Utility/Privacy Trade-off through the lens of Optimal Transport

Etienne Boursier
Université Paris-Saclay, ENS Paris-Saclay
CNRS, Centre Borelli, Cachan, France
eboursie@ens-paris-saclay.fr

Vianney Perchet
CREST, ENSAE Paris, Palaiseau, France
Criteo AI Lab, Paris, France
vianney.perchet@normalesup.org

Abstract

Strategic information is valuable either by remaining private (for instance if it is sensitive) or, on the other hand, by being used publicly to increase some utility. These two objectives are antagonistic and leaking this information might be more rewarding than concealing it. Unlike classical solutions that focus on the first point, we consider instead agents that optimize a natural trade-off between both objectives. We formalize this as an optimization problem where the objective mapping is regularized by the amount of information revealed to the adversary (measured as a divergence between the prior and posterior on the private knowledge). Quite surprisingly, when combined with the entropic regularization, the Sinkhorn loss naturally emerges in the optimization objective, making it efficiently solvable. We apply these techniques to preserve some privacy in online repeated auctions.

1 Introduction

In many economic mechanisms and strategic games involving different agents, asymmetries of information (induced by a private type, some knowledge on the hidden state of Nature, etc.) can and should be leveraged to increase one's utility. When these interactions between agents are repeated over time, preserving some asymmetry (i.e., not revealing private information) can be crucial to guarantee a larger utility in the long run. Indeed, the small short term utility of publicly using information can be overwhelmed by the long term effect of revealing it (Aumann et al., 1995).

Informally speaking, an agent should use, and potentially reveal, some private information only if she gets a subsequent utility increase in return. Keeping this information private is no longer a constraint (as in other classical privacy concepts such as differential privacy Dwork et al., 2006) but becomes part of the objective, which is then to decide how and when to use it. For instance, it might happen that revealing everything is optimal or, on the contrary, that a non-revealing policy is the best one. This is roughly similar to a poker player deciding whether to bluff or not. In some situations, it might be interesting to focus solely on the utility even if it implies losing the whole knowledge advantage, while in other situations, the immediate profit for using this advantage is so small that playing independently of it (or bluffing) is better.

After a rigorous mathematical formulation of this utility vs. privacy trade-off, it appears that this problem can be recast as a regularized optimal transport minimization. In the specific case of entropic regularization, this problem has received a lot of interest in the recent years as it induces a computationally tractable way to approximate an optimal transport distance between distributions and has thus been used in many applications (Cuturi, 2013). Our work showcases how the new Privacy Regularized Policy problem benefits in practice from this theory.

Private Mechanisms. Differential privacy is the most widely used private learning framework (Dwork, 2011; Dwork et al., 2006; Reed and Pierce, 2010) and ensures that the output of an algorithm does not significantly depend on a single element of the whole dataset. These privacy constraints are often too strong for economic applications (as illustrated before, it is sometimes optimal to disclose publicly some private information). f-divergence privacy costs have thus been proposed in recent literature as a promising alternative (Chaudhuri et al., 2019). These f-divergences, such as Kullback-Leibler, are also used by economists to measure the cost of information from a Bayesian perspective, as in the rational inattention literature (Sims, 2003; Matějka and McKay, 2015; Maćkowiak and Wiederholt, 2015). It was only recently that this ap-

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

proach has been considered to measure “privacy losses” in economic mechanisms (Eilat et al., 2019). This model assumes that the designer of the mechanism has some prior belief on the unobserved and private information. After observing the action of the player, this belief is updated and the cost of information corresponds to the KL between the prior and posterior distributions of this private information.

Optimal privacy preserving strategies with privacy constraints have been recently studied in this setting under specific conditions (Eilat et al., 2019). Loss of privacy can however be directly considered as a cost in the overall objective and an optimal strategy reveals information only if it actually leads to a significant increase in utility, whereas constrained strategies systematically reveal as much as allowed by the constraints, without incorporating the additional cost of this revelation.

Optimal Transport. Finding an appropriate way to compare probability distributions is a major challenge in learning theory. Optimal Transport manages to provide powerful tools to compare distributions in metric spaces (Villani, 2008). As a consequence, it has received an increasing interest these past years (Santambrogio, 2015), especially for generative models (Arjovsky et al., 2017; Genevay et al., 2018; Salimans et al., 2018). However, such powerful distances often come at the expense of heavy and intractable computations, which might not be suitable to learning algorithms. It was recently showcased that adding an entropic regularization term enables fast computations of approximated distances using Sinkhorn algorithm (Sinkhorn, 1967; Cuturi, 2013). Since then, the Sinkhorn loss has also shown promising results for applications such as generative models (Genevay et al., 2016, 2018), domain adaptation (Courty et al., 2014) and supervised learning (Frogner et al., 2015), besides having nice theoretical properties (Peyré and Cuturi, 2019; Feydy et al., 2019; Genevay et al., 2019).

Contributions and Organization of the paper. The new framework of Privacy Regularized Policy is motivated by several applications, presented in Section 2 and is formalized in Section 3. This problem is mathematically formulated as some optimization problem (yet eventually in an infinite dimensional space), which is convex if the privacy cost is an f-divergence, see Section 4. Also, if the private information space is discrete, this problem admits an optimal discrete distribution. The minimization problem then becomes dimensionally finite, but non-convex.

If the Kullback-Leibler divergence between the prior and the posterior is considered for the cost of information, the problem becomes a Sinkhorn loss minimization problem. Optimal transport techniques are developed in Section 5 (based on recent machinery) to compute partially revealing policies. Finally, with a linear utility cost, the problem is equivalent to the minimization of the difference of two convex functions. Using the theories of these specific problems, different optimization methods can be compared, which illustrates the practical aspect of our new model. This is done in Section 6 where we also compute partially revealing strategies for repeated auctions.

