



HAL
open science

On the Statistical Complexity of Estimation and Testing under Privacy Constraints

Clément Lalanne, Aurélien Garivier, Rémi Gribonval

► **To cite this version:**

Clément Lalanne, Aurélien Garivier, Rémi Gribonval. On the Statistical Complexity of Estimation and Testing under Privacy Constraints. Transactions on Machine Learning Research Journal, 2023. hal-03794374v4

HAL Id: hal-03794374

<https://hal.science/hal-03794374v4>

Submitted on 12 Sep 2024

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.

On the Statistical Complexity of Estimation and Testing under Privacy Constraints

Clément Lalanne

Univ. Lyon, ENS Lyon, UCBL, CNRS, Inria, LIP, F-69342, Lyon Cedex 07, France

clement.lalanne@ens-lyon.fr

Aurélien Garivier

Univ. Lyon, ENS Lyon, UMPA UMR 5669, 46 allée d'Italie, F-69364, Lyon cedex 07

aurelien.garivier@ens-lyon.fr

Rémi Gribonval

Univ. Lyon, ENS Lyon, UCBL, CNRS, Inria, LIP, F-69342, Lyon Cedex 07, France

remi.gribonval@inria.fr

Abstract

The challenge of producing accurate statistics while respecting the privacy of the individuals in a sample is an important area of research. We study minimax lower bounds for classes of differentially private estimators. In particular, we show how to characterize the power of a statistical test under differential privacy in a *plug-and-play* fashion by solving an appropriate transport problem. With specific coupling constructions, this observation allows us to derive Le Cam-type and Fano-type inequalities not only for regular definitions of differential privacy but also for those based on Renyi divergence. We then proceed to illustrate our results on three simple, fully worked out examples. In particular, we show that the problem class has a huge importance on the provable degradation of utility due to privacy. In certain scenarios, we show that maintaining privacy results in a noticeable reduction in performance only when the level of privacy protection is very high. Conversely, for other problems, even a modest level of privacy protection can lead to a significant decrease in performance. Finally, we demonstrate that the DP-SGLD algorithm, a private convex solver, can be employed for maximum likelihood estimation with a high degree of confidence, as it provides near-optimal results with respect to both the size of the sample and the level of privacy protection. This algorithm is applicable to a broad range of parametric estimation procedures, including exponential families.

1 Introduction

The ever-increasing data collection on individuals and their sometimes hazardous use has led to numerous threats to privacy and serious concerns have emerged (Narayanan & Shmatikov, 2006; Backstrom et al., 2007; Fredrikson et al., 2015; Dinur & Nissim, 2003; Homer et al., 2008; Loukides et al., 2010; Narayanan & Shmatikov, 2008; Sweeney, 2000; Wagner & Eckhoff, 2018; Sweeney, 2002). *Differential privacy* (Dwork & Roth, 2014) offers a future-proof solution to this problem by ensuring that individuals are protected from the result of an estimation procedure. It enables the inference of global statistics on a dataset while bounding each sample's influence and ensuring that the presence or absence of an individual in the dataset cannot be deduced from the result. In the last decade, research results have multiplied and nowadays, it is possible to build complex data pipelines under privacy constraints (Dwork et al., 2006; Kairouz et al., 2015; Dong et al., 2019; 2020; Abadi et al., 2016). Notably, differential privacy is now used in production by the US Census Bureau (Abowd, 2018), Google (Erlingsson et al., 2014), Apple (Thakurta et al., 2017) and Microsoft (Ding et al., 2017) among others. Differential privacy constrains the class of usable *stochastic* functions defined on a dataset and as a result it degrades the utility of an estimation. To quantify the loss due to privacy, many *utility bounds* are often used (McSherry & Talwar, 2007; McSherry, 2010). When the target of the estimation is not the pointwise evaluation of a function on a given dataset but rather a hidden quantity defined “at the scale of the population”, i.e., of the underlying data distribution, statistical problems arise and an important topic is to quantify the utility of private estimators (Dwork & Lei,

2009; Wasserman & Zhou, 2010; Hall et al., 2011; Smith, 2011; Chaudhuri et al., 2011; Rubinstein et al., 2012; Lalanne et al., 2022; Ryffel et al., 2022; Karwa & Vadhan, 2018; Du et al., 2020; Biswas et al., 2020; Diakonikolas et al., 2015; Bun et al., 2019; Bun & Steinke, 2019; Ben-Eliezer et al., 2022).

Part of the privacy literature focuses on lower bounds (Asi & Duchi, 2020a;b; Farhadi et al., 2022; Tao et al., 2022), i.e. bounds for which we know we "cannot" do better under certain hypotheses. Most of them are problem-specific and are to be understood as a worst case among all instances of a given problem. Some others, such as the ones based on the theory of *inverse sensitivity* (Asi & Duchi, 2020a;b), only consider a "local" worst case and thus give tighter results. In contrast, in statistics, it is possible to measure the probabilities of occurrence of each instance of a problem. It is then natural to consider lower bounds that are probabilistic in nature (i.e. with a certain probability or in expectation). In particular, the classical *minimax* theory looks at the best uniform risk of convergence of the estimators in order to estimate a quantity defined at the scale of a population and has a vast literature on lower bounds (Assouad, 1983; Ibragimov & Khasminskii, 1981; Bickel & Ritov, 1988; Giraud, 2021; Devroye, 1987; Han & Verdú, 1994; Cover & Thomas, 2006; Scarlett & Cevher, 2019; Rigollet & Hütter, 2015; Tsybakov, 2009; Györfi et al., 2002). Under privacy conditions, some work is still to be done on the minimax risk, both on lower bounds and on some matching upper bounds. Our work sits in this line of research.

At the time of writing, to the best of our knowledge, two series of work have already met some success in this task. The first one is due to Duchi et al. (2013; 2014) and looks at so-called *local* privacy (Bebensee, 2019; Yang et al., 2020; Cormode et al., 2018) which is a stronger definition of privacy than the ones that will be investigated in this article. To put it simply, it requires that *each piece of data is anonymized before collection*, whereas *global* privacy only requires *the aggregation* to be private. It paved the way for a consequent line of work on estimation under local privacy or more generally under communication constraints (Acharya et al., 2021a;b;c;d; Barnes et al., 2020a;b; 2019). The second one is due to Acharya et al. (2021e; 2018). It adapts Le Cam's, Fano's and Assouad's methods to differential privacy. Their work will be our main point of comparison, since we believe that our work nicely complements theirs by providing a somewhat *plug-and-play* framework, notably allowing to establish results for other types of privacy.

Indeed, the present work extends classical techniques like Le Cam and Fano for proving minimax lower bounds to the setting where the estimator must additionally be differentially private. While this was previously done by Acharya et al. (2021e; 2018) in the specific context of (ϵ, δ) -differential privacy, our contribution improves upon that work by getting quantitative bounds that are both tighter (e.g. by large constants in the exponent) and somewhat *plug-and-play* in that they apply to a variety of differential privacy variants, and notably to the setting of *zero-concentrated differential privacy (zCDP)*, bridging the theoretical gap with very recent work in that field (Kamath et al., 2022). As applications, we extend known rates for estimating the parameter of a Bernoulli distribution and for estimating the mean of a high-dimensional spherical Gaussian to the zCDP setting. We also show that for the problem of estimating the support of a uniform distribution over an unknown interval inside $[0,1]$, the minimax risk is *uniformly degraded* by the privacy constraint. This contrasts with Bernoulli and Gaussian mean estimation, and is perhaps somewhat surprising, as one usually expects bounds for which there are parameter regimes under which the non-private minimax risk dominates.

Lastly, we show a private minimax lower bound for maximum likelihood estimation, when the log likelihood is concave and smooth in the parameter being estimated and satisfies some additional non-degeneracy assumptions, and show that private SGLD qualitatively matches this lower bound in terms of sample size and privacy parameter dependence.

1.1 The Minimax Risk and Private Estimators

We start by defining the minimax risk. Given $n \in \mathbb{N}_*$ and a Polish feature space \mathcal{X} , \mathcal{X}^n may be viewed as a set of datasets containing n elements from \mathcal{X} . We consider a family of probability distributions $(\mathbb{P}_\theta)_{\theta \in \Theta}$ on \mathcal{X}^n where Θ is equipped with a semi-metric¹ $d_\Theta : \Theta^2 \rightarrow \mathbb{R}_+$. Often, for all $\theta \in \Theta$, $\mathbb{P}_\theta = \mathbb{p}_\theta^{\otimes n}$ where $(\mathbb{p}_\theta)_{\theta \in \Theta}$ is a family of probability distributions on \mathcal{X} . This corresponds to the classical statistical setup where we observe n i.i.d. random variables. The general setup allows capturing phenomena that are not i.i.d., for instance Markov processes. Given an estimator $\hat{\theta} : \mathcal{X}^n \rightarrow \Theta$ one might look at its uniform risk of estimation

¹i.e. that is positive, symmetric, that satisfies the triangular inequality and $d_\Theta(\theta, \theta) = 0, \forall \theta \in \Theta$

over Θ for a loss function $\Phi : [0, +\infty) \rightarrow [0, +\infty)$ that is non-decreasing and such that $\Phi(0) = 0$ which is

$$\sup_{\theta \in \Theta} \int_{\mathcal{X}^n} \Phi(d_{\Theta}(\hat{\theta}(\mathbf{X}), \theta)) d\mathbb{P}_{\theta}(\mathbf{X}) .$$

The best achievable uniform risk defines what is called the *minimax* risk

$$\mathfrak{M}_n((\mathbb{P}_{\theta})_{\theta \in \Theta}, d_{\Theta}, \Phi) := \inf_{\hat{\theta}} \sup_{\theta \in \Theta} \int_{\mathcal{X}^n} \Phi(d_{\Theta}(\hat{\theta}(\mathbf{X}), \theta)) d\mathbb{P}_{\theta}(\mathbf{X}) . \quad (1)$$

Here, the infimum over $\hat{\theta}$ is taken among all possible measurable functions of the samples.

When an estimator $\hat{\theta} = \hat{\theta}(\mathbf{X})$ is to be made public or is to be shared with some untrustworthy agents and when the records of \mathbf{X} are *sensitive* (on a privacy standpoint), the disclosure of $\hat{\theta}$ may reveal a lot of information about the records of \mathbf{X} . Against this background, *differential privacy* (Dwork et al., 2006) offers strong privacy guarantees. Given a (randomized) mechanism \mathfrak{M} , $\text{dom}(\mathfrak{M})$ refers to its domain (i.e. the set of admissible inputs) and $\text{codom}(\mathfrak{M})$ refers to its codomain (i.e. the set of admissible outputs). A differentially private mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ ensures that limited information can be inferred on the records of $\mathbf{X} \in \mathcal{X}^n$ from the sole observation of the output $\mathfrak{M}(\mathbf{X})$. Given $\epsilon \in \mathbb{R}_{+*}$ and $\delta \in [0, 1)$, a randomized mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ is (ϵ, δ) -differentially private (or (ϵ, δ) -DP) if for all $\mathbf{X}, \mathbf{Y} \in \mathcal{X}^n$ and all measurable $S \subseteq \text{codom}(\mathfrak{M})$ we have

$$d_{\text{ham}}(\mathbf{X}, \mathbf{Y}) \leq 1 \implies \mathbb{P}_{\mathfrak{M}}(\mathfrak{M}(\mathbf{X}) \in S) \leq e^{\epsilon} \mathbb{P}_{\mathfrak{M}}(\mathfrak{M}(\mathbf{Y}) \in S) + \delta .$$

Note that $d_{\text{ham}}(\cdot, \cdot)$ denotes the Hamming distance on \mathcal{X}^n . There is however no consensus yet on the "correct" definition of privacy and a few other useful definitions have emerged. For instance, more recent definitions of privacy are due to the need to sharply count the privacy of a composition of many Gaussian mechanisms. At first it was done implicitly via the so-called moment accountant method (Abadi et al., 2016) before being formalized under the name of Renyi differential privacy (Mironov, 2017). Nowadays, it seems that all these notions tend to converge towards the definition of zero concentrated differential privacy (Dwork & Rothblum, 2016; Bun & Steinke, 2016). This is the one that we will investigate in this article in addition to the (ϵ, δ) -differential privacy, but the results and the proofs can easily be adapted to other definitions that are based on Renyi divergences. Given $\rho \in (0, +\infty)$, a randomized mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ is ρ -zero concentrated differentially private (ρ -zCDP) if for all $\mathbf{X}, \mathbf{Y} \in \mathcal{X}^n$,

$$d_{\text{ham}}(\mathbf{X}, \mathbf{Y}) \leq 1 \implies D_{\alpha}(\mathfrak{M}(\mathbf{X}) \parallel \mathfrak{M}(\mathbf{Y})) \leq \rho\alpha, \forall 1 < \alpha < +\infty .$$

The Renyi divergence of level α , $D_{\alpha}(\cdot \parallel \cdot)$, is properly defined in Appendix A. There exist links between (ϵ, δ) -DP and ρ -zCDP. For instance, Bun & Steinke (2016, Proposition 3) states that if a mechanism is ρ -zCDP, it is (ϵ, δ) -DP for a collection of (ϵ, δ) 's that depend on ρ . In particular, finding minimax lower bounds for ρ -zCDP mechanisms can be done by taking the supremum of lower bounds on (ϵ, δ) -DP mechanisms. But as we will see later, we can do better directly. Conversely, if a mechanism is $(\epsilon, 0)$ -DP, it is also $\epsilon^2/2$ -zCDP (see Bun & Steinke (2016, Proposition 4)).

In order to factorize the results, we will use the abstract formulation that a randomized mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \Theta$ satisfies a certain condition \mathcal{C} rather than fixing the class in which it belongs. We define the *private minimax* risk as the best achievable uniform risk with mechanisms that satisfy the privacy condition \mathcal{C} (using the set convention, \mathcal{C} can alternatively refer to the set of estimators that satisfy this condition)

$$\mathfrak{M}_n(\mathcal{C}, (\mathbb{P}_{\theta})_{\theta \in \Theta}, d_{\Theta}, \Phi) := \inf_{\mathfrak{M} \in \mathcal{C}} \sup_{\theta \in \Theta} \int_{\mathcal{X}^n} \mathbb{E}_{\mathbb{P}_{\mathfrak{M}}}(\Phi(d_{\Theta}(\mathfrak{M}(\mathbf{X}), \theta))) d\mathbb{P}_{\theta}(\mathbf{X}) . \quad (2)$$

Because of their similarities, we use \mathfrak{M}_n to refer to both the non-private minimax risk and its private counterpart. With four arguments, \mathfrak{M}_n should be understood as the private minimax risk and when equipped with three arguments it is simply the regular minimax risk.

1.2 Introducing example

As a warmup we discuss here the simplest possible example on which we can present the questions that this article addresses and the flavor of the developed approaches. Let $p_1 < p_2$ be two parameters in $(0, 1)$ and let U_1, \dots, U_n , n be independent and identically distributed uniform random variables on $[0, 1]$. The random variables $Z_i := (X_i^{(1)}, X_i^{(2)}) \in \mathbb{R}^2$, $1 \leq i \leq n$, defined by

$$(X_i^{(1)}, X_i^{(2)}) = (1_{[0, p_1]}(U_i), 1_{[0, p_2]}(U_i))$$

are independent and identically distributed with marginal distributions Bernoulli $\mathcal{B}(p_1)$ and $\mathcal{B}(p_2)$. In the sequel we note $\mathbf{X}^{(j)} = (X_1^{(j)}, \dots, X_n^{(j)})$, $j = 1, 2$, $\mathbf{U} = (U_1, \dots, U_n)$, $S_1 := [0, (p_1 + p_2)/2]$. and $S_2 := [(p_1 + p_2)/2, 1]$. Given any $(\epsilon, 0)$ -DP mechanism $\mathfrak{M} : [0, 1]^n \rightarrow [0, 1]$ (where $\epsilon > 0$) to estimate the Bernoulli parameter, the risk satisfies

$$\begin{aligned} & \sup_{p \in [0, 1]} \mathbb{E}_{\mathbf{X} \sim \mathcal{B}(p)^{\otimes n}} ((\mathfrak{M}(\mathbf{X}) - p)^2) \\ & \geq (\mathbb{E}_{\mathbf{X} \sim \mathcal{B}(p_1)^{\otimes n}} ((\mathfrak{M}(\mathbf{X}) - p_1)^2) + \mathbb{E}_{\mathbf{X} \sim \mathcal{B}(p_2)^{\otimes n}} ((\mathfrak{M}(\mathbf{X}) - p_2)^2)) / 2 \\ & \stackrel{\text{Coupling}}{=} (\mathbb{E}_{\mathbf{U}, \mathfrak{M}} ((\mathfrak{M}(\mathbf{X}^{(1)}) - p_1)^2) + \mathbb{E}_{\mathbf{U}, \mathfrak{M}} ((\mathfrak{M}(\mathbf{X}^{(2)}) - p_2)^2)) / 2 \\ & \stackrel{\text{Conditioning}}{=} \mathbb{E}_{\mathbf{U}} (\mathbb{E}_{\mathfrak{M}} ((\mathfrak{M}(\mathbf{X}^{(1)}) - p_1)^2) + \mathbb{E}_{\mathfrak{M}} ((\mathfrak{M}(\mathbf{X}^{(2)}) - p_2)^2)) / 2 \\ & \geq \left(\frac{p_2 - p_1}{2}\right)^2 \mathbb{E}_{\mathbf{U}} (\mathbb{P}_{\mathfrak{M}} (\mathfrak{M}(\mathbf{X}^{(1)}) \in S_2) + \mathbb{P}_{\mathfrak{M}} (\mathfrak{M}(\mathbf{X}^{(2)}) \in S_1)) / 2. \end{aligned} \quad (3)$$

This is where the DP property yields a lower bound on the second factor as

$$\begin{aligned} & \mathbb{E}_{\mathbf{U}} \left(e^{-\epsilon d_{\text{ham}}(\mathbf{X}^{(1)}, \mathbf{X}^{(2)})} \mathbb{P}_{\mathfrak{M}} (\mathfrak{M}(\mathbf{X}^{(2)}) \in S_2) + \mathbb{P}_{\mathfrak{M}} (\mathfrak{M}(\mathbf{X}^{(2)}) \in S_1) \right) \\ & \stackrel{d_{\text{ham}}(\cdot, \cdot) \geq 0}{\geq} \mathbb{E}_{\mathbf{U}} \left(e^{-\epsilon d_{\text{ham}}(\mathbf{X}^{(1)}, \mathbf{X}^{(2)})} (\mathbb{P}_{\mathfrak{M}} (\mathfrak{M}(\mathbf{X}^{(2)}) \in S_2) + \mathbb{P}_{\mathfrak{M}} (\mathfrak{M}(\mathbf{X}^{(2)}) \in S_1)) \right) \\ & = \mathbb{E}_{\mathbf{U}} \left(e^{-\epsilon d_{\text{ham}}(\mathbf{X}^{(1)}, \mathbf{X}^{(2)})} \right) \stackrel{\text{Jensen}}{\geq} e^{-n\epsilon |p_2 - p_1|}, \end{aligned} \quad (4)$$

which overall yields the lower bound $\frac{(p_2 - p_1)^2}{8} e^{-n\epsilon |p_2 - p_1|}$. A good lower bound on the minimax risk is then provided by optimizing over p_1 and p_2 . For instance, when $n \geq \frac{2}{\epsilon}$, $p_1 = \frac{1}{2}$ and $p_2 = \frac{1}{2} + \frac{1}{n\epsilon}$ leads to

$$\sup_{p \in [0, 1]} \mathbb{E}_{\mathbf{X} \sim \mathcal{B}(p)^{\otimes n}} ((\mathfrak{M}(\mathbf{X}) - p)^2) \geq \frac{1}{8} \frac{1}{(n\epsilon)^2}.$$

The idea behind the first inequality in (3) is classical in the minimax literature and is recalled in Section 1.3 using the notion of *packing*. The *coupling construction* can be generalized and tailored to other settings and has a critical impact on the deduced lower bounds, as we present in Section 3. The minoration involving differential privacy is a special case of the techniques that we formalize under the notion of *admissible similarity functions* in Section 2, which are adapted to various types of privacy constraints. Practical implications of such generalizations are described in Section 1.4.

1.3 From Minimax Lower Bounds to Hypothesis Testing

A classical technique (see Duchi et al. (2013)) for finding lower bounds on $\mathfrak{M}_n((\mathbb{P}_\theta)_{\theta \in \Theta}, d_\Theta, \Phi)$ is to replace the parameter set Θ by a much ‘‘simpler’’ set $\Theta' \subseteq \Theta$ and to use the trivial lower bound

$$\mathfrak{M}_n((\mathbb{P}_\theta)_{\theta \in \Theta}, d_\Theta, \Phi) \geq \inf_{\hat{\theta}} \sup_{\theta \in \Theta'} \int_{\mathcal{X}^n} \Phi(d_\Theta(\hat{\theta}(\mathbf{X}), \theta)) d\mathbb{P}_\theta(\mathbf{X}).$$

Usually Θ' is chosen as an Ω -*packing* of Θ , for some real number $\Omega > 0$: it is a countable family $\Theta' := \{\theta_i, i \in \mathbb{N}_*\}$ $(\theta_i)_{i \in \mathbb{N}_*}$ (and most of the time, including in this article, it is taken to be finite) such that: a) $\theta_i \in \Theta$ for all i ; b) $i \neq j \implies d_\Theta(\theta_i, \theta_j) \geq 2\Omega$; and c) there is a well-defined function $\Psi_{\Theta'}$ satisfying

$$\Psi_{\Theta'}(\theta) \in \arg \min_{i \geq 1} d_\Theta(\theta_i, \theta)$$

for each $\theta \in \Theta$. Under such hypotheses, any estimator $\hat{\theta}$ satisfies (Duchi et al., 2013)

$$\sup_{\theta \in \Theta'} \int_{\mathcal{X}^n} \Phi(d_{\Theta}(\hat{\theta}(\mathbf{X}), \theta)) d\mathbb{P}_{\theta}(\mathbf{X}) \geq \Phi(\Omega) \sup_{i \in \{1, \dots, \#(\Theta')\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_{\theta_i}} \left(\Psi_{\Theta'} \left(\hat{\theta}(\mathbf{X}) \right) \neq i \right). \quad (5)$$

The mapping $\hat{\Psi} := \Psi_{\Theta'} \circ \hat{\theta} : \mathcal{X}^n \rightarrow \{1, \dots, \#(\Theta')\}$ may be viewed as a *test function* (that selects the model number) and thus

$$\mathfrak{M}_n((\mathbb{P}_{\theta})_{\theta \in \Theta}, d_{\Theta}, \Phi) \geq \Phi(\Omega) \inf_{\Psi: \mathcal{X}^n \rightarrow \{1, \dots, \#(\Theta')\}} \sup_{i \in \{1, \dots, \#(\Theta')\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_{\theta_i}} (\Psi(\mathbf{X}) \neq i). \quad (6)$$

Finding minimax lower bounds is thus done by finding a suitable Ω -packing of the parameter space and then by providing lower bounds on

$$\inf_{\Psi: \mathcal{X}^n \rightarrow \{1, \dots, \#(\Theta')\}} \sup_{i \in \{1, \dots, \#(\Theta')\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_{\theta_i}} (\Psi(\mathbf{X}) \neq i). \quad (7)$$

Two powerful tools to find such lower bounds come from information theory: Le Cam's lemma (see Fact 1) can be used when Θ' only contains two elements, while Fano's lemma (see Fact 2) is applicable when Θ' contains $N \geq 2$ elements.

