



HAL
open science

Rényi Pufferfish Privacy: General Additive Noise Mechanisms and Privacy Amplification by Iteration via Shift Reduction Lemmas

Clément Pierquin, Aurélien Bellet, Marc Tommasi, Matthieu Bousard

► **To cite this version:**

Clément Pierquin, Aurélien Bellet, Marc Tommasi, Matthieu Bousard. Rényi Pufferfish Privacy: General Additive Noise Mechanisms and Privacy Amplification by Iteration via Shift Reduction Lemmas. 2023. hal-04363020v1

HAL Id: hal-04363020

<https://inria.hal.science/hal-04363020v1>

Preprint submitted on 23 Dec 2023 (v1), last revised 13 Jun 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Rényi Pufferfish Privacy: General Additive Noise Mechanisms and Privacy Amplification by Iteration via Shift Reduction Lemmas

Clément Pierquin
Craft AI, Univ. Lille, Inria

Aurélien Bellet
Inria Montpellier

Marc Tommasi
Univ. Lille, Inria

Matthieu Bousard
Craft AI

Abstract

Pufferfish privacy is a flexible generalization of differential privacy that allows to model arbitrary secrets and adversary’s prior knowledge about the data. Unfortunately, designing general and tractable Pufferfish mechanisms that do not compromise utility is challenging. Furthermore, this framework does not provide the composition guarantees needed for a direct use in iterative machine learning algorithms. To mitigate these issues, we introduce a Rényi divergence-based variant of Pufferfish and show that it allows us to extend the applicability of the Pufferfish framework. We first generalize the Wasserstein mechanism to cover a wide range of noise distributions and introduce several ways to improve its utility. We also derive stronger guarantees against out-of-distribution adversaries. Finally, as an alternative to composition, we prove privacy amplification results for contractive noisy iterations and showcase the first use of Pufferfish in private convex optimization. A common ingredient underlying our results is the use and extension of shift reduction lemmas.

1 Introduction

Differential privacy (DP) (Dwork and Roth, 2014) is now considered as the gold standard for privacy-preserving data analysis. However, despite its many desirable properties, DP does not suit all types of data effectively. Specifically, the guarantees it offers are based on the underlying assumption that individuals in the dataset being analyzed are statistically independent. In reality, data often exhibit correlations, and when two correlated individuals are present in a dataset, performing the same analysis with and with-

out one of these individuals could leak more knowledge about the individual than the conventional differential privacy framework assumes (Humphries et al., 2023).

To address these situations, specialized privacy definitions have been designed. Certain direct extensions of DP, like group privacy (Dwork and Roth, 2014) or entry privacy (Hardt and Roth, 2013), protect entire instances or groups, which results in strong privacy guarantees but often much poorer utility. More flexible frameworks allow to tailor the privacy definition to a set of distributions which could have plausibly generated the dataset, and thereby allow a tighter privacy analysis. In this work, we focus on the general framework of Pufferfish privacy (Kifer and Machanavajjhala, 2014), which is closely related to other similar definitions like Blowfish privacy (He et al., 2014) and distribution privacy (Kawamoto and Murakami, 2019; Chen and Ohrimenko, 2023).

Pufferfish privacy however comes with new challenges, first and foremost in the design of general and computationally tractable Pufferfish private mechanisms. Indeed, the sensitivity of the query, which is critical in DP to design additive noise mechanisms, has no direct use in Pufferfish privacy. Moreover, while various ways to measure and efficiently track the privacy loss have been proposed for DP, see for instance Rényi differential privacy (RDP) (Mironov, 2017), this flexibility is lacking in Pufferfish privacy. As a result, previous work on the design of Pufferfish mechanisms has focused on specific noise distributions and applications (Kifer and Machanavajjhala, 2014; Ou et al., 2018; Kessler et al., 2015; Niu et al., 2019; Song et al., 2017). For instance, Song et al. (2017) proposed the Wasserstein mechanism for the Laplace noise, which relies on the computation of ∞ -Wasserstein distances. Another recent work proposes an exponential mechanism-based approach which provides a more computationally tractable approach but relies on (potentially loose) sufficient conditions for Pufferfish privacy (Ding, 2022). The Pufferfish framework thus lacks a unified theory that subsumes the original worst-

case definition and allows for the design of general additive mechanisms compatible with a wide range of noise distributions.

Another key limitation of Pufferfish privacy is that (sequential) composition results exist only for some Pufferfish instantiations and mechanisms (Kifer and Machanavajjhala, 2014). This is the case for instance of the Markov Quilt Mechanism (Song et al., 2017), but it is limited to Bayesian networks. The lack of a universal composition theorem currently makes Pufferfish privacy unfit for the analysis of iterative algorithms such as those used in differentially private machine learning (Abadi et al., 2016).

In this paper, we mitigate the above limitations of Pufferfish privacy by making the following contributions:

- We define the Rényi Pufferfish privacy framework and describe its basic properties.
- We introduce the General Wasserstein Mechanism (GWM), a generalization of the Wasserstein mechanism of Song et al. (2017). Our mechanism allows to derive (Rényi) Pufferfish privacy guarantees for all additive noise distributions that are absolutely continuous with respect to the Lebesgue measure.
- We propose two ways to improve the utility of GWM by relaxing the ∞ -Wasserstein distance used to calibrate the noise. Our first approach relies on a δ -approximation allowing the tail of the distribution of the mechanism to be disregarded, similar to what has been proposed by Chen and Ohrimenko (2023) for the distribution privacy framework. Incidentally, we demonstrate an equivalence between Pufferfish privacy and distribution privacy. Our second approach enables the use of p -Wasserstein distances.
- We generalize the guarantees against “close adversaries” of Song et al. (2017) to our Rényi Pufferfish privacy framework for greater robustness and applicability.
- Inspired by Feldman et al. (2018), we prove privacy amplification by iteration results for Pufferfish, allowing to bypass the use of composition in the analysis of contractive noisy iterations. This technique is particularly useful to analyze convex optimization with stochastic gradient descent, and thus constitutes a first step towards the integration of Pufferfish privacy in machine learning.

One of our key technical contributions lies in the novel use and generalization of shift reduction lemmas (Feldman et al., 2018; Altschuler and Talwar,

2022) in the context of Pufferfish privacy. We argue that shift reduction is the right tool to analyze Pufferfish privacy, and believe this view may yield more results in the future.

All proofs and some additional content can be found in the supplementary material.

2 Rényi Pufferfish Privacy

We start by recalling the definitions of Rényi differential privacy and Pufferfish privacy. Rényi differential privacy relies on Rényi divergences, which are defined as follows.

Definition 2.1. Let μ and ν be two distributions on a measurable space (E, \mathcal{A}) and $\alpha > 1$. We define the Rényi divergence of order α between μ and ν as:

$$D_\alpha(\mu, \nu) = \frac{1}{\alpha - 1} \log \mathbb{E}_{x \sim \nu} \left[\left(\frac{\mu(x)}{\nu(x)} \right)^\alpha \right].$$

The definition extends to the case $\alpha = +\infty$ by continuity.

Definition 2.2 (Rényi differential privacy, RDP (Mironov, 2017)). Let $\alpha > 1$ and $\varepsilon \geq 0$. A randomized algorithm $\mathcal{M}: \mathcal{D} \rightarrow \mathcal{R}$ satisfies (α, ε) -Rényi differential privacy if for any two adjacent datasets $X_1, X_2 \in \mathcal{D}$ differing by one element, it holds:

$$D_\alpha(P(\mathcal{M}(X_1)), P(\mathcal{M}(X_2))) \leq \varepsilon.$$

RDP ensures that an adversary cannot gain too much knowledge about whether an individual point is in the dataset or not by observing the output of the mechanism. In this definition, it is implied that the elements of the dataset are statistically independent.

A more general framework, Pufferfish privacy, has been designed to handle possibly correlated data and other types of secrets than the presence of an individual in a dataset (Kifer and Machanavajjhala, 2014). In a Pufferfish instantiation, we denote by \mathcal{S} the set of possible secrets to be protected, and by $\mathcal{Q} \subseteq \mathcal{S}^2$ the specific pairs of secrets we aim to make indistinguishable. In contrast to differential privacy, the variable X representing the dataset is not deterministic in Pufferfish privacy. Instead, it is sampled from a certain distribution $\theta \in \Theta$. The set Θ represents the possible prior knowledge of an adversary.

Definition 2.3 (Pufferfish privacy, PP (Kifer and Machanavajjhala, 2014; Ding, 2022)). Let $\varepsilon \geq 0$ and $\delta \in (0, 1)$. A privacy mechanism \mathcal{M} is said to be (ε, δ) -Pufferfish private in a framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ if for all $\theta \in \Theta$, for all secret pairs $(s_i, s_j) \in \mathcal{Q}$, and for all $w \in \text{Range}(\mathcal{M})$, we have:

$$\begin{aligned} P(\mathcal{M}(X) = w \mid s_i, \theta) &\leq e^\varepsilon P(\mathcal{M}(X) = w \mid s_j, \theta) + \delta, \\ P(\mathcal{M}(X) = w \mid s_j, \theta) &\leq e^\varepsilon P(\mathcal{M}(X) = w \mid s_i, \theta) + \delta, \end{aligned}$$

where $X \sim \theta$ and (s_i, s_j) is such that $P(s_i | \theta) \neq 0, P(s_j | \theta) \neq 0$. If $\delta = 0$, \mathcal{M} satisfies ε -Pufferfish privacy.

In this work, we introduce a Rényi divergence-based version of Pufferfish privacy. Using Rényi divergences in privacy definitions has several advantages. Especially relevant to our work will be the quantification of privacy guarantees by bounding certain moments of the exponential of the privacy loss (Mironov, 2017), and the ability to leverage a large body of results on Rényi divergences such as shift reduction lemmas (Feldman et al., 2018; Altschuler and Chewi, 2023).

Definition 2.4 (Rényi Pufferfish privacy, RPP). Let $\alpha > 1$ and $\varepsilon \geq 0$. A privacy mechanism \mathcal{M} is said to be (α, ε) -Rényi Pufferfish private in a framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ if for all $\theta \in \Theta$ and for all secret pairs $(s_i, s_j) \in \mathcal{Q}$, we have:

$$\begin{aligned} D_\alpha(P(\mathcal{M}(X) | s_i, \theta), P(\mathcal{M}(X) | s_j, \theta)) &\leq \varepsilon, \\ D_\alpha(P(\mathcal{M}(X) | s_j, \theta), P(\mathcal{M}(X) | s_i, \theta)) &\leq \varepsilon, \end{aligned}$$

where $X \sim \theta$ and (s_i, s_j) is such that $P(s_i | \theta) \neq 0, P(s_j | \theta) \neq 0$.

Rényi Pufferfish privacy upholds the post-processing inequality, which is a key attribute for any effective privacy framework.

Proposition 2.1 (Post-processing). *Let \mathcal{M}_1 be a randomized algorithm and \mathcal{M} be (α, ε) -RPP. Then,*

$$\begin{aligned} D_\alpha(P(\mathcal{M}_1(\mathcal{M}(X)) | s_i, \theta), P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)) \\ \leq D_\alpha(P(\mathcal{M}(X) | s_i, \theta), P(\mathcal{M}(X) | s_j, \theta)) \leq \varepsilon. \end{aligned}$$

It is easy to see that (∞, ε) -RPP corresponds to ε -PP. Furthermore, (α, ε) -RPP can be converted to (ε, δ) -PP.

Proposition 2.2 (RPP implies PP). *If \mathcal{M} is (α, ε) -RPP, it also satisfies $(\varepsilon + \frac{\log(1/\delta)}{\alpha-1}, \delta)$ -PP $\forall \delta \in (0, 1)$.*

Running examples. We introduce here some examples of RPP instantiations which we will use throughout the paper to illustrate our private mechanisms. Let $n > 0$ be the total number of participants in a study. Let \mathcal{X} be the potential values of an individual’s private features. Let $X = (X_1, \dots, X_n) \in \mathcal{X}^n$ describing the private properties of the n individuals. An adversary anticipates correlations among individuals within the study with a prior $\theta \in \Theta$. We define the set of secrets for this adversary as $\mathcal{S} = \left\{ s_i^a \stackrel{\text{def}}{=} \{X_i = a\}; a \in \mathcal{X}, i \in \llbracket 1, n \rrbracket \right\}$ and define $\mathcal{Q} = \{(s_i^a, s_j^b); a, b \in \mathcal{X}, i, j \in \llbracket 1, n \rrbracket\}$. Consider the following simple instantiations of this setting for datasets of size 2:

- **Example 1** (Counting query with correlation). Each individual i holds a binary value $X_i \in \{0, 1\}$ and we consider a counting query $f(X) = X_1 + X_2$. For $p \in (0, 1), \rho \in [-1, 1]$, the adversary has the following prior: $P(X_1 = 1) = P(X_2 = 1) = p$, where X_1 and X_2 are drawn with correlation ρ .
- **Example 2** (Average salary query). Each individual i holds her salary $X_i \geq 0$ and we consider an average query $f(X) = \frac{1}{2}(X_1 + X_2)$. The adversary has the following prior for the marginals: for $i \in \{1, 2\}$,

$$X_i = \begin{cases} 1 & \text{with prob. } 1/2 \\ 2 & \text{with prob. } 499/1000, \text{ for } i \in \{1, 2\} \\ 100 & \text{with prob. } 1/1000 \end{cases}$$

Here, X_1 and X_2 are thus considered independent.

- **Example 3** (Sum query with user-dependent prior). We consider $\mathcal{X} = (0, r)$ and a sum query $f(X) = X_1 + X_2$. The adversary has an arbitrary prior about the distribution of (X_1, X_2) but assumes that each individual i holds a different value $X_i \in (0, r_i)$ with $0 < r_i \leq r$.

3 A General Additive Mechanism for Rényi Pufferfish Privacy

In this section, we present a general approach to obtain Rényi Pufferfish privacy guarantees. Specifically, we introduce the General Wasserstein Mechanism (GWM), a generalization of the Laplacian-based Wasserstein mechanism of Song et al. (2017) to a wide range of noise distributions, and derive the corresponding RPP guarantees. We also highlight that the shift reduction lemma and its variants, introduced by Feldman et al. (2018) in the context of privacy amplification by iteration, provide the right framework for analyzing Rényi Pufferfish privacy.

We first introduce ∞ -Wasserstein distances and couplings.

Definition 3.1 (Couplings). Let μ and ν be two distributions on a measurable space $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ with $\mathcal{B}(\mathbb{R}^d)$ the Borel σ -algebra. A coupling π is a joint distribution on the product space $(\mathbb{R}^{d \times 2}, \mathcal{B}(\mathbb{R}^d)^2)$ with marginals μ and ν , where $\mathcal{B}(\mathbb{R}^d)^2$ is the product σ -algebra.

Definition 3.2 (∞ -Wasserstein distance). Let μ and ν be two distributions on \mathbb{R}^d . We note Γ the set of the couplings between μ and ν . We define the ∞ -Wasserstein distance between μ and ν as:

$$W_\infty(\mu, \nu) = \inf_{\pi \in \Gamma(\mu, \nu)} \sup_{(x, y) \in \text{supp}(\pi)} \|x - y\|.$$

Throughout the paper, $\|\cdot\|$ represents a norm of \mathbb{R}^d . When necessary, in later results, the type of norm will be specified.

We now recall the shift reduction lemma, a result that allows to split the Rényi divergence between two noised distributions into two distinct components: one involving the two original distributions, and one involving the noise. Let μ, ν, ζ be three distributions on \mathbb{R}^d and $z, a \geq 0$. We define the following quantities:

$$D_\alpha^{(z)}(\mu, \nu) = \inf_{W_\infty(\mu, \mu') \leq z} D_\alpha(\mu', \nu),$$

$$R_\alpha(\zeta, z) = \sup_{\|x\| < z} D_\alpha(\zeta_{-x}, \zeta),$$

where $\zeta_{-x} : y \mapsto \zeta(y - x)$, and denote by $*$ the convolution product.