2 Some Applications

Our model is motivated by different applications described in this section: online repeated auctions and learning models on external servers.

2.1 Online repeated auctions

When a website wants to sell an advertisement slot, firms such as Google or Criteo take part in an auction to buy this slot for one of their customer, a process illustrated in Figure 1. As this interaction happens each time a user lands on the website, this is no longer a one-time auction problem, but repeated auctions where the seller and/or the competitor might observe not just one bid, but a distribution of bids. As a consequence, if a firm were bidding truthfully, seller and other bidders would have access to its true value distribution μ. This has two possible downsides.

First, if the value distribution μ was known to the auctioneer, she could maximize her revenue at the expense of the bidder utility (Amin et al., 2013, 2014; Feldman et al., 2016; Golrezaei et al., 2019), for instance with personalized reserve prices. Second, the auctioneer can sometimes take part in the auction and becomes a direct concurrent of the bidder (this might be a unique characteristic of online repeated auctions for ads). For instance, Google is both running some auction platforms and bidding on some ad slots for their client. As a consequence, if the distribution μ was perfectly known to some concurrent bidder, he could use it in the future, by bidding more or less aggressively or by trying to conquer new markets.

It is also closely related to online pricing or repeated posted price auctions. When a user wants to buy a flight ticket (or any other good), the selling company can learn the value distribution of the buyer and then dynamically adapts its prices in order to increase its revenue. The user can prevent this behavior in order to maximize her long term utility, even if it means refusing some apparently good offers in the short term (in poker lingo, she would be “bluffing”).

As explained in Section 3.1 below, finding the best possible long term strategy is intractable, as the auctioneer could always adapt to the bidding strategy, leading to

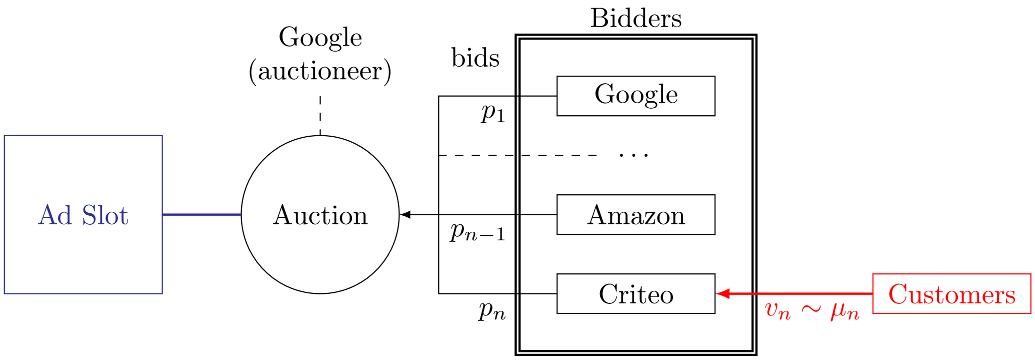

Figure 1: Online advertisement auction system.

an arm race where each bidder and auctioneer successively adapts to the other one's strategy. Such an arm race is instead avoided by trading-off between the best possible response to the auctioneer's fixed strategy as well as the leaked quantity of information. The privacy loss then aims at bounding the incurred loss in bidder's utility if the auctioneer adapts her strategy using the revealed information.

2.2 Learning through external servers

Nowadays, several servers or clusters allow their clients to perform heavy computations remotely, for instance to learn some model parameters (say a deep neural net) for a given training set. The privacy concern when querying a server can sometimes be handled using homomorphic encryption (Gilad-Bachrach et al., 2016; Bourse et al., 2018; Sanyal et al., 2018), if the cluster is designed in that way (typically a public model has been learned on the server). In this case, the client sends an encrypted testing set to the server, receives encrypted predictions and locally recovers the accurate ones. This technique, when available, is powerful, but requires heavy local computations.

Consider instead a client wanting to learn a new model (say, a linear/logistic regression or any neural net) on a dataset that has some confidential component. Directly sending the training set would reveal the whole data to the server owner, besides the risk of someone else observing it. The agent might instead prefer to send perturbed datasets, so that the computed model remains close to the accurate one, while keeping secret the true data. If the data contain sensitive information on individuals, then differential privacy is an appropriate solution. However, it is often the case that the private part is just a single piece of information of the client itself (say, its margin, its current wealth or its total number of users for instance) that is crucial to the final learned model but should not be totally revealed to a competitor. Then differential privacy is no longer the solution, as there is only a single element to protect and/or to use. Indeed, some privacy leakage is allowed and can lead to much more accurate parameters returned by the server and a higher utility at the end; the Privacy Regularized Policy aims at computing the best dataset to send to the server, in order to maximize the utility-privacy trade-off.

3 Model

We first introduce a simple toy example in Section 3.1 giving insights into the more general problem, whose formal and general formulation is given in Section 3.2.

3.1 Toy Example

Suppose an agent is publicly playing an action $x \in \mathcal{X}$ to minimize a loss $x^\top c_k$, where c_k is some loss vector. The true type $k \in [K]$ is only known to the agent and drawn from a prior p_0. Without privacy concern, the agent would then solve for every k: $\min_{x\in\mathcal{X}} x^\top c_k$. Let us denote by x_k^* the optimal solution of that problem. Besides maximizing her reward, the agent actually wants to protect the secret type k. After observing the action x taken by the agent, an adversary can update her posterior distribution of the hidden type p_x.

If the agent were to play deterministically x_k^* when her type is k, then the adversary could infer the true value of k based on the played action. The agent should instead choose her action randomly to hide her true type to the adversary. Given a type k, the strategy of the agent is then a probability distribution μ_k over $\mathcal{X}$ and her expected reward is $\mathbb{E}_{x\sim\mu_k}\left[x^\top c_k\right]$. In this case, the posterior distribution after playing the action x is computed using Bayes rule and if the different μ_k have overlapping supports, then the posterior distribution is no longer a Dirac mass, i.e., some asymmetry of information is maintained.