Fact 1 (Neyman-Pearson & Le Cam's lemma (Rigollet & Hütter, 2015, Lemma 5.3)). *Let $\mathbb{P}_1, \mathbb{P}_2$ be two probability distributions on a measure space \mathcal{E} , then*

$$\begin{aligned} \inf_{\Psi: \mathcal{E} \rightarrow \{1, 2\}} \max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi(\mathbf{X}) \neq i) &\geq \frac{1}{2} \inf_{\Psi: \mathcal{E} \rightarrow \{1, 2\}} \sum_{i=1}^2 \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi(\mathbf{X}) \neq i) \\ &= \frac{1}{2} (1 - \text{TV}(\mathbb{P}_1, \mathbb{P}_2)). \end{aligned} \quad (8)$$

The Total Variation (TV) is rigorously defined in Appendix A.

Fact 2 (Fano's lemma (Giraud, 2021, Theorem 3.1)). *Let $(\mathbb{P}_i)_{i \in \{1, \dots, N\}}$ be a family of probability distributions on a measure space \mathcal{E} . For any probability distribution \mathbb{Q} on \mathcal{E} such that $\mathbb{P}_i \ll \mathbb{Q}$ for all i , and for any test function $\Psi : \mathcal{X}^n \rightarrow \{1, \dots, N\}$,*

$$\begin{aligned} \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi(\mathbf{X}) \neq i) &\geq \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi(\mathbf{X}) \neq i) \\ &\geq 1 - \frac{1 + \frac{1}{N} \sum_{i=1}^N \text{KL}(\mathbb{P}_i \| \mathbb{Q})}{\ln(N)}. \end{aligned} \quad (9)$$

The KL divergence and the absolute continuity (\ll) are rigorously defined in Appendix A. Often \mathbb{Q} is set to $\frac{1}{N} \sum_{i=1}^N \mathbb{P}_i$.

With the same reasoning used (Duchi et al., 2013) to establish (5), with $\Theta' = (\theta_i)_{i \in \{1, \dots, \#(\Theta')\}}$ an Ω -packing of Θ , we can lower-bound the *private* minimax risk:

$$\begin{aligned} \mathfrak{M}_n(\mathcal{C}, (\mathbb{P}_{\theta})_{\theta \in \Theta}, d_{\Theta}, \Phi) \\ \geq \Phi(\Omega) \inf_{\mathfrak{M} \in \mathcal{C}} \inf_{\Psi: \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, \#(\Theta')\}} \sup_{i \in \{1, \dots, \#(\Theta')\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_{\theta_i}, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i). \end{aligned} \quad (10)$$

Consequently, finding private minimax lower bounds is done analogously to the non-private setting by finding an appropriate Ω -packing and a lower bound on

$$\inf_{\Psi: \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, \#(\Theta')\}} \sup_{i \in \{1, \dots, \#(\Theta')\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_{\theta_i}, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \quad (11)$$

that is independent on the mechanism \mathfrak{M} but only depends on the privacy condition \mathcal{C} .

1.4 Contributions

The main contribution of this work, presented in Section 2, is to propose a generic framework for the derivation of lower bounds on the minimax risk under various privacy conditions. Here and in the sequel, the symbols $o(\cdot)$, $O(\cdot)$, $\Theta(\cdot)$ and $\Omega(\cdot)$ are used without ambiguity as classical comparison operators for sequences, as recalled in Appendix A. Technically, the techniques of Le Cam and Fano are extended to the private context, reducing the distributional test problem (11) to a Kantorovich problem (Santambrogio, 2016; Peyré & Cuturi, 2019) of the form

$$\sup_{\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)} \int_{(\mathcal{X}^n)^N} s_{\mathcal{C}}(\mathbf{X}_1, \dots, \mathbf{X}_N) d\mathbb{Q}(\mathbf{X}_1, \dots, \mathbf{X}_N) . \quad (12)$$

Here, $\Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$ is the set of *couplings* between the considered distributions and $s_{\mathcal{C}}$ is an *admissible similarity function* depending on the nature of the constraint \mathcal{C} and the number of hypotheses (Theorem 6 and Theorem 7). For instance, regarding (ϵ, δ) -differential privacy, similarity functions are obtained by comparing datasets to a common *anchor*. This result is summarized in Theorem 5.

Unlike the prior work of Acharya et al. (2021e), the proposed framework allows us to consider joint couplings across all instances rather than just pairwise couplings. Additionally, the level of generality of our proofs leaves room for subsequent work to build upon this framework.

The general idea behind the proofs is as follows. In classical Fano's, one considers the decoding error probability: on average over a family of instances, what is the probability that the estimator, given samples from a given instance, fails to identify that the samples came from that instance. In place of Fano's inequality, the present work lower bounds this by noting that, given datasets $\mathbf{X}_1, \dots, \mathbf{X}_N$ coming from each instance, as well as an "anchor" dataset Λ (or alternatively an anchor distribution), differential privacy implies that the probability that the estimator decides \mathbf{X}_i comes from instance i cannot differ by much from the probability it decides Λ comes from instance i , provided Λ and \mathbf{X}_i are similar. The decoding error probability can thus be lower bounded in terms of the maximum distance between Λ and any of $\mathbf{X}_1, \dots, \mathbf{X}_N$, averaged over the randomness of $\mathbf{X}_1, \dots, \mathbf{X}_N$, where there is freedom in choosing how to couple this randomness.

Section 3 includes various coupling constructions yielding quantitative lower bounds for the Kantorovich formulation (12). These constructions only depend on the number of hypotheses N , the sample size n , the privacy parameters ϵ, δ, ρ , and information theoretic quantities such as the pairwise total variations or KL divergences between the distributions.

Those results will be presented in Section 2 and in Section 3. We showcase now useful consequences, starting with the case $N = 2$: similarly to Acharya et al. (2021e), we extend Le Cam's lemma to the (ϵ, δ) -differentially private setting:

Theorem 1 (Le Cam for (ϵ, δ) -DP). *If a randomized mechanism \mathfrak{M} satisfies (ϵ, δ) -DP, then for any test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, 2\}$ and any probability distributions \mathbb{P}_1 and \mathbb{P}_2 on \mathcal{X}^n we have*

$$\begin{aligned} \max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) &\geq \frac{1}{2} \max \left\{ 1 - \text{TV}(\mathbb{P}_1, \mathbb{P}_2) , \right. \\ &\left. 1 - (1 - e^{-n\epsilon} + 2ne^{-\epsilon\delta}) \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right\} . \end{aligned}$$

Furthermore, when $\mathbb{P}_1 = \mathbb{p}_1^{\otimes n}$ and $\mathbb{P}_2 = \mathbb{p}_2^{\otimes n}$ are product distributions,

$$\begin{aligned} \max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ \geq \frac{1}{2} \left((1 - (1 - e^{-\epsilon}) \text{TV}(\mathbb{p}_1, \mathbb{p}_2))^n - 2ne^{-\epsilon\delta} \text{TV}(\mathbb{p}_1, \mathbb{p}_2) \right) . \end{aligned}$$

The proof can be found in Appendix C. The classical lower bound of Le Cam (8) allows for a tunable testing difficulty depending on $\text{TV}(\mathbb{P}_1, \mathbb{P}_2)$. However, in the regime $\epsilon, \delta = o(1/n)$, the private lower bound is $\Omega(1)$: it becomes arbitrarily hard to distinguish between any pair of distributions.

For i.i.d. observations ($\mathbb{P}_i = \mathbb{p}_i^{\otimes n}$), it follows by convexity that for any (ϵ, δ) -DP mechanism \mathfrak{M} and test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, 2\}$

$$\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \frac{1}{2} \left(e^{-\epsilon n \text{TV}(\mathbb{P}_1, \mathbb{P}_2)} - 2e^{-\epsilon} \delta n \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right).$$

This is to be compared to the state of the art lower bound of Acharya et al. (2021e, Theorem 1):

$$\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \frac{1}{2} \left(0.9e^{-10\epsilon n \text{TV}(\mathbb{P}_1, \mathbb{P}_2)} - 10\delta n \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right).$$

Theorem 1 gives tighter results with better constants, notably thanks to a different proof technique avoiding some convexity and concentration inequalities, but also because Acharya et al. (2021e) did not optimize the constants. Indeed, qualitative results and rates of convergence do usually not depend on them.

We also prove an equivalent for so-called ρ -zero concentrated differential privacy (or in short ρ -zCDP), which is, to the best of our knowledge, the first successful attempt at doing so.

Theorem 2 (Le Cam for ρ -zCDP). *If a randomized mechanism \mathfrak{M} satisfies ρ -zCDP, then for any test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$ and any probability distributions \mathbb{P}_1 and \mathbb{P}_2 on \mathcal{X}^n ,*

$$\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \frac{1}{2} \max \left\{ 1 - \text{TV}(\mathbb{P}_1, \mathbb{P}_2), \right. \\ \left. 1 - n\sqrt{\rho/2} \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right\}.$$

Furthermore, when $\mathbb{P}_1 = \mathbb{p}_1^{\otimes n}$ and $\mathbb{P}_2 = \mathbb{p}_2^{\otimes n}$ are product distributions,

$$\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \frac{1}{2} \left(1 - n\sqrt{\rho/2} \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right).$$

The proof of this result can be found in Appendix C. As above, any two distributions can no longer be distinguished in the regime $\rho \ll 1/n^2$. For more than two hypotheses and (ϵ, δ) -DP, we also get a private version of Fano's lemma.

Theorem 3 (Multiple Distributional Tests for (ϵ, δ) -DP). *If a randomized mechanism \mathfrak{M} satisfies (ϵ, δ) -DP, then for any test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$, any family of probability distributions $(\mathbb{P}_i)_{i \in \{1, \dots, N\}}$ on \mathcal{X}^n and any Q such that $\mathbb{P}_i \ll Q$ for all i ,*

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \max \left\{ 1 - \frac{1 + \frac{1}{N} \sum_{i=1}^N \text{KL}(\mathbb{P}_i \| Q)}{\ln(N)}, \right. \\ \frac{1}{2} - \frac{1 - e^{-n\epsilon} + 2ne^{-\epsilon}\delta}{2N^2} \sum_{i,j} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}, \\ \left. 1_{\delta=0} \times \left(1 - \frac{1 + \frac{n\epsilon}{N^2} \sum_{i,j} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}}{\ln(N)} \right) \right\}.$$

Furthermore, when $\mathbb{P}_1 = \mathbb{p}_1^{\otimes n}, \dots, \mathbb{P}_N = \mathbb{p}_N^{\otimes n}$ are product distributions,

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \max \left\{ \frac{1}{2N^2} \sum_{i,j} \left(\left(1 - (1 - e^{-\epsilon}) \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \right)^n \right. \right. \\ \left. \left. - 2ne^{-\epsilon}\delta \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \right), \right. \\ \left. 1_{\delta=0} \times \left(1 - \frac{1 + \frac{n\epsilon}{N^2} \sum_{i,j} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}}{\ln(N)} \right) \right\}.$$

The proof is given in Appendix C. When dealing with product distributions, the quantity

$$D := \frac{n}{N^2} \sum_{i,j} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}$$

can roughly be seen as an averaged hamming distance between pairs of marginals. An implication of Theorem 3 is then that

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq 1_{\delta=0} \times \left(1 - \frac{1 + \epsilon D}{\ln(N)}\right).$$

As the bound of Acharya et al. (2021e, Theorem 2)

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq 1_{\delta=0} \times 0.9 \times \min \left\{1, \frac{N}{e^{10\epsilon D}}\right\}, \quad (13)$$

the lower bound is $\Omega(1)$ in the regime $D = o(\ln(N)/\epsilon)$. In particular, both inequalities are expected to yield similar qualitative results for a broad range of applications. However, the quantitative consequences of Theorem 3 can again be orders of magnitude better. Another improvement of our result is that, contrary to previous work, our bound allows to handle asymmetric hypotheses. Indeed, prior work is based on a uniform upper-bound on the family $(\text{TV}(\mathbb{P}_i, \mathbb{P}_j))_{i,j}$, whereas our work uses only their mean value. As an illustration, if a statistician was to discriminate between a set of N distributions with for instance $N - 1$ distributions close to each other in total variation distance and one outlier far from all the others, the results of Acharya et al. (2021e) only tell that the problem will be at least as hard as discriminating distributions that are far from one another (which is easy). In contrast, our Theorem 3 shows that the true testing difficulty lies in discriminating the distributions that are similar (the outlier vanishes), thus resulting in lower bounds that are less over-optimistic.

Similarly, we obtain results for multiple hypotheses under ρ -zCDP.

Theorem 4 (Multiple Distributional Tests for ρ -zCDP). *If a randomized mechanism \mathfrak{M} satisfies ρ -zCDP, then for any test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$, any family of probability distributions $(\mathbb{P}_i)_{i \in \{1, \dots, N\}}$ on \mathcal{X}^n and any \mathbb{Q} such that $\mathbb{P}_i \ll \mathbb{Q}$ for all i ,*

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \max \left\{ 1 - \frac{1 + \frac{1}{N} \sum_{i=1}^N \text{KL}(\mathbb{P}_i \| \mathbb{Q})}{\ln(N)}, \right. \\ \left. 1 - \frac{1 + \frac{n^2 \rho}{N^2} \sum_{i,j} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}}{\ln(N)} \right\}.$$

Furthermore, when $\mathbb{P}_1 = \mathbb{P}_1^{\otimes n}, \dots, \mathbb{P}_N = \mathbb{P}_N^{\otimes n}$ are product distributions,

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq 1 - \frac{1 + \frac{n^2 \rho}{N^2} \sum_{i,j} \frac{1}{n} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)} + \left(\frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}\right)^2}{\ln(N)}.$$

The proof is to be found in Appendix C. This result recovers a recently published result in Kamath et al. (2022), with the advantage again of better handling asymmetrical hypotheses (i.e. with possible outliers). Another interesting observation is that our framework unifies the proofs of lower bounds under a general technique based on multiple marginals coupling and similarity functions.

A more detailed discussion about the different privacy regimes is less direct compared to Le Cam's method and is postponed to Appendix D, where we discuss three specific examples, recovering known facts and uncovering novel results.

Bernoulli model. We first recover that the classical minimax rate $\Theta(1/n)$ for the squared error on the estimation of the parameter of a Bernoulli distribution becomes $\Theta\left(\max\left\{\frac{1}{n}, \frac{1}{(n\epsilon)^2}\right\}\right)$ under ϵ -differential privacy. Furthermore, we exhibit a new rate in the case of ρ -zero concentrated differential privacy : $\Theta\left(\max\left\{\frac{1}{n}, \frac{1}{n^2 \rho}\right\}\right)$.

Gaussian Model. We allow each piece of data to have dimensionality d (the dimension of \mathcal{X}). We again recover that the minimax risk $\Theta\left(\frac{\sigma^2 d}{n}\right)$ (see Rigollet & Hütter (2015, Corollary 5.13)) becomes $\Omega\left(\max\left\{\frac{\sigma^2 d}{n}, \frac{\sigma^2 d^2}{(n\epsilon)^2}\right\}\right)$ under ϵ -DP, and we prove that it becomes $\Omega\left(\max\left\{\frac{\sigma^2 d}{n}, \frac{\sigma^2 d}{n^2 \rho}\right\}\right)$ under ρ -zCDP.

Uniform model. This example allows us to exhibit a new behavior under privacy. Indeed, we show that the usual minimax risk $\Theta\left(\frac{1}{n^2}\right)$ becomes $\Omega\left(\max\left(\frac{1}{n^2}, \frac{1}{(n\epsilon)^2}\right)\right)$ under ϵ -DP and $\Omega\left(\max\left(\frac{1}{n^2}, \frac{1}{n^2 \rho}\right)\right)$ under ρ -zCDP, proving a systematic degradation of utility due to privacy that in particular does not depend on a disjunction on the rates at which the privacy-tuning parameters decrease.

Finally, in Section 4, we study the private parametric estimation of distributions that are log-concave with respect to the parameter. In particular, under some mild hypotheses, we exhibit some lower bounds on the minimax risk of estimation. We then show, based on existing upper bounds, that under mild hypotheses, Differentially Private Stochastic Gradient Langevin Dynamics (DP-SGLD, see Ryffel et al. (2022)), a private optimization solver, is near-minimax optimal when used to perform private maximum likelihood estimation in this class of distributions. Here "near-minimax optimal" means that the ratio between the risk of private estimation and the private minimax risk is upper-bounded by a quantity that only depends on the regularity of the log-likelihood, such as the eigenvalues of its Hessian. In particular, it is independent of the sample size n or of the constants that tune the privacy.

2 From Testing to a Transport Problem

This section presents our main theorem, which states that finding lower bounds on (11) can be done by solving a transport problem (Santambrogio, 2016; Peyré & Cuturi, 2019). In some sense, this view is close to the coupling of Acharya et al. (2021e) which considers couplings between pairs of marginals and controls the variations of the hamming distance compared to its expected value with Markov's inequality. However, the high level view that our result allows to obtain numerically sharper results because it allows to skip approximations such as those involving Jensen or Markov inequalities and more importantly, it allows handling divergence-based definitions of privacy which do not fit in the framework of Acharya et al. (2021e). Furthermore, a key difference is that the theory of Acharya et al. (2021e) only requires to build couplings between pairs of marginals, whereas our theory requires building couplings between all the marginals at the same time. This is both a drawback because it requires to use more complex coupling constructions, and an advantage because it allows to give results that are easier to use when there are more than two hypotheses.

Our analysis is based on the notion of *similarity* functions.

Definition 1. Given a condition \mathcal{C} , we say that a similarity function $s_{\mathcal{C}} : (\mathcal{X}^n)^N \rightarrow \mathbb{R}$ is admissible for \mathcal{C} if for any mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ that satisfies \mathcal{C} , for any test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$, and for any $\mathbf{X}_1, \dots, \mathbf{X}_N \in \mathcal{X}^n$, the following inequality holds:

$$\frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq s_{\mathcal{C}}(\mathbf{X}_1, \dots, \mathbf{X}_N) .$$

Theorem 5. If a randomized mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ satisfies the privacy condition \mathcal{C} , for any $N \geq 2$, if $s_{\mathcal{C}} : (\mathcal{X}^n)^N \rightarrow \mathbb{R}$ is an admissible similarity function for \mathcal{C} , for any distributions $\mathbb{P}_1, \dots, \mathbb{P}_N$ over \mathcal{X}^n we have

$$\begin{aligned} & \inf_{\Psi: \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}} \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ & \geq \sup_{\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)} \int_{(\mathcal{X}^n)^N} s_{\mathcal{C}}(\mathbf{X}_1, \dots, \mathbf{X}_N) d\mathbb{Q}(\mathbf{X}_1, \dots, \mathbf{X}_N) . \end{aligned} \tag{14}$$

Proof. Given a test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$ and a coupling $\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$,

$$\begin{aligned} \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) &\geq \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ &= \int_{(\mathcal{X}^n)^N} \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) d\mathbb{Q}(\mathbf{X}_1, \dots, \mathbf{X}_N) \\ &\geq \int_{(\mathcal{X}^n)^N} s_C(\mathbf{X}_1, \dots, \mathbf{X}_N) d\mathbb{Q}(\mathbf{X}_1, \dots, \mathbf{X}_N) . \end{aligned}$$

□

In particular, under (ϵ, δ) -DP, similarity functions are built using a technique that we call *anchoring* which will be introduced in Appendix B.1, where the proof of the following theorem is given.