Lemma 3.1 (Shift reduction (Feldman et al., 2018)). *Let μ, ν, ζ be three distributions on \mathbb{R}^d and $z, a \geq 0$. Then,*

$$D_\alpha^{(a)}(\mu * \zeta, \nu * \zeta) \leq D_\alpha^{(z+a)}(\mu, \nu) + R_\alpha(\zeta, z).$$

We now show that the shift reduction lemma allows to obtain a unified approach for RPP analysis. In fact, it gives a closed formula for the privacy guarantees of releasing a query with additive noise. This yields our General Wasserstein Mechanism (GWM) and its associated privacy guarantees.

Theorem 3.1 (General Wasserstein mechanism, GWM). *Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query and denote:*

$$\Delta_G = \max_{\substack{(s_i, s_j) \in \mathcal{S} \\ \theta \in \Theta}} W_\infty(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta)).$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the dataset X . Then, $\mathcal{M}(X) = f(X) + N$ satisfies $(\alpha, R_\alpha(\zeta, \Delta_G))$ -RPP for all $\alpha \in (1, +\infty)$ and $R_\infty(\zeta, \Delta_G)$ -PP.

While Theorem 3.1 is very general, we can easily derive explicit results for specific choices of noise distributions. In particular, for some distributions we can simplify the formula as a Rényi divergence, as shown by the following proposition.

Proposition 3.1. *Let $z > 0$. If ζ is defined on \mathbb{R}^d , radial, decreasing as the norm of its argument increases, then $R_\alpha(\zeta, z) = D_\alpha(\zeta_{-(z, 0, \dots, 0)}, \zeta)$. We refer to these functions as radial decreasing functions.*

Using this result to instantiate GWM with Laplacian noise, we recover the results of Song et al. (2017) for PP as a special case where $d = 1$. More interestingly, we also directly obtain a novel Gaussian mechanism and a novel Laplacian mechanism for RPP.

Corollary 3.1 (Privacy guarantees for usual noise distributions). *We note I_d the identity matrix of size d . Plugging the expressions of $R_\infty(\zeta, z)$ and $R_\alpha(\zeta, z)$ for Laplacian and Gaussian distributions, we obtain:*

- $\mathcal{M}(X) = f(X) + N$ with $N \sim \mathcal{N}(0, \frac{\alpha \Delta_G^2}{2\varepsilon} I_d)$ and Δ_G computed w.r.t. the l_2 norm is (α, ε) -RPP.
- $\mathcal{M}(X) = f(X) + L$ with $L \sim \text{Lap}(0, \rho I_d)$ and Δ_G computed w.r.t. the l_1 norm is $\left(\alpha, \frac{1}{\alpha-1} \log\left(\frac{\alpha}{2\alpha-1} e^{\Delta_G(\alpha-1)/\rho} + \frac{\alpha-1}{2\alpha-1} e^{-\Delta_G \alpha/\rho}\right)\right)$ -RPP.
- $\mathcal{M}(X) = f(X) + L$ with $L \sim \text{Lap}\left(0, \frac{\Delta_G}{\varepsilon} I_d\right)$ with Δ_G computed w.r.t. the l_1 norm is ε -PP.

The results of Corollary 3.1 are analogous to the results of Mironov (2017) for RDP, where the sensitivity of the query is replaced by Δ_G . It enables us to directly compare the utility of a RDP mechanism in the group privacy setting and the GWM in RPP. Considering Example 3, we have $\Delta_G \leq r_1 + r_2$, which is smaller than $\Delta_{\text{GROUP}} = 2r$. Therefore, GWM achieves better utility than group RDP in this case. This observation can be generalized to other settings as the utility guarantees of the Wasserstein mechanism of Song et al. (2017) extend to the GWM.

Proposition 3.2 (Utility of the GWM, informal). *Under mild conditions, an additive mechanism offers better utility in the GWM setting than in the group privacy setting (see Appendix A.2.4 for details).*

One drawback of GWM is that in some cases, Δ_G may be large, as it depends on ∞ -Wasserstein distances. In Example 1, $\Delta_G = \Delta_{\text{GROUP}} = 2$, thus GWM gives no utility advantage compared to group RDP. In Example 2, $\Delta_G = 98$ is large although the event $X_i = 100$ is rare. We deal with this issue in the next section.

4 Improving Utility by Relaxing the W_∞ Constraint

In this section, we propose two ways to improve the utility of GWM by relaxing the ∞ -Wasserstein constraint in the calibration of the noise.

4.1 δ -Approximation of (α, ε) -RPP

Our first approach is to define an approximation of Rényi Pufferfish Privacy that allows a low probability set of values to be disregarded.

Definition 4.1 (Approximate Rényi Pufferfish privacy). *A privacy mechanism \mathcal{M} is said to be $(\alpha, \varepsilon, \delta)$ -approximate Rényi Pufferfish private in a framework*

$(\mathcal{S}, \mathcal{Q}, \Theta)$ if for all $\theta \in \Theta$ and for all secret pairs $(s_i, s_j) \in \mathcal{Q}$, there exists E, E' such that $P(E) \geq 1 - \delta, P(E') \geq 1 - \delta$ and:

$$\begin{aligned} D_\alpha(P(\mathcal{M}(X) | s_i, \theta, E), P(\mathcal{M}(X) | s_j, \theta, E')) &\leq \varepsilon, \\ D_\alpha(P(\mathcal{M}(X) | s_j, \theta, E'), P(\mathcal{M}(X) | s_i, \theta, E)) &\leq \varepsilon, \end{aligned}$$

where $X \sim \theta$ and (s_i, s_j) is such that $P(s_i | \theta) \neq 0, P(s_j | \theta) \neq 0$.

This definition implies (ε, δ) -PP when $\alpha \rightarrow +\infty$.

Proposition 4.1. *If \mathcal{M} is $(+\infty, \varepsilon, \delta)$ -approximate RPP, then it is (ε, δ) -PP.*

We now design an approximate Wasserstein mechanism for Rényi Pufferfish privacy. To do so, we rely on the notion of (z, δ) -proximity (named *closeness* in (Chen and Ohrimenko, 2023)).

Definition 4.2 ((z, δ) -proximity). Let μ, ν two distributions on \mathbb{R}^d and $z \geq 0, \delta \in (0, 1)$. We say that μ and ν are (z, δ) -near if there exists a coupling π between μ and ν and $\mathcal{R} \subset \text{supp}(\pi)$ such that $\int_{\mathcal{R}} d\pi(x, y) \geq 1 - \delta$ and $\forall (x, y) \in \mathcal{R}, \|x - y\| \leq z$.

We also need to extend the shift reduction lemma of Feldman et al. (2018) to account for shifts that are (z, δ) -near to the original distribution μ , instead of shifts μ' such that $W_\infty(\mu, \mu') \leq z$. Our result relies on the following characterization of (z, δ) -proximity.

Lemma 4.1. *μ and ν are (z, δ) -near iff $\exists W \sim \mu, Y \sim \nu$ and $V \in \mathcal{P}(\mathbb{R}^d)$ such that $W + V = Y$ and $P(\|V\| > z) < \delta$.*

Lemma 4.2 (Approximate shift reduction). *Let μ, ν, ζ be three distributions on \mathbb{R}^d . We denote $D_\alpha^{(z, \delta)}(\mu, \nu) = \inf_{\mu', \mu' \text{ (z, } \delta)\text{-near}} D_\alpha(\mu', \nu)$. Then, for all $\delta \in (0, 1)$, there exists an event E such that $P(E) \geq 1 - \delta$ and:*

$$\begin{aligned} D_\alpha((\mu * \zeta)|_E, (\nu * \zeta)) \\ \leq D_\alpha^{(z, \delta)}(\mu, \nu) + R_\alpha(\zeta, z) + \frac{\alpha}{\alpha - 1} \log\left(\frac{1}{1 - \delta}\right). \end{aligned}$$

This approximate shift reduction lemma provides a general mechanism to achieve approximate RPP.

Theorem 4.1 (General approximate Wasserstein mechanism, GAWM). *Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query. For all $\delta \in (0, 1)$, let us denote:*

$$\begin{aligned} \Delta_{G, \delta} &> \inf\{z \in \mathbb{R}; \forall (s_i, s_j) \in \mathcal{S}, \forall \theta \in \Theta, \\ &(P((f(X)|_{s_i}, \theta), P(f(X)|_{s_j}, \theta)) \text{ are } (z, \delta)\text{-near}\}. \end{aligned}$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the dataset X . Then, $\mathcal{M} = f(X) + N$ satisfies $(\alpha, R_\alpha(\zeta, \Delta_{G, \delta}) + \frac{\alpha}{\alpha - 1} \log \frac{1}{1 - \delta}, \delta)$ -approximate RPP for all $\alpha \in (1, +\infty)$ and $(R_\infty(\zeta, \Delta_{G, \delta}) + \log \frac{1}{1 - \delta}, \delta)$ -PP.

From this general result, we can then design approximate RPP mechanisms for usual noise distributions. These results are similar to those of the general Wasserstein mechanism (see Corollary 3.1) but with an additive term that depends on δ . We refer to Appendix A.3.5 for details. Using this new mechanism, we can obtain better utility at a small privacy cost for queries that take large values with small probability. In Example 2, we have $\Delta_G = 98$ while for $\delta = 3 \cdot 10^{-3}$, $\Delta_{G, \delta} = 1$, which yields a major improvement in utility. This observation also holds in a more general case.

Proposition 4.2 (Utility of the GAWM, informal). *At a privacy cost of $\delta \in (0, 1)$, the GAWM offers more utility than the GWM (see Appendix A.3.7 for details).*

Remark (Relation to distribution privacy). A related result has been shown by Chen and Ohrimenko (2023) for the distribution privacy framework (see Appendix A.3.6 for the definition of distribution privacy and the result). The formulation of the results are similar, despite employing a different proof technique to get the conclusions. We prove a connection between the two results by establishing a formal equivalence between Pufferfish privacy and distribution privacy, which appears to be novel and could be of independent interest. In the interest of space, we refer to Appendix A.3.6 for the formal result and its proof. While our approximate shift reduction result (Lemma 4.2) induces an additional term which prevents us from recovering exactly the results of Chen and Ohrimenko (2023) in the particular case of the Laplace mechanism for PP, our result can be used with a wide range of noise distributions and in the RPP framework, which is more general than PP (and thus more general than distribution privacy).

4.2 Leveraging p -Wasserstein Metrics

As another way to improve the utility of the GWM, we propose to use shifts constrained by p -Wasserstein metrics instead of ∞ -Wasserstein metrics, thereby replacing the worst case transportation cost between $P(f(X)|_{s_i}, \theta)$ and $P(f(X)|_{s_j}, \theta)$ by moments of the transportation cost. This idea was explored in a different context by Altschuler and Chewi (2023), who considered Orlicz-Wasserstein shifts for Gaussian noise and identified a dependency between the noise distribution and the selected Wasserstein shift constraint. They argue that the Orlicz-Wasserstein metric is the “right” metric to use for the shifted Rényi analysis because the original shift reduction lemma fails for weaker shifts. Inspired by these considerations, we broaden the applicability of the Orlicz-Wasserstein shift reduction lemma of Altschuler and Chewi (2023) by adapting their result to a wider range of noise distributions.

Lemma 4.3 (Generalized shift reduction for radial decreasing noises). *Let ζ be a radial decreasing noise distribution: $\zeta(z) = \zeta_0(\|z\|)$. Let $z, p, q > 0$ such that $1/p + 1/q = 1$. We note: $D_{\alpha, \alpha', \zeta}^{(z)}(\mu, \nu) = \inf_{\xi: W \sim \xi} \inf_{[\exp((\alpha'-1)D_{\alpha'}(\zeta_0 * \|W\|, \zeta_0))] \leq z} D_{\alpha}(\mu * \xi, \nu)$.*

Then, we have:

$$D_{\alpha}(\mu * \zeta, \nu * \zeta) \leq D_{p(\alpha-1)+1, q(\alpha-1)+1, \zeta}^{(z)}(\mu, \nu) + \frac{\log(z)}{q(\alpha-1)}.$$

In the case $q = 1$:

$$D_{\alpha}(\mu * \zeta, \nu * \zeta) \leq D_{\infty, \alpha, \zeta}^{(z)}(\mu, \nu) + \frac{\log(z)}{\alpha-1}.$$

This lemma yields a general Wasserstein mechanism that incorporates the noise distribution within the shift.

Theorem 4.2 (Distribution Aware General Wasserstein Mechanism, DAGWM). *Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query and ζ a radial decreasing noise distribution: $\zeta(z) = \zeta_0(\|z\|)$. Let $q \geq 1$. For $(s_i, s_j) \in \mathcal{Q}, \theta \in \Theta$, we note $\mu_i^{\theta} = P(f(X)|s_i, \theta)$. We denote:*

$$\Delta_G^{\zeta, q, \alpha} = \max_{\substack{(s_i, s_j) \in \mathcal{S} \\ \theta \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^{\theta}, \mu_j^{\theta})} \mathbb{E} \left[e^{q(\alpha-1)D_{q(\alpha-1)+1}(\zeta_0 * \|X-Y\|, \zeta_0)} \right].$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the dataset X . Then, $\mathcal{M}(X) = f(X) + N$ satisfies $(\alpha, \frac{\log(\Delta_G^{\zeta, q, \alpha})}{q(\alpha-1)})$ -RPP for all $\alpha \in (1, +\infty)$ and $\lim_{\alpha \rightarrow +\infty} \frac{\log(\Delta_G^{\zeta, q, \alpha})}{q(\alpha-1)}$ -PP.

Leveraging this result allows for the design of mechanisms with sensitivity constrained by p -Wasserstein distances (W_p). In particular, we will consider noise drawn from generalized Cauchy distributions, originally introduced by Rider (1957).

Definition 4.3 (Generalized Cauchy Distributions). Let $k \geq 2, \lambda > 0$. We say that the real random variable $V \sim \text{GCauchy}(\lambda, k)$ if it has the following density: $\zeta_{k, \lambda}(x) = \frac{\beta_{k, \lambda}}{(1+(\lambda x)^2)^{k/2}}, x \in \mathbb{R}$ and $\int \zeta_{k, \lambda}(x) dx = 1$. The Cauchy distribution is the special case $k = 2$.

Using generalized Cauchy noise enables to consider W_p shifts while ensuring the existence of moments for large values of k .

Corollary 4.1 (Cauchy Mechanism). *We denote Q_{α} the Legendre polynomial of integer index $\alpha > 1$. Let $k \geq 2$ and $q \geq 1$ such that $kq(\alpha-1)/2$ is an integer. We note: $\Delta_G^{kq(\alpha-1)} = \max_{\substack{(s_i, s_j) \in \mathcal{S} \\ \theta \in \Theta}} W_{kq(\alpha-1)}(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta))$,*

with $W_{kq(\alpha-1)}$ computed with the l_2 norm. Then, $\mathcal{M}(X) = f(X) + V$ with $V \sim \text{GCauchy}(\lambda, k)$ is $\left(\alpha, \frac{\log \frac{\beta_{k, \lambda} \pi}{\lambda} Q_{kq(\alpha-1)/2} \left(1 + \left(\frac{\Delta_G^{kq(\alpha-1)}}{\lambda} \right)^2 \right)}{q(\alpha-1)} \right)$ -RPP.

In Example 1, for $q = 1$ and $\alpha = k = 2$, we have $\Delta_G^{\zeta, 2, 2} = \sqrt{1+3\rho}$ and noising with $V \sim \text{Cauchy}(\lambda)$ in DAGWM ensures $\left(\alpha, \frac{\log(1+\frac{1+3\rho}{\lambda^2})}{\alpha-1} \right)$ -RPP, while the GWM for the same noise distribution gives $\left(\alpha, \frac{\log(1+\frac{4}{\lambda^2})}{\alpha-1} \right)$ -RPP. Hence, in this case DAGWM is better than GWM, as it allows to capture the correlation between the attributes. In the general case, DAGWM consistently outperforms GWM.