The agent aims at simultaneously minimizing both the utility loss and the amount of information given to the adversary. A common way to measure the latter is given by the Kullback-Leibler (KL) divergence between the prior and the posterior (Sims, 2003): $\mathrm{KL}(p_x, p_0) =$

$\sum_{k=1}^K \log\left(\frac{p_x(k)}{p_0(k)}\right) p_x(k)$, where $p_x(k) = \frac{p_0(k)\mu_k(x)}{\sum_{l=1}^K p_0(l)\mu_l(x)}$. If the information cost scales in utility with $\lambda > 0$, the regularized loss of the agent is then $x^\top c_k + \lambda \mathrm{KL}(p_x, p_0)$ instead of $x^\top c_k$. Overall, the global objective of the agent is the following minimization:

$$\min_{\mu_1,\ldots,\mu_K} \sum_{k=1}^K p_0(k)\mathbb{E}_{x\sim\mu_k}\left[x^\top c_k + \lambda \mathrm{KL}(p_x, p_0)\right].$$

In the limit case $\lambda = 0$, the agent follows a totally revealing strategy and deterministically plays x_k^* given k. When $\lambda = \infty$, the agent focuses on perfect privacy and looks for the best action chosen independently of the type: $x \perp\!\!\!\perp k$. It corresponds to a so called non-revealing strategy in game theory and the best strategy is then to play $\arg\min_x x^\top c[p_0]$ where $c[p_0] = \sum_{k=1}^K p_0(k) c_k$. For a positive λ, the behavior of the player will then interpolate between these two extreme strategies.

This problem is related to repeated games with incomplete information (Aumann et al., 1995), where players have private information affecting their utility functions. Playing some action leaks information to the other players, who then change their strategies in consequence. The goal is then to control the amount of information leaked to the adversaries in order to maximize one's own utility. In practice, it can be impossible to compute the best adversarial strategy, e.g., the player is unaware of how the adversaries would adapt. The utility loss caused by adversarial actions is then modeled as a function of the amount of revealed information.

3.2 General model

We now introduce formally the general model sketched by the previous toy example. The agent (or player) has a private type $y \in \mathcal{Y}$ drawn according to a prior p_0 whose support can be infinite. She then chooses an action $x \in \mathcal{X}$ to maximize her utility, which depends on both her action and her type. Meanwhile, she wants to hide the true value of her type y. A strategy is thus a mapping $\mathcal{Y} \to \mathcal{P}(\mathcal{X})$, where $\mathcal{P}(\mathcal{X})$ denotes the set of distributions over $\mathcal{X}$; for the sake of conciseness, we denote by $X|Y \in \mathcal{P}(\mathcal{X})^{\mathcal{Y}}$ such a strategy. In the toy example, this mapping was given by $k \mapsto \mu_k$. The adversary observes her action x and tries to infer the type of the agent. We assume a perfect adversary, i.e., she can exactly compute the posterior distribution p_x.

Let $c(x, y)$ be the utility loss for playing $x \in \mathcal{X}$ with the type $y \in \mathcal{Y}$. The cost of information is $c_{\mathrm{priv}}(X, Y)$ where (X, Y) is the joint distribution of the action and the type. In the toy example given in Section 3.1, the utility cost was given by $c(x, k) = x^\top c_k$ and the privacy cost was the expected KL divergence between p_x and p_0. The previous frameworks aimed at minimizing the utility loss with a privacy cost below some threshold $\varepsilon > 0$, i.e., minimize $\mathbb{E}_{(x,y)\sim(X,Y)}\left[c(x, y)\right]$ such that $c_{\mathrm{priv}}(X, Y) \leq \varepsilon$. Here, this privacy loss has some utility scaling with $\lambda > 0$, which can be seen as the value of information. The final objective of the agent is then to minimize the following loss:

$$\inf_{X|Y \in \mathcal{P}(\mathcal{X})^{\mathcal{Y}}} \mathbb{E}_{(x,y)\sim(X,Y)}\left[c(x, y)\right] + \lambda\, c_{\mathrm{priv}}(X, Y). \qquad (1)$$

As mentioned above, the cost of information is here defined as a measure between the posterior p_x and the prior distribution p_0 of the type, i.e., $c_{\mathrm{priv}}(X, Y) = \mathbb{E}_{x\sim X} D(p_x, p_0)$ for some function D[1]. In the toy example of Section 3.1, $D(p_x, p_0) = \mathrm{KL}(p_x, p_0)$, which is a classical cost of information in economics.

For a distribution $\gamma \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$, we denote by $\pi_{1\#}\gamma$ (resp. $\pi_{2\#}\gamma$) the marginal distribution of X (resp. Y): $\pi_{1\#}\gamma(A) = \gamma(A \times \mathcal{Y})$ and $\pi_{2\#}\gamma(B) = \gamma(\mathcal{X} \times B)$. In order to have a simpler formulation of the problem, we remark that instead of defining a strategy by the conditional distribution $X|Y$, it is equivalent to see it as a joint distribution γ of (X, Y) with a marginal over the type equal to the prior: $\pi_{2\#}\gamma = p_0$. The remaining of the paper focuses on the problem below, which we call **Privacy Regularized Policy.** With the privacy cost defined as above, the minimization problem (1) is equivalent to

$$\inf_{\substack{\gamma \in \mathcal{P}(\mathcal{X}\times\mathcal{Y}) \\ \pi_{2\#}\gamma = p_0}} \int_{\mathcal{X}\times\mathcal{Y}} \left[c(x, y) + \lambda\, D(p_x, p_0)\right] \mathrm{d}\gamma(x, y). \qquad \text{(PRP)}$$

4 A convex minimization problem

In this section, we study some theoretical properties of the Problem (PRP). We first recall the definition of an f-divergence.

