of how we generated the improved "optimal" conversion rule in Figure 1 based on an extension of the technique of Balle et al. [2020].

Specifically, Balle et al. [2020] shows that (α, ϵ)-RDP implies f-DP for any tradeoff function f that lower bounds the following tradeoff region:

$$\left\{ (x, y) \in [0, 1]^2 \middle| \begin{array}{l} x^\alpha (1-y)^{1-\alpha} + (1-x)^\alpha y^{1-\alpha} \leq e^{(\alpha-1)\epsilon(\alpha)} \\ y^\alpha (1-x)^{1-\alpha} + (1-y)^\alpha x^{1-\alpha} \leq e^{(\alpha-1)\epsilon(\alpha)} \end{array} \right\}.$$

Then one can further convert this f-DP to (ϵ, δ)-DP according to our formula in Section B (originally due to [Dong et al., 2021].)

The main improvement that we propose is to consider the mechanism specific version of the same conversion rule, which converts the RDP function $\epsilon_{\mathcal{M}}(\cdot)$ satisfied by a mechanism $\mathcal{M}$ to an f-DP of $\mathcal{M}$, which involves taking the pointwise maximum of all f functions implied by each $(\alpha, \epsilon_{\mathcal{M}}(\alpha))$-RDP. The key to obtain the conversion that we have shown in Figure 1 was to consider and extended version of RDP that also includes $0 < \alpha < 1$, which we find to have a nontrivial effect in the resulting tradeoff function.

In the remainder of the section, we will first explain how the two-stage conversion works in Section F.1 and Section F.2 and then comment on whether the rule can be improved in Section F.3.

F.1 RDP to f-DP

The hypothesis testing interpretation (f-DP) of RDP is not entirely tight. Here we present a simplified derivation of the tradeoff function f implied by RDP via closedness to post-processing.

Consider any hypothesis testing procedure $h : \mathcal{O} \to \{0, 1\}$ that uses the output o of an (α, ϵ)-RDP mechanism. "1" denotes "Rejecting the null hypothesis" that individual z is not in the dataset, indicating that h predicts that z is in the dataset. "0" denotes the complement event of "Failing to reject the null hypothesis", indicating that h predicts that z is not in the dataset. By the closure to post-processing property, $h(o)$ satisfies (α, ϵ)-RDP. By definition of RDP,

$$\frac{1}{\alpha - 1} \log \mathbb{E}_{h(o)\sim q} \left[(\frac{p}{q})^\alpha \right] \leq \epsilon \tag{13}$$

for all pairs of neighboring datasets that induce distributions p, q.

Let q be the distribution where individual z is not in the dataset and p otherwise. Let x denote the probability of false positive (Type I error) — z is not in the dataset but the prediction is 1; and y denote the probability of false negative (Type II error) — z is in the dataset but the prediction is 0.

For $\alpha > 1$, (13) is equivalent to

$$(1 - y)^\alpha x^{1-\alpha} + y^\alpha (1 - x)^{1-\alpha} \leq e^{(\alpha-1)\epsilon},$$
$$x^\alpha (1 - y)^{1-\alpha} + (1 - x)^\alpha y^{1-\alpha} \leq e^{(\alpha-1)\epsilon},$$

where the first constraint follows by taking the moments of the density ratio of the binary random variable $h(o)$, by noting that the event for prediction 1 is false positive under q but true positive under q. The second constraint follows from swapping p, q.

When $\alpha = 1$, by the definition of KL-divergence, (13) is equivalent to

$$x \log(\frac{x}{1-y}) + (1 - x) \log(\frac{1-x}{y}) \leq \epsilon,$$
$$y \log(\frac{y}{1-x}) + (1 - y) \log(\frac{1-y}{x}) \leq \epsilon.$$

Finally, when $0 < \alpha < 1$, (13) is equivalent to

$$(1 - y)^\alpha x^{1-\alpha} + y^\alpha (1 - x)^{1-\alpha} \geq e^{(\alpha-1)\epsilon}$$
$$x^\alpha (1 - y)^{1-\alpha} + (1 - x)^\alpha y^{1-\alpha} \geq e^{(\alpha-1)\epsilon}.$$

Note that the only difference from the case when $\alpha > 1$ is the direction of the inequality. Also note that the symmetry of the two inequalities ensures that it suffices to consider $\alpha \geq 0.5$[10].