Proposition 4.3 (Utility of the DAGWM, informal). *The DAGWM always offers more utility than the GWM at no privacy cost (see Appendix A.4.4 for details).*

5 Protection Against Close Adversaries

In Pufferfish, the set Θ represents the possible beliefs of the adversary. It needs to be large enough to prevent harmful privacy leaks, but there is also a no free lunch theorem that states that if Θ is too large then the resulting mechanism will have poor utility (Kifer and Machanavajjhala, 2014). Hence, it is important to quantify the privacy protection offered by a mechanism \mathcal{M} when the belief θ' of the adversary is not in Θ .

This question has been addressed for a Pufferfish private mechanism by Song et al. (2017). The theorem derived by Song et al. (2017), which we recall in Appendix A.5 for completeness, shows that if θ' is Δ -close to some $\theta \in \Theta$, then \mathcal{M} retains its Pufferfish privacy guarantees for θ' up to an additive penalty 2Δ . However, Δ is measured in ∞ -Rényi divergence, which corresponds to a worst-case scenario, and can thus be very large. We extend this result to our RPP framework, allowing the use of α -Rényi divergences.

Theorem 5.1 (RPP protection against close adversaries). *Let $p, q, r > 0$ such that $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$, and let \mathcal{M} be a mechanism that satisfies \mathcal{M} is $(q(\alpha-1/p), \varepsilon)$ -RPP in a framework $(\mathcal{S}, \mathcal{Q}, \Theta)$. Let $\theta' \notin \Theta$ and*

$$\Delta_p^1 = \inf_{\theta \in \Theta} \sup_{s_i \in \mathcal{S}} D_{\alpha p}(P(X|s_i, \theta'), P(X|s_i, \theta)),$$

$$\Delta_r^2 = \inf_{\theta \in \Theta} \sup_{s_i \in \mathcal{S}} D_{(\alpha-1)r+1}(P(X|s_i, \theta), P(X|s_i, \theta')).$$

Then, for all $\alpha \in (1, \infty)$, \mathcal{M} satisfies:

$$\left(\alpha, \left(1 + \frac{1}{r(\alpha-1)} \right) \varepsilon + \left(1 + \frac{\frac{1}{r} + \frac{1}{q}}{\alpha-1} \right) \Delta_p^1 + \Delta_r^2 \right)\text{-RPP}$$

for $(\mathcal{S}, \mathcal{Q}, \Theta')$ with $\Theta' = \Theta \cup \{\theta'\}$.

This theorem employs α -Rényi divergences and can be viewed as a generalization of the result of Song et al. (2017), which we recover as a special case for $\alpha = +\infty$. Our result can provide better privacy guarantees in situations where the original one gives poor guarantees. Note that neither Theorem 5.1 nor the original result of Song et al. (2017) exploit the characteristics of the particular mechanism \mathcal{M} of interest in the quantification of the additional privacy loss. As a matter of fact, it is likely that a mechanism with large variance would yield more robust guarantees. Interestingly, we can address this issue by refining our result to additive noise mechanisms using the shift reduction lemma, see Appendix A.5 for details.

A possible application of the above result is to analyze the privacy guarantees of differentially private mechanisms under weakly-correlated data.

6 Privacy Amplification by Iteration

Analyzing the privacy guarantees of Pufferfish privacy under composition is known to be challenging (Kifer and Machanavajjhala, 2014). While Pufferfish satisfies a form of parallel composition (see Appendix A.6 for the result in RPP), to our knowledge there does not exist any theorem providing mechanism-agnostic guarantees for sequential composition in Pufferfish privacy. As an alternative to composition, we show in this section that RPP is amenable to privacy amplification by iteration, providing a way to analyze iterative gradient descent algorithms for convex optimization.

6.1 Theoretical Results

In differential privacy, privacy amplification by iteration (PABI) allows to evaluate the privacy loss of applying multiple contractive noisy iterations to a dataset and releasing only the output of the last iteration (Feldman et al., 2018; Altschuler and Talwar, 2022). PABI has often been employed in private machine learning to analyze the privacy cost of projected noisy stochastic gradient descent (DP-SGD), bypassing the use of composition (Feldman et al., 2018). However, existing PABI results for differential privacy cannot be used in Pufferfish privacy. These results consider the distribution shift between two processes performed on two neighboring datasets (equal up to

one element) and how this additional shift propagates through the rest of the iterations. In Pufferfish, privacy is obtained by conditioning over secrets and the dataset is sampled from an adversary’s prior. This means that two datasets with different secrets might share no common elements. Hence, the original worst case PABI analysis must be adapted to account for shifts at each iteration, while measuring these shifts based on the dataset distribution conditioned by the secrets.

We start by defining contractive noisy iterations.

Definition 6.1 (Contractive noisy iteration (CNI)). Let $\mathcal{Z} \subset \mathbb{R}^d$. Given an initial random state $W_0 \in \mathcal{Z}$, a sequence of random variables $\{X_t\}$, a sequence of contractive maps in their first argument $\psi_t : \mathcal{Z} \times \mathcal{D} \rightarrow \mathcal{Z}$ and a sequence of noise distributions $\{\zeta_t\}$, we define the Contractive Noisy Iteration (CNI) by the following update rule:

$$W_{t+1} = \psi_{t+1}(W_t, X_{t+1}) + N_{t+1},$$

where $N_{t+1} \sim \zeta_{t+1}$. For brevity, we refer to the result W_T of the CNI at the time step T by $CNI_T(W_0, \{X_t\}, \{\psi_t\}, \{\zeta_t\})$.

As opposed to the work of Feldman et al. (2018), we make an explicit reference to the dataset distribution modeled by the random sequence $\{X_t\}$ in the CNI definition. The original PABI analysis leverages a contraction lemma that we need to adapt to the Pufferfish setting. We prove a new contraction lemma which incorporates the ∞ -Wasserstein distance to take into account the dataset distribution.

Lemma 6.1 (Dataset Dependent Contraction lemma). Let ψ be a contractive map in its first argument on $(\mathcal{Z}, \|\cdot\|)$. Let X, X' be two r.v.’s. Suppose that $\sup_w W_\infty(\psi(w, X), \psi(w, X')) \leq s$. Then, for $z > 0$:

$$D_\alpha^{(z+s)}(\psi(W, X), \psi(W', X')) \leq D_\alpha^{(z)}(W, W').$$

Coupled with the original shift reduction lemma (Lemma 3.1), this contraction lemma yields a relaxation of the original PABI bounds, allowing take into account the dataset distribution in the measurement of the shifts.

Theorem 6.1 (Dataset Dependent PABI). Let X_T and X'_T denote the output of $CNI_T(W_0, \{\psi_t\}, \{\zeta_t\}, X)$ and $CNI_T(W_0, \{\psi_t\}, \{\zeta_t\}, X')$. Let $s_t = \sup_w W_\infty(\psi(w, X), \psi(w, X'))$. Let a_1, \dots, a_T be a sequence of reals and let $z_t = \sum_{i \leq t} s_i - \sum_{i \leq t} a_i$. If $z_t \geq 0$ for all t , then, we have:

$$D_\alpha^{(z_T)}(X_T, X'_T) \leq \sum_{t=1}^T R_\alpha(\zeta_t, a_t).$$

6.2 Application to Convex Optimization

Our new PABI bounds allow for an analysis of the projected noisy gradient descent in the RPP framework.

Let $m, d, T > 0$. Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish framework. We note \mathcal{X} the set of values taken by the elements of the dataset. Let $(s_i, s_j) \in \mathcal{Q}, \theta \in \Theta$. We note $X = (X_1, \dots, X_T) \sim P(X|s_i, \theta)$ and $X' = (X'_1, \dots, X'_T) \sim P(X|s_j, \theta)$. We assume that $\mathcal{X} \subset \mathbb{R}^m$. Let $f : \mathbb{R}^d \times \mathcal{X} \rightarrow \mathbb{R}$ be an objective function. We assume that f is convex and make the following additional assumptions:

- f is L -Lipschitz in its first argument: there exists $L > 0$ such that $\forall x \in \mathcal{X}, w_1, w_2 \in \mathbb{R}^d$,

$$\|f(w_1, x) - f(w_2, x)\| \leq L\|w_1 - w_2\|.$$

- f is β -smooth in its first argument: there exists $\beta > 0$ such that $\forall x \in \mathcal{X}, w_1, w_2 \in \mathbb{R}^d$,

$$\|\nabla_w f(w_1, x) - \nabla_w f(w_2, x)\| \leq \beta\|w_1 - w_2\|.$$

- f satisfies the following condition: $\forall x_1, x_2 \in \mathcal{X}, w_1 \in \mathbb{R}^d, \exists C_{w_1} > 0$ such as :

$$\|\nabla_w f(w_1, x_1) - \nabla_w f(w_1, x_2)\| \leq C_{w_1}\|x_1 - x_2\|.$$

The last assumption, which is used in the adversarial training literature (see e.g., Liu et al., 2020), is satisfied in certain simple settings as linear regression. It enables to take into account the distribution of the gradients as a function of the distribution of the data in our PABI analysis. From these conditions, we can show that $\forall x_1, x_2 \in \mathcal{X}, w_1 \in \mathbb{R}^d, \exists C_{w_1} > 0$ such that:

$$\|\nabla_w f(w_1, x_1) - \nabla_w f(w_1, x_2)\| \leq \min(2L, C_{w_1}\|x_1 - x_2\|).$$

Let $\Pi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a projection over a compact $\mathcal{K} \subset \mathbb{R}^d$ and $\eta > 0$ such that $\eta < 2/\beta$. By Proposition 18 of Feldman et al. (2018), the weight update function:

$$\begin{aligned} \psi : \mathbb{R}^d \times \mathcal{X} &\rightarrow \mathbb{R}^d \\ (v, x) &\mapsto \Pi(v - \eta \nabla_w f(v, x)) \end{aligned}$$

is contractive. Let $W_0 = W'_0 \in \mathcal{K}$ be the initial weight and ζ_1, \dots, ζ_T be noise distributions. We note $(N_1, \dots, N_T) \sim \otimes_{k=1}^T \zeta_k$ and for all $t \in \llbracket 1, T \rrbracket$, $W_t = \psi(W_{t-1}, X_t) + N_t$, $W'_t = \psi(W'_{t-1}, X'_t) + N_t$. Then, we note:

$$\begin{aligned} s_t &= \eta \sup_{v \in \mathcal{K}} W_\infty(\nabla_w f(v, X_t), \nabla_w f(v, X'_t)) \\ &\leq \eta \min(2L, \sup_{v \in \mathcal{K}} C_v W_\infty(X_t, X'_t)) \end{aligned}$$

As an example of application of Theorem 6.1, in the case where $N_t \sim \mathcal{N}(0, \sigma^2 I_d)$ as in DP-SGD, and taking $(a_t) = (s_t)$, we have:

$$D_\alpha(W_T, W'_T) \leq \frac{\alpha \eta^2}{2\sigma^2} \sum_{t=1}^T \min(2L, \sup_{v \in \mathcal{K}} C_v W_\infty(X_t, X'_t))^2.$$

To interpret this formula, we can look at some extreme cases. Let the secrets $s_t^a = \{X_t = a\}$, $s_t^b = \{X_t = b\}$, $t \in \llbracket 1, T \rrbracket, a, b \in \mathcal{X}$. If the adversary has a prior of high correlations, such as for example $X_1 = \dots = X_t$, $X'_1 = \dots = X'_t$, we get:

$$D_\alpha(W_T, W'_T) \leq \frac{T \alpha \eta^2}{2\sigma^2} \min(2L, \|a - b\| \sup_{v \in \mathcal{K}} C_v)^2,$$

which is no better than the group privacy analysis. On the other hand, when data points are independent as in differential privacy, we get:

$$D_\alpha(W_T, W'_T) \leq \frac{\alpha \eta^2}{2\sigma^2} \min(2L, \|a - b\| \sup_{v \in \mathcal{K}} C_v)^2.$$

In this case, the upper bound is independent of T and we thus obtain much better results than with group privacy. In fact, our result in Theorem 6.1 is general enough to recover the original results of Feldman et al. (2018) for DP-SGD as a special case.

Remark (DP as a special case, informal). Theorem 6.1 allows to recover the same privacy bounds as Theorem 23 of Feldman et al. (2018) (see Appendix A.6.4 for details).

7 Conclusion

We presented a new framework, called Rényi Pufferfish privacy, which extends the original Pufferfish privacy definition. We designed general additive noise mechanisms for achieving (approximate) Rényi Pufferfish privacy and discussed their utility and robustness to close adversaries. As a way to use Pufferfish privacy to analyze sequential algorithms, we derived a privacy amplification by iteration result which allows to bypass the lack of sequential composition theorems. We put forward a first application of this analysis for convex optimization with gradient descent. We believe that our results are a first step towards the integration of Pufferfish in machine learning algorithms. Potential areas for future work include a tighter PABI analysis with other shift reduction lemmas, and a numerical analysis of Rényi Pufferfish privacy mechanisms to optimize utility in practical use-cases.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. (2016). Deep

- learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, page 308–318.
- Altschuler, J. M. and Chewi, S. (2023). Faster high-accuracy log-concave sampling via algorithmic warm starts. arXiv:2302.10249.
- Altschuler, J. M. and Talwar, K. (2022). Privacy of noisy stochastic gradient descent: More iterations without more privacy loss. In *NeurIPS*.
- Bun, M. and Steinke, T. (2016). Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography*, pages 635–658. Springer Berlin Heidelberg.
- Chen, M. and Ohrimenko, O. (2023). Protecting global properties of datasets with distribution privacy mechanisms. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 7472–7491. PMLR.
- Ding, N. (2022). Kantorovich mechanism for pufferfish privacy. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 5084–5103. PMLR.
- Dwork, C. and Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407.
- Feldman, V., Mironov, I., Talwar, K., and Thakurta, A. (2018). Privacy amplification by iteration. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 521–532.
- Hardt, M. and Roth, A. (2013). Beyond worst-case analysis in private singular vector computation. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, page 331–340.
- He, X., Machanavajjhala, A., and Ding, B. (2014). Blowfish privacy: Tuning privacy-utility trade-offs using policies. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, page 1447–1458. Association for Computing Machinery.
- Humphries, T., Oya, S., Tulloch, L., Rafuse, M., Goldberg, I., Hengartner, U., and Kerschbaum, F. (2023). Investigating membership inference attacks under data dependencies. In *2023 IEEE 36th Computer Security Foundations Symposium (CSF) (CSF)*, pages 194–209. IEEE Computer Society.
- Kawamoto, Y. and Murakami, T. (2019). Local obfuscation mechanisms for hiding probability distributions. In *Computer Security – ESORICS 2019*, pages 128–148, Cham. Springer International Publishing.
- Kessler, S., Buchmann, E., and Böhm, K. (2015). Deploying and evaluating pufferfish privacy for smart meter data. In *2015 IEEE 12th Intl Conf on Ubiquitous Intelligence and Computing and 2015 IEEE 12th Intl Conf on Autonomic and Trusted Computing and 2015 IEEE 15th Intl Conf on Scalable Computing and Communications and Its Associated Workshops (UIC-ATC-ScalCom)*, pages 229–238.
- Kifer, D. and Machanavajjhala, A. (2014). Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems*, 39(1).
- Liu, C., Salzmann, M., Lin, T., Tomioka, R., and Süsstrunk, S. (2020). On the loss landscape of adversarial training: Identifying challenges and how to overcome them. In *Advances in Neural Information Processing Systems*, pages 21476–21487.
- Mironov, I. (2017). Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275.
- Niu, C., Zheng, Z., Tang, S., Gao, X., and Wu, F. (2019). Making big money from small sensors: Trading time-series data under pufferfish privacy. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, pages 568–576.
- Ou, L., Qin, Z., Liao, S., Yin, H., and Jia, X. (2018). An optimal pufferfish privacy mechanism for temporally correlated trajectories. *IEEE Access*, 6:37150–37165.
- Rider, P. R. (1957). Generalized cauchy distributions. *Annals of the Institute of Statistical Mathematics*, 9:215–223.
- Song, S., Wang, Y., and Chaudhuri, K. (2017). Pufferfish privacy mechanisms for correlated data. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD '17, page 1291–1306.
- Verdú, S. (2023). The cauchy distribution in information theory. *Entropy*, 25(2).