Definition 1. *D is an f-divergence if for all distributions P, Q such that P is absolutely continuous w.r.t. Q, $D(P, Q) = \int_{\mathcal{Y}} f\left(\frac{\mathrm{d}P(y)}{\mathrm{d}Q(y)}\right) \mathrm{d}Q(y)$ where f is a convex function defined on $\mathbb{R}_+^*$ with $f(1) = 0$.*

The set of f-divergences includes common divergences such as the Kullback-Leibler divergence, the reverse Kullback-Leibler or the Total Variation distance.

Also, the min-entropy defined by $D(P, Q) = \log\left(\operatorname{ess\,sup} \mathrm{d}P/\mathrm{d}Q\right)$ is widely used for privacy (Tóth et al., 2004; Smith, 2009). It corresponds to the limit of the Renyi divergence $\ln\left(\sum_{i=1}^n p_i^\alpha q_i^{1-\alpha}\right)/(\alpha - 1)$, when $\alpha \to +\infty$ (Rényi, 1961; Mironov, 2017). Although it is not an f-divergence, the Renyi divergence derives from the f-divergence associated to the convex function $t \mapsto (t^\alpha - 1)/(\alpha - 1)$. f-divergence costs have been recently considered in the computer science literature in a non-Bayesian case and then present the good properties of convexity, composition and post-processing invariance (Chaudhuri et al., 2019).

In the remaining of the paper, D is an f-divergence. (PRP) then becomes a convex minimization problem.

[1] We here favor ex-ante costs as they suggest that the value of information can be heterogeneous among types.

Theorem 1. *If D is an f-divergence,* (PRP) *is a convex problem in $\gamma \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$*[2]*.*

The proof is given in Appendix A. Although $\mathcal{P}(\mathcal{X} \times \mathcal{Y})$ has generally an infinite dimension, it is dimensionally finite if both sets $\mathcal{X}$ and $\mathcal{Y}$ are discrete. A minimum can then be found using classical optimization methods such as gradient descent. In the case of low dimensional spaces $\mathcal{X}$ and $\mathcal{Y}$, they can be approximated by finite grids. However, the size of the grid grows exponentially with the dimension and another approach is needed for large dimensions of $\mathcal{X}$ and $\mathcal{Y}$.

4.1 Discrete type space

We assume here that $\mathcal{X}$ is an infinite action space and $\mathcal{Y}$ is of cardinality K (or equivalently, that p_0 is a discrete prior of size K), so that $p_0 = \sum_{k=1}^K p_0^k \delta_{y_k}$. For a fixed joint distribution γ, let the measure μ_k be defined for any $A \subset \mathcal{X}$ by $\mu_k(A) = \gamma(A \times \{y_k\})$ and $\mu = \sum_{k=1}^K \mu_k = \pi_{1\#}\gamma$. The function $p^k(x) = \frac{\mathrm{d}\mu_k(x)}{\mathrm{d}\mu(x)}$, defined over the support of μ by absolute continuity, is the posterior probability of having the type k when playing x. In this specific setting, the tuple $(\mu, (p^k)_k)$ exactly determines γ. (PRP) is then equivalent to:

$$\inf_{\substack{\mu,(p^k(\cdot))_{1\leq k \leq K} \\ p^k \geq 0, \sum_{l=1}^K p^l(\cdot)=1}} \sum_k \int_{\mathcal{X}} \left[p^k(x)c(x,y_k) + \lambda p_0^k f\left(\frac{p^k(x)}{p_0^k}\right)\right] \mathrm{d}\mu(x)$$
$$\text{such that } \forall k \leq K, \int_{\mathcal{X}} p^k(x)\mathrm{d}\mu(x) = p_0^k.$$

For fixed posterior distributions p^k, this is a generalized moment problem on the distribution μ (Lasserre, 2001). The same types of arguments can then be used for the existence and the form of optimal solutions.

Theorem 2. *If the prior is dicrete of size K, for all $\varepsilon > 0$,* (PRP) *has an ε-optimal solution such that $\pi_{1\#}\gamma = \mu$ has a finite support of at most $K+2$ points. Furthermore, if $\mathcal{X}$ is compact and $c(\cdot, y_k)$ is lower semi-continuous for every k, then it also holds for $\varepsilon = 0$.*

The proof is delayed to Appendix A. If the support of γ is included in $\{(x_i, y_k) \mid 1 \leq i \leq K+2,\ 1 \leq k \leq K\}$, we will denote it as a matrix $\gamma_{i,k} := \gamma(\{(x_i, y_k)\})$.

Corollary 1. *In the case of a discrete prior,* (PRP) *is equivalent to:*

$$\inf_{(\gamma,x)\in\mathbb{R}_+^{(K+2)\times K}\times\mathcal{X}^{K+2}} \sum_{i,k} \gamma_{i,k}\ c(x_i,y_k) + \lambda \sum_{i,k} \gamma_{i,k} D(p_{x_i}, p_0)$$
$$\textit{such that } \forall k \leq K,\ \sum_i \gamma_{i,k} = p_0^k.$$

Although it seems easier to consider the dimensionally finite problem given by Corollary 1, it is not jointly convex in (γ, x). No general algorithms exist to efficiently minimize non-convex problems. We refer the reader to (Horst et al., 2000) for an introduction to non-convex optimization.

The remaining of the paper reformulates the problem to better understand its structure, which then leads to better local minima. Computing global minima of Problem (PRP) is yet left open for future work.

5 Sinkhorn Loss minimization

Formally, (PRP) is expressed as Optimal Transport Minimization for the utility cost c with a regularization given by the privacy cost. In this section, we focus on the case where this privacy cost is the Kullback-Leibler divergence. In this case, the problem becomes a Sinkhorn loss minimization, which presents computationally tractable schemes (Peyré and Cuturi, 2019). If the privacy cost is the KL divergence between the posterior and the prior, i.e., $f(t) = t\log(t)$, then the regularization term corresponds to the mutual information $I(X;Y)$. As explained above, this is the classical cost of information in economics.