The f-DP of a mechanism satisfying $\epsilon(\alpha)$-RDP for a family α is therefore the pointwise maximum of the resulting f function for all α.

F.2 f-DP to (ϵ, δ)-DP

fDP is related to (ϵ, δ)-DP in the following lemma.

Lemma 45. *Let f be the lower bound of Type II error given Type I error, then mechanisms that satisfy fDP with function f also obeys a family of $(\epsilon(x), \delta(x))$-DP for all $x \in [0, 1]$ such that*

$$\delta(x) = 1 - f(x) - (-\partial f(x))x$$
$$\epsilon(x) = \log(-\partial f(x))$$

where $\partial f(x)$ is any subgradient of f at x. Recall that by definition of subgradient, g is a subgradient of f at x if for all $y \in \mathrm{Dom}(f)$, $f(y) \geq f(x) + g \cdot (y - x)$ (note that $\mathrm{Dom}(f) = [0, 1]$ for trade-off function f).

[10] Let $0 < \alpha < 0.5$ and $\alpha' = 1 - \alpha$. Notice that $(1 - y)^\alpha x^{1-\alpha} + y^\alpha (1 - x)^{1-\alpha} = x^{\alpha'} (1 - y)^{1-\alpha'} + (1 - x)^{\alpha'} y^{1-\alpha}$. Next, check that for all positive ϵ and $0 < \alpha < 0.5$, $e^{(\alpha-1)\epsilon} < e^{(\alpha'-1)\epsilon}$. In other words, the bound with α in $(0, 0.5)$ is never active.

Numerically stable computation of ϵ given δ or δ given ϵ involves working with ∂f and $1-f$ in logarithmic scale. Specifically, we find x such that there exists subgradient $g \in \partial f(x)$ such that

$$\log(\delta) \geq \log(e^{\log(1-f(x))} - e^{\log(-g)+\log(x)}). \tag{14}$$

Then it is true that $\epsilon(\delta) = \log(-g)$. In fact, a more general procedure finds an x (why any feasible x works is left as an exercise) such that $\partial f(x)$ contains g satisfying (14). It then follows that

$$\epsilon(\delta) = \min_{g \in \partial f(x)} \log(-g).$$

When $f(x)$ is differentiable everywhere, we can solve (14) as a nonlinear equation and then we can write $\epsilon(\delta) = \log(-f'(x(\delta)))$ if $f'(x(\delta)) < -1$, and $\epsilon(\delta) = 0$ otherwise.

Similarly, $\delta(\epsilon)$ can be found by solving the nonlinear equation $\log(-\partial f(x)) = \epsilon$ for x and then plug into (14):

$$\delta(\epsilon) = e^{\log(1-f(x(\epsilon)))} - e^{\log(-\partial f(x(\epsilon)))+\log(x(\epsilon))}.$$

In other word, provided that $\log(1-f(x))$ and $\log(-\partial f(x))$ admit an analytical implementation, we can convert f-DP to (ϵ, δ)-DP in a numerically stable fashion.

F.3 Optimality of this conversion rule?

A natural question to ask is that whether the aforementioned conversion rule from RDP to (ϵ, δ)-DP is optimal. The answer is yes and no. Let us explain.

From Figure 1, we can clearly see that for the randomized response mechanism, the resulting conversion matches exactly with the exact (ϵ, δ)-DP obtained via the privacy-profile.

Moreover, observe that the converted f-function of the Gaussian mechanism touches that of the randomized response mechanism which satisfies the same RDP bound for all $\alpha > 0$. For this reason, we know that for Gaussian mechanism, the conversion rule cannot be improved in any ways that strictly improves the stated conversion rule at all inputs (Type I error).