A Appendix

This appendix provides some useful background, as well as more detailed versions of our results, along with their proofs.

A.1 Properties of Rényi Pufferfish Privacy (Section 2)

A.1.1 Proof of Proposition 2.1

Proposition 2.1 (Post-processing). *Let \mathcal{M}_1 be a randomized algorithm and \mathcal{M} be (α, ε) -RPP. Then,*

$$\begin{aligned} D_\alpha(P(\mathcal{M}_1(\mathcal{M}(X)) | s_i, \theta), P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)) \\ \leq D_\alpha(P(\mathcal{M}(X) | s_i, \theta), P(\mathcal{M}(X) | s_j, \theta)) \leq \varepsilon. \end{aligned}$$

Proof. Let $\mathcal{S}, \mathcal{Q}, \Theta$ be a Pufferfish framework. Let $(s_i, s_j) \in \mathcal{Q}, \theta \in \Theta, \alpha > 1$ and $\varepsilon > 0$. Let \mathcal{M}_1 be a randomized algorithm and \mathcal{M} satisfying (α, ε) -RPP. Then,

$$\begin{aligned} D_\alpha(P(\mathcal{M}(X) | s_i, \theta), P(\mathcal{M}(X) | s_j, \theta)) &= \mathbb{E}_{Z \sim P(\mathcal{M}(X) | s_j, \theta)} \left[\left(\frac{P(\mathcal{M}(X) = Z | s_i, \theta)}{P(\mathcal{M}(X) = Z | s_j, \theta)} \right)^\alpha \right] \\ &= \mathbb{E}_{(Z', Z) \sim (P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta), P(\mathcal{M}(X) | s_j, \theta))} \left[\left(\frac{P(\mathcal{M}(X) = Z | s_i, \theta)}{P(\mathcal{M}(X) = Z | s_j, \theta)} \right)^\alpha \right] \\ &= \mathbb{E}_{(Z', Z) \sim (P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta), P(\mathcal{M}(X) | s_j, \theta))} \left[\left(\frac{P(\mathcal{M}(X) = Z | s_i, \theta) P(\mathcal{M}_1(\mathcal{M}(X)) = Z' | \mathcal{M}(X) = Z)}{P(\mathcal{M}(X) = Z | s_j, \theta) P(\mathcal{M}_1(\mathcal{M}(X)) = Z' | \mathcal{M}(X) = Z)} \right)^\alpha \right] \\ &= \mathbb{E}_{(Z', Z) \sim (P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta), P(\mathcal{M}(X) = s_j, \theta))} \left[\left(\frac{P(\mathcal{M}_1(\mathcal{M}(X)) = Z', \mathcal{M}(X) = Z | s_i, \theta)}{P(\mathcal{M}_1(\mathcal{M}(X)) = Z', \mathcal{M}(X) = Z | s_j, \theta)} \right)^\alpha \right] \\ &= \mathbb{E}_{Z' \sim P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)} \left[\mathbb{E}_{Z \sim P(\mathcal{M}(X) | \mathcal{M}_1(\mathcal{M}(X)), s_j, \theta)} \left[\left(\frac{P(\mathcal{M}_1(\mathcal{M}(X)) = Z', \mathcal{M}(X) = Z | s_i, \theta)}{P(\mathcal{M}_1(\mathcal{M}(X)) = Z', \mathcal{M}(X) = Z | s_j, \theta)} \right)^\alpha \right] \right] \\ &\geq \mathbb{E}_{Z' \sim P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)} \left[\left(\mathbb{E}_{Z \sim P(\mathcal{M}(X) | \mathcal{M}_1(\mathcal{M}(X)), s_j, \theta)} \left[\frac{P(\mathcal{M}_1(\mathcal{M}(X)) = Z', \mathcal{M}(X) = Z | s_i, \theta)}{P(\mathcal{M}_1(\mathcal{M}(X)) = Z', \mathcal{M}(X) = Z | s_j, \theta)} \right] \right)^\alpha \right] \quad \text{Jensen inequality} \\ &= \mathbb{E}_{Z' \sim P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)} \left[\left(\mathbb{E}_{Z \sim P(\mathcal{M}(X) | \mathcal{M}_1(\mathcal{M}(X)), s_j, \theta)} \left[\frac{P(\mathcal{M}_1(\mathcal{M}(X)) = Z' | s_i, \theta)}{P(\mathcal{M}_1(\mathcal{M}(X)) = Z' | s_j, \theta)} \right] \right)^\alpha \right] \\ &= \mathbb{E}_{Z' \sim P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)} \left[\left(\frac{P(\mathcal{M}_1(\mathcal{M}(X)) = Z | s_i, \theta)}{P(\mathcal{M}_1(\mathcal{M}(X)) = Z | s_j, \theta)} \right)^\alpha \right] \\ &= D_\alpha(P(\mathcal{M}_1(\mathcal{M}(X)) | s_i, \theta), P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)). \end{aligned}$$

Thus,

$$D_\alpha(P(\mathcal{M}_1(\mathcal{M}(X)) | s_i, \theta), P(\mathcal{M}_1(\mathcal{M}(X)) | s_j, \theta)) \leq D_\alpha(P(\mathcal{M}(X) | s_i, \theta), P(\mathcal{M}(X) | s_j, \theta)) \leq \varepsilon. \quad \square$$

A.1.2 Proof of Proposition 2.2

Proposition 2.2 (RPP implies PP). *If \mathcal{M} is (α, ε) -RPP, it also satisfies $(\varepsilon + \frac{\log(1/\delta)}{\alpha-1}, \delta)$ -PP for all $\delta \in (0, 1)$.*

Proof. The proof technique of Mironov (2017) remains applicable in the context of Rényi Pufferfish privacy. For clarity and completeness, we showcase it here. Let $\varepsilon \geq 0, \alpha > 1$. Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish privacy framework and \mathcal{M} an (α, ε) -RPP mechanism. Let $\delta \in (0, 1), \theta \in \Theta, (s_i, s_j) \in \mathcal{Q}$ and $z \in \text{Range}(\mathcal{M})$. Then, we have:

$$P(\mathcal{M}(X) = z | s_i, \theta)^\alpha \leq e^{(\alpha-1)D_\alpha(P(\mathcal{M}(X) | s_i, \theta), P(\mathcal{M}(X) | s_j, \theta))} P(\mathcal{M}(X) = z | s_j, \theta)^{\alpha-1} \leq e^{\varepsilon(\alpha-1)} P(\mathcal{M}(X) = z | s_j, \theta)^{\alpha-1},$$

where the first inequality is obtained by Hölder inequality applied to the functions $\left(\frac{f^\alpha}{g^{\alpha-1}}\right)^{\frac{1}{\alpha}}$ and $g^{\frac{\alpha-1}{\alpha}}$. We then consider two cases:

- Case 1: $e^\varepsilon P(\mathcal{M}(X) = z | s_j, \theta) \leq \delta^{\frac{\alpha}{\alpha-1}}$. Then, $P(\mathcal{M}(X) = z | s_i, \theta) \leq \delta \leq P(\mathcal{M}(X) = z | s_j, \theta) + \delta$.

- Case 2: $e^\varepsilon P(\mathcal{M}(X) = z|s_j, \theta) > \delta^{\frac{\alpha}{\alpha-1}}$. Then,

$$\begin{aligned} P(\mathcal{M}(X) = z|s_i, \theta) &\leq (e^\varepsilon P(\mathcal{M}(X) = z|s_j, \theta)) (e^\varepsilon P(\mathcal{M}(X) = z|s_j, \theta))^{\frac{-1}{\alpha}} \\ &\leq e^\varepsilon P(\mathcal{M}(X) = z|s_j, \theta) \delta^{\frac{-1}{\alpha-1}} \\ &\leq e^\varepsilon P(\mathcal{M}(X) = z|s_j, \theta) + \delta. \end{aligned} \quad \square$$

A.2 General Wasserstein Mechanism (Section 3)

A.2.1 Proof of Proposition 3.1

Proposition 3.1. *We refer to radial decreasing functions ζ defined on \mathbb{R}^d as the functions which have the following properties:*

- ζ is a radial function: there exists a real valued function ζ_0 such that for all $x \in \mathbb{R}^d$, $\zeta(x) = \zeta_0(\|x\|)$.
- ζ_0 is decreasing.

Let ζ defined on \mathbb{R}^d be a radial decreasing function. Then, $x \in \mathbb{R}^d \mapsto D_\alpha(\zeta_{-x}, \zeta)$ is radial decreasing and for $z \in \mathbb{R}$, $R_\alpha(\zeta, z) = D_\alpha(\zeta_{0-z}, \zeta_0) = D_\alpha(\zeta_{(-z, 0, \dots, 0)}, \zeta)$.

Proof. We show that $x \in \mathbb{R}^d \mapsto D_\alpha(\zeta_{-x}, \zeta)$ is radial decreasing. Let $x \in \mathbb{R}^d$. We note: $\zeta_1(z) = \int \frac{\zeta_0(\|u - (z, 0, \dots, 0)\|)^\alpha}{\zeta_0(\|u\|)^\alpha} du$. We note ρ the rotation which rotates x around the origin onto the first coordinate: $\rho(x) = (\|x\|, 0, \dots, 0)$. We have:

$$\begin{aligned} \exp((\alpha - 1)D_\alpha(\zeta_{-x}, \zeta)) &= \int \frac{\zeta(u - x)^\alpha}{\zeta(x)^{\alpha-1}} du \\ &= \int \frac{\zeta_0(\|\rho(u - x)\|)^\alpha}{\zeta_0(\|\rho(u)\|)^\alpha} du && (\rho \text{ is norm-preserving}) \\ &= \int \frac{\zeta_0(\|\rho(u) - (\|x\|, 0, \dots, 0)\|)^\alpha}{\zeta_0(\|\rho(u)\|)^\alpha} du && (\rho \text{ is linear}) \\ &= \int \frac{\zeta_0(\|v - (\|x\|, 0, \dots, 0)\|)^\alpha}{\zeta_0(\|v\|)^\alpha} |\det J_{\rho^{-1}}(v)| dv && (\text{by change of variable } v = \rho(u)) \\ &= \zeta_1(\|x\|) && (\text{because } \rho \text{ is a rotation: } |\det J_{\rho^{-1}}(v)| = 1) \end{aligned}$$

Then, $x \mapsto D_\alpha(\zeta_{-x}, \zeta)$ is radial decreasing.

Now we prove that, for $z > 0$, $R_\alpha(\zeta, z) = D_\alpha(\zeta_{0-z}, \zeta_0)$. We have: $R_\alpha(\zeta, z) = \sup_{\|x\| < z} D_\alpha(\zeta_{-x}, \zeta) = \sup_{0 \leq y \leq z} D_\alpha(\zeta_{0-y}, \zeta_0)$ and the application $z \mapsto \zeta_0(|z|)$ is defined on \mathbb{R} , symmetric and decreasing on \mathbb{R}^+ .

Let $0 \leq z \leq a$, $N \sim \zeta$. We have:

$$\begin{aligned} e^{(\alpha-1)D_\alpha(N+(a,0,\dots,0),N)} - e^{(\alpha-1)D_\alpha(N+(z,0,\dots,0),N)} &= \int_{-\infty}^{+\infty} \frac{\zeta_0(x-a)^\alpha}{\zeta_0(x)^{\alpha-1}} dx - \int_{-\infty}^{+\infty} \frac{\zeta_0(x-z)^\alpha}{\zeta_0(x)^{\alpha-1}} dx \\ &= \int_0^{+\infty} \frac{\zeta_0(x)^\alpha - \zeta_0(x-z+a)^\alpha}{\zeta_0(x+a)^{\alpha-1}} dx - \int_{-a}^{+\infty} \frac{\zeta_0(-x-a)^\alpha - \zeta_0(-x-z)^\alpha}{\zeta_0(-x)^{\alpha-1}} dx \\ &= \int_0^{+\infty} \frac{\zeta_0(x)^\alpha - \zeta_0(x-z+a)^\alpha}{\zeta_0(x+a)^{\alpha-1}} + \frac{\zeta_0(x+z-a)^\alpha - \zeta_0(x)^\alpha}{\zeta_0(x-a)^{\alpha-1}} dx, \end{aligned}$$

using the symmetry of ζ_0 and multiple changes of variables.

Now, we show that $\int_0^{+\infty} \frac{\zeta_0(x)^\alpha - \zeta_0(x-z+a)^\alpha}{\zeta_0(x+a)^{\alpha-1}} + \frac{\zeta_0(x+z-a)^\alpha - \zeta_0(x)^\alpha}{\zeta_0(x-a)^{\alpha-1}} \geq 0$. First, we know that $\zeta_0(x+a) \leq \zeta_0(x-a)$. If $x \geq a$, it comes from the fact that ζ_0 is decreasing on \mathbb{R}^+ . If $0 \leq x \leq a$, by symmetry $\zeta_0(x-a) = \zeta_0(a-x) \geq \zeta_0(a+x)$. By the same argument, $\zeta_0(x+a-z) \leq \zeta_0(x+z-a)$. Since for any $x \geq 0$, $\zeta_0(x) \geq \zeta_0(x+a-z)$, we

have:

$$\begin{aligned} \frac{\zeta_0(x)^\alpha - \zeta_0(x-z+a)^\alpha}{\zeta_0(x+a)^{\alpha-1}} + \frac{\zeta_0(x+z-a)^\alpha - \zeta_0(x)^\alpha}{\zeta_0(x-a)^{\alpha-1}} &\geq \frac{\zeta_0(x)^\alpha - \zeta_0(x+a-z)^\alpha}{\zeta_0(x+a)^{\alpha-1}} + \frac{\zeta_0(x+a-z)^\alpha - \zeta_0(x)^\alpha}{\zeta_0(x-a)^{\alpha-1}} \\ &= (\zeta_0(x)^\alpha - \zeta_0(x+a-z)^\alpha) \left(\frac{1}{\zeta_0(x+a)^{\alpha-1}} - \frac{1}{\zeta_0(x-a)^{\alpha-1}} \right) \geq 0 \end{aligned}$$

□

A.2.2 Proof of Theorem 3.1

Theorem 3.1 (General Wasserstein mechanism, GWM). *Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query and denote:*

$$\Delta_G = \max_{\substack{(s_i, s_j) \in \mathcal{S} \\ \theta \in \Theta}} W_\infty(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta)).$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the data X . Then, $\mathcal{M}(X) = f(X) + N$ satisfies $(\alpha, R_\alpha(\zeta, \Delta_G))$ -RPP for all $\alpha \in (1, +\infty)$ and $R_\infty(\zeta, \Delta_G)$ -PP.

Proof. Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish privacy instance. Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query and denote:

$$\Delta_G = \max_{\substack{(s_i, s_j) \in \mathcal{S} \\ \theta_i \in \Theta}} W_\infty(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta)).$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the data X . We use the abuse of notation $D_\alpha(X|E, Y|E) = D_\alpha(P(X|E), P(Y|E))$. Let $\alpha > 1$, $z > 0$, $(s_i, s_j) \in \mathcal{Q}$ and $\theta \in \Theta$. By the shift reduction lemma (Lemma 3.1), we have:

$$D_\alpha \left((f(X) + N)|_{s_i, \theta}, (f(X) + N)|_{s_j, \theta} \right) \leq D_\alpha^{(z)}(f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta}) + R_\alpha(\zeta, z).$$

By definition,

$$D_\alpha^{(z)}(f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta}) = \inf_{W \in \mathcal{P}(\mathbb{R}^d); W_\infty(W, f(X)|_{s_i, \theta}) \leq z} D_\alpha(W, f(X)|_{s_j, \theta}),$$

and

$$D_\alpha^{(W_\infty(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta)))}(f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta}) = 0.$$

Then,

$$D_\alpha \left((f(X) + N)|_{s_i, \theta}, (f(X) + N)|_{s_j, \theta} \right) \leq R_\alpha(\zeta, W_\infty(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta))) \leq R_\alpha(\zeta, \Delta_G). \quad \square$$

A.2.3 Proof of Corollary 3.1

Corollary 3.1 (Privacy guarantees for usual noise distributions). *We note I_d the identity matrix of size d . Plugging the expressions of $R_\infty(\zeta, z)$ and $R_\alpha(\zeta, z)$ for Laplacian and Gaussian distributions, we obtain:*

- $\mathcal{M}(X) = f(X) + N$ with $N \sim \mathcal{N}(0, \frac{\alpha \Delta_G^2}{2\varepsilon} I_d)$ and Δ_G computed on the l_2 norm is (α, ε) -RPP.
- $\mathcal{M}(X) = f(X) + L$ with $L \sim \text{Lap}(0, \rho I_d)$ and Δ_G computed on the l_1 norm is $\left(\alpha, \frac{1}{\alpha-1} \log \left(\frac{\alpha}{2\alpha-1} e^{\Delta_G(\alpha-1)/\rho} + \frac{\alpha-1}{2\alpha-1} e^{-\Delta_G \alpha/\rho} \right) \right)$ -RPP.
- $\mathcal{M}(X) = f(X) + L$ with $L \sim \text{Lap}(0, \frac{\Delta_G}{\varepsilon} I_d)$ with Δ_G computed on the l_1 norm is ε -PP.