Recall that the Sinkhorn loss for given distributions $(\mu, \nu) \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})$ is defined by

$$\begin{aligned}\mathrm{OT}_{c,\lambda}(\mu,\nu) := \min_{\gamma\in\Pi(\mu,\nu)} &\int c(x,y)\mathrm{d}\gamma(x,y) \\ &+ \lambda \int \log\left(\frac{\mathrm{d}\gamma(x,y)}{\mathrm{d}\mu(x)\mathrm{d}\nu(y)}\right)\mathrm{d}\gamma(x,y),\end{aligned} \tag{2}$$

where $\Pi(\mu,\nu) = \{\gamma \in \mathcal{P}(\mathcal{X} \times \mathcal{Y}) \mid \pi_{1\#}\gamma = \mu \text{ and } \pi_{2\#}\gamma = \nu\}$. Problem (PRP) with $D = \mathrm{KL}$ can then be rewritten as the following Optimal Transport minimization:

$$\inf_{\mu\in\mathcal{P}(\mathcal{X})} \mathrm{OT}_{c,\lambda}(\mu, p_0).$$

Indeed, observe that $\frac{\mathrm{d}\gamma(x,y)}{\mathrm{d}\mu(x)}$ is the posterior probability $\mathrm{d}p_x(y)$, thanks to Bayes rule. The regularization term in equation (2) then corresponds to $D(p_x, p_0)$ as $p_0 = \nu$ and $D = \mathrm{KL}$ here. The minimization problem given by equation (2) is thus equivalent to equation (PRP) with the additional constraint $\pi_{1\#}\gamma = \mu$. Minimizing without this constraint is thus equivalent to minimizing the Sinkhorn loss over all action distributions μ.

While the regularization term is usually only added to speed up the computations, it here directly appears in the cost of the original problem since it corresponds to the privacy cost! An approximation of $\mathrm{OT}_{c,\lambda}(\mu,\nu)$ can then be quickly computed for discrete distributions using Sinkhorn algorithm (Cuturi, 2013).

Notice that the definition of Sinkhorn loss sometimes differs in the literature and instead considers $\int \log(\mathrm{d}\gamma(x,y))\,\mathrm{d}\gamma(x,y)$ for the regularization term. When μ and ν are both fixed, the optimal transport plan γ remains the same. As μ is varying here, these

[2] It is convex in a usual sense and not geodesically here.

notions yet become different. For this alternative definition, a minimizing distribution μ would actually be easy to compute. It is much more complex in our problem because of the presence of μ in the denominator of the logarithmic term.

In the case of discrete support, we can then look for a distribution $\mu = \sum_{j=1}^{K+2} \alpha_j \delta_{x_j}$. In case of continuous distributions, they could still be approximated using sampled discrete distributions as previously done for generative models (Genevay et al., 2019, 2018).

Besides being a new interpretation of Sinkhorn loss, this reformulation mainly allows to better understand the problem structure and reduce the support size of the distribution in the minimization problem.

5.1 Minimization algorithm

We now consider the following minimization problem over the tuple (α, x):

$$\inf_{(\alpha,x)\in\Delta_{K+2}\times\mathcal{X}^{K+2}} \mathrm{OT}_{c,\lambda}\Big(\sum_{i=1}^{K+2}\alpha_i\delta_{x_i}, p_0\Big). \qquad (3)$$

The main difficulties come from the computation of the objective function and its gradient to use classical gradient based methods.

Sinkhorn algorithm. It was recently suggested to use the Sinkhorn algorithm, which has a linear convergence rate, to compute $\mathrm{OT}_{c,\lambda}(\mu,\nu)$ for distributions $\mu = \sum_{i=1}^{n}\alpha_i\delta_{x_i}$ and $\nu = \sum_{j=1}^{m}\beta_j\delta_{y_j}$ (Knight, 2008; Cuturi, 2013). Let K be the exponential cost matrix defined by $K_{i,j} = e^{-\frac{c(x_i,y_j)}{\lambda}}$. In the discrete case, the unique matrix γ solution of the Problem (2) has the form $\mathrm{diag}(u)K\mathrm{diag}(v)$. The Sinkhorn algorithm then updates $(u,v) \leftarrow (\alpha/Kv, \beta/K^\top u)$ (with componentwise division) for n iterations or until convergence.

Gradient computation. Computing $\nabla\mathrm{OT}_{c,\lambda}$ is a known difficult task (Feydy et al., 2019; Luise et al., 2018; Genevay et al., 2018). A simple solution consists in using automatic differentiation, i.e., computing the gradient using the chain rule over the simple successive operations computed during the Sinkhorn algorithm.

The gradient can also be computed from the dual solution of Problem (2). This method is faster as it does not need to store all the Sinkhorn iterations in memory and *backpropagate* through them afterwards. Convergence of Sinkhorn algorithm has yet to be guaranteed to provide an accurate approximation of the gradient (see Peyré and Cuturi, 2019, for an extended discussion). Automatic differentiation is used in the experiments because of this last reason.

6 Experiments and particular cases

In this section, the case of linear utility cost is first considered and shown to have relations with DC programming, which allows efficient algorithms. The performances of different optimization schemes are then compared on a simple example. Simulations based on the Sinkhorn scheme are then run for the real problem of online repeated auctions. The code is available at `github.com/eboursier/regularized_private_learning`.