This example, however, does not rule out the possibility of improving the RDP-implied f-DP elsewhere for Gaussian mechanism. Therefore, it is unclear whether the proposed RDP-to-DP conversion rule is optimal in the strong sense:

Is it optimal for all RDP functions and all input (type I error) at the same time?

Our conjecture is positive, but it is beyond the scope of the current paper to formally prove this.

A promising direction is to compare our approach to the optimal conversion rule proposed by Asoodeh et al. [2021] (which uses a very different approach to derive almost the same formula as that in [Balle et al., 2020]).

G Omitted proofs in Appendix B

Lemma 46.

$$H_\alpha(Q\|P) = \alpha H_{\alpha^{-1}}(P\|Q) - \alpha + 1.$$

Proof of Lemma 46. Let p, q be the Radon–Nikodym derivative of P and Q with respect to a common dominating measure μ (say $\frac{1}{2}P + \frac{1}{2}Q$). Then $H_\alpha(P\|Q) = \int (p - \alpha q)_+ \, \mathrm{d}\mu$ and $H_\alpha(Q\|P) = \int (q - \alpha p)_+ \, \mathrm{d}\mu$. We will drop the $\mathrm{d}\mu$ notation in the integrals for convenience.

$$\begin{aligned} H_\alpha(Q\|P) &= \int (q - \alpha p)_+ = \alpha \int (\alpha^{-1} q - p)_+ \\ &= \alpha \int (\alpha^{-1} q - p) + \alpha \int (p - \alpha^{-1} q)_+ \\ &= \alpha(\alpha^{-1} - 1) + \alpha H_{\alpha^{-1}}(P\|Q) = \alpha H_{\alpha^{-1}}(P\|Q) - \alpha + 1 \end{aligned}$$

where we used the fact that $a_+ - (-a)_+ = a$ on the second line. □

Proof of Lemma 19. Let $A = \{\omega : \frac{\mathrm{d}P}{\mathrm{d}Q}(\omega) \geqslant \alpha\} = \{\omega : \log \frac{\mathrm{d}Q}{\mathrm{d}P}(\omega) \leqslant \log \alpha^{-1}\}$. We have

$$\begin{aligned} H_\alpha(P\|Q) &= \mathbb{E}_{\omega\sim Q}[(\frac{\mathrm{d}P}{\mathrm{d}Q}(\omega) - \alpha)_+] \\ &= \mathbb{E}_{\omega\sim Q}[(\frac{\mathrm{d}P}{\mathrm{d}Q}(\omega) - \alpha)\cdot 1_A] \\ &= P[A] - \alpha Q[A] = F(-\log\alpha) - \alpha G(-\log\alpha). \end{aligned}$$

The second identity can be obtained in a similar fashion. □

Proof of Lemma 22. Follow directly from Lemma 48. □

Proof of Lemma 23. The first expression is the direct consequence of $f'\big(1-F(x)\big) = -\mathrm{e}^x$ and the well-known fact in convex analysis that for a convex function f, f' and $(f^*)'$ are inverse functions of each other.

Let $t = -\mathrm{e}^x$. Then from Lemma 48 we have $H_{\mathrm{e}^x}(Q\|P) = 1 + f^*(t)$ and hence $\frac{\mathrm{d}}{\mathrm{d}t}H_{\mathrm{e}^x}(Q\|P) = (f^*)'(t)$. Therefore,

$$\begin{aligned} F(x) &= 1 - (f^*)'(t) = 1 - \frac{\mathrm{d}}{\mathrm{d}t}H_{\mathrm{e}^x}(Q\|P) = 1 - \frac{\mathrm{d}x}{\mathrm{d}t}\cdot\frac{\mathrm{d}}{\mathrm{d}x}H_{\mathrm{e}^x}(Q\|P) \\ &= 1 - \left(\frac{\mathrm{d}t}{\mathrm{d}x}\right)^{-1}\cdot\frac{\mathrm{d}}{\mathrm{d}x}H_{\mathrm{e}^x}(Q\|P) = 1 + \mathrm{e}^{-x}\cdot\frac{\mathrm{d}}{\mathrm{d}x}H_{\mathrm{e}^x}(Q\|P) \end{aligned}$$