Proof. The result is directly obtained by plugging Rényi divergences of shifts into the GWM and using Proposition 3.1. Let $\alpha > 1, z \geq 0$.

- $\text{Lap}(0, \rho I_d)$ is radial decreasing for the l_1 norm and for $L \sim \text{Lap}(0, \rho I_d)$,

$$D_\alpha(L + z, L) = \frac{1}{\alpha - 1} \log \left(\frac{\alpha}{2\alpha - 1} e^{\Delta_G(\alpha-1)/\rho} + \frac{\alpha - 1}{2\alpha - 1} e^{-\Delta_G \alpha/\rho} \right).$$

- $\mathcal{N}(0, \sigma^2 I_d)$ is radial decreasing for the l_2 norm and for $N \sim \mathcal{N}(0, \sigma^2 I_d)$, $D_\alpha(N + z, N) = \frac{\alpha z^2}{2\sigma^2}$.

□

A.2.4 Utility of the GWM (Proposition 3.2)

Below, we make the informal result of Proposition 3.2 precise and provide its proof.

Proposition A.1 (Utility of the GWM). *Let $n, d_1, \dots, d_n \in \mathbb{N}^*$. Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish framework such that, for each $\theta \in \Theta$, $\theta = \otimes_{k=1}^n \theta_k$, with $\theta_k \in \mathcal{P}(\mathbb{R}^{d_k})$. We note $X = (X_1^1, \dots, X_{d_1}^1, \dots, X_{d_n}^n) \sim \theta$. We assume that $s_{i,k}^a = \{X_i^k = a\} \in \mathcal{S}$ and $\mathcal{Q} = \{(s_{i,k}^a, s_{i,k}^b); k \in \{1, \dots, n\}, i \in \{1, \dots, d_k\}, a, b \in \mathbb{R}\}$. Following Song et al. (2017), we define the corresponding group differential privacy of the Pufferfish framework as: $G_k = (x_1^k, \dots, x_{d_k}^k) \in \mathbb{R}^{d_k}$ and $D_k = \{(x, x') \in \mathbb{R}^{d_k} \text{ such that } x \text{ and } x' \text{ only differ in } G_k\}$.*

$$\Delta_{\text{GROUP}}(f) = \max_{k \in \{1, \dots, n\}} \max_{(x, x') \in D_k} \|f(x) - f(x')\|.$$

Then, $\Delta_G \leq \Delta_{\text{GROUP}}(f)$.

Proof. Let $(s_{i,l}^a, s_{i,l}^b) \in \mathcal{Q}, \theta \in \Theta$, with $\theta = \otimes_{k=1}^n \theta_k$. Let $Y \sim P(f(X)|s_{i,l}^a, \theta)$. Let $Z \sim \theta_l|s_{i,l}^b$ drawn independently from Y . For $k \in \llbracket 1, n \rrbracket, i \in \llbracket 1, d_k \rrbracket$. We define $Y_i^k = \begin{cases} Y_i^k & \text{if } k \neq l \\ Z_i & \text{else} \end{cases}$ and $Y' = (Y_1^1, \dots, Y_{d_1}^1, \dots, Y_{d_n}^n)$.

Then, $(Y, Y') \in D_l, Y' \sim P(f(X)|s_{i,l}^b)$ and:

$$\|Y - Y'\| \leq \max_{(x, x') \in D_l} \|f(x) - f(x')\| \leq \Delta_{\text{GROUP}}(f).$$

Then, $W_\infty(P(f(X)|s_{i,l}^a), P(f(X)|s_{i,l}^b)) \leq \Delta_{\text{GROUP}}(f)$ and $\Delta_G \leq \Delta_{\text{GROUP}}(f)$. □

A.3 Approximate General Wasserstein Mechanism (Section 4.1)

A.3.1 Proof of Lemma 4.1

Lemma 4.1. *μ and ν are (z, δ) -near iff $\exists X \sim \mu, Y \sim \nu$ and $V \in \mathcal{P}(\mathbb{R}^d)$ such that $X + V = Y$ and $P(\|V\| > z) < \delta$.*

Proof. Let $z \geq 0, \delta \in (0, 1), \mu, \nu$ two distributions on \mathbb{R}^d such that μ and ν are (z, δ) -near. Then, there exists π a coupling between μ and ν such that $\int_{\mathcal{R}} d\pi(x, y) \geq 1 - \delta$ and $\forall (x, y) \in \mathcal{R}, \|x - y\| \leq z$. We note $V = Y - W$ where (W, Y) is drawn from the coupling π . We observe that $\mathcal{R} \subset \{(x, y); \|x - y\| \leq z\}$.

Then, $P(\|V\| > z) \leq P((W, Y) \notin \mathcal{R}) = \int_{\mathbb{R}^d \setminus \mathcal{R}} d\pi(x, y) < \delta$.

For the opposite side, consider the coupling π of the pair (W, Y) such that $W \sim \mu, Y \sim \nu$ and $W + V = Y$ with $P(\|V\| > z) < \delta$.

Then, $P(\|V\| \leq z) = \int_{\|x-y\| \leq z} d\pi(x, y) \geq 1 - \delta$. □

A.3.2 Proof of Lemma 4.2

Lemma 4.2 (Approximate shift reduction). *Let μ, ν, ζ be three distributions on \mathbb{R}^d . We denote $D_\alpha^{(z, \delta)}(\mu, \nu) = \inf_{\mu', \nu' \text{ } (z, \delta)\text{-near}} D_\alpha(\mu', \nu')$. Then, for all $\delta \in (0, 1)$, there exists an event E such that $P(E) \geq 1 - \delta$ and:*

$$D_\alpha((\mu * \zeta)|_E, (\nu * \zeta)) \leq D_\alpha^{(z, \delta)}(\mu, \nu) + R_\alpha(\zeta, z) + \frac{\alpha}{\alpha - 1} \log \left(\frac{1}{1 - \delta} \right).$$

Proof. Let $\alpha > 1$, $z > 0$, $X \sim \mu, Y \sim \nu, N \sim \zeta$ and $W \sim \xi \in \mathcal{P}(\mathbb{R}^d)$ such that $P(\|W\| \geq z) = \delta$ and N is independent of X, Y and W . We use the abuse of notation $D_\alpha(\mu, \nu) = D_\alpha(X, Y)$, with $X \sim \mu, Y \sim \nu$. We consider the event $E = \{\|W\| \leq z\}$. Like in the original proof of the shift reduction lemma of Feldman et al. (2018), we have:

$$D_\alpha((X + N)|_E, Y + N) = D_\alpha((X + W + N - W)|_E, Y + N) \leq D_\alpha((X + W, N - W)|_E, (Y, N)).$$

by post-processing (Proposition 2.1) for $\mathcal{M}_1(x, y) = x + y$. Then, we have:

$$\begin{aligned} & D_\alpha((X + W, N - W)|_E, (Y, N)) \\ &= \frac{1}{\alpha - 1} \log \left(\int \frac{P_{(X+W, N-W)|_E}(x, y)^\alpha}{P_{Y, N}(x, y)^{\alpha-1}} dx dy \right) \\ &= \frac{1}{\alpha - 1} \log \left(\int \frac{P_{X+W|E}(x)^\alpha P_{N-W|E, X+W=x}(y)^\alpha}{\nu(x)^{\alpha-1} \zeta(y)^{\alpha-1}} dx dy \right) \\ &= \frac{1}{\alpha - 1} \log \left(\int \frac{P_{X+W|E}(x)^\alpha}{\nu(x)^{\alpha-1}} \left(\int \frac{P_{N-W|E, X+W=x}(y)^\alpha}{\zeta(y)^{\alpha-1}} dy \right) dx \right) \\ &= \frac{1}{\alpha - 1} \log \int \frac{P_{X+W|E}(x)^\alpha}{\nu(x)^{\alpha-1}} \left(\int \frac{\left(\int_{\|u\| \leq z} P_{N-W|X+W=x, W=u}(y) \xi(u) du \right)^\alpha}{\zeta(y)^{\alpha-1}} dy \right) dx \\ &\leq \frac{1}{\alpha - 1} \log \left(\int \frac{P_{X+W|E}(x)^\alpha}{\nu(x)^{\alpha-1}} \left(\int_{\|u\| \leq z} \frac{\zeta(y+u)^\alpha}{\zeta(y)^{\alpha-1}} \xi(u) du dy \right) dx \right) \\ &\leq \frac{1}{\alpha - 1} \log \left(\int \frac{P_{X+W|E}(x)^\alpha}{\nu(x)^{\alpha-1}} dx \right) + R_\alpha(\zeta, z). \end{aligned}$$

Yet,

$$P_{X+W|E}(x)^\alpha = \left(\frac{P_{X+W}(x) - P(\bar{E})P_{X+W|\bar{E}}(x)}{P(E)} \right)^\alpha \leq \frac{P_{X+W}(x)^\alpha}{(1 - \delta)^\alpha}.$$

Thus:

$$\begin{aligned} & D_\alpha((X + W, N - W)|_E, (Y, N)) \\ &\leq D_\alpha(X + W, Y) + R_\alpha(\zeta, z) - \frac{\alpha}{\alpha - 1} \log(1 - \delta). \end{aligned} \quad \square$$

A.3.3 Proof of Proposition 4.1

Proposition 4.1. *If \mathcal{M} is $(+\infty, \varepsilon, \delta)$ -approximate RPP, then it is (ε, δ) -PP.*

Proof. The proof uses the same approach as the Lemma 8.8 of Bun and Steinke (2016). Let $(s_i, s_j) \in \mathcal{Q}$, $\theta \in \Theta$. Without loss of generality, we assume that there exists E, E' such that $P(E) = 1 - \delta, P(E') = 1 - \delta$ and we have: $D_\infty(P(\mathcal{M}(X) = w | s_i, \theta, E), P(\mathcal{M}(X) = w | s_j, \theta, E')) \leq \varepsilon$. Then,

$$\sup_{w \in \text{Range}(\mathcal{M})} \log \frac{P(\mathcal{M}(X) = w | s_i, \theta, E)}{P(\mathcal{M}(X) = w | s_j, \theta, E')} \leq \varepsilon.$$

$$\begin{aligned} P(\mathcal{M}(X) = w | s_j, \theta) &= P(E')P(\mathcal{M}(X) = w | s_j, \theta, E') + P(\bar{E}')P(\mathcal{M}(X) = w | s_j, \theta, \bar{E}') \\ &\geq (1 - \delta)P(\mathcal{M}(X) = w | s_j, \theta, E'), \\ P(\mathcal{M}(X) = w | s_i, \theta) &= P(E)P(\mathcal{M}(X) = w | s_i, \theta, E) + P(\bar{E})P(\mathcal{M}(X) = w | s_i, \theta, \bar{E}) \\ &\leq (1 - \delta)P(\mathcal{M}(X) = w | s_i, \theta, E) + \delta \\ &\leq (1 - \delta)P(\mathcal{M}(X) = w | s_j, \theta, E') e^\varepsilon + \delta \\ &\leq P(\mathcal{M}(X) = w | s_j, \theta) e^\varepsilon + \delta. \end{aligned} \quad \square$$

A.3.4 Proof of Theorem 4.1

Theorem 4.1 (General approximate Wasserstein mechanism). *Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query. For all $\delta \in (0, 1)$, let us denote:*

$$\Delta_{G,\delta} > \inf\{z \in \mathbb{R}; \forall (s_i, s_j) \in S, \forall \theta \in \Theta, \\ (P((f(X)|_{s_i}, \theta), P(f(X)|_{s_j}, \theta)) \text{ are } (z, \delta)\text{-near}\}.$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the dataset X . Then, $\mathcal{M} = f(X) + N$ satisfies $(\alpha, R_\alpha(\zeta, \Delta_{G,\delta}) + \frac{\alpha}{\alpha-1} \log \frac{1}{1-\delta}, \delta)$ -approximate RPP for all $\alpha \in (1, +\infty)$ and $(R_\infty(\zeta, \Delta_{G,\delta}) + \log \frac{1}{1-\delta}, \delta)$ -PP.

Proof. This proof is similar to Theorem 3.1 but we use the approximate shift reduction lemma (Lemma 4.2). We use the abuse of notation $D_\alpha(X|_E, Y|_E) = D_\alpha(P(X|E), P(Y|E))$. Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query and $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the data X . Let $\delta \in (0, 1)$. Let us denote:

$$\Delta_{G,\delta} > \inf\{z \in \mathbb{R}; \forall (s_i, s_j) \in S, \forall \theta \in \Theta, (P(f(X)|_{s_i}, \theta), P(f(X)|_{s_j}, \theta)) \text{ are } (z, \delta)\text{-near}\}.$$

By the approximate shift reduction lemma (Lemma 4.2), there exists E such that $P(E) \geq 1 - \delta$ and:

$$D_\alpha \left((f(X) + N)|_{E, s_i, \theta}, (f(X) + N)|_{s_j, \theta} \right) \leq D_\alpha^{(z, \delta)} (f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta}) + R_\alpha(\zeta, z) - \frac{\alpha}{\alpha-1} \log(1 - \delta).$$

By definition,

$$D_\alpha^{(z, \delta)} (f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta}) = \inf_{\mu \in \mathcal{P}(\mathbb{R}^d); \mu, P(f(X)|_{s_i, \theta}) \text{ are } (z, \delta)\text{-near}} D_\alpha(\mu, P(f(X)|_{s_j, \theta})),$$

and

$$D_\alpha^{(\Delta_{G,\delta}, \delta)} (f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta}) = 0.$$

Then,

$$D_\alpha \left((f(X) + N)|_{E, s_i, \theta}, (f(X) + N)|_{s_j, \theta} \right) \leq R_\alpha(\zeta, \Delta_{G,\delta}) - \frac{\alpha}{\alpha-1} \log(1 - \delta). \quad \square$$

A.3.5 Result for Usual Noise Distributions

We provide below a corollary of Theorem 4.1 that gives closed formula for usual noise distributions to get approximate RPP guarantees.

Proposition A.2 (Approximate Wasserstein mechanism). *We note I_d the identity matrix of size d . The results are similar to those of the general Wasserstein mechanism (Corollary 3.1), but with an additive term which depends on δ :*

- $\mathcal{M}(X) = X + N$ with $N \sim \mathcal{N}\left(0, \frac{\alpha \Delta_{G,\delta}^2}{2(\varepsilon + \frac{\alpha}{\alpha-1} \log(1-\delta))} I_d\right)$ is $(\alpha, \varepsilon, \delta)$ -approximate RPP.
- $\mathcal{M}(X) = X + L$ with $L \sim \text{Lap}(0, \rho I_d)$ is $(\alpha, \frac{1}{\alpha-1} (\log(b) - \alpha \log(1-\delta)), \delta)$ -approximate RPP for $b = \frac{\alpha}{2\alpha-1} e^{\Delta_{G,\delta}(\alpha-1)/\rho} + \frac{\alpha-1}{2\alpha-1} e^{-\Delta_{G,\delta}\alpha/\rho}$.
- $\mathcal{M}(X) = X + L$ with $L \sim \text{Lap}\left(0, \frac{\Delta_{G,\delta}}{\varepsilon + \log(1-\delta)} I_d\right)$ is (ε, δ) -PP.