6.1 Linear utility cost

Section 4 described a general optimization scheme for Problem (PRP) with a discrete type prior. It used a dimensionally finite, non-convex problem. An objective is then to find a local minimum. Local minima can be found using classical techniques of gradient descent (Wright, 2015). However in some particular cases, better schemes are possible as claimed in Section 5 for the particular case of entropic regularization. In the case of a linear utility for any privacy cost, it is related to DC programming (Horst et al., 2000). A standard DC program is of the form $\min_{x\in\mathcal{X}} f(x) - g(x)$, where both f and g are convex functions. Specific optimization schemes are then possible (Tao and An, 1997; Horst and Thoai, 1999; Horst et al., 2000). In the case of linear utility costs over a hyperrectangle, (PRP) can be reformulated as a DC program stated in Theorem 3. Its proof is delayed to Appendix A.

Theorem 3. *If $\mathcal{X} = \prod_{l=1}^{d}[a_l,\ b_l]$ and $c(x,y) = x^\top y$, then* (PRP) *is equivalent to the following DC program:*

$$\min_{\gamma\in\mathbb{R}_+^{(K+2)\times K}} \lambda\sum_{i,k} p_0^k h_k(\gamma_i) - \sum_{i=1}^{K+2}\Big\|\sum_{k=1}^{K}\gamma_{i,k}\phi(y_k)\Big\|_1,$$

$$\textit{such that } \forall k \le K,\ \sum_{i=1}^{K+2}\gamma_{i,k} = p_0^k,$$

with $\phi(y)^l := (b_l - a_l)y^l/2$
and $h_k(\gamma_i) := \big(\sum_{m=1}^{K}\gamma_{i,m}\big) f\big(\frac{\gamma_{i,k}}{p_0^k\sum_{m=1}^{K}\gamma_{i,m}}\big)$.

More generally, if the cost c is concave and the action space $\mathcal{X}$ is a polytope, optimal actions are located on the vertices of $\mathcal{X}$. In that case, we can therefore replace $\mathcal{X}$ by the set of its vertices and the problem becomes a dimensionally finite convex problem as already claimed in Section 3.2. Unfortunately, for some polytopes such as hyperrectangles, the number of vertices grows exponentially with the dimension and the optimization scheme is no longer tractable in large dimensions.

6.2 Comparing methods on the toy example

We consider the linear utility loss $c(x,y) = x^\top y$ over the space $\mathcal{X} = [-1,1]^d$ and the Kullback-Leibler divergence for privacy cost, so that both DC and Sinkhorn schemes are possible. Different methods exist for DC programming and they compute either a local or a global minimum. We here choose the DCA algorithm (Tao and An, 1997) as it computes a local minimum

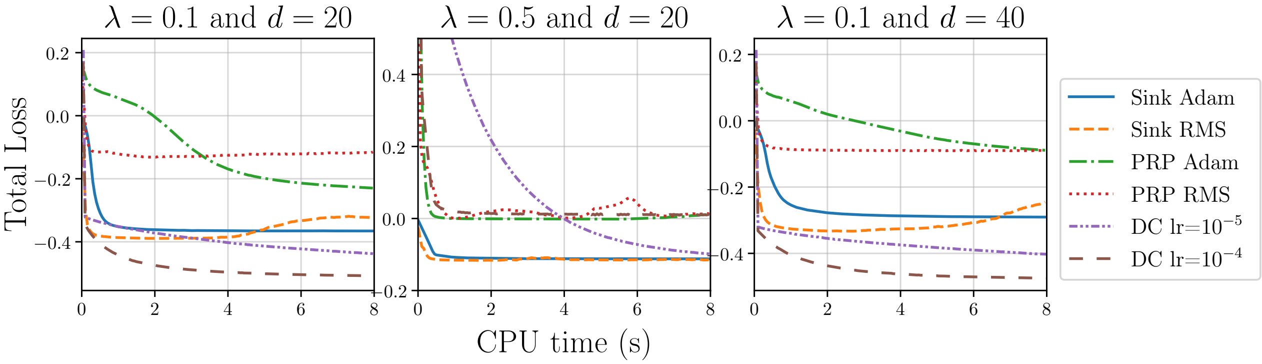

Figure 2: Comparison of optimization schemes.

and is thus comparable to the other considered schemes. Figure 2 compares, for different problem parameters, the convergence rates of usual non-convex optimization methods (ADAM and RMS), as well as DCA. The former methods are used on different minimizations given by Corrolary 1 and Equation (3) (resp. PRP and Sink).

We optimized using projected gradient descent for well tuned learning rates. The prior p_0^k is chosen proportional to e^{Z_k} for any $k \in [K]$, where Z_k is drawn uniformly at random in $[0, 1]$. Each y_i^k is taken uniformly at random in $[-1, 1]$ and is rescaled so that $\|y_i\|_1 = 1$. The values are averaged over 200 runs.

The DC method finds better local minima than the other ones. This was already observed in practice (Tao and An, 1997) and confirms that it is more adapted to the structure of the problem, despite being only applicable in very specific cases such as linear cost on hyperrectangles. Also, the PRP method converges to worse spurious local minima as it optimizes in higher dimensional spaces than the Sinkhorn method. We also observed in our experiments that PRP method is more sensitive to problem parameters than Sinkhorn method.

The Sinkhorn method seems to perform better for larger values of λ. Indeed, given the actions, the Sinkhorn method computes the best joint distribution for each iteration and thus performs well when the privacy cost is predominant, while DCA computes the best actions given a joint distribution and thus performs well when the utility cost is predominant. It is thus crucial to choose the method which is most adapted to the problem structure as it can lead to significant improvement in the solution.

6.3 Utility-privacy in repeated auctions

For repeated second price auctions following a precise scheme (Leme et al., 2016), there exist numerical methods to implement an optimal strategy for the bidder (Nedelec et al., 2019). However, if the auctioneer knows that the bidder plays this strategy, he can still infer the bidder's type and adapt to it. We thus require to add a privacy cost to avoid this kind of behavior from the auctioneer.