Now the first identity about G is a direct consequence of Lemma 22, and the second one is a direct consequence of Lemma 19. □

Lemma 47. $G(x) = \int_{-\infty}^{x} \mathrm{e}^t \,\mathrm{d}F(t)$.

Proof of Lemma 47. Let $h : \mathbb{R} \to \mathbb{R}$ be a Borel measurable function. Then

$$\int h(x)\,\mathrm{d}G(x) = \int h\left(\log\tfrac{\mathrm{d}Q}{\mathrm{d}P}(\omega)\right)\,\mathrm{d}Q(\omega) = \int h\left(\log\tfrac{\mathrm{d}Q}{\mathrm{d}P}(\omega)\right)\cdot\tfrac{\mathrm{d}Q}{\mathrm{d}P}(\omega)\,\mathrm{d}P(\omega) = \int h(x)\mathrm{e}^x\,\mathrm{d}F(x)$$

The choice of h as the indicator function of $(-\infty, x]$ yields the second identity. □

Lemma 48. *The functions f, F, G and the hockey-stick divergence have the following relations:*

$$\begin{aligned} f(1-F(x)) &= G(x) \\ f'\big(1-F(x)\big) &= -\mathrm{e}^x \\ H_{\mathrm{e}^\varepsilon}(Q\|P) &= 1 + f^*(-\mathrm{e}^\varepsilon) \end{aligned}$$

Proof of Lemma 48. The first identity follows from the definition of the trade-off function and Neyman–Pearson lemma. In fact, $1-F(x)$ and $G(x)$ are the type I and type II errors of the likelihood ratio test with threshold at x. Taking derivative with respect to x on both sides of the first identity, we have

$$f'(1-F(x))\cdot(-F'(x)) = G'(x).$$

Now the second identity follows by plugging in Lemma 47. □

Proof of Lemma 25. It is known that hockey-stick divergence is determined by the trade-off function, namely $H_{\mathrm{e}^\varepsilon}(Q\|P) = 1 + f^*(-\mathrm{e}^\varepsilon)$. By Lemma 23, we know the distributions of $\log\frac{\mathrm{d}P}{\mathrm{d}Q}$ (under P and Q respectively) are also determined by the trade-off function f. Therefore, ϕ and ϕ' are determined by f. In particular, if we can find P', Q' such that $f = T[P', Q']$, then we can use $\phi(t) = \mathbb{E}_{P'}\mathrm{e}^{it\log\frac{\mathrm{d}P'}{\mathrm{d}Q'}}$ and $\phi(t) = \mathbb{E}_{Q'}\mathrm{e}^{it\log\frac{\mathrm{d}Q'}{\mathrm{d}P'}}$ to compute ϕ and ϕ'.

From the proof of Proposition 2.2 of Dong et al. [2021], we know that we can pick P' as the uniform distribution over $[0,1]$ and Q' has density $-f'(1-x) = |f'(1-x)|$ on $[0,1]$. Therefore,

$$\begin{aligned}\phi(t) = \mathbb{E}_{P'} \mathrm{e}^{it \log \frac{\mathrm{d}P'}{\mathrm{d}Q'}} &= \int_0^1 \mathrm{e}^{-it \log \frac{\mathrm{d}Q'}{\mathrm{d}P'}} \,\mathrm{d}x \\ &= \int_0^1 \mathrm{e}^{-it \log |f'(1-x)|} \,\mathrm{d}x = \int_0^1 \mathrm{e}^{-it \log |f'(x)|} \,\mathrm{d}x\end{aligned}$$

The second identity can be proved similarly. □