A.3.6 Relationship with distribution privacy results of Chen and Ohrimenko (2023)

We start by recalling the definition of distribution privacy.

Definition A.1 (Distribution privacy Chen and Ohrimenko (2023)). A mechanism \mathcal{M} satisfies (ε, δ) -distribution privacy with respect to a set of distribution pairs $\Psi \subset \Theta \times \Theta$ if for all pairs $(\psi_i, \psi_j) \in \Psi$ and all subsets $S \subset \text{Range}(\mathcal{M})$,

$$P(\mathcal{M}(X) \in S | \psi_i) \leq e^\varepsilon P(\mathcal{M}(X) \in S | \psi_j) + \delta,$$

where the expression $P(\mathcal{M}(X) \in S | \psi)$ denotes the probability that $\mathcal{M}(X)$ given $X \sim \psi$.

For completeness, we recall the original approximate Wasserstein mechanism Theorem for distribution privacy from Chen and Ohrimenko (2023).

Theorem A.1 (Approximate Wasserstein mechanism for distribution privacy Chen and Ohrimenko (2023)). *Let (Ψ, Θ) be a distribution privacy framework. Let $W > 0$, $\delta \in (0, 1)$. Suppose that for all $(\psi_i, \psi_j) \in \Psi$, $P(X | \psi_i)$ and $P(X | \psi_j)$ are (W, δ) -near. Then $\mathcal{M}(X) = X + L$ where $L \sim \text{Lap}(0, \frac{W}{\epsilon}I)$ is (ϵ, δ) -distribution private.*

We now formally state and prove the equivalence between Pufferfish privacy and distribution privacy.

Proposition A.3. *Let $(E, \mathcal{B}(E))$ be a measurable space, where $|E| \leq \aleph_1$ is a topological space with its Borel σ -algebra $\mathcal{B}(E)$ and \aleph_1 is the cardinality of \mathbb{R} . Let $\Theta \subset \mathcal{P}(\mathcal{B}(E))$. Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish privacy instance and \mathcal{M} a randomized mechanism. Then, there exists a distribution privacy instance (Ψ, Θ') such that \mathcal{M} is (ϵ, δ) -PP iff \mathcal{M} is (ϵ, δ) -distribution private. Conversely, let (Ψ, Θ) be a distribution privacy instance. Then, there exists a Pufferfish privacy instance $(\mathcal{S}, \mathcal{Q}, \Theta')$ such that \mathcal{M} is (ϵ, δ) -PP iff \mathcal{M} is (ϵ, δ) -distribution private.*

Remark. The condition $|E| \leq \aleph_1$ is quite general. In particular, it allows the data space to be (a subset of) \mathbb{R}^d , thus covering typical data domains found in fields like data analysis, machine learning, text processing, computer vision, and database management.

Proof. We show the equivalence between the Pufferfish privacy framework and the distribution privacy framework. Let $\Theta \subset \mathcal{P}(\mathcal{B}(E))$, where $|E| \leq \aleph_1$.

- Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish privacy instance. We consider:

$$\Psi = \{(P(X|s_i, \theta), P(X|s_j, \theta)) \text{ such that } (s_i, s_j) \in \mathcal{Q}, \theta \in \Theta \text{ and } P(s_i | \theta) \neq 0, P(s_j | \theta) \neq 0\}.$$

Then,

$$\begin{aligned} & \forall w \in \text{Range}(\mathcal{M}), \forall (\psi_i, \psi_j) \in \Psi, \\ & P(\mathcal{M}(X) = w | \psi_i) \leq e^\epsilon P(\mathcal{M}(X) = w | \psi_j) + \delta \\ & \iff \\ & \forall w \in \text{Range}(\mathcal{M}), \forall (s_i, s_j) \in \mathcal{Q}, \theta \in \Theta \text{ such that } P(s_i | \theta) \neq 0, P(s_j | \theta) \neq 0, \\ & P(\mathcal{M}(X) = w | s_i, \theta) \leq e^\epsilon P(\mathcal{M}(X) = w | s_j, \theta) + \delta. \end{aligned}$$

- Let (Ψ, Θ) be a distribution privacy instance. First, we consider the case where each $\psi \in \Theta$ is parametrized by a vector $\rho \in \mathbb{R}^d$, which means that there exists a bijection between a subset of \mathbb{R}^d and Θ . For $\rho \in \mathbb{R}^d$, if it exists, we denote $\psi_\rho \in \Theta$ the corresponding distribution. Then, we denote $\Phi = \{\rho \in \mathbb{R}^d \text{ such that } \exists \psi \in \Psi; (\psi_\rho, \psi) \in \Psi \vee (\psi, \psi_\rho) \in \Psi\}$ and $\Omega = \{(\rho_1, \rho_2) \in \Phi \times \Phi \text{ such that } (\psi_{\rho_1}, \psi_{\rho_2}) \in \Psi\} \subset \mathbb{R}^{n \times 2}$ and $\Pi = \{\pi \in \mathcal{P}(\mathbb{R}^d) \text{ such that } \text{supp}(\pi) = \Phi\}$. We consider:

$$\begin{aligned} \mathcal{S} &= \{(s_\rho = \text{“}X \text{ has been generated from the distribution } \psi_\rho\text{”}), \forall \rho \in \Phi\}, \\ \mathcal{Q} &= \{(s_{\rho_1}, s_{\rho_2}) \text{ such that } (\rho_1, \rho_2) \in \Omega\}, \\ \Theta' &= \left\{ \theta_\pi \in \mathcal{P}(\mathcal{B}(E)) \text{ such that } \pi \in \Pi \wedge P(X|\theta_\pi) = \int_{\Phi} \pi(\rho) P(X|\psi_\rho) d\rho \right\}. \end{aligned}$$

Then, $\forall w \in \text{Range}(\mathcal{M}), \forall (s_i, s_j) \in \mathcal{Q}, \theta_\pi \in \Theta', P(X|\theta_\pi, s_i) = P(X|\psi_i)$. Thus, we have:

$$\begin{aligned} & \forall w \in \text{Range}(\mathcal{M}), \forall (\psi_i, \psi_j) \in \Psi, \\ & P(\mathcal{M}(X) = w | \psi_i) \leq e^\epsilon P(\mathcal{M}(X) = w | \psi_j) + \delta \\ & \iff \\ & \forall w \in \text{Range}(\mathcal{M}), \forall (s_i, s_j) \in \mathcal{Q}, \theta \in \Theta \text{ such that } P(s_i | \theta) \neq 0, P(s_j | \theta) \neq 0, \\ & P(\mathcal{M}(X) = w | s_i, \theta) \leq e^\epsilon P(\mathcal{M}(X) = w | s_j, \theta) + \delta. \end{aligned}$$

In this proof, the case $|\Theta| = n \in \mathbb{N}^*$ is a case where $\psi \in \Theta$ can be parameterized. One such parameterization is to define $\Theta = \{\psi_1, \dots, \psi_n\}$ and the mapping $i \in \mathbb{N} \mapsto \psi_i \in \Theta$.

The second part of the proof relies on the fact that the distributions of $\mathcal{P}(\mathcal{B}(E))$ are parameterizable. The hypothesis $|E| \leq \aleph_1$ allows us to reduce to the case $E = \mathbb{R}$, up to a bijection. Yet, every distribution of

$\mathcal{B}(\mathbb{R})$ is entirely defined by its values taken on open intervals of \mathbb{R} and each open interval of \mathbb{R} is a countable union of open intervals with rational endpoints. Therefore, $|\mathcal{P}(\mathcal{B}(\mathbb{R}))| \leq 2^{\aleph_0} = \aleph_1$, where the notation \aleph_0 denotes the cardinal of \mathbb{N} and we can map every distribution of \mathbb{R} with elements of \mathbb{R} . □

Remark. The proof shows how to transition from the Pufferfish privacy framework to the distribution privacy framework. Thus, it is possible to use Pufferfish private mechanisms to achieve distribution privacy guarantees (and vice versa).

This equivalence result allows us to precisely compare our result (Theorem 4.1) to the result of Chen and Ohrimenko (2023). Our approximate shift reduction result (Lemma 4.2) induces an additional term which prevents us from recovering exactly the results of Chen and Ohrimenko (2023) in the particular case of the Laplace mechanism for PP. However, we believe that our analysis can be improved and lead to better results. More generally, our result can be used with a wide range of noise distributions and in the RPP framework, which is more general than PP (and thus more general than distribution privacy).

A.3.7 Utility of the GAWM (Proposition 4.2)

Below, we make the informal result of Proposition 4.2 precise and provide its proof.

Proposition A.4 (Utility of the GAWM). *Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish framework, $\delta \in (0, 1)$, $\alpha > 1$ and let $\mathcal{M}(X) = f(X) + N$, where $X \sim \theta \in \Theta$, $N \sim \zeta$ and f is a numerical query. Then, Δ_G as defined in Theorem 3.1 is greater or equal than $\Delta_{G,\delta}$ defined in Theorem 4.1. Moreover, if $R_\alpha(\Delta_{G,\delta}, \zeta) \leq R_\alpha(\Delta_G, \zeta) + \frac{\alpha}{\alpha-1} \log(1-\delta)$ then the GAWM achieves better utility than the GWM with $(\alpha, \varepsilon, \delta)$ -RPP, without additional privacy cost on the ε . It happens when Δ_G is sufficiently larger than $\Delta_{G,\delta}$, which happens when there exists $(s_i, s_j) \in \mathcal{Q}, \theta \in \Theta$ and $(Y, Y') \sim \pi \in \Gamma(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta))$ such that $\|Y - Y'\|$ is large with small probability.*

Proof. Let $(s_i, s_j) \in \mathcal{Q}, \theta \in \Theta$. Then, there exists $(Y, Y') \sim \pi \in \Gamma(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta))$ such that $P(\|Y - Y'\| > \Delta_G) = 0$. Then, for any $\delta \in (0, 1)$, $P(\|Y - Y'\| > \Delta_G) < \delta$ and by Lemma 4.1, Y and Y' are (Δ_G, δ) -near. Finally, $\Delta_{G,\delta} \leq \Delta_G$. □

A.4 Leveraging W_p metrics (Section 4.2)

A.4.1 Proof of Lemma 4.3

Lemma 4.3 (Generalized shift reduction for radial decreasing noises). *Let ζ be a radial decreasing noise distribution: $\zeta(z) = \zeta_0(\|z\|)$. Let $z, p, q > 0$ such that $1/p + 1/q = 1$. We note :*

$$D_{\alpha, \alpha', \zeta}^{(z)}(\mu, \nu) = \inf_{\xi; \mathbb{E}_{W \sim \xi} [\exp((\alpha'-1)D_{\alpha'}(\zeta_0 * \|W\|, \zeta_0))] \leq z} D_\alpha(\mu * \xi, \nu).$$

Then, we have :

$$D_\alpha(\mu * \zeta, \nu * \zeta) \leq D_{p(\alpha-1)+1, q(\alpha-1)+1, \zeta}^{(z)}(\mu, \nu) + \frac{\log(z)}{q(\alpha-1)}.$$

In the case $q = 1$:

$$D_\alpha(\mu * \zeta, \nu * \zeta) \leq D_{\infty, \alpha, \zeta}^{(z)}(\mu, \nu) + \frac{\log(z)}{\alpha-1}.$$

Proof. The proof construction is similar to the one developed in Chen and Ohrimenko (2023). We do not apply Jensen inequality at the last step of the proof to obtain Orlicz-Wasserstein metrics, and keep the result general and working for a broader range of distributions. We use the abuse of notation $D_\alpha(\mu, \nu) = D_\alpha(X, Y)$, with $X \sim \mu, Y \sim \nu$. Let $z > 0, X \sim \mu, Y \sim \nu, N \sim \zeta$ be a radial decreasing noise and $W \sim \xi \in \mathcal{P}(\mathbb{R})$ such that:

$$\mathbb{E}_W[\exp(q(\alpha-1)D_{q(\alpha-1)+1}(\zeta_0 * \|W\|, \zeta_0))] \leq z.$$

Let $p, q > 0$ such that $\frac{1}{p} + \frac{1}{q} = 1$. We want to compute : $D_\alpha(X + N, Y + N)$. By the post processing theorem applied on the map $f : (x, y) \rightarrow x + y$, and the fact that $X + N = X + W - W + N$, we have :

$$D_\alpha(X + N, Y + N) \leq D_\alpha((X + W, N - W), (Y, N)).$$

We have:

$$\begin{aligned}
 D_\alpha((X+W, N-W), (Y, N)) &= \frac{1}{\alpha-1} \log \left(\int \frac{P_{(X+W, N-W)}(x, y)^\alpha}{P_{Y, N}(x, y)^{\alpha-1}} dx dy \right) \\
 &= \frac{1}{\alpha-1} \log \left(\int \frac{P_{X+W}(x)^\alpha P_{N-W|X+W=x}(y)^{\alpha-1}}{\nu(x)^{\alpha-1} \zeta(y)^\alpha} dx dy \right) \\
 &= \frac{1}{\alpha-1} \log \mathbb{E}_{\substack{U \sim X+W \\ V \sim N-W|X+W=U}} \left[\left(\frac{P_{X+W}(U)}{\nu(U)} \right)^{\alpha-1} \left(\frac{P_{N-W|X+W=x}(V)}{\zeta(V)} \right)^{\alpha-1} \right] \\
 &\leq \frac{1}{p(\alpha-1)} \log \mathbb{E}_{U \sim X+W} \left[\left(\frac{P_{X+W}(U)}{\nu(U)} \right)^{p(\alpha-1)} \right] \quad (1) \\
 &\quad + \frac{1}{q(\alpha-1)} \log \mathbb{E}_{\substack{U \sim X+W \\ V \sim N-W|X+W=U}} \left[\left(\frac{P_{N-W|X+W=x}(V)}{\zeta(V)} \right)^{q(\alpha-1)} \right] \quad (2) \text{ by Hölder inequality}
 \end{aligned}$$

Immediately (1) = $D_{p(\alpha-1)+1}(X+W, Y)$ and, given that

$$\begin{aligned}
 P_{N-W|X+W=x}(y)^{q(\alpha-1)+1} &= \left(\int P_{N-W|W=z}(y) \xi(z) dz \right)^{q(\alpha-1)+1} \\
 &= \mathbb{E}_W [\zeta(y+W)]^{q(\alpha-1)+1} \\
 &\leq \mathbb{E}_W [\zeta(y+W)^{q(\alpha-1)+1}],
 \end{aligned}$$

we have:

$$\begin{aligned}
 (2) &= \frac{1}{q(\alpha-1)} \log \int \left(\frac{P_{N-W|X+W=x}(y)}{\zeta(y)} \right)^{q(\alpha-1)} P_{X+W}(x) P_{N-W|X+W=x}(y) dx dy \\
 &\leq \frac{1}{q(\alpha-1)} \log \int \frac{\zeta(y+u)^{q(\alpha-1)+1}}{\zeta(y)^{q(\alpha-1)}} \xi(u) P_{X+W}(x) dx dy du \\
 &\leq \frac{1}{q(\alpha-1)} \log \mathbb{E}_{W \sim \xi} [\exp(q(\alpha-1) D_{q(\alpha-1)+1}(\zeta_0 * \|W\|, \zeta_0))] \\
 &\leq \frac{\log(z)}{q(\alpha-1)}.
 \end{aligned}$$

In the case $p = +\infty$, let $W \sim \xi \in \mathcal{P}(\mathbb{R})$ such that:

$$\mathbb{E}_W [\exp((\alpha-1) D_\alpha(\zeta_0 * \|W\|, \zeta_0))] \leq z.$$

$$\begin{aligned}
 D_\alpha((X+W, N-W), (Y, N)) &\leq \sup_{U \sim X+W} \frac{1}{\alpha-1} \log \left(\frac{P_{X+W}(U)}{\nu(U)} \right)^{\alpha-1} \quad (3) \\
 &\quad + \frac{1}{(\alpha-1)} \log \mathbb{E}_{\substack{U \sim X+W \\ V \sim N-W|X+W=U}} \left[\left(\frac{P_{N-W|X+W=x}(V)}{\zeta(V)} \right)^{\alpha-1} \right] \quad (4)
 \end{aligned}$$

Yet, (3) = $D_\infty(P_{X+W}, \nu)$ and:

$$\begin{aligned}
 (4) &= \frac{1}{\alpha-1} \log \int \left(\frac{P_{N-W|X+W=x}(y)}{\zeta(y)} \right)^{\alpha-1} P_{X+W}(x) P_{N-W|X+W=x}(y) dx dy \\
 &\leq \frac{1}{\alpha-1} \log \mathbb{E}_{W \sim \xi} [\exp((\alpha-1) D_\alpha(\zeta_0 * \|W\|, \zeta_0))] \\
 &\leq \frac{\log(z)}{\alpha-1}.
 \end{aligned}$$

□

A.4.2 Proof of Theorem 4.2

Theorem 4.2 (Distribution Aware General Wasserstein Mechanism). *Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query and ζ a radial decreasing noise distribution. Let $q \geq 1$. For $(s_i, s_j) \in \mathcal{Q}, \theta \in \Theta$, we note $\mu_i^\theta = P(f(X)|s_i, \theta)$. We denote:*

$$\Delta_G^{\zeta, q, \alpha} = \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \mathbb{E} \left[e^{q(\alpha-1)D_{q(\alpha-1)+1}(\zeta_0 * \|X-Y\|, \zeta_0)} \right].$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the data X . Then, $\mathcal{M}(X) = f(X) + N$ satisfies $(\alpha, \frac{\log(\Delta_G^{\zeta, q, \alpha})}{q(\alpha-1)})$ -RPP for all $\alpha \in (1, +\infty)$ and $\lim_{\alpha \rightarrow +\infty} \frac{\log(\Delta_G^{\zeta, q, \alpha})}{q(\alpha-1)}$ -PP.