For simplicity, bidder's valuations are assumed to be exponential distributions, so that the private type y corresponds to the only parameter of this distribution, i.e., its expectation: $y = \mathbb{E}_{v\sim\mu_y}[v]$. Moreover, we assume that the prior p_0 over y is the discretized uniform distribution on $[0, 1]$ with a support of size $K = 10$; let $\{y_j\}_{j=1,\dots,K}$ be the support of p_0.

In repeated auctions, values v are repeatedly sampled from the distribution μ_{y_j} and a bidder policy is a mapping $\beta(\cdot)$ from values to bids, i.e., she bids $\beta(v)$ if her value is v. So a type y_j and a policy $\beta(\cdot)$ generate the bid distribution $\beta_{\#}\mu_{y_j}$, which corresponds to an action in $\mathcal{X}$ in our setting. As a consequence, the set of actions of the agent are the probability distributions over $\mathbb{R}_+$ and an action ρ_i is naturally generated from the valuation distribution via the optimal monotone transport map denoted by β_j^i, i.e., $\rho_i = \beta^i_{j\#}\mu_{y_j}$ (Santambrogio, 2015). In the particular case of exponential distributions, this implies that $\beta_i^j(v) = \beta_i(v/y_j)$ where β_i is the unique monotone transport map from Exp(1) to ρ_i. The revenue of the bidder is then deduced for exponential distributions (Nedelec et al., 2019) as

$$r(\beta_i, y_j) = 1 - c(\beta_i, y_j)$$
$$= \mathbb{E}_{v\sim\text{Exp}(1)}\left[\left(y_j v - \beta_i(v) + \beta_i'(v)\right)G\left(\beta_i(v)\right)\mathbb{1}_{\beta_i(v)-\beta_i'(v)\geq 0}\right],$$

where G is the *c.d.f.* of the maximum bid of the other bidders. We here consider a single truthful opponent with a uniform value distribution on $[0, 1]$, so that $G(x) = \min(x, 1)$. This utility is averaged over 10^3 values drawn from the corresponding distribution at each training step and 10^6 values for the final evaluation.

Considering the KL for privacy cost, we compute a strategy (γ, β) using the Sinkhorn scheme described in Section 5. Every action β_i is parametrized as a single layer neural network of 100 ReLUs. Figure 3a represents both utility and privacy as a function of the regularization factor λ.

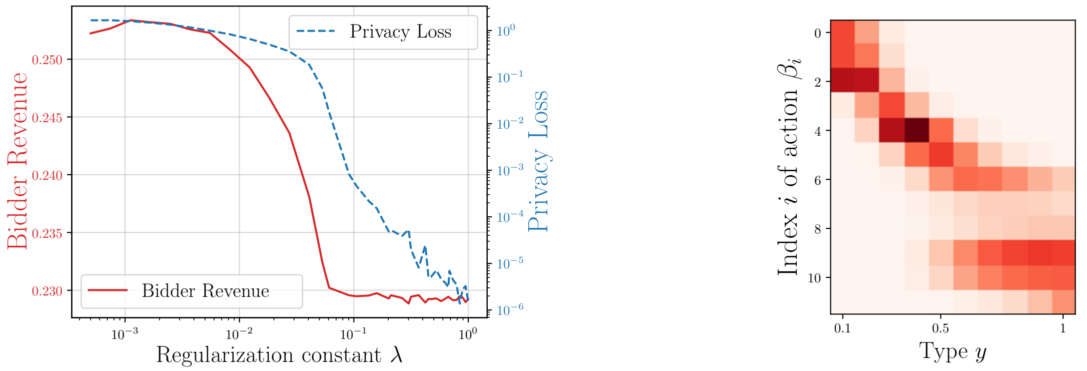

(a) Evolution of privacy-utility with λ.

(b) Joint distribution map for $\lambda = 0.01$. The intensity of a point (i, j) corresponds to the value of $\gamma(\beta_i, y_j)$.

Figure 3: Privacy-utility trade-off in online repeated auctions.

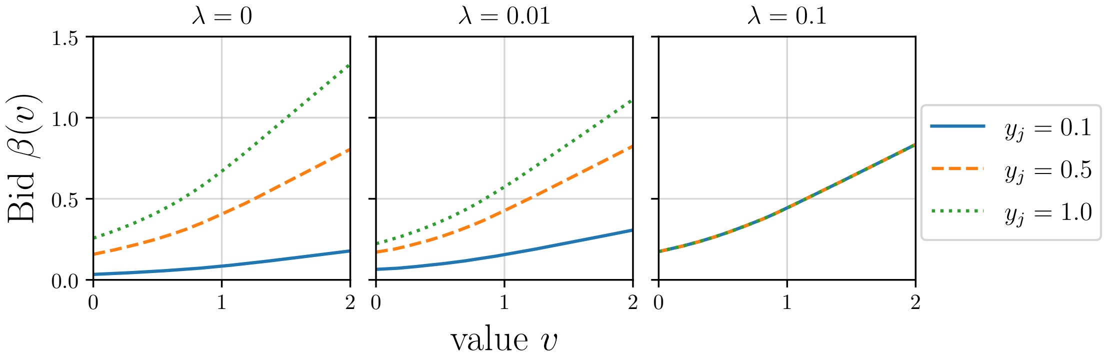

Figure 4: Evolution of the bidding strategy with the type and the regularization constant.

Naturally, both the bidder revenue and the privacy loss decrease with λ, going from revealing strategies for $\lambda \simeq 10^{-3}$ to non-revealing strategies for larger λ. They significantly drop at a critical point near 0.05, which can be seen as the cost of information here. There is a 8% revenue difference[3] between the non revealing strategy and the partially revealing strategy shown in Figure 3b. The latter randomizes the type over the neighboring types and reveals more information when the revenue is sensible to the action, i.e., for low types y_j here. This strategy thus takes advantage from the fact that the value of information is here heterogeneous among types, as desired in the design of our model.