Theorem 6 (Admissible similarity functions for (ϵ, δ) -DP). *When \mathcal{C} is (ϵ, δ) -differential privacy, the following approaches yield admissible similarity functions.*

- **Global anchoring.** Consider any anchor function $\Lambda : (\mathcal{X}^n)^N \rightarrow \mathcal{X}^n$, and define the admissible similarity function as

$$s_C(\mathbf{X}_1, \dots, \mathbf{X}_N) := \frac{N-1}{N} e^{-\epsilon \max_i (d_{\text{ham}}(\mathbf{X}_i, \Lambda(\mathbf{X}_1, \dots, \mathbf{X}_N)))} - e^{-\epsilon \delta \max_i (d_{\text{ham}}(\mathbf{X}_i, \Lambda(\mathbf{X}_1, \dots, \mathbf{X}_N)))} .$$

- **Projection anchoring.** In particular, for any $j \in \{1, \dots, N\}$, consider the projection anchor $\Lambda_j(\mathbf{X}_1, \dots, \mathbf{X}_N) := \mathbf{X}_j$, and define the corresponding admissible similarity function

$$s_C(\mathbf{X}_1, \dots, \mathbf{X}_N) := \frac{N-1}{N} e^{-\epsilon \max_i (d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j))} - e^{-\epsilon \delta \max_i (d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j))}$$

- **(ϵ, δ) -DP Le Cam matching.** When $N = 2$, there is a global anchor function yielding the admissible similarity function

$$s_C(\mathbf{X}_1, \mathbf{X}_2) := \frac{1}{2} e^{-\epsilon \lceil d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2) / 2 \rceil} - e^{-\epsilon \delta \lceil d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2) / 2 \rceil} .$$

- **Pairwise anchoring.** An admissible similarity function is

$$s_C(\mathbf{X}_1, \dots, \mathbf{X}_N) := \frac{1}{2N^2} \sum_{i=1}^N \sum_{j=1}^N e^{-\epsilon \lceil d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j) / 2 \rceil} - 2e^{-\epsilon \delta \lceil d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j) / 2 \rceil} .$$

- **$(\epsilon, 0)$ -DP Fano matching.** When $\delta = 0$, an admissible similarity function is

$$s_C(\mathbf{X}_1, \dots, \mathbf{X}_N) := 1 - \frac{1 + \frac{\epsilon}{N^2} \sum_{i=1}^N \sum_{j=1}^N d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)}{\ln(N)} .$$

When working under ρ -zCDP, admissible similarity functions are built using classical information theoretic inequalities directly. It can be seen as a form of anchoring on the distributions rather than on the observed random variables (i.e. all the distributions are compared to a common distribution directly that is not necessarily a pushforward by \mathfrak{M}). The following result is proved in Appendix B.2.

Theorem 7 (Admissible similarity functions for ρ -zCDP). *When \mathcal{C} is the ρ -zero concentrated-differential privacy, the two following quantities are admissible similarity functions:*

- ρ -zCDP Fano matching

$$s_C(\mathbf{X}_1, \dots, \mathbf{X}_N) := 1 - \frac{1 + \frac{\rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)^2}{\ln(N)} .$$

- ρ -zCDP Le Cam matching When $N = 2$

$$s_C(\mathbf{X}_1, \mathbf{X}_2) := \frac{1}{2} \left(1 - \sqrt{\rho/2} d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2) \right) .$$

Note that similarity functions can also be easily built for the more general notion of (ξ, ρ) -concentrated differential privacy by swapping the group privacy property for its correct variant (see Bun & Steinke (2016)). We do not include the results about (ξ, ρ) -concentrated differential in this article because our objective is more to illustrate the versatility of our framework rather than to build a complete catalogue.

3 Lower Bound via Couplings

The transport problem (12) can be studied either theoretically (Santambrogio, 2016) or numerically (Peyré & Cuturi, 2019) in order to give the best lower bounds that our technique permits. However, identifying an optimal coupling is out of the scope of this article. We here provide coupling constructions that are sufficient to exhibit useful lower bounds.

Most of the similarity functions expressed in Theorem 5 yield lower bounds that are based on or further lower-bounded by expressions involving the quantities

$$\mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathbb{Q}} (g(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)))$$

for a coupling $\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$ where g is a fixed non-increasing function. Hence, finding reasonably good lower bounds can be achieved by finding a coupling that *minimizes* the expected pairwise Hamming distance between the marginals.

As a proxy, we first aim at maximizing the probabilities of pairwise equalities between all the marginals simultaneously. We then control the Hamming distance by observing that when $\mathbf{X}_i = \mathbf{X}_j$, $d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j) = 0$ and otherwise, $d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j) \leq n$. It is known (Lindvall, 2003) that if $\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$, the disagreement probabilities (i.e. the probability that two marginal random variables are not equal) between the marginals satisfy

$$\forall i, j, \quad \text{TV}(\mathbb{P}_i, \mathbb{P}_j) \leq \mathbb{P}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathbb{Q}} (\mathbf{X}_i \neq \mathbf{X}_j) . \quad (15)$$

A natural question is whether this lower bound is achievable by a coupling simultaneously for all pairs of marginals. When there are only two marginals (i.e. $N = 2$), a classical construction (see Lindvall (2003)) answers this question positively:

Fact 3 (Maximal coupling). *Let \mathbb{P}_1 and \mathbb{P}_2 be two probability distributions on \mathcal{X}^n that share the same σ -algebra. There exists a coupling $\pi^\infty(\mathbb{P}_1, \mathbb{P}_2) \in \Pi(\mathbb{P}_1, \mathbb{P}_2)$ (which is a distribution on $(\mathcal{X}^n)^2$), called a maximal coupling, such that*

$$\begin{aligned} \mathbb{P}_{(X_1, X_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)} (X_1 \neq X_2) &= \text{TV}(\mathbb{P}_1, \mathbb{P}_2) , \\ \forall S \text{ measurable} , \quad \mathbb{P}_{(X_1, X_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)} (X_1 \in S) &= \mathbb{P}_1(X_1 \in S) , \\ \forall S \text{ measurable} , \quad \mathbb{P}_{(X_1, X_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)} (X_2 \in S) &= \mathbb{P}_2(X_2 \in S) . \end{aligned}$$

This construction unfortunately does not generically scale to more than two marginals, even though on simple examples, couplings can be built that still match the lower bound (15) for any pair of marginals.

Example 1 (Bernoulli optimal coupling). *Given $\mathbb{P}_i = \mathcal{B}(p_i)$, $1 \leq i \leq N$ a family of Bernoulli distributions and $U \sim \mathcal{U}([0, 1])$ a uniformly distributed variable on $[0, 1]$, the random vector (X_1, \dots, X_N) defined by $X_i := 1_{[0, p_i)}(U)$ is distributed according to a coupling $\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$, and for every i, j*

$$\mathbb{P}(X_i \neq X_j) = |p_i - p_j| = \text{TV}(\mathcal{B}(p_i), \mathcal{B}(p_j)) .$$

There are however examples for which it is provably impossible to build couplings that match the lower bound (15) for any pair of marginals.

Example 2 (A counterexample). Let $X_1 \sim \mathcal{U}(\{-1, 0\})$, $X_2 \sim \mathcal{U}(\{0, 1\})$ and $X_3 \sim \mathcal{U}(\{1, -1\})$, and let \mathbb{P} be a coupling between X_1, X_2 and X_3 . We have that,

$$1_{X_1 \neq X_2} + 1_{X_2 \neq X_3} + 1_{X_3 \neq X_1} \geq 2$$

and as a consequence on \mathbb{P} ,

$$\begin{aligned} \mathbb{P}(X_1 \neq X_2) + \mathbb{P}(X_2 \neq X_3) + \mathbb{P}(X_3 \neq X_1) &\geq 2 \\ &> \text{TV}(X_1, X_2) + \text{TV}(X_2, X_3) + \text{TV}(X_3, X_1), \end{aligned}$$

which proves that at least one of the disagreement probabilities is strictly bigger than the corresponding total variation.

Recent constructions based on Poisson point processes allow in general, for any number of marginals N , to match the lower bound (15) up to a factor 2.

Fact 4 (Near optimal coupling of multiple distributions (Angel & Spinka, 2021)). Let $\mathbb{P}_1, \dots, \mathbb{P}_N$ be N distributions on the same measurable set. There exists a coupling $\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$ such that

$$\forall i, j \in \{1, \dots, N\}, \quad \mathbb{P}_{(X_1, \dots, X_N) \sim \mathbb{Q}}(X_i \neq X_j) \leq \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}.$$

We give a small note about this coupling in Appendix E that presents a simplified view on its construction.

In the rest of this article, the notation $\pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)$ refers to a coupling that satisfies this condition. When there are only two distributions, it refers to the construction of Fact 3. This factor 2 is not a problem for minimax theory, since it is a common practice to overlook the constants by looking at rates of convergence. However, for some more precise applications, working on more specific couplings may improve our results. With either coupling constructions, the lower bounds of Theorem 5 can be controlled with the following straightforward lemma:

Lemma 1. Let $\mathbb{P}_1, \dots, \mathbb{P}_N$ be N distributions on \mathcal{X}^n and $\mathbb{Q} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$. Consider $1 \leq i, j \leq N$ and denote $\Delta_{i,j} := \mathbb{P}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathbb{Q}}(\mathbf{X}_i \neq \mathbf{X}_j)$. We have

$$\begin{aligned} \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathbb{Q}}(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)) &\leq n\Delta_{i,j} \\ \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathbb{Q}}(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)^2) &\leq n^2\Delta_{i,j} \\ \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathbb{Q}}(e^{-\epsilon d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)}) &\geq 1 - (1 - e^{-n\epsilon})\Delta_{i,j}. \end{aligned}$$

Note that $\Delta_{i,j}$ directly depends on the coupling construction, but that with any of the ones presented above, we always have $\forall i, j, \Delta_{i,j} \leq 2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)$.

When the distributions that we are trying to couple are product distributions (i.e. $\mathbb{P}_1 = \mathbb{p}_1^{\otimes n}, \dots, \mathbb{P}_N = \mathbb{p}_N^{\otimes n}$), we can notice that any coupling $\mathfrak{q} \in \Pi(\mathbb{p}_1, \dots, \mathbb{p}_N)$ induces a coupling $\mathfrak{q}^{\otimes n} \in \Pi(\mathbb{P}_1, \dots, \mathbb{P}_N)$. Under this coupling, the Hamming distances between the pairs of marginals follow binomial distributions. For the rest of this article, we define the product (near) optimal coupling

$$\pi^\otimes(\mathbb{p}_1^{\otimes n}, \dots, \mathbb{p}_N^{\otimes n}) := \pi^\infty(\mathbb{p}_1, \dots, \mathbb{p}_N)^{\otimes n}.$$

Straightforward computations yield the following lemma.

Lemma 2. Let $\mathbb{P}_1 = \mathbb{p}_1^{\otimes n}, \dots, \mathbb{P}_N = \mathbb{p}_N^{\otimes n}$ be N product distributions on \mathcal{X}^n and $\mathfrak{q} \in \Pi(\mathbb{p}_1, \dots, \mathbb{p}_N)$. Consider any $1 \leq i, j \leq N$ and denote² $\delta_{i,j} := \mathbb{P}_{(X_1, \dots, X_N) \sim \mathfrak{q}}(X_i \neq X_j)$. We have:

$$\begin{aligned} \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathfrak{q}^{\otimes n}}(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)) &= n\delta_{i,j} \\ \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathfrak{q}^{\otimes n}}(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)^2) &= n^2\delta_{i,j}^2 + n\delta_{i,j}(1 - \delta_{i,j}) \leq n^2\delta_{i,j}^2 + n\delta_{i,j} \\ \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \mathfrak{q}^{\otimes n}}(e^{-\epsilon d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)}) &= (1 - (1 - e^{-\epsilon})\delta_{i,j})^n \geq e^{-n\epsilon\delta_{i,j}}. \end{aligned}$$

²not to be confused with the Kronecker symbol.

Note that $\delta_{i,j}$ directly depends on the coupling construction, but that with any of the ones presented above (applied to $\mathbb{P}_1, \dots, \mathbb{P}_N$), we always have $\forall i, j, \delta_{i,j} \leq 2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)$.

Putting the pieces together. Each of the coupling constructions presented above has its own merits and can be used to establish the quantitative lower-bounds establishing Theorem 1, Theorem 2, Theorem 3 and Theorem 4 (see details in Appendix C) as well as the fully worked-out examples given at the end of the introduction (see details in Appendix D, including a discussion on different regimes where privacy induces an estimation overhead).

Discussion: beyond the i.i.d. structure? Note that the techniques presented in this article can be used in non i.i.d. setups as well, by emulating the structure of the probability space in the coupling construction. For instance, consider an $m \times m$ stochastic matrix K , the kernel of a Markov chain kernel K on m states $\{1, \dots, m\}$. The column i represents the vector $P(\cdot|i)$ that gives the conditional probabilities of ending in the different states, knowing that the current state is i . We assume that initial distribution of the chain is uniform on the states. The objective is to build a differentially private test of K based on the observation $(X_t)_{1 \leq t \leq n+1}$, $X_t \in \{1, \dots, m\}$ of a trajectory of length $n+1$. Let us illustrate this with $m=2$: consider the two kernels $K = \begin{pmatrix} k_{1,1} & k_{1,2} \\ k_{2,1} & k_{2,2} \end{pmatrix}$ and $L = \begin{pmatrix} l_{1,1} & l_{1,2} \\ l_{2,1} & l_{2,2} \end{pmatrix}$, and their associated Markov chains M_K and M_L . We build a Markov kernel Q on the set of pairs of states of M_K and M_L , such that for any x_K state of M_K and any x_L state of M_L , $Q((\cdot, \cdot)|(x_K, x_L))$ is a coupling between $K(\cdot|x_K)$ and $L(\cdot|x_L)$. Let us take $Q((1,1)|(x_K, x_L)) = \min(k_{1,x_K}, l_{1,x_L})$, $Q((2,2)|(x_K, x_L)) = \min(k_{2,x_K}, l_{2,x_L})$, $Q((1,2)|(x_K, x_L)) = k_{1,x_K} - l_{1,x_L}$ if $k_{1,x_K} > l_{1,x_L}$ or 0 otherwise, and $Q((2,1)|(x_K, x_L)) = l_{1,x_L} - k_{1,x_K}$ if $l_{1,x_L} > k_{1,x_K}$ or 0 otherwise. We consider M_Q , the Markov chain that starts with the uniform distribution on the pairs of states of M_K and M_L , and has transition kernel Q . We observe that the probability distribution over pairs of sequences of length $n+1$ generated by M_Q is a coupling between the corresponding distributions over single sequences associated to M_1 and M_2 . Furthermore, in general, the structure of the probability space is Markovian and is not a product one (meaning that the generated trajectories are not i.i.d.).

By integrating Le Cam matching similarity function (Theorem 6 and Theorem 7) against this distribution on the pairs of trajectories, we obtain the following lower-bounds: Any ϵ -DP test that tries to discriminate M_K from M_L must have a type 1 or a type 2 error at least equal to $\frac{1}{2}(1 - (1 - e^{-\epsilon})\alpha)^n$, where $\alpha = \max\{|k_{1,1} - l_{1,1}|, |k_{1,1} - l_{1,2}|, |k_{2,1} - l_{1,1}|, |k_{2,1} - l_{2,1}|\}$. Similarly, any ρ -zCDP test that tries to discriminate M_L from M_K must have a type 1 or a type 2 error at least equal to $\frac{1}{2}(1 - n\alpha\sqrt{\rho/2})$.

When there are more than two Markov chains to test (say N), a similar coupling can be built by building a Markov chain on the N -tuples of states of the different Markov chains. The technicality is that one has to use Fact 4 instead of Fact 3 for coupling the transition probabilities. When there are more than two states, the construction is the same. The expressions of the total variations between the pairs of transition probabilities can however be more difficult.

4 Near Optimal Private Maximum Likelihood

In the different models presented in Appendix D and for many other parametric models, the statistician typically would like to consider the maximum likelihood estimator. Given X_1, \dots, X_n i.i.d. random variables of distribution \mathbb{P}_{θ^*} , the maximum likelihood estimator has value

$$\hat{\theta}_{\text{ML}} \in \arg \max_{\theta \in \Theta} \left\{ l(\theta) := \frac{1}{n} \sum_{i=1}^n f(X_i, \theta) \right\}, \quad (16)$$

where f is the log-likelihood. The parametric model with respect to a reference measure μ is thus

$$\forall X, \quad \frac{d\mathbb{P}_\theta}{d\mu}(X) := e^{f(X, \theta)}, \quad \theta \in \Theta,$$

where $\frac{d\mathbb{P}_\theta}{d\mu}$ is the Radon-Nikodym density of \mathbb{P}_θ with respect to μ and Θ is often a closed, convex subset of \mathbb{R}^d with nonempty interior. This setup covers, in particular, exponential families (Van der Vaart, 1998) with $f(X, \theta) = \theta^T T(X) - \ln(Z(\theta))$ associated with some statistic T and normalization factor $Z(\theta)$. This section

first presents a lower bound on the minimax risk for the private estimation in such parametric models and then studies the optimality properties of the Differentially Private Stochastic Gradient Langevin Dynamics (DP-SGLD) of Ryffel et al. (2022) for this specific task based on the existing upper bounds for this private convex optimizer.

4.1 On the regularity of f and the estimation complexity

First, we may assume that the parametric model is not degenerate in the sense that f satisfies

$$\forall \theta \in \Theta, \quad \int \nabla_{\theta} f(X, \theta) d\mathbb{P}_{\theta}(X) = 0. \quad (17)$$

This hypothesis is for instance satisfied in the Gaussian model presented previously. Indeed, in this case $\forall X, \nabla_{\theta} f(\theta + X, \theta) + \nabla_{\theta} f(\theta - X, \theta) = 0$ and $\forall X, \frac{d\mathbb{P}_{\theta}}{d\mu}(\theta + X) = \frac{d\mathbb{P}_{\theta}}{d\mu}(\theta - X)$. This hypothesis is more generally satisfied in the broader model of the exponential families (see Boucheron et al. (2019, Théorème 4.10)). Under such hypothesis, we have the following lemma which will allow to leverage Proposition 1:

Lemma 3. *If $(\mathbb{P}_{\theta})_{\theta \in \Theta}$ satisfies the property (17) and if f is concave and β -smooth in its second argument, then*

$$\forall \theta_1, \theta_2 \in \Theta, \text{KL}(\mathbb{P}_{\theta_1} \| \mathbb{P}_{\theta_2}) \leq \frac{\beta}{2} \|\theta_2 - \theta_1\|^2.$$

Note that the family $(\mathbb{P}_{\theta})_{\theta \in \Theta}$ directly depends on f . In particular, for the Gaussian Model, $\beta = \frac{1}{\sigma^2}$, we recover the classical upper bound on the KL divergence between multivariate normal distributions, which is in fact in this case, an equality.

Proof. Because of concavity in the second argument of f and the fact that it is β -smooth, we have the following result:

$$\forall \theta_1, \theta_2 \in \Theta, \forall x, \quad f(x, \theta_1) + \nabla_{\theta} f(x, \theta_1)^T (\theta_2 - \theta_1) \leq f(x, \theta_2) + \frac{\beta}{2} \|\theta_1 - \theta_2\|^2.$$

As a consequence,

$$\begin{aligned} \text{KL}(\mathbb{P}_{\theta_1} \| \mathbb{P}_{\theta_2}) &= \int \ln \left(\frac{d\mathbb{P}_{\theta_1}}{d\mathbb{P}_{\theta_2}} \right) d\mathbb{P}_{\theta_1} = \int (f(X, \theta_1) - f(X, \theta_2)) d\mathbb{P}_{\theta_1}(X) \\ &\leq \int \left(-\nabla_{\theta} f(X, \theta_1)^T (\theta_2 - \theta_1) + \frac{\beta}{2} \|\theta_1 - \theta_2\|^2 \right) d\mathbb{P}_{\theta_1}(X) \\ &\stackrel{(17)}{=} \int \frac{\beta}{2} \|\theta_1 - \theta_2\|^2 d\mathbb{P}_{\theta_1}(X) = \frac{\beta}{2} \|\theta_1 - \theta_2\|^2. \end{aligned}$$

□

We may apply Proposition 1 with $\gamma = \beta/2$ and we obtain that

$$\mathfrak{M}_n(\rho\text{-zCDP}, (\mathbb{P}_{\theta}^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2) = \Omega \left(\max \left\{ \frac{d}{n^2 \beta \rho}, \frac{d}{n \beta} \right\} \right) \quad (18)$$

Under the hypotheses of Proposition 1: d is big enough, ρ is small enough and the interior of the parameter space is big enough. In particular, this gives us a lower bound to compare any private estimator to.

4.2 Private maximum likelihood

In general, $\hat{\theta}_{\text{ML}}$ has no closed form formula. Even when it has some, the closed form formula usually does not respect differential privacy.

The problem (16) is typically addressed via numerical optimization: instead of considering its explicit maximum, a provably converging sequence is constructed. This requires some assumptions on the log-likelihood f . A convenient combination of hypotheses is that f is λ -strongly concave, β -smooth and L -Lipschitz in its

second argument: then, the stochastic gradient ascend algorithm converges rapidly to $\hat{\theta}_{\text{ML}}$ (Beck, 2017). Exponential families typically obey those requirements with $\beta := \sup_{\theta \in \Theta} \lambda_{\max}(C_\theta)$ and, $\lambda := \inf_{\theta \in \Theta} \lambda_{\min}(C_\theta)$ where $C_\theta := \text{Cov}_{X \sim \mathbb{P}_\theta}(T(X))$ and $\lambda_{\min}(C)$ (resp. $\lambda_{\max}(C)$) denotes the smallest (resp. largest) eigenvalue of a matrix C (see Boucheron et al. (2019, Théorème 4.10)).

The issue of privacy can be addressed directly in the optimization procedure. DP-SGD (Abadi et al., 2016) is an adaptation of the Stochastic Gradient Descent method where the gradient is first clipped and then noised. The privacy guarantees are based on the moment accountant method or on the composition of Renyi differential privacy (Mironov, 2017). The results are obtained under very general hypotheses on the objective function, but are based on a pessimistic scenario where an adversary may observe every gradient in the optimizer. Recent work based on Langevin diffusion (Chourasia et al., 2021; Ryffel et al., 2022) has adapted the Gradient Descent algorithm and the Stochastic Gradient Descent algorithm in order to have privacy guarantees with tighter utility bounds at the price of stronger hypotheses on the objective function which is required to have a compact domain and to be strongly convex.

Building on DP-SGLD by Ryffel et al. (2022), we consider its adaptation for maximum likelihood DP-SGML (Algorithm 1). For a batch $\mathcal{B} \subseteq \{1, \dots, n\}$, the batch log-likelihood is defined as

$$l_{\mathcal{B}}(\theta) := \frac{1}{\#\mathcal{B}} \sum_{i \in \mathcal{B}} f(X_i, \theta).$$

For a closed convex set Θ , Π_Θ refers to the projection onto Θ .