Proof. The proof is similar to Theorem 3.1 but we use the generalized shift reduction lemma (Lemma 4.3). Let (S, \mathcal{Q}, Θ) be a Pufferfish privacy instance. Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query and denote:

$$\Delta_G^{\zeta, q, \alpha} = \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \mathbb{E} \left[e^{q(\alpha-1)D_{q(\alpha-1)+1}(\zeta_0 * \|X-Y\|, \zeta_0)} \right].$$

Let $N = (N_1, \dots, N_d) \sim \zeta$, where N_1, \dots, N_d are iid real random variables independent of the data X . Let $\alpha > 1, z > 0, (s_i, s_j) \in \mathcal{Q}$ and $\theta \in \Theta$. We use the abuse of notation $D_\alpha(X|_E, Y|_E) = D_\alpha(P(X|E), P(Y|E))$. By the shift reduction lemma (Lemma 4.3), we have:

$$D_\alpha \left((f(X) + N)|_{s_i, \theta}, (f(X) + N)|_{s_j, \theta} \right) \leq D_{p(\alpha-1)+1, q(\alpha-1)+1, \zeta}^{(z)} \left(f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta} \right) + \frac{\log(z)}{q(\alpha-1)}.$$

By definition,

$$D_{p(\alpha-1)+1, q(\alpha-1)+1, \zeta}^{(z)} \left(f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta} \right) = \inf_{W \in \mathcal{P}(\mathbb{R}^d); \mathbb{E}_{W \sim \xi} [\exp(q(\alpha-1)D_{q(\alpha-1)+1}(\zeta_0 * \|W - f(X)|_{s_i, \theta}\|, \zeta_0))] \leq z} D_\alpha(W, f(X)|_{s_j, \theta}),$$

and

$$D_{p(\alpha-1)+1, q(\alpha-1)+1, \zeta}^{(\exp(q(\alpha-1)D_{q(\alpha-1)+1}(\zeta_0 * \|f(X)|_{s_i, \theta} - f(X)|_{s_j, \theta}\|, \zeta_0))} \left(f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta} \right) = 0.$$

Then,

$$D_\alpha \left((f(X) + N)|_{s_i, \theta}, (f(X) + N)|_{s_j, \theta} \right) \leq \frac{\log(\Delta_G^{\zeta, q, \alpha})}{q(\alpha-1)}. \quad \square$$

A.4.3 Proof of Corollary 4.1

Divergences of shifts in Cauchy distributions have been discussed in Verdú (2023). We generalize their results for certain types of generalized Cauchy distributions in the following lemma.

Lemma A.1 (Shifts of generalized Cauchy distributions). *Let $k \in \mathbb{N}^*, \alpha > 1, \lambda > 0$ and $\beta_{k, \lambda} > 0$ such that $\zeta_{k, \lambda} : x \mapsto \beta_{k, \lambda} \left(\frac{1}{1+(\lambda x)^2} \right)^{\frac{k}{2}}$ verifies $\int \zeta_{k, \lambda}(x) dx = 1$. Let $X \sim \zeta_{k, \lambda}$ and $z \geq 0$. Then,*

$$D_\alpha(X + z, X) \leq \frac{1}{\alpha-1} \log \frac{\beta_{k, \lambda} \pi}{\lambda} Q_{k(\alpha-1)/2} \left(1 + \frac{z^2}{\lambda^2} \right),$$

where $Q_{k(\alpha-1)/2}$ is the Legendre function of the first kind of index $k(\alpha-1)/2$.

Proof. Let $\lambda, z > 0, k \in \mathbb{N}^*$. We have:

$$\begin{aligned}
 \int \frac{\zeta_{k,\lambda}(x+z)^\alpha}{\zeta_{k,\lambda}(x)^{\alpha-1}} dx &= \beta_{k,\lambda} \int \frac{(1 + (\lambda(x-z))^2)^{(\alpha-1)k/2}}{(1 + (\lambda x)^2)^{\alpha k/2}} dx \\
 &= \frac{\beta_{k,\lambda}}{\lambda} \int \frac{(1 + (u - \lambda z)^2)^{(\alpha-1)k/2}}{(1 + u^2)^{\alpha k/2}} du \\
 &= \frac{\beta_{k,\lambda}}{\lambda} \int_{-\pi/2}^{\pi/2} \frac{(1 + (\tan(t) - \lambda z)^2)^{(\alpha-1)k/2}}{(1 + \tan^2(t))^{\alpha k/2}} (1 + \tan^2(t)) dt \\
 &= \frac{\beta_{k,\lambda}}{\lambda} \int_{-\pi/2}^{\pi/2} (1 + \tan^2(t) - 2 \tan(t)\lambda z + \lambda^2 z^2)^{(\alpha-1)k/2} (\cos^2(t))^{\alpha k/2 - 1} dt \\
 &= \frac{\beta_{k,\lambda}}{\lambda} \int_{-\pi/2}^{\pi/2} (\cos^2(t)(1 + \tan^2(t) - 2 \tan(t)\lambda z + \lambda^2 z^2))^{(\alpha-1)k/2} (\cos^2(t))^{k/2 - 1} dt \\
 &\leq \frac{\beta_{k,\lambda}}{\lambda} \int_{-\pi/2}^{\pi/2} (1 - 2 \sin(t) \cos(t)\lambda z + \cos^2(t)\lambda^2 z^2)^{(\alpha-1)k/2} dt \\
 &\leq \frac{\beta_{k,\lambda}}{2\lambda} \int_{-\pi}^{\pi} (1 - \sin(t)\lambda z + (\cos(t) + 1)\lambda^2 z^2/2)^{(\alpha-1)k/2} dt \\
 &\leq \frac{\beta_{k,\lambda}}{2\lambda} \int_{-\pi}^{\pi} (1 + \lambda^2 z^2/2 - \sin(t)\lambda z + \cos(t)\lambda^2 z^2/2)^{(\alpha-1)k/2} dt \\
 &\leq \frac{\beta_{k,\lambda}}{2\lambda} \int_{-\pi}^{\pi} \left(1 + \lambda^2 z^2/2 + \sqrt{\lambda z + \lambda^2 z^2/2} \cos(t)\right)^{(\alpha-1)k/2} dt,
 \end{aligned}$$

And $Q_\alpha(z)$ is defined by:

$$Q_\alpha(z) = \frac{1}{\pi} \int_0^\pi \left(z + \sqrt{z^2 - 1} \cos(t)\right)^\alpha dt. \quad \square$$

We are now ready to prove Corollary 4.1.

Corollary 4.1 (Cauchy Mechanism). *We denote Q_α the Legendre polynomial of integer index $\alpha > 1$. Let $k \geq 2$ and $q \geq 1$ such that $kq(\alpha - 1)/2$ is an integer. We note:*

$$\Delta_G^{kq(\alpha-1)} = \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} W_{kq(\alpha-1)}(P(f(X)|s_i, \theta), P(f(X)|s_j, \theta)).$$

Then, $\mathcal{M}(X) = f(X) + V$ with $V \sim \text{GCauchy}(0, \lambda, k)$ is $\left(\alpha, \frac{\log \frac{\beta_{k,\lambda}\pi}{\lambda} Q_{kq(\alpha-1)/2} \left(1 + \left(\frac{\Delta_G^{kq(\alpha-1)}}{\lambda}\right)^2\right)}{q(\alpha-1)}\right)$ -RPP.

Proof. By Lemma A.1, we have:

$$\begin{aligned}
 \Delta_G^{\zeta, q, \alpha} &= \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \mathbb{E} \left[e^{q(\alpha-1)D_{q(\alpha-1)+1}(\zeta_0 * \|X-Y\|, \zeta_0)} \right] \\
 &\leq \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \mathbb{E} \left[\frac{\beta_{k, \lambda} \pi}{\lambda} Q_{kq(\alpha-1)/2} \left(1 + \frac{\|X-Y\|^2}{\lambda^2} \right) \right] \\
 &\leq \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \sum_{i=0}^{kq(\alpha-1)/2} \frac{\beta_{k, \lambda} \pi a_i}{\lambda} \mathbb{E} \left[\left(1 + \frac{\|X-Y\|^2}{\lambda^2} \right)^i \right] \quad P_{kq(\alpha-1)/2} \text{ is a polynomial} \\
 &\leq \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \sum_{i=0}^{kq(\alpha-1)/2} \sum_{l=0}^i \binom{i}{l} \frac{\beta_{k, \lambda} \pi a_i}{\lambda} \mathbb{E} \left[\frac{\|X-Y\|^{2i}}{\lambda^{2i}} \right] \\
 &\leq \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \sum_{i=0}^{kq(\alpha-1)/2} \sum_{l=0}^i \binom{i}{l} \frac{\beta_{k, \lambda} \pi a_i}{\lambda} \frac{\mathbb{E} [\|X-Y\|^{kq(\alpha-1)}]^{2i/kq(\alpha-1)}}{\lambda^{2i}} \quad \text{Jensen inequality } (2i \leq kq(\alpha-1)) \\
 &\leq \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \sum_{i=0}^{kq(\alpha-1)/2} \sum_{l=0}^i \binom{i}{l} \frac{\beta_{k, \lambda} \pi a_i}{\lambda} \frac{W_{kq(\alpha-1)}(\mu_i^\theta, \mu_j^\theta)^{2i}}{\lambda^{2i}} \quad \text{by definition of } W_{kq(\alpha-1)} \\
 &\leq \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \frac{\beta_{k, \lambda} \pi}{\lambda} Q_{kq(\alpha-1)/2} \left(1 + \frac{W_{kq(\alpha-1)}(\mu_i^\theta, \mu_j^\theta)^2}{\lambda^2} \right). \quad \square
 \end{aligned}$$

A.4.4 Utility of the DAGWM (Proposition 4.3)

Below, we make the informal result of Proposition 4.3 precise and provide its proof.

Proposition A.5 (Utility of the DAGWM). *Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish framework, and let $\mathcal{M}(X) = f(X) + N$, where $X \sim \theta \in \Theta$, $N \sim \zeta$ a radial decreasing distribution and f is a numerical query. Let $\alpha > 1$. Then, $R_\alpha(\zeta, \Delta_G)$ as defined in Theorem 3.1 is greater or equal to $\frac{\log(\Delta_G^{\zeta, 1, \alpha})}{\alpha-1}$ defined in Theorem 4.2.*

Proof. By definition: for $(s_i, s_j) \in \mathcal{Q}, \theta \in \Theta$, we note $(\mu_i^\theta, \mu_j^\theta) = (P(f(X)|_{s_i}, \theta), P(f(X)|_{s_j}, \theta))$, and if $(f(X)|_{s_i, \theta}, f(X)|_{s_j, \theta}) \sim \pi^* \in \Gamma(\mu_i^\theta, \mu_j^\theta)$ realises the optimal transport plan for $W_\infty(\mu_i^\theta, \mu_j^\theta)$:

$$\|f(X)|_{s_i, \theta} - f(X)|_{s_j, \theta}\| \leq W_\infty(\mu_i^\theta, \mu_j^\theta) \text{ a.s.}$$

Given that ζ is radial decreasing, $\zeta(z) = \zeta_0(\|z\|)$ and $x \in \mathbb{R} \rightarrow D_\alpha(\zeta_0 * x, \zeta_0)$ is increasing:

$$e^{(\alpha-1)D_{(\alpha-1)}(\zeta_0 * \|f(X)|_{s_i, \theta} - f(X)|_{s_j, \theta}\|, \zeta_0)} \leq e^{(\alpha-1)D_\alpha(\zeta_0 * W_\infty(\mu_i^\theta, \mu_j^\theta), \zeta_0)} \text{ a.s.}$$

Then,

$$\mathbb{E} \left[e^{(\alpha-1)D_{(\alpha-1)}(\zeta_0 * \|f(X)|_{s_i, \theta} - f(X)|_{s_j, \theta}\|, \zeta_0)} \right] \leq \mathbb{E} \left[e^{(\alpha-1)D_\alpha(\zeta_0 * W_\infty(\mu_i^\theta, \mu_j^\theta), \zeta_0)} \right].$$

It follows:

$$\begin{aligned}
 \Delta_G^{\zeta, 1, \alpha} &= \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} \inf_{P(X, Y) \in \Gamma(\mu_i^\theta, \mu_j^\theta)} \mathbb{E} \left[e^{(\alpha-1)D_\alpha(\zeta_0 * \|X-Y\|, \zeta_0)} \right] \\
 &\leq \mathbb{E} \left[e^{(\alpha-1)D_{(\alpha-1)}(\zeta_0 * \|f(X)|_{s_i, \theta} - f(X)|_{s_j, \theta}\|, \zeta_0)} \right] \\
 &\leq e^{(\alpha-1)D_\alpha(\zeta_0 * W_\infty(\mu_i^\theta, \mu_j^\theta), \zeta_0)}.
 \end{aligned}$$

Finally :

$$\frac{\log(\Delta_G^{\zeta, 1, \alpha})}{\alpha-1} \leq \max_{\substack{(s_i, s_j) \in S \\ \theta_i \in \Theta}} D_\alpha(\zeta_0 * W_\infty(\mu_i^\theta, \mu_j^\theta), \zeta_0) = R_\alpha(\zeta, \Delta_G). \quad \square$$

A.5 Guarantees Against Close Adversaries (Section 5)

A.5.1 Original result from Song et al. (2017)

For completeness, we recall here the original theorem from Song et al. (2017) on the robustness of the Pufferfish privacy framework.

Theorem A.2 (Protection against close adversaries Song et al. (2017)). *Let \mathcal{M} be a mechanism that satisfies ε -PP in a framework $(\mathcal{S}, \mathcal{Q}, \Theta)$. Let $\theta' \notin \Theta$ and*

$$\Delta = \inf_{\theta \in \Theta} \sup_{s_i \in \mathcal{Q}} \max\{D_\infty(P(X|s_i, \theta), P(X|s_i, \theta')), D_\infty(P(X|s_i, \theta), P(X|s_i, \theta'))\}.$$

Then, \mathcal{M} is $(\varepsilon + 2\Delta)$ -PP for the framework $(\mathcal{S}, \mathcal{Q}, \Theta')$ with $\Theta' = \Theta \cup \{\theta'\}$.