Figure 4 shows the most used action for different types and λ. In the revealing strategy ($\lambda = 0$), the action significantly scales with the type. But as λ grows, this rescaling shrinks so that the actions perform for several types, until having a single action in the non-revealing strategy. This shrinkage is also more important for large values of y_j. This confirms the observation made above: the player loses less by hiding her type for large values than for low values and she is thus more willing to hide her type when it is large.

Besides confirming expected results, this illustrates how the Privacy Regularized Policy is adapted to complex utility costs and action spaces, such as distributions or function spaces.

7 Conclusion

We formalized a new utility-privacy trade-off problem to compute strategies revealing private information only if it induces a significant increase in utility. For classical costs in economics, it benefits from recent advances of Optimal Transport. It yet leads to a hard non-convex minimization problem and future work includes designing efficient algorithms computing global minima for this problem.

We believe that this work is a step towards the design of optimal utility vs. privacy trade-offs in economic mechanisms as well as for other applications. Its many connections with recent topics of interest motivate a better understanding of them as future work.

[3] Which is significant for large firms such as those presented in Figure 1 besides the revenue difference brought by considering non truthful strategies (Nedelec et al., 2019).

Bibliography

K. Amin, A. Rostamizadeh, and U. Syed. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pages 1169–1177, 2013.

K. Amin, A. Rostamizadeh, and U. Syed. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pages 622–630, 2014.

M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.

R. Aumann, M. Maschler, and R. Stearns. *Repeated games with incomplete information*. MIT press, 1995.

F. Bourse, M. Minelli, M. Minihold, and P. Paillier. Fast homomorphic evaluation of deep discretized neural networks. In *Annual International Cryptology Conference*, pages 483–512, 2018.

S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

K. Chaudhuri, J. Imola, and A. Machanavajjhala. Capacity bounded differential privacy. In *Advances in Neural Information Processing Systems*, pages 3469–3478, 2019.

N. Courty, R. Flamary, and D. Tuia. Domain adaptation with regularized optimal transport. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 274–289, 2014.

M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.

C. Dwork. Differential privacy. *Encyclopedia of Cryptography and Security*, pages 338–340, 2011.

C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.

R. Eilat, K. Eliaz, and X. Mu. Optimal Privacy-Constrained Mechanisms. Technical report, C.E.P.R. Discussion Papers, 2019.

M. Feldman, T. Koren, R. Livni, Y. Mansour, and A. Zohar. Online pricing with strategic and patient buyers. In *Advances in Neural Information Processing Systems*, pages 3864–3872, 2016.

J. Feydy, T. Séjourné, F.-X. Vialard, S. i. Amari, A. Trouve, and G. Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2681–2690, 2019.

C. Frogner, C. Zhang, H. Mobahi, M. Araya-Polo, and T. Poggio. Learning with a Wasserstein loss. In *Advances in Neural Information Processing Systems*, 2015.

A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems*, pages 3440–3448, 2016.

A. Genevay, G. Peyre, and M. Cuturi. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617, 2018.

A. Genevay, L. Chizat, F. Bach, M. Cuturi, and G. Peyré. Sample complexity of sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1574–1583, 2019.

R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning*, pages 201–210, 2016.

N. Golrezaei, A. Javanmard, and V. Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. In *Advances in Neural Information Processing Systems*, pages 9756–9766, 2019.

R. Horst and N. Thoai. DC programming: overview. *Journal of Optimization Theory and Applications*, 103(1):1–43, 1999.

R. Horst, P. Pardalos, and N. V. Thoai. *Introduction to global optimization*. Springer Science & Business Media, 2000.

P. Knight. The sinkhorn–knopp algorithm: convergence and applications. *SIAM Journal on Matrix Analysis and Applications*, 30(1):261–275, 2008.

J. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.

R. P. Leme, M. Pal, and S. Vassilvitskii. A field guide to personalized reserve prices. In *Proceedings of the 25th international conference on world wide web*, pages 1093–1102, 2016.

G. Luise, A. Rudi, M. Pontil, and C. Ciliberto. Differential properties of sinkhorn approximation for learning with wasserstein distance. In *Advances in Neural Information Processing Systems*, pages 5859–5870, 2018.

B. Maćkowiak and M. Wiederholt. Business cycle dynamics under rational inattention. *The Review of Economic Studies*, 82(4):1502–1532, 2015.

F. Matějka and A. McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105 (1):272–98, 2015.

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

T. Nedelec, N. E. Karoui, and V. Perchet. Learning to bid in revenue-maximizing auctions. In *International Conference on Machine Learning*, pages 4781–4789, 2019.

G. Peyré and M. Cuturi. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

J. Reed and B. Pierce. Distance makes the types grow stronger: a calculus for differential privacy. In *ACM Sigplan Notices*, volume 45, pages 157–168, 2010.

A. Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, 1961.

T. Salimans, D. Metaxas, H. Zhang, and A. Radford. Improving gans using optimal transport. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.

F. Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55:58–63, 2015.

A. Sanyal, M. Kusner, A. Gascon, and V. Kanade. Tapas: Tricks to accelerate (encrypted) prediction as a service. In *International Conference on Machine Learning*, pages 4497–4506, 2018.

C. Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690, 2003.

R. Sinkhorn. Diagonal equivalence to matrices with prescribed row and column sums. *The American Mathematical Monthly*, 74(4):402–405, 1967.

G. Smith. On the foundations of quantitative information flow. In *International Conference on Foundations of Software Science and Computational Structures*, pages 288–302, 2009.

P. Tao and L. An. Convex analysis approach to DC programming: Theory, algorithms and applications. *Acta mathematica vietnamica*, 22(1):289–355, 1997.

G. Tóth, Z. Hornák, and F. Vajda. Measuring anonymity revisited. In *Proceedings of the Ninth Nordic Workshop on Secure IT Systems*, pages 85–90, 2004.

C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

S. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015.