Data: X_1, \dots, X_n , f , step sizes $(\eta_k)_{k \geq 0}$, batch size m , noise variance σ^2 , initial parameter θ_0 , stopping time K .

for $k = 0, \dots, K - 1$ **do**

Sample batch \mathcal{B}_k from X_1, \dots, X_n with replacement of size m ;
 Compute $\nabla l_{\mathcal{B}_k}(\theta_k) = \frac{1}{\#\mathcal{B}_k} \sum_{i \in \mathcal{B}_k} \nabla_\theta f(X_i, \theta_k)$;
 Update parameter $\theta_{k+1} = \Pi_\Theta(\theta_k + \eta_k \nabla l_{\mathcal{B}_k}(\theta_k) + \sqrt{2\eta_k} \mathcal{N}(0, \sigma^2 I_d))$.

end

return θ_K

Algorithm 1: DP-SGML: Differentially Private Stochastic Gradient Maximum Likelihood

A choice of the parameters $(\eta_k)_{k \geq 0}$, σ^2 , θ_0 and K is suggested by the privacy-utility theorem Fact 5 which is a direct corollary of Ryffel et al. (2022).

Fact 5 (Utility and Privacy of Algorithm 1, Fixed Step Size). *Assume that f is λ -strongly concave, β -smooth and L -Lipschitz in its second argument on Θ . Consider any $\rho > 0$, an integer $n \geq 1$, a batch size m and set*

$$\sigma^2 := \frac{4L^2}{\rho\lambda n^2}, \quad K := \frac{2\beta}{\lambda} \ln\left(\frac{\rho n^2}{d}\right), \quad \xi^2 := \mathbb{E}_{\mathcal{B}}(\|\nabla l_{\mathcal{B}}(\theta_{\text{ML}})\|^2)$$

Given a collection \mathbf{X} of n arbitrary samples, consider $\mathfrak{M}(\mathbf{X}) = \theta_K$ obtained using DP-SGML with $\theta_0 \sim \Pi_\Theta\left(\mathcal{N}(0, \frac{2\sigma^2}{\lambda} I_d)\right)$ and constant step size $\eta = \frac{1}{2\beta}$. This mechanism satisfies ρ -zCDP. Moreover, if \mathbf{X} is such that the solution θ_{ML} of (16) is in the interior of Θ , then

$$\mathbb{E}(\|\theta_{\text{ML}} - \theta_K\|^2) = O\left(\frac{\beta d L^2}{\rho \lambda^3 n^2}\right) + \frac{\xi^2}{2\lambda^2}$$

where the expectation is with respect to initialization, random batch sampling, and noise addition in the parameter update step.

Indeed, the direct application of (Ryffel et al., 2022, Theorem 4.1) gives the privacy guarantee, and that

$$\mathbb{E}(l(\theta_{\text{ML}}) - l(\theta_K)) = O\left(\frac{\beta d L^2}{\rho \lambda^2 n^2}\right) + \frac{\xi^2}{4\lambda}.$$

Furthermore, by λ -strong concavity of l ,

$$l(\theta_{\text{ML}}) - l(\theta_K) \geq \nabla l(\theta_{\text{ML}})(\theta_K - \theta_{\text{ML}}) + \frac{\lambda}{2} \|\theta_K - \theta_{\text{ML}}\|^2$$

and since θ_{ML} is in the interior of Θ , $\nabla l(\theta_{\text{ML}}) = 0$ which concludes the proof. The term $\xi^2 := \mathbb{E}_{\mathcal{B}}(\|\nabla l_{\mathcal{B}}(\theta_{\text{ML}})\|^2)$ is due to the stochastic noise of the batch sampling. Indeed, even though $\nabla l(\theta_{\text{ML}}) = 0$, this is not necessarily the case when working on batches. This term depends on the batch size m and can be made arbitrarily small by choosing m large enough.

4.3 About minimax optimality

The quadratic risk of any (private or not) solver \mathfrak{M} can be decomposed (by the triangle inequality and since $(a+b)^2 \leq 2a^2 + 2b^2, \forall a, b \geq 0$) as:

$$\mathbb{E}(\|\theta^* - \mathfrak{M}(\mathbf{X})\|^2) \leq 2 \left(\mathbb{E}(\|\theta^* - \theta_{\text{ML}}\|^2) + \mathbb{E}(\|\theta_{\text{ML}} - \mathfrak{M}(\mathbf{X})\|^2) \right) \quad (19)$$

where the expectation is over the draw of \mathbf{X} and, in the case of a private solver, on the intrinsic randomness of \mathfrak{M} .

The first term in the right hand side of (19) only depends on the properties of the “ideal” maximum likelihood estimator in this parametric model. Under mild assumptions, it is asymptotically normal – for example, in exponential families (see Van der Vaart (1998, Theorem 4.6)): we have

$$\sqrt{n}(\theta^* - \theta_{\text{ML}}) \rightsquigarrow \mathcal{N}(0, C_{\theta^*}^{-1}),$$

where \rightsquigarrow refers to the convergence in distribution, and

$$\mathbb{E}(\|\theta^* - \theta_{\text{ML}}\|^2) = O\left(\frac{d}{n\lambda}\right). \quad (20)$$

The second term in (19) depends on the solver, which here can be controlled with Fact 5. As a consequence, the ratio between the error of estimation and the minimax risk which is lower-bounded in (18) can be bounded as follows:

$$\begin{aligned} & \frac{\mathbb{E}(\|\theta^* - \mathfrak{M}(\mathbf{X})\|^2)}{\mathfrak{M}_n(\rho\text{-zCDP}, (\mathbb{P}_{\theta}^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2)} \\ & \stackrel{(19) \& (18)}{=} O\left(\frac{\mathbb{E}(\|\theta^* - \theta_{\text{ML}}\|^2) + \mathbb{E}(\|\theta_{\text{ML}} - \mathfrak{M}(\mathbf{X})\|^2)}{\max\left\{\frac{d}{n^2\beta\rho}, \frac{d}{n\beta}\right\}}\right) \\ & = O\left(\frac{n\beta}{d}\mathbb{E}(\|\theta^* - \theta_{\text{ML}}\|^2) + \frac{n^2\beta\rho}{d}\mathbb{E}(\|\theta_{\text{ML}} - \mathfrak{M}(\mathbf{X})\|^2)\right). \end{aligned}$$

In particular, for the fixed step-size (see Fact 5), when θ_{ML} is in the interior of Θ and when the variance term due to the clipped gradient is negligible (i.e., when the batch size is big enough to have $\frac{\xi^2}{4\lambda} = O\left(\frac{\beta d L^2}{\rho \lambda^2 n^2}\right)$), the second term is $O\left(\frac{\beta^2 L^2}{\lambda^3}\right)$.

All in all, the ratio between the risk of DP-SGML for maximum likelihood in exponential families when the maximum likelihood estimator is in the interior of the search set is

$$\frac{\mathbb{E}(\|\theta^* - \mathfrak{M}(\mathbf{X})\|^2)}{\mathfrak{M}_n(\rho\text{-zCDP}, (\mathbb{P}_{\theta}^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2)} = O\left(\frac{\beta}{\lambda} + \frac{\beta^2 L^2}{\lambda^3}\right).$$

DP-SGML optimally captures the variation in the sample size n , in the privacy parameter ρ , and to some extent, in the dimensionality d (to some extent because even if d vanishes in the expressions, L , β and λ may vary with d). This proves what we call the near-minimax optimality of DP-SG(L)D for performing inference via maximum likelihood in a broad class of parametric models.

Acknowledgement

Aurélien Garivier acknowledges the support of the Project IDEXLYON of the University of Lyon, in the framework of the Programme Investissements d’Avenir (ANR-16-IDEX-0005), and Chaire SeqALO (ANR-20-CHIA-0020-01). This project was supported in part by the AllegroAssai ANR project ANR-19-CHIA-0009. Additionally, we thank the anonymous reviewers for their precious inputs and suggestions.

References

- Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (eds.), *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pp. 308–318. ACM, 2016. doi: 10.1145/2976749.2978318. URL <https://doi.org/10.1145/2976749.2978318>.
- John M Abowd. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2867–2867, 2018.
- Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private testing of identity and closeness of discrete distributions. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 6879–6891, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/7de32147a4f1055bed9e4faf3485a84d-Abstract.html>.
- Jayadev Acharya, Clement Canonne, Aditya Vikram Singh, and Himanshu Tyagi. Optimal rates for nonparametric density estimation under communication constraints. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 26754–26766. Curran Associates, Inc., 2021a. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/e1021d43911ca2c1845910d84f40aeae-Paper.pdf.
- Jayadev Acharya, Clément L. Canonne, Prathamesh Mayekar, and Himanshu Tyagi. Information-constrained optimization: can adaptive processing of gradients help? *CoRR*, abs/2104.00979, 2021b. URL <https://arxiv.org/abs/2104.00979>.
- Jayadev Acharya, Clément L. Canonne, Zuteng Sun, and Himanshu Tyagi. Unified lower bounds for interactive high-dimensional estimation under information constraints. *CoRR*, abs/2010.06562, 2021c. URL <https://arxiv.org/abs/2010.06562>.
- Jayadev Acharya, Clément L. Canonne, Cody Freitag, Ziteng Sun, and Himanshu Tyagi. Inference under information constraints iii: Local privacy constraints. *IEEE Journal on Selected Areas in Information Theory*, 2(1):253–267, 2021d. doi: 10.1109/JSAIT.2021.3053569. URL <https://doi.org/10.1109/JSAIT.2021.3053569>.
- Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private assouad, fano, and le cam. In Vitaly Feldman, Katrina Ligett, and Sivan Sabato (eds.), *Algorithmic Learning Theory, 16-19 March 2021, Virtual Conference, Worldwide*, volume 132 of *Proceedings of Machine Learning Research*, pp. 48–78. PMLR, 2021e. URL <http://proceedings.mlr.press/v132/acharya21a.html>.
- Omer Angel and Yinon Spinka. Pairwise optimal coupling of multiple random variables, 2021. URL <https://arxiv.org/pdf/1903.00632.pdf>.
- Hilal Asi and John C. Duchi. Near instance-optimality in differential privacy. *CoRR*, abs/2005.10630, 2020a. URL <https://arxiv.org/abs/2005.10630>.
- Hilal Asi and John C. Duchi. Instance-optimality in differential privacy via approximate inverse sensitivity mechanisms. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020b. URL <https://proceedings.neurips.cc/paper/2020/hash/a267f936e54d7c10a2bb70dbe6ad7a89-Abstract.html>.

-
- Patrice Assouad. Deux remarques sur l'estimation. *Comptes rendus des séances de l'Académie des sciences. Série 1, Mathématique*, 296(23):1021–1024, 1983.
- Lars Backstrom, Cynthia Dwork, and Jon M. Kleinberg. Wherefore art thou r3579x?: anonymized social networks, hidden patterns, and structural steganography. In Carey L. Williamson, Mary Ellen Zurko, Peter F. Patel-Schneider, and Prashant J. Shenoy (eds.), *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, pp. 181–190. ACM, 2007. doi: 10.1145/1242572.1242598. URL <https://doi.org/10.1145/1242572.1242598>.
- Leighton Pate Barnes, Yanjun Han, and Ayfer Ozgur. Fisher information for distributed estimation under a blackboard communication protocol. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 2704–2708, 2019. doi: 10.1109/ISIT.2019.8849821.
- Leighton Pate Barnes, Wei-Ning Chen, and Ayfer Ozgur. Fisher information under local differential privacy. *IEEE Journal on Selected Areas in Information Theory*, 1(3):645–659, 2020a. doi: 10.1109/JSAIT.2020.3039461. URL <https://doi.org/10.1109/JSAIT.2020.3039461>.
- Leighton Pate Barnes, Yanjun Han, and Ayfer Özgür. Lower bounds for learning distributions under communication constraints via fisher information. *Journal of Machine Learning Research*, 21:Paper No. 236, 30, 2020b. ISSN 1532-4435. URL <https://jmlr.csail.mit.edu/papers/volume21/19-737/19-737.pdf>.
- Björn Bebensee. Local differential privacy: a tutorial. *arXiv preprint arXiv:1907.11908*, 2019.
- Amir Beck. *First-Order Methods in Optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017. doi: 10.1137/1.9781611974997. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611974997>.
- Omri Ben-Eliezer, Dan Mikulincer, and Ilias Zadik. Archimedes meets privacy: On privately estimating quantiles in high dimensions under minimal assumptions. *arXiv preprint arXiv:2208.07438*, 2022.
- Peter J Bickel and Yaacov Ritov. Estimating integrated squared density derivatives: sharp best order of convergence estimates. *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 381–393, 1988.
- Sourav Biswas, Yihe Dong, Gautam Kamath, and Jonathan R. Ullman. Coinpress: Practical private mean and covariance estimation. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/a684ecccc76fc522773286a895bc8436-Abstract.html>.
- Adam Block, Ali Jadbabaie, Daniel Pfrommer, Max Simchowitz, and Russ Tedrake. Provable guarantees for generative behavior cloning: Bridging low-level stability and high-level behavior. *Advances in Neural Information Processing Systems*, 36, 2024.
- Boucheron, Boyer, and Ryder. Cours de statistiques, ens ulm, fimfa ens. *Lectures at the École Normale Supérieure*, 2019. URL <https://stephane-v-boucheron.fr/files/stats/notes-17-18.pdf>.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In Martin Hirt and Adam D. Smith (eds.), *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pp. 635–658, 2016. doi: 10.1007/978-3-662-53641-4_24. URL https://doi.org/10.1007/978-3-662-53641-4_24.
- Mark Bun and Thomas Steinke. Average-case averages: Private algorithms for smooth sensitivity and mean estimation. *Advances in Neural Information Processing Systems*, 32, 2019.
- Mark Bun, Gautam Kamath, Thomas Steinke, and Steven Z Wu. Private hypothesis selection. *Advances in Neural Information Processing Systems*, 32, 2019.

-
- Mark Bun, Marco Gaboardi, Max Hopkins, Russell Impagliazzo, Rex Lei, Toniann Pitassi, Satchit Sivakumar, and Jessica Sorrell. Stability is stable: Connections between replicability, privacy, and adaptive generalization. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pp. 520–527, 2023.
- Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *J. Mach. Learn. Res.*, 12:1069–1109, 2011. doi: 10.5555/1953048.2021036. URL <https://dl.acm.org/doi/10.5555/1953048.2021036>.
- Rishav Chourasia, Jiayuan Ye, and Reza Shokri. Differential privacy dynamics of langevin diffusion and noisy gradient descent. *Advances in Neural Information Processing Systems*, 34:14771–14781, 2021.
- Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. Privacy at scale: Local differential privacy in practice. In Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein (eds.), *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, pp. 1655–1658. ACM, 2018. doi: 10.1145/3183713.3197390. URL <https://doi.org/10.1145/3183713.3197390>.
- Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006. ISBN 978-0-471-24195-9. URL <http://www.elementsofinformationtheory.com/>.
- Luc Devroye. *A course in density estimation*. Birkhauser Boston Inc., 1987.
- Ilias Diakonikolas, Moritz Hardt, and Ludwig Schmidt. Differentially private learning of structured discrete distributions. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 2566–2574, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/2b3bf3eee2475e03885a110e9acaab61-Abstract.html>.
- Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 3571–3580, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/253614bbac999b38b5b60cae531c4969-Abstract.html>.
- Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In Frank Neven, Catriel Beeri, and Tova Milo (eds.), *Proceedings of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 9-12, 2003, San Diego, CA, USA*, pp. 202–210. ACM, 2003. doi: 10.1145/773153.773173. URL <https://doi.org/10.1145/773153.773173>.
- Jinshuo Dong, Aaron Roth, and Weijie J. Su. Gaussian differential privacy. *CoRR*, abs/1905.02383, 2019. URL <http://arxiv.org/abs/1905.02383>.
- Jinshuo Dong, David Durfee, and Ryan Rogers. Optimal differential privacy composition for exponential mechanisms. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 2597–2606. PMLR, 2020. URL <http://proceedings.mlr.press/v119/dong20a.html>.
- Wenxin Du, Canyon Foot, Monica Moniot, Andrew Bray, and Adam Groce. Differentially private confidence intervals. *CoRR*, abs/2001.02285, 2020. URL <http://arxiv.org/abs/2001.02285>.
- John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Local privacy and statistical minimax rates. In *51st Annual Allerton Conference on Communication, Control, and Computing, Allerton 2013, Allerton Park & Retreat Center, Monticello, IL, USA, October 2-4, 2013*, pp. 1592. IEEE, 2013. doi: 10.1109/Allerton.2013.6736718. URL <https://doi.org/10.1109/Allerton.2013.6736718>.
- John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Local privacy, data processing inequalities, and statistical minimax rates, 2014. URL <https://arxiv.org/abs/1302.3203>.

-
- Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In Michael Mitzenmacher (ed.), *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pp. 371–380. ACM, 2009. doi: 10.1145/1536414.1536466. URL <https://doi.org/10.1145/1536414.1536466>.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014. doi: 10.1561/04000000042. URL <https://doi.org/10.1561/04000000042>.
- Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin (eds.), *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pp. 265–284. Springer, 2006. doi: 10.1007/11681878_14. URL https://doi.org/10.1007/11681878_14.
- Niva Elkin-Koren, Uri Hacoheh, Roi Livni, and Shay Moran. Can copyright be reduced to privacy? *arXiv preprint arXiv:2305.14822*, 2023.
- Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: randomized aggregatable privacy-preserving ordinal response. In Gail-Joon Ahn, Moti Yung, and Ninghui Li (eds.), *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pp. 1054–1067. ACM, 2014. doi: 10.1145/2660267.2660348. URL <https://doi.org/10.1145/2660267.2660348>.
- Alireza Farhadi, MohammadTaghi Hajiaghayi, and Elaine Shi. Differentially private densest subgraph. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pp. 11581–11597. PMLR, 2022. URL <https://proceedings.mlr.press/v151/farhadi22a.html>.
- Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In Indrajit Ray, Ninghui Li, and Christopher Kruegel (eds.), *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, pp. 1322–1333. ACM, 2015. doi: 10.1145/2810103.2813677. URL <https://doi.org/10.1145/2810103.2813677>.
- Christophe Giraud. *Introduction to high-dimensional statistics*. Chapman and Hall/CRC, 2021. ISBN 9781003158745. doi: 10.1201/9781003158745.
- László Györfi, Michael Kohler, Adam Krzyzak, and Harro Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer series in statistics. Springer, 2002. ISBN 978-0-387-95441-7. doi: 10.1007/b97848. URL <https://doi.org/10.1007/b97848>.
- Rob Hall, Alessandro Rinaldo, and Larry A. Wasserman. Random differential privacy. *CoRR*, abs/1112.2680, 2011. URL <http://arxiv.org/abs/1112.2680>.
- Te Sun Han and Sergio Verdú. Generalizing the fano inequality. *IEEE Trans. Inf. Theory*, 40(4):1247–1251, 1994. doi: 10.1109/18.335943. URL <https://doi.org/10.1109/18.335943>.
- Nils Homer, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V Pearson, Dietrich A Stephan, Stanley F Nelson, and David W Craig. Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS Genet*, 4(8):e1000167, 2008.
- I. A. Ibragimov and R. Z. Khasminskii. *Statistical estimation : asymptotic theory / I. A. Ibragimov, R. Z. Hasminskii ; translated by Samuel Kotz*. Springer-Verlag New York, 1981. ISBN 3540905235 0387905235.

-
- Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In Francis R. Bach and David M. Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1376–1385. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/kairouz15.html>.
- Alkis Kalavasis, Amin Karbasi, Shay Moran, and Grigoris Velegkas. Statistical indistinguishability of learning algorithms. In *International Conference on Machine Learning*, pp. 15586–15622. PMLR, 2023.
- Gautam Kamath, Xingtu Liu, and Huanyu Zhang. Improved rates for differentially private stochastic convex optimization with heavy-tailed data. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 10633–10660. PMLR, 2022. URL <https://proceedings.mlr.press/v162/kamath22a.html>.
- Amin Karbasi, Grigoris Velegkas, Lin Yang, and Felix Zhou. Replicability in reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Vishesh Karwa and Salil P. Vadhan. Finite sample differentially private confidence intervals. In Anna R. Karlin (ed.), *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pp. 44:1–44:9. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi: 10.4230/LIPIcs.ITCS.2018.44. URL <https://doi.org/10.4230/LIPIcs.ITCS.2018.44>.
- Clément Lalanne, Clément Gastaud, Nicolas Grislain, Aurélien Garivier, and Rémi Gribonval. Private quantiles estimation in the presence of atoms. *CoRR*, abs/2202.08969, 2022. URL <https://arxiv.org/abs/2202.08969>.
- Torgny Lindvall. *Lectures on the Coupling Method*. Dover Publications Inc., 2003. ISBN 978-0486421452.
- Grigorios Loukides, Joshua C. Denny, and Bradley A. Malin. The disclosure of diagnosis codes can breach research participants’ privacy. *J. Am. Medical Informatics Assoc.*, 17(3):322–327, 2010. doi: 10.1136/jamia.2009.002725. URL <https://doi.org/10.1136/jamia.2009.002725>.
- Frank McSherry. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. *Commun. ACM*, 53(9):89–97, 2010. doi: 10.1145/1810891.1810916. URL <https://doi.org/10.1145/1810891.1810916>.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pp. 94–103. IEEE Computer Society, 2007. doi: 10.1109/FOCS.2007.41. URL <https://doi.org/10.1109/FOCS.2007.41>.
- Ilya Mironov. Rényi differential privacy. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pp. 263–275. IEEE Computer Society, 2017. doi: 10.1109/CSF.2017.11. URL <https://doi.org/10.1109/CSF.2017.11>.
- Arvind Narayanan and Vitaly Shmatikov. How to break anonymity of the netflix prize dataset. *CoRR*, abs/cs/0610105, 2006. URL <http://arxiv.org/abs/cs/0610105>.
- Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (S&P 2008), 18-21 May 2008, Oakland, California, USA*, pp. 111–125. IEEE Computer Society, 2008. doi: 10.1109/SP.2008.33. URL <https://doi.org/10.1109/SP.2008.33>.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. ISSN 1935-8237. doi: 10.1561/22000000073. URL <http://dx.doi.org/10.1561/22000000073>.
- Phillippe Rigollet and Jan-Christian Hütter. High dimensional statistics. *MIT lecture notes for course 18S997*, 2015. URL <https://math.mit.edu/~rigollet/PDFs/RigNotes17.pdf>.