A.5.2 Proof of Theorem 5.1

Theorem 5.1 (RPP protection against close adversaries). *Let $p, q, r > 0$ such that $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$, and let \mathcal{M} be a mechanism that satisfies \mathcal{M} is $(q(\alpha - 1/p), \varepsilon)$ -RPP in a framework $(\mathcal{S}, \mathcal{Q}, \Theta)$. Let $\theta' \notin \Theta$ and*

$$\begin{aligned} \Delta_p^1 &= \inf_{\theta \in \Theta} \sup_{s_i \in \mathcal{S}} D_{\alpha p}(P(X|s_i, \theta'), P(X|s_i, \theta)), \\ \Delta_r^2 &= \inf_{\theta \in \Theta} \sup_{s_i \in \mathcal{S}} D_{(\alpha-1)r+1}(P(X|s_i, \theta), P(X|s_i, \theta')). \end{aligned}$$

Then, for all $\alpha \in (1, \infty)$, \mathcal{M} satisfies:

$$\left(\alpha, \left(1 + \frac{1}{r(\alpha-1)}\right)\varepsilon + \left(1 + \frac{\frac{1}{r} + \frac{1}{q}}{\alpha-1}\right)\Delta_p^1 + \Delta_r^2 \right)\text{-RPP}$$

for $(\mathcal{S}, \mathcal{Q}, \Theta')$ with $\Theta' = \Theta \cup \{\theta'\}$.

Proof. Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish privacy instance and \mathcal{M} a randomized mechanism. Let $\theta' \notin \Theta$ and $p, q, r > 0$ such that $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$. Let $s_i, s_j \in \mathcal{Q}$. We have:

$$\begin{aligned} & \exp(\alpha-1)D_\alpha(P(\mathcal{M}(X)|s_i, \theta'), P(\mathcal{M}(X)|s_j, \theta')) \\ &= \int \frac{P(\mathcal{M}(X) = z|s_i, \theta')^\alpha}{P(\mathcal{M}(X) = z|s_j, \theta')^{\alpha-1}} dz \\ &= \int \frac{P(\mathcal{M}(X) = z|s_i, \theta')^\alpha}{P(\mathcal{M}(X) = z|s_i, \theta)^{\alpha-1/p}} \frac{P(\mathcal{M}(X) = z|s_i, \theta)^{\alpha-1/p}}{P(\mathcal{M}(X) = z|s_i, \theta)^{\alpha-1/p-1/q}} \frac{P(\mathcal{M}(X) = z|s_j, \theta)^{\alpha-1/p-1/q}}{P(\mathcal{M}(X) = z|s_j, \theta')^{\alpha-1}} dz \\ &\leq \left(\int \frac{P(\mathcal{M}(X) = z|s_i, \theta')^{\alpha p}}{P(\mathcal{M}(X) = z|s_i, \theta)^{\alpha p-1}} dz \right)^{\frac{1}{p}} \cdot \left(\int \frac{P(\mathcal{M}(X) = z|s_i, \theta)^{q(\alpha-1/p)}}{P(\mathcal{M}(X) = z|s_i, \theta)^{q(\alpha-1/p)-1}} dz \right)^{\frac{1}{q}} \\ &\quad \cdot \left(\int \frac{P(\mathcal{M}(X) = z|s_j, \theta)^{\alpha-1/p-1/q}}{P(\mathcal{M}(X) = z|s_j, \theta')^{\alpha-1}} dz \right)^{\frac{1}{r}} \\ &\leq \exp(\alpha-1/p)D_{\alpha p}(P(\mathcal{M}(X)|s_i, \theta'), P(\mathcal{M}(X)|s_i, \theta)) \\ &\quad + \exp((\alpha-1+1/r)D_{q(\alpha-1/p)}(P(\mathcal{M}(X)|s_i, \theta), P(\mathcal{M}(X)|s_j, \theta))) \\ &\quad + \exp((\alpha-1)D_{(\alpha-1)r+1}(P(\mathcal{M}(X)|s_j, \theta), P(\mathcal{M}(X)|s_j, \theta'))) \end{aligned}$$

by using the generalized Hölder inequality: for $p, q, r, t > 0$ such that $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = \frac{1}{t}$ and $f \in L^p, g \in L^q, h \in L^r$,

$$\|fgh\|_t \leq \|f\|_p \|g\|_q \|h\|_r.$$

Then, the post-processing property of RPP (Proposition 2.1) gives the result. \square

A.5.3 Refinement of Theorem 5.1 for Additive Mechanisms

Leveraging the shift reduction lemma (Lemma 3.1), we refine Theorem 5.1 for additive mechanisms.

Theorem A.3 (RPP protection against close adversaries for additive noise mechanisms). *Let $p, q, r > 0$ such that $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$. Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query. Let $\mathcal{M}(X) = f(X) + N$ with $N \sim \zeta$ be an additive noise mechanism that satisfies $(q(\alpha - 1/p), \varepsilon)$ -RPP for $(\mathcal{S}, \mathcal{Q}, \Theta)$. Let $\theta' \notin \Theta$ and*

$$\Delta_{\theta'} = \inf_{\theta \in \Theta} \sup_{s_i \in \mathcal{S}} W_{\infty}(P(f(X)|s_i, \theta'), P(f(X)|s_i, \theta)).$$

Then, for all $\alpha \in (1, \infty)$ and denoting

$$K = \left(1 + \frac{\frac{1}{r} + \frac{1}{q}}{\alpha - 1}\right) R_{\alpha p}(\zeta, \Delta_{\theta'}) + R_{(\alpha-1)r+1}(\zeta, \Delta_{\theta'}),$$

\mathcal{M} satisfies:

$$\left(\alpha, \left(1 + \frac{1}{r(\alpha - 1)}\right) \varepsilon + K\right)\text{-RPP}$$

for $(\mathcal{S}, \mathcal{Q}, \Theta')$ with $\Theta' = \Theta \cup \{\theta'\}$.

This theorem enables us to take into account the characteristics of the mechanism when examining the robustness of a RPP instance. We illustrate this below with the Gaussian mechanism.

Corollary A.1 (RPP protection against close adversaries for the Gaussian mechanism). *We note I_d the identity matrix of size d . Let $p, q, r > 0$ such that $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$. Let $f : \mathcal{D} \rightarrow \mathbb{R}^d$ be a numerical query. Let $\mathcal{M}(X) = f(X) + N$ with $N \sim \mathcal{N}(0, \frac{q(\alpha-1/p)\Delta_G^2}{2\varepsilon} I_d)$, where Δ_G is defined in Theorem 3.1. Let $\theta' \notin \Theta$.*

Then, for all $\alpha \in (1, \infty)$, \mathcal{M} satisfies:

$$\left(\alpha, \left(\left(1 + \frac{1}{r(\alpha - 1)}\right) + \left(\alpha \left(p + \frac{p-1}{\alpha-1}\right) + (\alpha-1)r + 1\right) \frac{\Delta_{\theta'}^2}{\Delta_G^2} \frac{1}{q(\alpha-1/p)}\right) \varepsilon\right)\text{-RPP}$$

for $(\mathcal{S}, \mathcal{Q}, \Theta')$ with $\Theta' = \Theta \cup \{\theta'\}$.

One can see that the additive penalty vanishes proportionally to $\frac{1}{\Delta_G^2}$. It establishes a trade-off between the utility of the mechanism and the robustness of the Pufferfish privacy framework when designing Θ . Remarkably, this consideration could not have been derived from our Theorem 5.1 for RPP nor from the original result from Song et al. (2017) (Theorem A.2).

A.6 Privacy Amplification by Iteration (Section 6)

A.6.1 Parallel Composition

Assessing the privacy guarantees of composition in RPP may be challenging. As a matter of fact, there does not exist, to our knowledge, any theorem stating the mechanism-agnostic privacy guarantees of sequential composition in Pufferfish privacy. However, we can recover a straightforward result of parallel composition for the RPP framework.

Proposition A.6 (RPP parallel composition for queries performed over independent datasets). *Let $m > 0$ and $(\mathcal{S}, \mathcal{Q}, \Theta_k)$ be Pufferfish frameworks corresponding to each dataset $X_k \sim P(\cdot | s_i^k, \theta_k)$. We assume that each secret s_i^k is independent of the distributions θ_l , for $l \neq k$ and that \mathcal{Q} only contains pairs of the form (s_i^k, s_j^k) . For all $k \in \{1, \dots, n\}$, let $\mathcal{M}_k(X_k)$ be mechanisms that satisfy (α, ε_k) -RPP. Let $\Theta = \{\otimes_{k=1}^m \theta_k; \forall k \in \{1, \dots, m\}, \theta_k \in \Theta_k\}$. Then, the mechanism $(\mathcal{M}_1, \dots, \mathcal{M}_m)$ satisfies $(\alpha, \max_k \varepsilon_k)$ -RPP for $(\mathcal{S}, \mathcal{Q}, \Theta)$.*

Proof. Let $s_i^l, s_j^l \in \mathcal{Q}, \theta = \otimes_{k=1}^m \theta_k \in \Theta$.

$$\begin{aligned} D_{\alpha}(P(\mathcal{M}(X)|s_i^l, \theta), P(\mathcal{M}(X)|s_j^l, \theta)) &= D_{\alpha}(P((\mathcal{M}_1(X_1), \dots, \mathcal{M}_n(X_n))|s_i^l, \otimes_{k=1}^m \theta_k), P((\mathcal{M}_1(X_1), \dots, \mathcal{M}_n(X_n))|s_j^l, \otimes_{k=1}^m \theta_k)) \\ &= \sum_{k=1}^n D_{\alpha}(P(\mathcal{M}_k(X_k)|s_i^l, \theta_l), P(\mathcal{M}_k(X_k)|s_j^l, \theta_k)) \\ &= D_{\alpha}(P(\mathcal{M}_l(X_l)|s_i^l, \theta_l), P(\mathcal{M}_l(X_l)|s_j^l, \theta_l)) \leq \varepsilon_l. \end{aligned}$$

□

This theorem states that if an adversary assumes that the dataset can be split into independent parts and if the secrets have some form of separability, such as in our Example 2, it is possible to apply a different RPP mechanisms to each independent part while paying only for the maximum privacy loss, similar to the parallel composition result for differential privacy.

A.6.2 Proof of Lemma 6.1

Lemma 6.1 (Dataset Dependent Contraction lemma). *Let ψ be a contractive map in its first argument on $(\mathcal{Z}, \|\cdot\|)$. Let X, X' be two r.v.'s. Suppose that $\sup_w W_\infty(\psi(w, X), \psi(w, X')) \leq s$. Then, for $z > 0$:*

$$D_\alpha^{(z+s)}(\psi(W, X), \psi(W', X')) \leq D_\alpha^{(z)}(W, W').$$

Proof. This proof is similar to the contraction lemma of Feldman et al. (2018). Let $s > 0$ such that $\sup_w W_\infty(\psi(W, X), \psi(W, X')) \leq s$, we have, for Y a v.a. such that $D_\alpha^{(z)}(W, W') = D_\alpha(Y, W')$ and $W_\infty(W, Y) \leq z$:

$$\begin{aligned} W_\infty(\psi(W, X), \psi(W', X')) &\leq W_\infty(\psi(W, X), \psi(W, X')) + W_\infty(\psi(W, X'), \psi(W', X')) \\ &\leq s + W_\infty(W, W') \\ &\leq s + z. \end{aligned}$$

It follows that:

$$D_\alpha^{(z+s)}(\psi(W, X), \psi(W', X')) \leq D_\alpha(\psi(Y, X'), \psi(W', X')) \leq D_\alpha(Y, W') = D_\alpha^{(z)}(W, W'). \quad \square$$

A.6.3 Proof of Theorem 6.1

Theorem 6.1 (Dataset Dependent PABI). *Let X_T and X'_T denote the output of $CNI_T(W_0, \{\psi_t\}, \{\zeta_t\}, X)$ and $CNI_T(W_0, \{\psi_t\}, \{\zeta_t\}, X')$. Let $s_t = \sup_w W_\infty(\psi(w, X), \psi(w, X'))$. Let a_1, \dots, a_T be a sequence of reals and let $z_t = \sum_{i \leq t} s_i - \sum_{i \leq t} a_i$. If $z_t \geq 0$ for all t , then, we have:*

$$D_\alpha^{(z_T)}(X_T, X'_T) \leq \sum_{t=1}^T R_\alpha(\zeta_t, a_t).$$

Proof. The proof is similar to the original PABI proof of Feldman et al. (2018). It is obtained by induction by replacing in the original PABI proof $s_t = \sup_{w \in \mathbb{R}^d, x, x' \in \mathcal{X}} \|\psi(w, x) - \psi(w, x')\|$ by $s_t = \sup_w W_\infty(\psi(w, X), \psi(w, X'))$ and using the dataset dependent contraction lemma (Lemma 6.1). \square

A.6.4 Application to DP

Lemma A.2 (Example: DP as a special case). *In the case of DP, each distribution $\theta \in \Theta$ corresponds to a prior of independence between the elements of the dataset. Let $\beta, \eta, \sigma, L, T > 0, \alpha > 1$ such that $\eta > 2/\beta$. We set the secrets $\mathcal{S} = \left\{ s_i^a \stackrel{\text{def}}{=} \{X_i = a\}; a \in \mathcal{X} \right\}$ and the pairs of secrets : $\mathcal{Q} = \{(s_i^a, s_i^b); a, b \in \mathcal{X}\}$. Let $(X, X') \sim \pi \in \Gamma(P(X|s_i^a), P(X|s_i^b))$. Let f be an objective function which is convex, β -smooth and L -Lipschitz. Let $\mathcal{K} \subset \mathbb{R}^d$ be a compact set. Let $W_0 = W'_0 \in \mathcal{K}$ be the original weight of the stochastic gradient descent and ψ the update function of the projected noisy stochastic gradient descent of learning rate η . Let $\zeta = \mathcal{N}(0, \sigma^2 \eta^2 I_d)$ be the noising distribution. For $t \in \llbracket 0, T \rrbracket$, we define $W_t = CNI_t(W_0, \psi, \zeta, X)$, $W'_t = CNI_t(W'_0, \psi, \zeta, X')$. Then, Theorem 6.1 allows to obtain:*

$$D_\alpha^{(z_T)}(X_T, X'_T) \leq \frac{2\alpha L^2}{\sigma^2(T-i+1)}.$$

This recovers the results of Feldman et al. (2018) for the case of DP-SGD.

Proof. Let $\sigma > 0$. Let $(s_i^a, s_i^b) \in \mathcal{Q}, \theta \in \Theta$, with θ representing a prior of independence. Then, for $t \in \llbracket 1, T \rrbracket$, $(X, X') \sim \pi \in \Gamma(P(X|s_i^a), P(X|s_i^b))$, $s_t = \sup_w W_\infty(\psi(w, X_t), \psi(w, X'_t)) =$

$\begin{cases} \sup_w \|\psi(w, a) - \psi(w, b)\| & \text{if } t = i \\ 0 & \text{else} \end{cases}$, $\zeta_t = \mathcal{N}(0, (\eta\sigma)^2 I_d)$. Then, setting $a_t = \begin{cases} \frac{s_i}{T-i+1} & \text{if } t \geq i \\ 0 & \text{else} \end{cases}$, we get:

$$\begin{aligned} D_\alpha^{(z^T)}(X_T, X'_T) &\leq \sum_{t=i}^T R_\alpha \left(\zeta_t, \frac{\sup_w \|\psi(w, a) - \psi(w, b)\|}{T-i+1} \right) \\ &\leq \sum_{t=i}^T \frac{\alpha \sup_w \|\psi(w, a) - \psi(w, b)\|}{2\eta^2 \sigma^2 (T-i+1)^2} \\ &\leq \frac{2\alpha L^2}{\sigma^2 (T-i+1)}, \end{aligned}$$

which is the bound of Theorem 23 of Feldman et al. (2018). □