-
- Benjamin I. P. Rubinstein, Peter L. Bartlett, Ling Huang, and Nina Taft. Learning in a large function space: Privacy-preserving mechanisms for SVM learning. *J. Priv. Confidentiality*, 4(1), 2012. doi: 10.29012/jpc.v4i1.612. URL <https://doi.org/10.29012/jpc.v4i1.612>.
- Théo Ryffel, Francis R. Bach, and David Pointcheval. Differential privacy guarantees for stochastic gradient langevin dynamics. *CoRR*, abs/2201.11980, 2022. URL <https://arxiv.org/abs/2201.11980>.
- Filippo Santambrogio. *Optimal Transport for Applied Mathematicians*. Birkhäuser Cham/Springer, 2016. doi: 10.1007/978-3-319-20828-2.
- Jonathan Scarlett and Volkan Cevher. An introductory guide to fano’s inequality with applications in statistical estimation. *CoRR*, abs/1901.00555, 2019. URL <http://arxiv.org/abs/1901.00555>.
- Adam D. Smith. Privacy-preserving statistical estimation with optimal convergence rates. In Lance Fortnow and Salil P. Vadhan (eds.), *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pp. 813–822. ACM, 2011. doi: 10.1145/1993636.1993743. URL <https://doi.org/10.1145/1993636.1993743>.
- Latanya Sweeney. Simple demographics often identify people uniquely. *Health (San Francisco)*, 671(2000): 1–34, 2000.
- Latanya Sweeney. k-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl. Based Syst.*, 10(5):557–570, 2002. doi: 10.1142/S0218488502001648. URL <https://doi.org/10.1142/S0218488502001648>.
- Youming Tao, Yulian Wu, Peng Zhao, and Di Wang. Optimal rates of (locally) differentially private heavy-tailed multi-armed bandits. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pp. 1546–1574. PMLR, 2022. URL <https://proceedings.mlr.press/v151/tao22a.html>.
- Abhradeep Guha Thakurta, Andrew H Vyrros, Umesh S Vaishampayan, Gaurav Kapoor, Julien Freudiger, Vivek Rangarajan Sridhar, and Doug Davidson. Learning new words. *Granted US Patents*, 9594741, 2017.
- Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer series in statistics. Springer, 2009. ISBN 978-0-387-79051-0. doi: 10.1007/b13794. URL <https://doi.org/10.1007/b13794>.
- A. W. Van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998. doi: 10.1017/CBO9780511802256.
- Tim van Erven and Peter Harremoës. Rényi divergence and kullback-leibler divergence. *IEEE Trans. Inf. Theory*, 60(7):3797–3820, 2014. doi: 10.1109/TIT.2014.2320500. URL <https://doi.org/10.1109/TIT.2014.2320500>.
- Isabel Wagner and David Eckhoff. Technical privacy metrics: A systematic survey. *ACM Comput. Surv.*, 51(3):57:1–57:38, 2018. doi: 10.1145/3168389. URL <https://doi.org/10.1145/3168389>.
- Larry A. Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010. doi: 10.1198/jasa.2009.tm08651. URL <https://doi.org/10.1198/jasa.2009.tm08651>.
- Mengmeng Yang, Lingjuan Lyu, Jun Zhao, Tianqing Zhu, and Kwok-Yan Lam. Local differential privacy and its applications: A comprehensive survey. *CoRR*, abs/2008.03686, 2020. URL <https://arxiv.org/abs/2008.03686>.

A Notation

In this article, the term *Fact* will be reserved for results that are directly borrowed from the literature. In contrast, *Lemma*, *Proposition*, *Theorem* and *Corollary* will be used as soon as the result requires some work from the existing literature. \mathbb{N} is the set of natural numbers starting at 0. $\#(S) \in \mathbb{N} \cup \{+\infty\}$ refers to the cardinality of a set S . Bold letters will be used for vectors, and their non-bold counterparts with $i \in \mathbb{N} \setminus \{0\}$ as subscript will refer to their i -th entry. For instance, $\mathbf{X} = (X_1, \dots, X_n)$. Considering any set \mathcal{X} and $n \in \mathbb{N} \setminus \{0\}$, $d_{\text{ham}}(\cdot, \cdot) : \mathcal{X}^n \times \mathcal{X}^n \rightarrow \mathbb{N}$ is the Hamming distance defined as

$$d_{\text{ham}}((X_1, \dots, X_n), (Y_1, \dots, Y_n)) := \sum_{i=1}^n 1_{X_i \neq Y_i}.$$

The notation $\Pi(\mathbb{P}_1, \dots, \mathbb{P}_n)$ is used to refer to the set of distributions that have $\mathbb{P}_1, \dots, \mathbb{P}_n$ as their marginals. Any such distribution is called a *coupling* between/of $\mathbb{P}_1, \dots, \mathbb{P}_n$. For two distributions $\mathbb{P}_1, \mathbb{P}_2$, the notation $\mathbb{P}_1 \ll \mathbb{P}_2$ means that \mathbb{P}_1 is absolutely continuous with respect to \mathbb{P}_2 , i.e. that for any S measurable,

$$\mathbb{P}_2(S) = 0 \implies \mathbb{P}_1(S) = 0.$$

For two probability distributions \mathbb{P} and \mathbb{Q} on the same space, the total variation between \mathbb{P} and \mathbb{Q} is

$$\text{TV}(\mathbb{P}, \mathbb{Q}) := \sup_{S \text{ measurable}} |\mathbb{P}(S) - \mathbb{Q}(S)|.$$

Note that it is a symmetric function of \mathbb{P} and \mathbb{Q} . For two probability distributions \mathbb{P} and \mathbb{Q} on the same space with \mathbb{P} that is absolutely continuous with respect to \mathbb{Q} , the KL divergence between \mathbb{P} and \mathbb{Q} is

$$\text{KL}(\mathbb{P} \parallel \mathbb{Q}) := \int \ln \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right) d\mathbb{P}.$$

It is always positive but may be infinite. Furthermore, whenever we have access to a σ -finite measure μ such that \mathbb{Q} is absolutely continuous with respect to μ (for instance $\mu = \mathbb{P} + \mathbb{Q}$), if we note $\pi_{\mathbb{P}}$ and $\pi_{\mathbb{Q}}$ the Radon-Nikodym densities of \mathbb{P} and \mathbb{Q} respectively with respect to μ , then

$$\text{KL}(\mathbb{P} \parallel \mathbb{Q}) = \int \ln \left(\frac{\pi_{\mathbb{P}}}{\pi_{\mathbb{Q}}} \right) \pi_{\mathbb{P}} d\mu.$$

For two probability distributions \mathbb{P} and \mathbb{Q} on the same space with \mathbb{P} that is absolutely continuous with respect to \mathbb{Q} , and $\alpha \in [1, +\infty)$, the Rényi divergence of level α between \mathbb{P} and \mathbb{Q} is $D_{\alpha}(\mathbb{P} \parallel \mathbb{Q}) := \text{KL}(\mathbb{P} \parallel \mathbb{Q})$ if $\alpha = 1$ or

$$D_{\alpha}(\mathbb{P} \parallel \mathbb{Q}) := \frac{1}{\alpha - 1} \ln \int \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right)^{\alpha - 1} d\mathbb{Q}$$

when $\alpha > 1$. Furthermore, whenever we have access to a σ -finite measure μ such that \mathbb{Q} is absolutely continuous with respect to μ (for instance $\mu = \mathbb{P} + \mathbb{Q}$), if we note $\pi_{\mathbb{P}}$ and $\pi_{\mathbb{Q}}$ the Radon-Nikodym densities of \mathbb{P} and \mathbb{Q} respectively with respect to μ ,

$$D_{\alpha}(\mathbb{P} \parallel \mathbb{Q}) = \frac{1}{\alpha - 1} \ln \int \left(\frac{\pi_{\mathbb{P}}}{\pi_{\mathbb{Q}}} \right)^{\alpha - 1} \pi_{\mathbb{Q}} d\mu$$

when $\alpha > 1$. For brevity of notations, when applied to random variables, information theoretic quantities such as $\text{TV}(\cdot, \cdot)$, $\text{KL}(\cdot \parallel \cdot)$, \dots are to be understood as applied to the *probability distributions* of these random variables. For two series $(a_n)_{n \in \mathbb{N}} \in \mathbb{R}^{\mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}} \in \mathbb{R}^{\mathbb{N}}$, we use the following notations for asymptotic comparisons:

- $a_n = o(b_n)$ or equivalently $a_n \ll b_n$ when there exists $(c_n)_{n \in \mathbb{N}} \in \mathbb{R}_{+,*}^{\mathbb{N}}$ such that $c_n \xrightarrow[n \rightarrow +\infty]{} 0$ and

$$\exists n_0 \in \mathbb{N}, \forall n \geq n_0, \quad a_n = c_n b_n.$$

- $a_n = \Omega(b_n)$ when there exist $M > 0$ and $(c_n)_{n \in \mathbb{N}} \in [M, +\infty)^{\mathbb{N}}$ such that

$$\exists n_0 \in \mathbb{N}, \forall n \geq n_0, \quad a_n = c_n b_n .$$

- $a_n = \Theta(b_n)$ when $a_n = \Omega(b_n)$ and $b_n = \Omega(a_n)$.

These notations may be used in addition to other notations of the problem when there is no ambiguity, for instance when Θ also refers to the parameter space later. All the different parameters are considered as series of n , the sample size. For $x \in (0, +\infty)$, $\lceil x \rceil := \inf_{y \in \mathbb{N} \cap [x, +\infty)} y$. For a differentiable $f : E \rightarrow \mathbb{R}$ where E a Euclidean space, $\nabla f(x)$ refers to its gradient at the point x . When sub-scripted by a subset of variables, $\nabla_{\theta} f(x)$ refers to the gradient at the point x of f restricted to the variables in θ . We say that a function is L -smooth if it is differentiable and if its gradient is L -Lipschitz. A function $f : \Theta \rightarrow \mathbb{R}$ where Θ is convex is said to be λ -strongly convex if

$$\forall x, y \in \Theta, f(y) \geq f(x) + \nabla f(x)^T (y - x) + \frac{\lambda}{2} \|y - x\|^2 .$$

Furthermore, a function is said to be λ -strongly concave if its opposite is λ -strongly convex.

B Similarity functions

B.1 The case of (ϵ, δ) -differential privacy

(ϵ, δ) -differential privacy allows to compare conditional distributions for datasets depending on their Hamming distance. In particular, characterizing the pushforward of a distribution by a private mechanism is not an easy task. We overtake that difficulty with a technique that we call *anchoring*. Informally, an anchor is a function that, given multiple datasets, decides a common dataset to exploit so called *group privacy* of (ϵ, δ) -DP mechanisms and to give numerically tractable results.

Fact 6 ((ϵ, δ) -DP Group Privacy). *Given $\epsilon \in \mathbb{R}_{+*}$ and $\delta \in [0, 1)$, if a randomized mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ is (ϵ, δ) -differentially private, then, for all $\mathbf{X}, \mathbf{Y} \in \mathcal{X}^n$ and all measurable $S \subseteq \text{codom}(\mathfrak{M})$, we have*

$$\mathbb{P}_{\mathfrak{M}}(\mathfrak{M}(\mathbf{X}) \in S) \leq e^{\epsilon d_{\text{ham}}(\mathbf{X}, \mathbf{Y})} \mathbb{P}_{\mathfrak{M}}(\mathfrak{M}(\mathbf{Y}) \in S) + \delta d_{\text{ham}}(\mathbf{X}, \mathbf{Y}) e^{\epsilon(d_{\text{ham}}(\mathbf{X}, \mathbf{Y}) - 1)} .$$

B.1.1 Global Anchoring

The first type of anchor is a global anchor, where all the marginal datasets are compared to the same one.

Lemma 4 (Global Anchoring). *Consider an (ϵ, δ) -DP mechanism \mathfrak{M} , a test function*

$\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$, *and datasets $\mathbf{X}_1, \dots, \mathbf{X}_N \in \mathcal{X}^n$. For any anchor function $\Lambda : (\mathcal{X}^n)^N \rightarrow \mathcal{X}^n$, we have*

$$\frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq \frac{N-1}{N} e^{-\epsilon \max_i d_{\text{ham}}(\mathbf{X}_i, \Lambda)} - e^{-\epsilon} \delta \max_i d_{\text{ham}}(\mathbf{X}_i, \Lambda)$$

where Λ is a shorthand for $\Lambda(\mathbf{X}_1, \dots, \mathbf{X}_N)$.

Proof. By the group privacy property (see Fact 6), we have for each $i \in \{1, \dots, N\}$

$$\mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq e^{-\epsilon d_{\text{ham}}(\mathbf{X}_i, \Lambda)} \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\Lambda)) \neq i) - e^{-\epsilon} \delta d_{\text{ham}}(\mathbf{X}_i, \Lambda) .$$

As a result,

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_i)) \neq i) \\
& \geq \frac{1}{N} \left(\sum_{i=1}^N e^{-\epsilon d_{\text{ham}}(\mathbf{X}_i, \Lambda)} \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\Lambda)) \neq i) - e^{-\epsilon} \delta d_{\text{ham}} (\mathbf{X}_i, \Lambda) \right) \\
& \geq \frac{1}{N} \left(e^{-\epsilon \max_i d_{\text{ham}}(\mathbf{X}_i, \Lambda)} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\Lambda)) \neq i) \right. \\
& \quad \left. - N e^{-\epsilon} \delta \max_i d_{\text{ham}} (\mathbf{X}_i, \Lambda) \right) \\
& = \frac{N-1}{N} e^{-\epsilon \max_i d_{\text{ham}}(\mathbf{X}_i, \Lambda)} - e^{-\epsilon} \delta \max_i d_{\text{ham}} (\mathbf{X}_i, \Lambda) ,
\end{aligned}$$

where we used $\sum_{i=1}^N \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\Lambda)) \neq i) = \sum_{i=1}^N (1 - \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\Lambda)) = i)) = N - 1$ to get the last equality. \square

Remark 1 ((ϵ, δ) -DP Le Cam Matching). *When we have to find an anchor between only two datasets, we can design it optimally. Considering any $\mathbf{X}_1, \mathbf{X}_2 \in \mathcal{X}^n$, by definition these datasets disagree on exactly $d_{\text{ham}} (\mathbf{X}_1, \mathbf{X}_2)$ entries. The projection anchor $\Lambda = \Lambda_j$, $j \in \{1, 2\}$ consists in anchoring both \mathbf{X}_1 and \mathbf{X}_2 to \mathbf{X}_j . Consequently, we have $\max \{d_{\text{ham}} (\mathbf{X}_1, \Lambda), d_{\text{ham}} (\mathbf{X}_2, \Lambda)\} = d_{\text{ham}} (\mathbf{X}_1, \mathbf{X}_2)$. If instead we allocate in the anchor Λ half of the disagreeing components to \mathbf{X}_1 and the other half to \mathbf{X}_2 , we get an anchor that satisfies*

$$\max \{d_{\text{ham}} (\mathbf{X}_1, \Lambda), d_{\text{ham}} (\mathbf{X}_2, \Lambda)\} = \lceil d_{\text{ham}} (\mathbf{X}_1, \mathbf{X}_2) / 2 \rceil .$$

Furthermore, one can check that no anchor can achieve a better bound. With this new anchor, the direct application of Lemma 4 yields

$$\begin{aligned}
& \frac{1}{2} (\mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_1)) \neq 1) + \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_2)) \neq 2)) \\
& \geq \frac{1}{2} e^{-\epsilon \lceil d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2) / 2 \rceil} - e^{-\epsilon} \delta \lceil d_{\text{ham}} (\mathbf{X}_1, \mathbf{X}_2) / 2 \rceil .
\end{aligned} \tag{21}$$

B.1.2 Pairwise Anchoring

The fact that one needs to control the maximum of the hamming distances between a single anchor and the marginals might be prohibitive. We give here a symmetrized version that only requires to control the hamming distances between the pairs of marginals.

Lemma 5 (Pairwise Anchoring). *Under the assumptions of Lemma 4 we have*

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_i)) \neq i) \\
& \geq \frac{1}{2N^2} \sum_{i=1}^N \sum_{j=1}^N \left(e^{-\epsilon \lceil d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j) / 2 \rceil} - 2e^{-\epsilon} \delta \lceil d_{\text{ham}} (\mathbf{X}_i, \mathbf{X}_j) / 2 \rceil \right) .
\end{aligned}$$

Proof. First we observe that

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_i)) \neq i) \\
& = \frac{1}{2N^2} \sum_{i=1}^N \sum_{j=1}^N (\mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_i)) \neq i) + \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_j)) \neq j)) .
\end{aligned}$$

We then consider the two-point anchor defined in Remark 1 and get using (21) that for every $1 \leq i, j \leq N$,

$$\begin{aligned} & \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) + \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_j)) \neq j) \\ & \geq e^{-\epsilon \lceil d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)/2 \rceil} - 2e^{-\epsilon \delta \lceil d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)/2 \rceil}. \end{aligned}$$

□

B.1.3 The special case of $(\epsilon, 0)$ -DP

The following lemma yields a bound on the KL divergence between the output distributions of an $(\epsilon, 0)$ -DP mechanism applied to different datasets.

Lemma 6. *If a randomized mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ is $(\epsilon, 0)$ -DP, then*

$$\forall \mathbf{X}, \mathbf{Y} \in \mathcal{X}^n, \quad \frac{d\mathbb{P}_{\mathfrak{M}(\mathbf{X})}}{d\mathbb{P}_{\mathfrak{M}(\mathbf{Y})}} \leq e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon}, \quad \mathbb{P}_{\mathfrak{M}(\mathbf{X})} - \text{almost surely}$$

where $\frac{d\mathbb{P}_{\mathfrak{M}(\mathbf{X})}}{d\mathbb{P}_{\mathfrak{M}(\mathbf{Y})}}$ is the Radon-Nikodym density of the distribution of the output of the mechanism with input \mathbf{X} , with respect to the distribution of the output of the mechanism with input \mathbf{Y} . As a consequence,

$$\forall \mathbf{X}, \mathbf{Y} \in \mathcal{X}^n, \quad \text{KL}(\mathfrak{M}(\mathbf{X}) \parallel \mathfrak{M}(\mathbf{Y})) \leq \epsilon d_{\text{ham}}(\mathbf{X}, \mathbf{Y}).$$

Proof. By the group privacy property (see Fact 6), it is clear that the measurable sets of null measure for $\mathbb{P}_{\mathfrak{M}(\mathbf{X})}$ are exactly the measurable sets of null measure for $\mathbb{P}_{\mathfrak{M}(\mathbf{Y})}$. In particular, $\mathbb{P}_{\mathfrak{M}(\mathbf{X})} \ll \mathbb{P}_{\mathfrak{M}(\mathbf{Y})}$ and hence $p := \frac{d\mathbb{P}_{\mathfrak{M}(\mathbf{X})}}{d\mathbb{P}_{\mathfrak{M}(\mathbf{Y})}}$ exists. By group privacy property again for each measurable set $S \subseteq \text{codom}(\mathfrak{M})$ we have

$$\begin{aligned} \mathbb{P}_{\mathfrak{M}(\mathbf{Y})}(S) & \geq e^{-\epsilon d_{\text{ham}}(\mathbf{X}, \mathbf{Y})} \mathbb{P}_{\mathfrak{M}(\mathbf{X})}(S) \\ & = e^{-\epsilon d_{\text{ham}}(\mathbf{X}, \mathbf{Y})} \int_S p d\mathbb{P}_{\mathfrak{M}(\mathbf{Y})} \\ & \geq e^{-\epsilon d_{\text{ham}}(\mathbf{X}, \mathbf{Y})} \left(\inf_S p \right) \mathbb{P}_{\mathfrak{M}(\mathbf{Y})}(S). \end{aligned}$$

So, for each measurable set S ,

$$\mathbb{P}_{\mathfrak{M}(\mathbf{Y})}(S) > 0 \implies \inf_S p \leq e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon}.$$

Furthermore, p is measurable for the Borel σ -algebra of \mathbb{R} . In particular, for any $n \in \mathbb{N}_*$, $p^{-1}([e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon} + \frac{1}{n}, +\infty))$ is measurable. As a consequence,

$$\forall n \in \mathbb{N}_*, \quad \mathbb{P}_{\mathfrak{M}(\mathbf{Y})} \left(p^{-1} \left(\left[e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon} + \frac{1}{n}, +\infty \right] \right) \right) = 0$$

and then

$$\begin{aligned} \mathbb{P}_{\mathfrak{M}(\mathbf{Y})} \left(p^{-1} \left(\left(e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon}, +\infty \right) \right) \right) & = \mathbb{P}_{\mathfrak{M}(\mathbf{Y})} \left(p^{-1} \left(\bigcup_{n \in \mathbb{N}_*} \left[e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon} + \frac{1}{n}, +\infty \right] \right) \right) \\ & = \mathbb{P}_{\mathfrak{M}(\mathbf{Y})} \left(\bigcup_{n \in \mathbb{N}_*} p^{-1} \left(\left[e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon} + \frac{1}{n}, +\infty \right] \right) \right) \\ & \leq \sum_{n \in \mathbb{N}_*} \mathbb{P}_{\mathfrak{M}(\mathbf{Y})} \left(p^{-1} \left(\left[e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon} + \frac{1}{n}, +\infty \right] \right) \right) \\ & = 0 \end{aligned}$$

which proves that $\frac{d\mathbb{P}_{\mathfrak{M}(\mathbf{X})}}{d\mathbb{P}_{\mathfrak{M}(\mathbf{Y})}} \leq e^{d_{\text{ham}}(\mathbf{X}, \mathbf{Y})\epsilon}$, $\mathbb{P}_{\mathfrak{M}(\mathbf{Y})}$ -almost surely, which is also the case $\mathbb{P}_{\mathfrak{M}(\mathbf{X})}$ -almost surely, thanks to the first remark of the proof. The result about the KL divergence is a direct consequence of this inequality. □

In particular, this result allows us to apply Fano's lemma in order to obtain a similarity function that is based on anchoring the conditional distributions rather than the marginals. I.e., given, $\mathbf{X}_1, \dots, \mathbf{X}_N \in \mathcal{X}^n$, $\mathbb{P}_{\mathfrak{M}(\mathbf{X}_1)}, \dots, \mathbb{P}_{\mathfrak{M}(\mathbf{X}_N)}$ are anchored to $\frac{1}{N} \sum_{j=1}^N \mathbb{P}_{\mathfrak{M}(\mathbf{X}_j)}$.

Lemma 7 ($(\epsilon, 0)$ -DP Fano Matching). *Let $\mathbf{X}_1, \dots, \mathbf{X}_N \in \mathcal{X}^n$ and $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$,*

$$\frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq 1 - \frac{1 + \frac{\epsilon}{N^2} \sum_{i=1}^N \sum_{j=1}^N d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)}{\ln(N)}.$$

Proof. By Fano's lemma (see Fact 2),

$$\frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq 1 - \frac{1 + \frac{1}{N} \sum_{i=1}^N \text{KL}\left(\mathbb{P}_{\mathfrak{M}(\mathbf{X}_i)} \parallel \frac{1}{N} \sum_{j=1}^N \mathbb{P}_{\mathfrak{M}(\mathbf{X}_j)}\right)}{\ln(N)}.$$

By convexity of the KL divergence with respect to its second argument (see van Erven & Harremoës (2014, Theorem 12)), it follows that

$$\frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq 1 - \frac{1 + \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \text{KL}\left(\mathbb{P}_{\mathfrak{M}(\mathbf{X}_i)} \parallel \mathbb{P}_{\mathfrak{M}(\mathbf{X}_j)}\right)}{\ln(N)}. \quad (22)$$

An application of Lemma 6 concludes the proof. \square

B.2 The case of ρ -zero concentrated differential privacy

For ρ -zero concentrated differential privacy, the fact that the definition uses information theoretic quantities makes things easier than with the traditional definition of privacy. In particular, the anchoring technique happens implicitly on the distributions rather than on the marginals (similarly as with the $(\epsilon, 0)$ -DP case). Again, the notion of *group privacy* is central in our proofs.

Fact 7 (ρ -zCDP Group Privacy (Bun & Steinke, 2016, Proposition 27)). *Let $\rho \in \mathbb{R}_{+*}$, if a randomized mechanism $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom}(\mathfrak{M})$ is ρ -zero concentrated differentially private, then, for any $\mathbf{X}, \mathbf{Y} \in \mathcal{X}^n$ and for any $\alpha \in (1, \infty)$,*

$$D_{\alpha}(\mathfrak{M}(\mathbf{X}) \parallel \mathfrak{M}(\mathbf{Y})) \leq \rho d_{\text{ham}}(\mathbf{X}, \mathbf{Y})^2 \alpha.$$

Lemma 8 (ρ -zCDP Le Cam Matching). *Consider a ρ -zCDP mechanism \mathfrak{M} , a test function $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, 2\}$, and two datasets $\mathbf{X}_1, \mathbf{X}_2 \in \mathcal{X}^n$. We have*

$$\frac{1}{2} \sum_{i=1}^2 \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq \frac{1}{2} \left(1 - \sqrt{\rho/2} d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2)\right).$$

Proof. By the Neyman-Pearson lemma (see Fact 1),

$$\frac{1}{2} \sum_{i=1}^2 \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq \frac{1}{2} (1 - \text{TV}(\mathfrak{M}(\mathbf{X}_1), \mathfrak{M}(\mathbf{X}_2))).$$

By Pinsker's lemma (see Tsybakov (2009, Lemma 2.5)), $\text{TV}(\mathbb{P}, \mathbb{Q}) \leq \sqrt{\text{KL}(\mathbb{P} \parallel \mathbb{Q})/2}$, and hence

$$\begin{aligned} \frac{1}{2} \sum_{i=1}^2 \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) &\geq \frac{1}{2} \left(1 - \sqrt{\text{KL}(\mathfrak{M}(\mathbf{X}_1) \parallel \mathfrak{M}(\mathbf{X}_2))/2}\right) \\ &= \frac{1}{2} \left(1 - \sqrt{D_1(\mathfrak{M}(\mathbf{X}_1) \parallel \mathfrak{M}(\mathbf{X}_2))/2}\right). \end{aligned}$$

Since the Renyi divergence between a given pair of distributions $D_{\alpha}(\cdot \parallel \cdot)$ is non-decreasing in α (see van Erven & Harremoës (2014, Theorem 3)), we obtain for any $\alpha \in (1, +\infty)$, s

$$\frac{1}{2} \sum_{i=1}^2 \mathbb{P}_{\mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X}_i)) \neq i) \geq \frac{1}{2} \left(1 - \sqrt{D_{\alpha}(\mathfrak{M}(\mathbf{X}_1) \parallel \mathfrak{M}(\mathbf{X}_2))/2}\right).$$

Eventually, we obtain using group privacy (see Fact 7) that

$$\frac{1}{2} \sum_{i=1}^2 \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_i)) \neq i) \geq \frac{1}{2} \left(1 - \sqrt{\rho\alpha/2} d_{\text{ham}} (\mathbf{X}_1, \mathbf{X}_2) \right) .$$

The supremum of the right hand side over $\alpha \in (1, +\infty)$ yields the result. \square

We also obtain a zero concentrated DP version of the Fano matching method that we introduced previously for $(\epsilon, 0)$ -DP.

Lemma 9 (ρ -zCDP Fano Matching). *Consider a ρ -zCDP mechanism \mathfrak{M} , a test function $\Psi := \text{codom} (\mathfrak{M}) \rightarrow \{1, \dots, N\}$, and datasets $\mathbf{X}_1, \dots, \mathbf{X}_N \in \mathcal{X}^n$. We have*

$$\frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_i)) \neq i) \geq 1 - \frac{1 + \frac{\rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N d_{\text{ham}} (\mathbf{X}_i, \mathbf{X}_j)^2}{\ln(N)} .$$

Proof. By the inequality (22) established in the proof of Lemma 7, and using again the fact that $D_\alpha (\cdot \| \cdot)$ is non-decreasing in α (see van Erven & Harremoës (2014, Theorem 3)), as well as the group privacy property (see Fact 7), we obtain that for any $\alpha \in (1, +\infty)$,

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \mathbb{P}_{\mathfrak{M}} (\Psi (\mathfrak{M} (\mathbf{X}_i)) \neq i) &\geq 1 - \frac{1 + \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \text{KL} (\mathbb{P}_{\mathfrak{M}(\mathbf{X}_i)} \| \mathbb{P}_{\mathfrak{M}(\mathbf{X}_j)})}{\ln(N)} \\ &\geq 1 - \frac{1 + \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N D_\alpha (\mathbb{P}_{\mathfrak{M}(\mathbf{X}_i)} \| \mathbb{P}_{\mathfrak{M}(\mathbf{X}_j)})}{\ln(N)} . \\ &\geq 1 - \frac{1 + \frac{\rho\alpha}{N^2} \sum_{i=1}^N \sum_{j=1}^N d_{\text{ham}} (\mathbf{X}_i, \mathbf{X}_j)^2}{\ln(N)} . \end{aligned}$$

The supremum of the right-hand side over $\alpha \in (1, +\infty)$ yields the result. \square

C Quantitative lower bounds

In this subsection, we finally put the pieces together in order to obtain quantitative lower bounds on (11). This subsection serves as a joint proof for Theorem 1, Theorem 2, Theorem 3 and Theorem 4.

Immediate results on the private minimax risk. A usual estimator (i.e. a measurable function of the samples) $\hat{\theta}$ may be viewed as randomized and almost surely constant to $\hat{\theta}$ (i.e. $\forall \mathbf{X}, \mathfrak{M}(\mathbf{X}) := \hat{\theta}(\mathbf{X})$ almost surely). As a result, it is clear that the private minimax risk is always bigger than the non-private one. For distributional tests, the result is not so obvious, and we give the following general purpose lemma that ensures that Fano's and Le Cam's regular inequalities still hold.

Lemma 10. *Let $(\mathbb{P}_i)_{i \in \{1, \dots, N\}}$ be a family of probability distributions on \mathcal{X}^n and let $\mathfrak{M} : \mathcal{X}^n \rightarrow \text{codom} (\mathfrak{M})$ be a randomized mechanism,*

$$\inf_{\Psi: \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}} \sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi (\mathfrak{M}(\mathbf{X})) \neq i) \geq \inf_{\Psi: \mathcal{X}^N \rightarrow \{1, \dots, N\}} \sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi (\mathbf{X}) \neq i) .$$

In particular, the inequalities in Le Cam's lemma (see Fact 1) or Fano's lemma (see Fact 2) still hold when the test function Ψ is fed with an input $\mathfrak{M}(\mathbf{X}) \in \text{codom} (\mathfrak{M})$ instead of an input $\mathbf{X} \in \mathcal{X}^n$.

Proof. Let $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$ be a test function. Then,

$$\begin{aligned}
\sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) &= \sum_{i=1}^N \int \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) d\mathbb{P}_{\mathfrak{M}} \\
&= \int \sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} ((\Psi \circ \mathfrak{M})(\mathbf{X}) \neq i) d\mathbb{P}_{\mathfrak{M}} \\
&\geq \int \inf_{\Psi': \mathcal{X}^N \rightarrow \{1, \dots, N\}} \sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi'(\mathbf{X}) \neq i) d\mathbb{P}_{\mathfrak{M}} \\
&= \inf_{\Psi': \mathcal{X}^N \rightarrow \{1, \dots, N\}} \sum_{i=1}^N \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i} (\Psi'(\mathbf{X}) \neq i) .
\end{aligned}$$

Taking the infimum over $\Psi : \text{codom}(\mathfrak{M}) \rightarrow \{1, \dots, N\}$ concludes the proof. \square

The case of two hypotheses ($N = 2$). At first, we look at the implications of couplings between pairs of distributions. Given \mathbb{P}_1 and \mathbb{P}_2 distributions on \mathcal{X}^n , a direct implication of Lemma 10 and of Le Cam's lemma (see Fact 1) is that independently on the privacy condition \mathcal{C} imposed on \mathfrak{M} ,

$$\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \frac{1}{2} (1 - \text{TV}(\mathbb{P}_1, \mathbb{P}_2)) .$$

This is the first ingredient in the proof of Theorem 1 and Theorem 2 that we now detail.

Proof of Theorem 1: When \mathfrak{M} is (ϵ, δ) -DP, the generic bound of Theorem 5 applied with the Le Cam matching technique described in Theorem 6 and the coupling $\pi^\infty(\mathbb{P}_1, \mathbb{P}_2)$ leads to

$$\begin{aligned}
\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) &\geq \frac{1}{2} \mathbb{E}_{(\mathbf{X}_1, \mathbf{X}_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)} \left(e^{-\epsilon d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2)} \right) \\
&\quad - e^{-\epsilon \delta} \mathbb{E}_{(\mathbf{X}_1, \mathbf{X}_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)} (d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2)) \\
&\stackrel{\text{Lemma 1}}{\geq} \frac{1}{2} \left(1 - (1 - e^{-n\epsilon}) \Delta_{1,2} \right) - e^{-\epsilon \delta} n \Delta_{1,2} \\
&= \frac{1}{2} \left(1 - (1 - e^{-n\epsilon} + 2ne^{-\epsilon \delta}) \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right) .
\end{aligned}$$

where in the second line we denote $\Delta_{1,2} := \mathbb{P}_{(\mathbf{X}_1, \mathbf{X}_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)} (\mathbf{X}_1 \neq \mathbf{X}_2)$ and in the last line we use that $\Delta_{1,2} = \text{TV}(\mathbb{P}_1, \mathbb{P}_2)$ with the chosen coupling. Similarly, in the case of product distributions, with the same matching but $\pi^\otimes(\mathbb{P}_1^{\otimes n}, \mathbb{P}_2^{\otimes n})$ we obtain as a consequence of Lemma 2

$$\begin{aligned}
&\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\
&\geq \frac{1}{2} \left((1 - (1 - e^{-\epsilon}) \text{TV}(\mathbb{P}_1, \mathbb{P}_2))^n - 2ne^{-\epsilon \delta} \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right) .
\end{aligned}$$

Proof of Theorem 2: When \mathfrak{M} is ρ -DP, the generic bound of Theorem 5 applied with the Le Cam matching technique described in Theorem 7 and the coupling $\pi^\infty(\mathbb{P}_1, \mathbb{P}_2)$ leads to

$$\begin{aligned}
\max_{i \in \{1, 2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) &\geq \frac{1}{2} \left(1 - \sqrt{\rho/2} \mathbb{E}_{(\mathbf{X}_1, \mathbf{X}_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)} (d_{\text{ham}}(\mathbf{X}_1, \mathbf{X}_2)) \right) \\
&\stackrel{\text{Lemma 1}}{\geq} \frac{1}{2} \left(1 - \sqrt{\rho/2} n \Delta_{1,2} \right) \\
&= \frac{1}{2} \left(1 - n \sqrt{\rho/2} \text{TV}(\mathbb{P}_1, \mathbb{P}_2) \right) .
\end{aligned}$$

where in the second line we denote $\Delta_{1,2} := \mathbb{P}_{(\mathbf{X}_1, \mathbf{X}_2) \sim \pi^\infty(\mathbb{P}_1, \mathbb{P}_2)}(\mathbf{X}_1 \neq \mathbf{X}_2)$ and in the last line we use that $\Delta_{1,2} = \text{TV}(\mathbb{P}_1, \mathbb{P}_2)$ with the chosen coupling. Similarly, in the case of product distributions, with the same matching but $\pi^\infty(\mathbb{P}_1^{\otimes n}, \mathbb{P}_2^{\otimes n})$ we obtain as a consequence of Lemma 2

$$\max_{i \in \{1,2\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq \frac{1}{2} \left(1 - n\sqrt{\rho/2} \text{TV}(\mathbb{P}_1, \mathbb{P}_2)\right).$$

The case of arbitrary many hypotheses ($N \geq 2$). Given $\mathbb{P}_1, \dots, \mathbb{P}_N$ distributions on \mathcal{X}^n , a direct implication of Lemma 10 and of Fano's lemma (see Fact 2) is that independently on the privacy condition \mathcal{C} imposed on \mathfrak{M} , for any \mathbb{Q} such that $\mathbb{P}_i \ll \mathbb{Q}$ for all i ,

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq 1 - \frac{1 + \frac{1}{N} \sum_{i=1}^N \text{KL}(\mathbb{P}_i \| \mathbb{Q})}{\ln(N)}.$$

Again, this serves as the first ingredient of the proof of Theorem 3 and Theorem 4 that we now detail.

Proof of Theorem 3: When \mathfrak{M} is (ϵ, δ) -DP, the generic bound of Theorem 5 applied with the pairwise anchoring technique described in Theorem 6 and the coupling $\pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)$ leads to (since $\lceil n/2 \rceil \leq n$ for each integer $n \geq 0$)

$$\begin{aligned} & \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ & \geq \frac{1}{2N^2} \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)} \left(\sum_{i=1}^N \sum_{j=1}^N e^{-\epsilon d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)} \right. \\ & \quad \left. - 2e^{-\epsilon} \delta d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j) \right) \\ & \stackrel{\text{Lemma 1}}{\geq} \frac{1}{2N^2} \left(\sum_{i=1}^N \sum_{j=1}^N (1 - (1 - e^{-n\epsilon}) \Delta_{i,j}) - 2e^{-\epsilon} \delta n \Delta_{i,j} \right) \\ & \geq \frac{1}{2} - \frac{1 - e^{-n\epsilon} + 2ne^{-\epsilon}\delta}{2N^2} \sum_{i,j} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \end{aligned}$$

where in the second line we denote $\Delta_{i,j} := \mathbb{P}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)}(\mathbf{X}_i \neq \mathbf{X}_j)$ and in the last line we use that $\Delta_{i,j} \leq \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}$ with the chosen coupling. Similarly, in the case of product distributions, with the same matching but the product coupling $\pi^\infty(\mathbb{P}_1^{\otimes n}, \dots, \mathbb{P}_N^{\otimes n})$ we obtain as a consequence of Lemma 2

$$\begin{aligned} & \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ & \geq \frac{1}{2N^2} \sum_{i,j} \left(\left(1 - (1 - e^{-\epsilon}) \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \right)^n \right. \\ & \quad \left. - 2ne^{-\epsilon}\delta \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \right). \end{aligned}$$

When $\delta = 0$, the generic bound of Theorem 5 applied with the Fano matching technique described in Theorem 6 and the coupling $\pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)$ leads to

$$\begin{aligned} & \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}}(\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ & \geq 1 - \frac{1 + \frac{\epsilon}{N^2} \sum_{i=1}^N \sum_{j=1}^N \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)}(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j))}{\ln N} \\ & \stackrel{\text{Lemma 1}}{\geq} 1 - \frac{1 + \frac{\epsilon}{N^2} \sum_{i=1}^N \sum_{j=1}^N n \Delta_{i,j}}{\ln N} \\ & \geq 1 - \frac{1 + \frac{n\epsilon}{N^2} \sum_{i=1}^N \sum_{j=1}^N \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}}{\ln N} \end{aligned}$$

where in the second line we denote $\Delta_{i,j} := \mathbb{P}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)} (\mathbf{X}_i \neq \mathbf{X}_j)$ and in the last line we use that $\Delta_{i,j} \leq \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1+\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}$ with the chosen coupling. Similarly, in the case of product distributions, with the same matching but the coupling $\pi^\otimes(\mathbb{P}_1^{\otimes n}, \dots, \mathbb{P}_N^{\otimes n})$ we obtain as a consequence of Lemma 2

$$\max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \geq 1 - \frac{1 + \frac{n\epsilon}{N^2} \sum_{i=1}^N \sum_{j=1}^N \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1+\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}}{\ln N}.$$

Proof of Theorem 4: When \mathfrak{M} is ρ -zCDP, the generic bound of Theorem 5 applied with the Fano matching technique described in Theorem 7 and the coupling $\pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)$ leads to

$$\begin{aligned} & \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ & \geq 1 - \frac{1 + \frac{\rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)} \left(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)^2 \right)}{\ln N} \\ & \stackrel{\text{Lemma 1}}{\geq} 1 - \frac{1 + \frac{\rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N n^2 \Delta_{i,j}}{\ln N} \\ & \geq 1 - \frac{1 + \frac{n^2 \rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1+\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}}{\ln N} \end{aligned}$$

where in the second line we denote $\Delta_{i,j} := \mathbb{P}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)} (\mathbf{X}_i \neq \mathbf{X}_j)$ and in the last line we use that $\Delta_{i,j} \leq \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1+\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}$ with the chosen coupling. Similarly, in the case of product distributions, with the same matching but with the product coupling $\pi^\otimes(\mathbb{P}_1^{\otimes n}, \dots, \mathbb{P}_N^{\otimes n})$ we obtain,

$$\begin{aligned} & \max_{i \in \{1, \dots, N\}} \mathbb{P}_{\mathbf{X} \sim \mathbb{P}_i, \mathfrak{M}} (\Psi(\mathfrak{M}(\mathbf{X})) \neq i) \\ & \geq 1 - \frac{1 + \frac{\rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N \mathbb{E}_{(\mathbf{X}_1, \dots, \mathbf{X}_N) \sim \pi^\otimes(\mathbb{P}_1^{\otimes n}, \dots, \mathbb{P}_N^{\otimes n})} \left(d_{\text{ham}}(\mathbf{X}_i, \mathbf{X}_j)^2 \right)}{\ln N} \\ & \stackrel{\text{Lemma 2}}{\geq} 1 - \frac{1 + \frac{\rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N (n^2 \delta_{i,j}^2 + n \delta_{i,j})}{\ln N} \\ & \geq 1 - \frac{1 + \frac{n^2 \rho}{N^2} \sum_{i=1}^N \sum_{j=1}^N \left(\left(\frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1+\text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \right)^2 + \frac{1}{n} \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1+\text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \right)}{\ln N} \end{aligned}$$

where in the second line we denote $\delta_{i,j} := \mathbb{P}_{(X_1, \dots, X_N) \sim \pi^\infty(\mathbb{P}_1, \dots, \mathbb{P}_N)} (X_i \neq X_j)$ and in the last line we use that $\delta_{i,j} \leq \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1+\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}$ with the chosen coupling.

D Examples of applications

This section presents three examples of applications of our lower bounds: The Bernoulli Model, the Gaussian Model and the Uniform Model.

In the first two applications, we show that the rate at which the privacy parameters vary has an importance. Namely, we show in both models that if the privacy parameters are too small, the private minimax risk becomes predominant compared to the non-private one, or to put it simply, we show that under strong privacy constraints, the performance of estimation *has* to be degraded.

Furthermore, we also show for the first model that above this threshold, the minimax risk is not degraded by privacy, essentially meaning that we then have privacy "for free".

In contrast, in the last example we prove that the minimax risk is systematically degraded by privacy as soon as we consider estimation procedures with increasing privacy requirements in the sense that the privacy parameters decrease as the sample size increases.

D.1 Bernoulli Model

The first application is the estimation of the proportion of a population that satisfies a certain property. It is a prime example of the application of Le Cam's lemma Fact 1 and its private counterparts Theorem 1 and Theorem 2. When we consider the parametric Bernoulli model

$$(\mathcal{B}(\theta))_{\theta \in \Theta}, \quad \Theta = (0, 1),$$

a classical and simple estimator for estimating the true parameter θ^* from i.i.d. samples X_1, \dots, X_n drawn according to $\mathcal{B}(\theta^*)$ is via the empirical average

$$\hat{\theta} := \frac{1}{n} \sum_{i=1}^n X_i.$$

The quadratic risk of this estimator is

$$\mathbb{E} \left((\theta^* - \hat{\theta})^2 \right) = \frac{\theta^*(1 - \theta^*)}{n} \leq \frac{1/4}{n}.$$

In order to find lower bounds on the minimax risk (with or without privacy constraints), let us investigate an $\Omega = \frac{\alpha}{4}$ -packing³ with $\theta_1 := \frac{1+\alpha}{2}$ and $\theta_2 := \frac{1}{2}$.

Regular Minimax Risk. By the master bound (6), Le Cam's lemma (Fact 1) and Pinsker's inequality (see Tsybakov (2009, Lemma 2.5)),

$$\begin{aligned} \mathfrak{M}_n \left((\mathcal{B}(\theta)^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) &\geq (\alpha/4)^2 \cdot \frac{1}{2} (1 - \text{TV}(\mathcal{B}(\theta_1)^{\otimes n}, \mathcal{B}(\theta_2)^{\otimes n})) \\ &\geq \frac{\alpha^2}{32} \left(1 - \sqrt{\text{KL}(\mathcal{B}(\theta_1)^{\otimes n} \| \mathcal{B}(\theta_2)^{\otimes n}) / 2} \right) \\ &= \frac{\alpha^2}{32} \left(1 - \sqrt{n \text{KL}(\mathcal{B}(\theta_1) \| \mathcal{B}(\theta_2)) / 2} \right). \end{aligned}$$

where we used the tensorization property of the KL divergence (see van Erven & Harremoës (2014, Theorem 28)). We can observe that when $\alpha \in [0, 1/2]$,

$$\text{KL}(\mathcal{B}(\theta_1) \| \mathcal{B}(\theta_2)) \leq \alpha^2.$$

Indeed, let us note $g(x) = \frac{1+x}{2} \ln(1+x) + \frac{1-x}{2} \ln(1-x) - x^2$. We have that $\frac{dg(x)}{dx}(x) = \frac{\ln(1+x) + \ln(1-x)}{2} - 2x$ and since $g(0) = 0$ and $x \mapsto \ln(1+x)$ is 2-Lipschitz on $[-1/2, 1/2]$, we have that $g(x) \leq 0$, $\forall x \in [0, 1/2]$. In particular, when $\alpha \in [0, 1/2]$,

$$\begin{aligned} \text{KL}(\mathcal{B}(\theta_1) \| \mathcal{B}(\theta_2)) &= \left(\theta_1 \ln \left(\frac{\theta_1}{\theta_2} \right) + (1 - \theta_1) \ln \left(\frac{1 - \theta_1}{1 - \theta_2} \right) \right) \\ &= \left(\frac{1 + \alpha}{2} \ln(1 + \alpha) + \frac{1 - \alpha}{2} \ln(1 - \alpha) \right) \\ &\leq \alpha^2. \end{aligned}$$

So, with $\alpha = \frac{1}{\sqrt{n}}$, as soon as $n \geq 4$, we obtain that

$$\mathfrak{M}_n \left((\mathcal{B}(\theta)^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) \geq \frac{\alpha^2}{32} \left(1 - \sqrt{n\alpha^2/2} \right) = \frac{1/160}{n} = \Omega \left(\frac{1}{n} \right),$$

which concludes that the non-private minimax rate is lower bounded by a quantity of the order of $\frac{1}{n}$ and in particular, that the empirical mean estimator $\hat{\theta}$ is minimax optimal in term of rates of convergence. Furthermore, any private minimax rate also has to be of the order of at least $\frac{1}{n}$.

³With $d(\cdot, \cdot) = |\cdot - \cdot|$, see Section 1.3: an Ω -packing must satisfy $d(\theta_i, \theta_j) \geq 2\Omega$, $i \neq j$.

Minimax Risk with ϵ -Differential Privacy. By the private master lower bound (10) and the product form of Le Cam's lemma for $(\epsilon, 0)$ -DP (see Theorem 1) combined with the last inequality in Lemma 2 we obtain

$$\begin{aligned} \mathfrak{M}_n \left(\epsilon\text{-DP}, (\mathcal{B}(\theta)^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) &\geq (\alpha/4)^2 \cdot \frac{1}{2} e^{-n\epsilon \text{TV}(\mathcal{B}(\theta_1), \mathcal{B}(\theta_2))} \\ &\geq \frac{\alpha^2}{32} e^{-n\epsilon \sqrt{\text{KL}(\mathcal{B}(\theta_1) \parallel \mathcal{B}(\theta_2))/2}} \\ &= \frac{\alpha^2}{32} e^{-\sqrt{(n\epsilon)^2 \alpha^2 / 2}} \end{aligned}$$

where we used again Pinsker's inequality.

So, with $\alpha = \frac{1}{n\epsilon}$, when $n\epsilon \geq 2$, we obtain that

$$\mathfrak{M}_n \left(\epsilon\text{-DP}, (\mathcal{B}(\theta)^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) \geq \frac{1/32}{(n\epsilon)^2} e^{-\sqrt{1/2}} \geq \frac{1/80}{(n\epsilon)^2} = \Omega \left(\frac{1}{(n\epsilon)^2} \right).$$

ρ -zero Concentrated Differential Privacy. Similarly, by the product form of Le Cam's lemma for ρ -zCDP (see Theorem 2), we get with $\alpha = \frac{1}{n\sqrt{\rho}}$ when $n\sqrt{\rho} \geq 2$,

$$\begin{aligned} \mathfrak{M}_n \left(\rho\text{-zCDP}, (\mathcal{B}(\theta)^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) &\geq \frac{\alpha^2}{32} \left(1 - n\sqrt{\rho/2} \text{TV}(\mathcal{B}(\theta_1), \mathcal{B}(\theta_2)) \right) \\ &\geq \frac{\alpha^2}{32} \left(1 - n\sqrt{\rho \text{KL}(\mathcal{B}(\theta_1) \parallel \mathcal{B}(\theta_2)) / 4} \right) \\ &= \frac{\alpha^2}{32} \left(1 - \sqrt{n^2 \rho \alpha^2 / 4} \right) = \frac{1/64}{n^2 \rho} \\ &= \Omega \left(\frac{1}{n^2 \rho} \right). \end{aligned}$$

Matching Upper Bounds. Consider the Laplace mechanism $\mathfrak{M}(\mathbf{X}) := \frac{1}{n} \sum_{i=1}^n X_i + \frac{1}{n\epsilon} \text{Lap}(1)$. It is an $(\epsilon, 0)$ -DP estimator \mathbf{X} (Dwork & Roth, 2014) and its quadratic risk is $O\left(\frac{1}{n} + \frac{1}{(n\epsilon)^2}\right)$. Likewise, the Gaussian mechanism $\mathfrak{M}(\mathbf{X}) = \frac{1}{n} \sum_{i=1}^n X_i + \frac{2}{n\sqrt{\rho}} \mathcal{N}(0, 1)$ is ρ -zCDP (Bun & Steinke, 2016) and its one is $O\left(\frac{1}{n} + \frac{1}{n^2 \rho}\right)$. Combined with the lower bounds established so far and with Lemma 10, this allows to conclude that in fact

$$\mathfrak{M}_n \left(\epsilon\text{-DP}, (\mathcal{B}(\theta)^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) = \Theta \left(\max \left\{ \frac{1}{n}, \frac{1}{(n\epsilon)^2} \right\} \right),$$

and that this optimal rate is achieved with the Laplace mechanism, while

$$\mathfrak{M}_n \left(\rho\text{-zCDP}, (\mathcal{B}(\theta)^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) = \Theta \left(\max \left\{ \frac{1}{n}, \frac{1}{n^2 \rho} \right\} \right),$$

which is an optimal rate achieved by the Gaussian mechanism.

The Cost of Privacy. An interesting observation for both definitions of privacy is that there exist regimes ($\epsilon \ll 1/\sqrt{n}$ or $\rho \ll 1/n$) for which the minimax rate of convergence is degraded compared to the non private one. In other words, privacy has an unavoidable cost on utility, no matter the mechanism used. Conversely, the order of magnitude of the minimax risk is not degraded otherwise.

D.2 Gaussian Model

The second application is the estimation of the unknown mean $\theta^* \in \mathbb{R}^d$ of multivariate normally distributed data with fixed covariance matrix $\sigma^2 I_d$. When we consider the parametric model $(\mathcal{N}(\theta, \sigma^2 I_d))_{\theta \in \Theta}$, $\Theta = \mathbb{R}^d$, a classical and simple estimator for estimating the mean θ^* from i.i.d. samples X_1, \dots, X_n is the empirical average $\hat{\theta} := \frac{1}{n} \sum_{i=1}^n X_i$. The quadratic risk of this estimator is

$$\mathbb{E} \left(\|\theta^* - \hat{\theta}\|^2 \right) = \frac{\sigma^2 d}{n}. \quad (23)$$

If we were to apply Le Cam's lemma Fact 1 or its private counterparts Theorem 1 and Theorem 2, the parameter that tunes the dimensionality d would not be captured by the resulting minimax lower bounds which would thus be overly optimistic. This example forces us to use Fano's lemma Fact 2 or its private counterparts Theorem 3 or Theorem 4 in order to have a chance to capture this phenomenon.

The total variation that appears in Fano's inequality is controlled via Pinsker's inequality in terms of a Kullback-Leibler divergence, which in the case of isotropic Gaussians is known to be proportional to the squared Euclidean distance.

$$\forall \theta_1, \theta_2 \in \Theta, \quad \text{KL}(\mathcal{N}(\theta_1, \sigma^2 I_d) \parallel \mathcal{N}(\theta_2, \sigma^2 I_d)) = \frac{\|\theta_2 - \theta_1\|^2}{2\sigma^2}. \quad (24)$$

This enables the use of packing results for the Euclidean norm, and minimax bounds valid in the more general case where the KL divergence is controlled by the Euclidean norm between parameters.

Packing Choice. In high dimension, the packing is chosen with an exponential number of hypotheses. A good way to obtain well-spread points is to use Varshamov–Gilbert's theorem

Fact 8 (Varshamov–Gilbert's theorem (Rigollet & Hütter, 2015, Lemma 5.12)). *For any $\zeta \in (0, \frac{1}{2})$ and for every dimension $d \geq 1$ there exist $N \geq e^{\frac{\zeta^2 d}{2}}$ and $w_1, \dots, w_N \in \{0, 1\}^d$ such that,*

$$i \neq j \implies d_{\text{ham}}(w_i, w_j) \geq \left(\frac{1}{2} - \zeta\right) d.$$

Minimax Lower Bounds. We obtain the following minimax lower bounds that we factorized in a single result:

Proposition 1. *Let $(\mathbb{P}_\theta)_{\theta \in \Theta}$ be a family of probability distributions on the same measurable space and Θ be a subset of \mathbb{R}^d with $d \geq 66$ that contains a ball of radius r_0 for the euclidean distance. Assume that $\gamma > 0$ is such that*

$$\forall \theta_1, \theta_2 \in \Theta, \quad \text{KL}(\mathbb{P}_{\theta_1} \parallel \mathbb{P}_{\theta_2}) \leq \gamma \|\theta_2 - \theta_1\|^2. \quad (25)$$

Then we have the following results on the minimax rates:

$$\begin{aligned} \mathfrak{M}_n \left((\mathbb{P}_\theta^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) &\geq \frac{\min \left(\frac{r_0}{\sqrt{d}}, \frac{1}{64\sqrt{n\gamma}} \right)^2 d}{32} = \Omega \left(\frac{d}{n\gamma} \right), \\ \mathfrak{M}_n \left(\epsilon\text{-DP}, (\mathbb{P}_\theta^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) &\geq \frac{\max \left(\min \left(\frac{r_0}{\sqrt{d}}, \frac{1}{64\sqrt{n\gamma}} \right), \min \left(\frac{r_0}{\sqrt{d}}, \frac{\sqrt{d}}{64^2 \sqrt{2n\epsilon\sqrt{\gamma}}} \right) \right)^2 d}{32} \\ &= \Omega \left(\max \left\{ \frac{d}{n\gamma}, \frac{d^2}{(n\epsilon)^2 \gamma} \right\} \right), \\ \mathfrak{M}_n \left(\rho\text{-zCDP}, (\mathbb{P}_\theta^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) &\geq \frac{\max \left(\min \left(\frac{r_0}{\sqrt{d}}, \frac{1}{64\sqrt{n\gamma}} \right), \min \left(\frac{r_0}{\sqrt{d}}, \frac{1}{64^2 2\sqrt{2n\sqrt{\rho\gamma}}} \right) \right)^2 d}{32} = \Omega \left(\max \left\{ \frac{d}{n\gamma}, \frac{d}{n^2 \rho \gamma} \right\} \right), \end{aligned}$$

when $\rho < 1$. Note that all the asymptotic expressions are taken when $r_0 > C\sqrt{d}$ for a positive constant C i.e. when the parameter space is not "too small".

Note that the constraint $d \geq 66$ can be relaxed to smaller constants by changing the ζ in the application of Varshamov–Gilbert's theorem at the cost of changing the constants in the minimax lower bounds. Likewise, the constraint $\rho < 1$ can be replaced by $\rho < M$ for any positive constant M at the cost again of worse constants. Since we aim to use this result in high dimension and with high privacy, those hypotheses are natural in order to simplify the expressions. Before giving the proof, we discuss some practical consequences.

The Cost of Privacy. For Gaussians, by (23) and Proposition 1 with $\gamma = \frac{1}{2\sigma^2}$ (cf (24)) the non-private minimax risk is

$$\mathfrak{M}_n \left((\mathcal{N}(\theta, \sigma^2 I_d))_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) = \Theta \left(\frac{\sigma^2 d}{n} \right),$$

hence Proposition 1 shows that there is a degradation of the private minimax rate over the non-private minimax rate in the regime $\epsilon \ll \sqrt{\frac{d}{n}}$ when working under ϵ -DP. Note that this shift in regime depends on the dimensionality. For ρ -zCDP, the minimax rate of convergence is degraded as soon as $\rho \ll \frac{1}{n}$. Compared to ϵ -DP, the rate at which we observe a degradation does not depend on the dimension d . We study an upper bound in Section 4.

Proof of Proposition 1. Without loss of generality, let us suppose that 0 is the center of the ball of radius r_0 (without loss of generality because we are going to work on a neighborhood of 0 but it can be translated to any point). Varshamov–Gilbert’s theorem (Fact 8) with $\zeta = \frac{1}{4}$ allows us to consider N and w_1, \dots, w_N and to define a packing of the form $\theta_1 := \alpha w_1, \dots, \theta_N := \alpha w_N$ such that

$$i \neq j \implies \frac{\alpha^2 d}{4} \leq \|\theta_i - \theta_j\|^2 \leq \alpha^2 d.$$

This yields an $\Omega = \alpha\sqrt{d}/4$ -packing with respect to the Euclidean metric. Since 0 is in the interior of Θ , all the θ_i ’s are in Θ provided that α is small enough. By the (non-private) master lower bound (6) and Fano’s lemma (Fact 2),

$$\begin{aligned} & \mathfrak{M}_n \left((\mathbb{P}_{\theta}^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) \\ & \geq (\alpha\sqrt{d}/4)^2 \cdot \left(1 - \frac{1 + \frac{1}{N} \sum_i \text{KL} \left(\mathbb{P}_{\theta_i}^{\otimes n} \parallel \frac{1}{N} \sum_j \mathbb{P}_{\theta_j}^{\otimes n} \right)}{\ln N} \right) \\ & \stackrel{\text{Jensen's inequality}}{\geq} (\alpha\sqrt{d}/4)^2 \cdot \left(1 - \frac{1 + \frac{1}{N^2} \sum_{i,j} \text{KL} \left(\mathbb{P}_{\theta_i}^{\otimes n} \parallel \mathbb{P}_{\theta_j}^{\otimes n} \right)}{\ln N} \right) \\ & = (\alpha\sqrt{d}/4)^2 \cdot \left(1 - \frac{1 + \frac{1}{N^2} \sum_{i,j} n \text{KL} \left(\mathbb{P}_{\theta_i} \parallel \mathbb{P}_{\theta_j} \right)}{\ln N} \right) \\ & \stackrel{(25)}{\geq} \frac{\alpha^2 d}{16} \left(1 - \frac{1 + \frac{1}{N^2} \sum_{i,j} n\gamma \|\theta_i - \theta_j\|^2}{\ln N} \right) \\ & \geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + n\gamma\alpha^2 d}{d/32} \right), \end{aligned}$$

where in the last line we used that $N \geq e^{d/32}$ and $\|\theta_i - \theta_j\|^2 \leq \alpha^2 d$. With $\alpha := \min \left(\frac{r_0}{\sqrt{d}}, \frac{1}{64\sqrt{n\gamma}} \right)$ when $d \geq 66$ leads to

$$\mathfrak{M}_n \left((\mathbb{P}_{\theta}^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) \geq \frac{\min \left(\frac{r_0}{\sqrt{d}}, \frac{1}{64\sqrt{n\gamma}} \right)^2 d}{32} = \Omega \left(\frac{d}{n\gamma} \right).$$

For ϵ -DP and ρ -zCDP, the first term in the max expressed in Proposition 1 is a direct consequence of the above bound and of Lemma 10 so we now concentrate on the other term. By the private master lower bound (10) and Fano’s lemma for product distributions and $(\epsilon, 0)$ -DP (see Theorem 3), arguments as above

show that

$$\begin{aligned}
\mathfrak{M}_n \left(\epsilon\text{-DP}, (\mathbb{P}_\theta^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) &\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + \frac{2n\epsilon}{N^2} \sum_{i,j} \text{TV}(\mathbb{P}_{\theta_i}, \mathbb{P}_{\theta_j})}{\ln N} \right) \\
&\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + \frac{2n\epsilon}{N^2} \sum_{i,j} \sqrt{\text{KL}(\mathbb{P}_{\theta_i} \parallel \mathbb{P}_{\theta_j}) / 2}}{\ln N} \right) \\
&\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + \frac{2n\epsilon}{N^2} \sum_{i,j} \sqrt{\gamma/2} \|\theta_i - \theta_j\|}{\ln N} \right) \\
&\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + 2n\epsilon\alpha\sqrt{\gamma/2\sqrt{d}}}{d/32} \right).
\end{aligned}$$

Again, setting $\alpha := \min\left(\frac{r_0}{\sqrt{d}}, \frac{\sqrt{d}}{64^2 \sqrt{2n\epsilon\sqrt{\gamma}}}\right)$ when $d \geq 66$ allows to conclude that

$$\begin{aligned}
\mathfrak{M}_n \left(\epsilon\text{-DP}, (\mathbb{P}_\theta^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) &\geq \frac{\min\left(\frac{r_0}{\sqrt{d}}, \frac{\sqrt{d}}{64^2 \sqrt{2n\epsilon\sqrt{\gamma}}}\right)^2 d}{32} \\
&= \Omega\left(\frac{d^2}{(n\epsilon)^2 \gamma}\right).
\end{aligned}$$

Similarly, by Fano's lemma for product distributions and ρ -zCDP (see Theorem 4),

$$\begin{aligned}
&\mathfrak{M}_n \left(\rho\text{-zCDP}, (\mathbb{P}_\theta^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) \\
&\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + \frac{4n^2 \rho}{N^2} \sum_{i,j} \frac{1}{2n} \text{TV}(\mathbb{P}_{\theta_i}, \mathbb{P}_{\theta_j}) + \text{TV}(\mathbb{P}_{\theta_i}, \mathbb{P}_{\theta_j})^2}{\ln N} \right) \\
&\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + \frac{4n^2 \rho}{N^2} \sum_{i,j} \frac{1}{2n} \sqrt{\text{KL}(\mathbb{P}_{\theta_i} \parallel \mathbb{P}_{\theta_j}) / 2} + \text{KL}(\mathbb{P}_{\theta_i} \parallel \mathbb{P}_{\theta_j}) / 2}{\ln N} \right) \\
&\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + \frac{4n^2 \rho}{N^2} \sum_{i,j} \frac{1}{2n} \sqrt{\gamma/2} \|\theta_i - \theta_j\| + \gamma \|\theta_i - \theta_j\|^2 / 2}{\ln N} \right) \\
&\geq \frac{\alpha^2 d}{16} \left(1 - \frac{1 + (2\sqrt{2}n\rho\alpha\sqrt{\gamma d} + 2n^2\rho\gamma\alpha^2 d)}{d/32} \right),
\end{aligned}$$

and setting $\alpha := \min\left(\frac{r_0}{\sqrt{d}}, \frac{1}{64^2 2\sqrt{2n\sqrt{\rho\gamma}}}\right)$ when $d \geq 66$ concludes that (because $\rho \leq 1$)

$$\begin{aligned}
\mathfrak{M}_n \left(\rho\text{-zCDP}, (\mathbb{P}_\theta^{\otimes n})_{\theta \in \Theta}, \|\cdot - \cdot\|, (\cdot)^2 \right) &\geq \frac{\min\left(\frac{r_0}{\sqrt{d}}, \frac{1}{64^2 2\sqrt{2n\sqrt{\rho\gamma}}}\right)^2 d}{32} \\
&= \Omega\left(\frac{d}{n^2 \rho \gamma}\right).
\end{aligned}$$

□

D.3 Support of Uniform Distributions

For the last example, we chose to investigate a statistical problem that has a non-private minimax rate faster than $\frac{1}{n}$. We consider the parametric model

$$(\mathbb{P}_\theta := \mathcal{U}([0, \theta]))_{\theta \in \Theta}, \quad \Theta = (0, 1].$$

To exploit Le Cam's lemma we will need to control the total variation between two distributions. In this model, it can be done explicitly. The total variation between $\mathbb{P}_{\theta_1}^{\otimes n}$ and $\mathbb{P}_{\theta_2}^{\otimes n}$ can be computed as

$$\text{TV}(\mathbb{P}_{\theta_1}^{\otimes n}, \mathbb{P}_{\theta_2}^{\otimes n}) = 1 - \int_{[0,1]^n} \min(\pi_{\mathbb{P}_{\theta_1}^{\otimes n}}, \pi_{\mathbb{P}_{\theta_2}^{\otimes n}}) = 1 - \left(\frac{\min(\theta_1, \theta_2)}{\max(\theta_1, \theta_2)} \right)^n.$$

Non-Private Minimax Risk. By the (non-private) master lower bound (6) and Le Cam's lemma (Fact 1), applied to the $\frac{1}{2n}$ -packing $\theta_1 = 1 - \frac{1}{n}$ and $\theta_2 = 1$, we have

$$\mathfrak{M}_n \left((\mathcal{U}([0, \theta])^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) \geq \frac{e^{-1}}{8n^2} = \Omega\left(\frac{1}{n^2}\right).$$

where we used that $1 - \text{TV}(\mathbb{P}_{\theta_1}^{\otimes n}, \mathbb{P}_{\theta_2}^{\otimes n}) = \left(1 - \frac{1}{n}\right)^n \geq e^{-1}$. Furthermore, as we now show, the estimator $\max \mathbf{X}$ achieves this rate of convergence when $X_1, \dots, X_n \sim \mathcal{U}([0, \theta^*])$ are independent. Indeed, for any $t \in [0, \theta^*]$,

$$\mathbb{P}(\max \mathbf{X} < t) = \prod_{i=1}^n \mathbb{P}(X_i < t) = \left(\frac{t}{\theta^*}\right)^n.$$

Hence, $\max \mathbf{X}$ has a density $\pi_{\max \mathbf{X}}$ with respect to the Lebesgue measure where

$$\forall t \in \mathbb{R}, \quad \pi_{\max \mathbf{X}}(t) = 1_{[0, \theta^*]}(t) \frac{nt^{n-1}}{\theta^{*n}},$$

so that

$$\begin{aligned} \mathbb{E}(\max \mathbf{X}) &= \int_0^{\theta^*} t \left(\frac{nt^{n-1}}{\theta^{*n}} \right) dt = \frac{n}{n+1} \theta^*, \\ \mathbb{V}(\max \mathbf{X}) &= \int_0^{\theta^*} t^2 \left(\frac{nt^{n-1}}{\theta^{*n}} \right) dt - [\mathbb{E}(\max \mathbf{X})]^2 = \theta^{*2} \left(\frac{n}{n+2} - \frac{n^2}{(n+1)^2} \right). \end{aligned}$$

By the bias-variance tradeoff, the quadratic risk of $\max \mathbf{X}$ is thus $O\left(\frac{\theta^{*2}}{n^2}\right)$. In particular, this proves that the non-private minimax rate of convergence is $\Theta\left(\frac{1}{n^2}\right)$ and that $\max \mathbf{X}$ achieves this minimax rate of convergence.

Minimax Risk with ϵ -Differential Privacy. By the private master lower bound (10) and the product form of Le Cam's private lemma for ϵ -DP on product distributions (see Theorem 1 with $\delta = 0$) with the $\frac{1}{2n\epsilon}$ -packing $\theta_1 = 1 - \frac{1}{n\epsilon}$ and $\theta_2 = 1$ we have when $n\epsilon > 1$

$$\mathfrak{M}_n \left(\epsilon\text{-DP}, (\mathcal{U}([0, \theta])^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) \geq \frac{e^{-1}}{8(n\epsilon)^2} = \Omega\left(\frac{1}{(n\epsilon)^2}\right),$$

In particular, the rate is degraded compared to the non-private one as soon as ϵ is decreasing.

Minimax Risk with ρ -zero Concentrated Differential Privacy. Similarly, using the product form of Le Cam's private lemma for ρ -zCDP on product distributions (see Theorem 2) and the $\frac{1}{2n\sqrt{\rho}}$ -packing $\theta_1 = 1 - \frac{1}{n\sqrt{\rho}}$ and $\theta_2 = 1$ gives that when $n\sqrt{\rho} > 1$,

$$\mathfrak{M}_n \left(\rho\text{-zCDP}, (\mathcal{U}([0, \theta])^{\otimes n})_{\theta \in \Theta}, |\cdot - \cdot|, (\cdot)^2 \right) \geq \frac{1 - \frac{1}{\sqrt{2}}}{8n^2\rho} = \Omega\left(\frac{1}{n^2\rho}\right).$$

In particular, the rate is degraded compared to the non-private one as soon as ρ is decreasing.

This example shows that when the stochastic noise due to sampling shrinks too fast (here $\max \mathbf{X}$ has quadratic risk $O(1/n^2)$), then the noise due to privacy becomes predominant. In particular, we do not observe a distinction on the rate at which ϵ or ρ tends to 0 in order to conclude to a degradation of the minimax risk. It is systematically degraded.

E A note on Fact 4

This section was added in the latest version of this article after noticing that the use of Poisson point processes in the proof of Fact 4 can make the results presented in this article harder to communicate on than it should. We thus present here a simplified view on the proof of Angel & Spinka (2021).

Let $(\mathcal{X}, \mathcal{B})$ be a measurable space. Let $(\mathbb{P}_i)_{i \in I}$ be a collection of probability measures that are compatible with $(\mathcal{X}, \mathcal{B})$. A coupling between $(\mathbb{P}_i)_{i \in I}$ is a collection of random variables $(X_i)_{i \in I}$ defined on a *common* probability space (with probability measure \mathbb{P}), and with values in $(\mathcal{X}, \mathcal{B})$, such that

$$\forall i \in I, \quad \mathbb{P}_{X_i} = \mathbb{P}_i .$$

Here, for any i , \mathbb{P}_{X_i} represents the push-forward probability of \mathbb{P} by X_i defined as

$$\forall B \in \mathcal{B}, \quad \mathbb{P}_{X_i}(B) = \mathbb{P}(X_i \in B) .$$

A folklore result Lindvall (2003) states that if $(X_i)_{i \in I}$ is a coupling between $(\mathbb{P}_i)_{i \in I}$, then

$$\forall i, j \in I, \quad \forall E_{i,j} \in \mathcal{B} \text{ s.t. } X_i(\omega) = X_j(\omega) \quad \forall \omega \in E_{i,j}, \quad \mathbb{P}(E_{i,j}) \leq 1 - \text{TV}(\mathbb{P}_i, \mathbb{P}_j) , \quad (26)$$

where we recall that $\text{TV}(\cdot, \cdot)$ is the total variation distance on the set on probability measures on $(\mathcal{X}, \mathcal{B})$ defined as

$$\text{TV}(\mathbb{P}_i, \mathbb{P}_j) = \sup_{B \in \mathcal{B}} |\mathbb{P}_i(B) - \mathbb{P}_j(B)| .$$

Indeed, for any $i, j \in I$, any $E_{i,j} \in \mathcal{B}$ s.t. $X_i(\omega) = X_j(\omega) \quad \forall \omega \in E_{i,j}$, and any $B \in \mathcal{B}$,

$$\begin{aligned} \mathbb{P}(E_{i,j}) &\leq \mathbb{P}((X_i \in B, X_j \in B) \cup (X_i \in \bar{B}, X_j \in \bar{B})) \\ &= \mathbb{P}(X_i \in B, X_j \in B) + \mathbb{P}(X_i \in \bar{B}, X_j \in \bar{B}) \\ &\leq \mathbb{P}(X_j \in B) + \mathbb{P}(X_i \in \bar{B}) \\ &= \mathbb{P}_j(B) + 1 - \mathbb{P}_i(B) \\ &= 1 - (\mathbb{P}_i(B) - \mathbb{P}_j(B)) , \end{aligned}$$

which gives the result after optimizing over B .

Note that in the previous expressions, we looked at events on which random variables are equal without directly referring to the exact sets of equalities between those random variables. This is because without further caution, such sets may not be measurable. However, mild topological hypotheses on $(\mathcal{X}, \mathcal{B})$ are enough to guarantee that for any random variables X and Y defined on a common probability space and taking values in $(\mathcal{X}, \mathcal{B})$ then $(X = Y)$ is measurable. This is for instance the case if $(\mathcal{X}, \mathcal{B})$ is a *Polish* space equipped with its Borel σ -algebra, a ubiquitous class of measurable spaces in applied mathematics and in theoretical computer science. Indeed, if \mathcal{X} is a Polish space, by noting δ its distance (up to homeomorphism) and D its countable dense part, one has that

$$(X = Y) = \bigcap_{n \in \mathbb{N}_*} \left(\bigcup_{d \in D} \left(\delta(X, d) < \frac{1}{n} \right) \cap \left(\delta(Y, d) < \frac{1}{n} \right) \right) .$$

E.1 Related work

A folklore coupling construction Lindvall (2003) states that when I is of cardinality two, i.e. when one is interested in finding a coupling between two marginal distributions only, and when $(\mathcal{X}, \mathcal{B})$ is a Polish space, then the upper-bound (26) is tight, which means that there exists a coupling (X_1, X_2) between \mathbb{P}_1 and \mathbb{P}_2 such that

$$\mathbb{P}(X_1 = X_2) = 1 - \text{TV}(\mathbb{P}_1, \mathbb{P}_2) .$$

In their article Angel & Spinka (2021), Omer Angel and Yinon Spinka investigate what happens when I is of cardinality bigger than two, and even possibly of infinite (non-countable) cardinality. First, they prove

that as soon as I is of cardinality at least three, the upper-bound (26) is unachievable, which means there exists at least an instance of the problem (i.e. a measurable space and a set of marginal distributions) on which no coupling may match the upper-bound. Then, they provide two different coupling constructions that achieve the following quantitative result : denoting by $(X_i)_{i \in I}$ the coupling between $(\mathbb{P}_i)_{i \in I}$, then

$$\mathbb{P}(X_i = X_j) \geq 1 - \frac{2\text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)} \geq 1 - 2\text{TV}(\mathbb{P}_i, \mathbb{P}_j), \quad \forall i, j \in I. \quad (27)$$

In this expression, $(X_i = X_j)$ should be understood in a weak sense (i.e. as a *measurable* set on which $X_i = X_j$, which may be smaller than the set $(X_i = X_j)$). (26) together with (27) show that the constructions of the authors are optimal up to a factor two. Both of their constructions work under different setups, use different construction tools and are qualitatively different. They can be summarized as follows :

1. (i) Their first construction is the most general, and is the one that is simplified here. It builds on inhomogeneous Poisson point processes and assumes the existence a σ -finite measure μ on $(\mathcal{X}, \mathcal{B})$ such that for any $i \in I$, \mathbb{P}_i is absolutely continuous w.r.t. μ . In particular, no assumption is made on the cardinality of I .
2. (ii) Their second construction only builds on elementary notions of probability theory but requires $(\mathcal{X}, \mathcal{B})$ being discrete. Note that since on discrete spaces, any probability measure is absolutely continuous w.r.t. the counting measure, the domination assumption in (i) is always satisfied in this setup. Again, no assumption is made on the cardinality of I .

E.2 About this note

This note aims to present a simplification of Angel and Spinka's first coupling construction ((i)), while still relying on the same core interesting idea, that only builds on elementary notions of probability theory, and that is still applicable to continuous distributions (contrary to (ii)).

This simplification will work under the following hypothesis. Let us assume that there exists a distribution of probability \mathbb{Q} on $(\mathcal{X}, \mathcal{B})$ such that there exists a $M > 1$ such that

$$\forall i \in I, \quad \forall B \in \mathcal{B}, \quad \mathbb{P}_i(B) \leq M\mathbb{Q}(B). \quad (28)$$

In particular, this hypothesis is stronger than the domination one in (i). Even if (28) does not have restrictive implications on the cardinality of I , an important observation is that (28) is always satisfied when I is finite. Indeed, one may take \mathbb{Q} as the mixture distribution

$$\mathbb{Q} = \frac{1}{\#(I)} \sum_{i \in I} \mathbb{P}_i,$$

and $M = \#(I)$.

Under this extra hypothesis, it is possible to present Angel and Spinka's first coupling construction under the light of traditional rejection sampling. In particular, the cornerstone of the construction, which is to add dependence between two marginals proportionally to their common area below their densities, is identical to Angel and Spinka's Poisson construction.

Since it presents Angel and Spinka's coupling construction with only elementary theory, we hope that this note will help communication and diffusion of such results that are interesting for both the theoretical computer science and the machine learning communities Elkin-Koren et al. (2023); Bun et al. (2023); Karbasi et al. (2024); Block et al. (2024); Kalavasis et al. (2023).

E.3 Coupling construction

Let $Y_1, U_1, Y_2, U_2, \dots, Y_n, U_n, \dots$ be a sequence of independent random variables⁴ such that, for any i , Z_i has distribution \mathbb{Q} on $(\mathcal{X}, \mathcal{B})$, and U_i is uniform on $[0, 1]$. Let $(\Omega, \mathcal{T}, \mathbb{P})$ be the probability space supporting this sequence.

⁴Formally, it is a stochastic process indexed by \mathbb{N} . The independence means that its finite-dimensional distributions are product distributions. Its existence, without topological assumptions of $(\mathcal{X}, \mathcal{B})$, is guaranteed by the Ionescu-Tulcea extension theorem.

From the domination assumption on the distributions ((28)), it is immediate that for any i , \mathbb{P}_i is absolutely continuous w.r.t. \mathbb{Q} , let p_i denote its Radon-Nikodym density. Let $\perp \notin \mathcal{X}$ be a "fresh" element. For any $i \in I$ and $n \in \mathbb{N}$, let

$$\begin{aligned} Z_n^{(i)} : \Omega &\longrightarrow \mathcal{X} \cup \{\perp\} \\ \omega &\longmapsto Y_n(\omega) && \text{if } U_n(\omega) \in [0, p_i(Y_n(\omega))/M] \\ \omega &\longmapsto \perp && \text{otherwise.} \end{aligned}$$

For any $i \in I$, let $X_i := Z_{n_0(i)}^{(i)}$ where $n_0(i)$ is the smallest index n such that $Z_n^{(i)} \neq \perp$. If such an index does not exist, simply assign to X_i a fixed value d in \mathcal{X} .

Then, $(X_i)_{i \in I}$ is a coupling between $(\mathbb{P}_i)_{i \in I}$ that satisfies (27). Furthermore, if it is possible to simulate independent sampling from \mathbb{Q} and if it is possible to evaluate $p_i(\cdot)$ \mathbb{Q} almost anywhere for any $i \in I$, then this construction gives a straightforward sampling algorithm for $(X_i)_{i \in I}$.

E.4 Proof of correctness

E.4.1 Measurability

As a safety check, we verify the measurability of the different quantities that were introduced.

For any $i \in I$ and $n \in \mathbb{N}$, $Z_n^{(i)}$ is measurable on $\mathcal{X} \cup \{\perp\}$ for the σ -algebra induced by $\mathcal{B} \cup \{\{\perp\}\}$ (it is thus a random variable). Indeed,

$$(Z_n^{(i)} = \perp) = \cup_{q \in \mathbb{Q} \cap [0,1]} \left(Y_n \in \left(\frac{p_i}{M} \right)^{-1}([0, q]) \right) \cap (U_n \in (q, +\infty)) ,$$

and for any $A \in \mathcal{B}$,

$$(Z_n^{(i)} \in A) = (Y_n \in A) \cap (Z_n^{(i)} \neq \perp) .$$

Here, \mathbb{Q} was locally used to refer to the set of rational numbers instead of the dominating measure.

As a consequence, for any i , X_i is measurable on $(\mathcal{X}, \mathcal{B})$ (and is thus also a random variable). Indeed, for any $i \in I$ and $A \in \mathcal{B}$,

$$(X_i \in A) = \cup_{n \in \mathbb{N}} (Z_0^{(i)} = \dots = Z_{n-1}^{(i)} = \perp) \cap (Z_n^{(i)} \in A)$$

if the default value d is not in A , and

$$(X_i \in A) = \left(\cap_{n \in \mathbb{N}} (Z_n^{(i)} = \perp) \right) \cup \left(\cup_{n \in \mathbb{N}} (Z_0^{(i)} = \dots = Z_{n-1}^{(i)} = \perp) \cap (Z_n^{(i)} \in A) \right)$$

otherwise.

E.4.2 $(X_i)_{i \in I}$ is a coupling between $(\mathbb{P}_i)_{i \in I}$

For any $i \in I$, the sequence $Z_1^{(i)}, Z_2^{(i)}, \dots$ is independent⁵. Finally, for any $i \in I$ and $n \in \mathbb{N}$, and any $A \in \mathcal{B}$, $1_A(Z_n^{(i)})$ is a Bernoulli random variable of probability of success

$$\begin{aligned} \mathbb{P}(Z_n^{(i)} \in A) &= \int_{\mathcal{X} \times [0,1]} 1_A(y) 1_{[0, p_i(y)/M]}(u) du d\mathbb{Q}(y) \\ &\stackrel{\text{Fubini}}{=} \frac{1}{M} \int_{\mathcal{X}} 1_A(y) p_i(y) d\mathbb{Q}(y) \\ &= \frac{1}{M} \int_{\mathcal{X}} 1_A(y) d\mathbb{P}_i(y) \\ &= \frac{\mathbb{P}_i(A)}{M} . \end{aligned} \tag{29}$$

In particular, for any $i \in I$ and $n \in \mathbb{N}$,

$$\mathbb{P}(Z_n^{(i)} = \perp) = \left(1 - \frac{1}{M} \right) \tag{30}$$

⁵Again, this must be understood as a stochastic process.

For a $A \in \mathcal{B}$, and $i \in I$,

$$\begin{aligned} (X_i \in A) &\supset (Z_0^{(i)} \in A) \\ &\cup (Z_0^{(i)} = \perp) \cap (Z_1^{(i)} \in A) \\ &\cup (Z_0^{(i)} = Z_1^{(i)} = \perp) \cap (Z_2^{(i)} \in A) \\ &\cup \dots \end{aligned}$$

Hence, for any $A \in \mathcal{B}$, and any $i \in I$,

$$\begin{aligned} \mathbb{P}(X_i \in A) &\geq \sum_{n \in \mathbb{N}} \mathbb{P} \left((Z_0^{(i)} = \dots = Z_{n-1}^{(i)} = \perp) \cap (Z_n^{(i)} \in A) \right) \\ &= \sum_{n \in \mathbb{N}} \mathbb{P}(Z_0^{(i)} = \perp) \dots \mathbb{P}(Z_{n-1}^{(i)} = \perp) \mathbb{P}(Z_n^{(i)} \in A) \\ &\stackrel{(29) \ \& \ (30)}{=} \sum_{n \in \mathbb{N}} \left(1 - \frac{1}{M}\right)^n \frac{\mathbb{P}_i(A)}{M} \\ &= \mathbb{P}_i(A). \end{aligned}$$

Since this is true for any $A \in \mathcal{B}$ and $i \in I$, $\mathbb{P}_{X_i} = \mathbb{P}_i$ for any $i \in I$. Hence, $(X_i)_{i \in I}$ is a coupling between $(\mathbb{P}_i)_{i \in I}$.

E.4.3 $(X_i)_{i \in I}$ satisfies (27)

For any pair $i, j \in I$, and any $n \in \mathbb{N}$, using the classical equality $1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j) = \int_{\mathcal{X}} (p_i \vee p_j)(y) d\mathbb{Q}(y)$,

$$\begin{aligned} \mathbb{P}(Z_n^{(i)} = Z_n^{(j)} = \perp) &= \int_{\mathcal{X} \times [0,1]} 1_{((p_i \vee p_j)(y)/M, 1]}(u) du d\mathbb{Q}(y) \\ &= \int_{\mathcal{X} \times [0,1]} 1 - 1_{[0, (p_i \vee p_j)(y)/M]}(u) du d\mathbb{Q}(y) \\ &\stackrel{\text{Fubini}}{=} \int_{\mathcal{X}} 1 - \frac{(p_i \vee p_j)(y)}{M} d\mathbb{Q}(y) \\ &= 1 - \frac{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{M}, \end{aligned} \tag{31}$$

and using the other classical equality $1 - \text{TV}(\mathbb{P}_i, \mathbb{P}_j) = \int_{\mathcal{X}} (p_i \wedge p_j)(y) d\mathbb{Q}(y)$,

$$\begin{aligned} \mathbb{P}(Z_n^{(i)} = Z_n^{(j)} \neq \perp) &= \int_{\mathcal{X} \times [0,1]} 1_{[0, (p_i \wedge p_j)(y)/M]}(u) du d\mathbb{Q}(y) \\ &\stackrel{\text{Fubini}}{=} \int_{\mathcal{X}} \frac{(p_i \wedge p_j)(y)}{M} d\mathbb{Q}(y) \\ &= \frac{1 - \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{M}. \end{aligned} \tag{32}$$

Finally, for any $i, j \in I$, defining

$$\begin{aligned} E_{i,j} &= (Z_0^{(i)} = Z_0^{(j)} \neq \perp) \\ &\cup (Z_0^{(i)} = Z_0^{(j)} = \perp) \cap (Z_1^{(i)} = Z_1^{(j)} \neq \perp) \\ &\cup (Z_0^{(i)} = Z_0^{(j)} = \perp) \cap (Z_1^{(i)} = Z_1^{(j)} = \perp) \cap (Z_2^{(i)} = Z_2^{(j)} \neq \perp) \\ &\cup \dots, \end{aligned}$$

one has

$$X_i(\omega) = X_j(\omega), \quad \forall \omega \in E_{i,j}$$

and

$$\begin{aligned}
\mathbb{P}(E_{i,j}) &= \sum_{n \in \mathbb{N}} \mathbb{P} \left(\bigcap_{k=0}^{n-1} (Z_k^{(i)} = Z_k^{(j)} = \perp) \cap (Z_n^{(i)} = Z_n^{(j)} \neq \perp) \right) \\
&= \sum_{n \in \mathbb{N}} \left(\prod_{k=0}^{n-1} \mathbb{P}(Z_k^{(i)} = Z_k^{(j)} = \perp) \right) \mathbb{P}(Z_n^{(i)} = Z_n^{(j)} \neq \perp) \\
&\stackrel{(31) \ \& \ (32)}{=} \sum_{n \in \mathbb{N}} \left(1 - \frac{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{M} \right)^n \frac{1 - \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{M} \\
&= \frac{1 - \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}{1 + \text{TV}(\mathbb{P}_i, \mathbb{P}_j)}.
\end{aligned}$$

This concludes the proof.