

Private Online Prediction from Experts: Separations and Faster Rates

Hilal Asi

Apple

HILAL.ASI94@GMAIL.COM

Vitaly Feldman

Apple

VITALY.EDU@GMAIL.COM

Tomer Koren

Tel Aviv University

TKOREN@TAUEX.TAU.AC.IL

Kunal Talwar

Apple

KUNAL@KUNALTALWAR.ORG

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

Online prediction from experts is a fundamental problem in machine learning and several works have studied this problem under privacy constraints. We propose and analyze new algorithms for this problem that improve over the regret bounds of the best existing algorithms for non-adaptive adversaries. For approximate differential privacy, our algorithms achieve regret bounds of $O(\sqrt{T \log d} + \log d/\epsilon)$ for the stochastic setting and $O(\sqrt{T \log d} + T^{1/3} \log d/\epsilon)$ for oblivious adversaries (where d is the number of experts). For pure DP, our algorithms are the first to obtain sub-linear regret for oblivious adversaries in the high-dimensional regime $d \geq T$. Moreover, we prove new lower bounds for adaptive adversaries. Our results imply that unlike the non-private setting, there is a strong separation between the optimal regret for adaptive and non-adaptive adversaries for this problem. Our lower bounds also show a separation between pure and approximate differential privacy for adaptive adversaries where the latter is necessary to achieve the non-private $O(\sqrt{T})$ regret.

1. Introduction

We study the problem of differentially private online prediction from experts (DP-OPE), where the algorithm interacts with an adversary for T rounds. In each round, the algorithm picks an expert $x_t \in [d]$ and the adversary chooses a loss function $\ell_t : [d] \rightarrow [0, 1]$. The algorithm incurs loss $\ell_t(x_t)$ at round t , and the goal is to design algorithms that minimize the regret, that is, the cumulative loss compared to the best fixed expert in hindsight, defined as

$$\text{Reg}_T = \sum_{t=1}^T \ell_t(x_t) - \min_{x^* \in [d]} \sum_{t=1}^T \ell_t(x^*).$$

Online prediction from experts is an important problem in machine learning with numerous applications (Arora et al., 2012). Without any privacy restrictions, the power of the adversary (oblivious adversary that picks the losses in advance versus adaptive adversary that picks the losses online in response to the algorithm) does not change the optimal rates for this problem (Cesa-Bianchi and Lugosi, 2006). This has perhaps led prior work in private online experts to focus on the strongest notion of adaptive adversaries (Jain et al., 2012; Smith and Thakurta, 2013; Jain and Thakurta, 2014; Agarwal and Singh, 2017). In this work, we study the problem of DP-OPE against

oblivious adversaries as well and show that, somewhat surprisingly, the power of the adversary can significantly affect the optimal rates for this problem.

We consider three types of adversaries: the strongest, *adaptive adversary*, can choose the loss sequence $(\ell_t)_{t=1}^T$ adversarially in an online manner, where the loss ℓ_t may depend arbitrarily on the choices made by the algorithm in previous time steps $(x_\ell)_{\ell=1}^{t-1}$. A slightly weaker notion is that of an *oblivious adversary*, which chooses a sequence of loss functions in advance. The weakest adversary is the *stochastic adversary* which chooses a distribution over loss functions and at each round samples a loss function i.i.d. from this distribution.

In the classical non-private setting, all of these adversarial models are equivalent, in the sense that they all induce an optimal rate of order $\sqrt{T \log d}$ (Cesa-Bianchi and Lugosi, 2006). We study algorithms that are required to be differentially private, where we view the sequence of loss functions as the dataset and adjacent datasets differ in a single loss function. We note that all our algorithms are private with respect to adaptive adversaries (see Section 2 for precise definitions) and only the utility bounds assume a weaker adversary.

Under the constraint of (ϵ, δ) -differential privacy, the best existing results obtain regret of order

$$\sqrt{T \log d} + \min \left(\frac{1}{\epsilon} \sqrt{T \log(1/\delta) \log d}, \frac{1}{\epsilon} \sqrt{d \log(1/\delta) \log d \log^2 T} \right)$$

for adaptive adversaries (Jain and Thakurta, 2014; Agarwal and Singh, 2017). For pure ϵ -differential privacy, the best existing regret (Jain and Thakurta, 2014) is of order

$$\sqrt{T \log d} + \frac{d \log d \log^2 T}{\epsilon}.$$

However, none of existing results prove any lower bounds (besides the trivial non-private lower bound $\sqrt{T \log d}$) to certify the optimality of these rates; thus, it is currently unclear whether these rates are optimal for adaptive adversaries, let alone for oblivious and stochastic adversaries.

1.1. Our contributions

We study DP-OPE for different types of adversaries and develop new algorithms and lower bounds. More precisely, we obtain the following results.

Faster rates for oblivious adversaries (Section 3.1): We develop new algorithms for DP-OPE with oblivious adversaries based on a lazy version of the multiplicative weights algorithm. For pure ϵ -DP, our algorithms obtain regret

$$\frac{\sqrt{T \log d}}{\epsilon}.$$

This is the first algorithm to achieve sub-linear regret for pure DP in the high-dimensional regime $d \geq T$. For approximate (ϵ, δ) -DP, our algorithm achieves regret

$$\sqrt{T \log d} + \frac{T^{1/3} \log^{1/3}(1/\delta) \log d}{\epsilon}.$$

This essentially demonstrates that the privacy cost for DP-OPE is negligible as long as $\epsilon \gg T^{-1/6}$. In contrast, previous work has established privacy cost $\epsilon^{-1} \sqrt{T \log d \log(1/\delta)}$ which is larger than the non-private cost even when ϵ is constant and $\delta = 1/T$. See Tables 1a and 1b for more details.

Separation between adaptive and oblivious adversaries (Section 4): We prove the first lower bounds for DP-OPE with adaptive adversaries that are stronger than the non-private lower bounds. These bounds show that any private algorithm must suffer linear regret if $\varepsilon \leq 1/\sqrt{T}$ for approximate DP and $\varepsilon \leq 1/10$ for pure DP. As our algorithms for oblivious adversaries obtain sub-linear regret in this regime of privacy parameters, this demonstrates that the oblivious model is significantly weaker than the adaptive model in the private setting (see Fig. 1). Moreover, these lower bounds show a separation between pure and approximate DP for DP-OPE with adaptive adversaries as the latter is necessary to obtain sub-linear regret.

Near-optimal rates for stochastic adversaries (Section 3.2): We design a general reduction that transforms any algorithm for private stochastic optimization in the offline setting into an algorithm for private online optimization with nearly the same rates (up to logarithmic factors). By building on algorithms for the offline setting (Asi et al., 2021b), we obtain regret $O(\sqrt{T \log d} + \varepsilon^{-1} \log d \log T)$ for DP-OPE with stochastic adversaries. Moreover, using this reduction with general algorithms for differentially private stochastic convex optimization (DP-SCO) (Feldman et al., 2020), we obtain near-optimal regret $O(\sqrt{T} + \varepsilon^{-1} \sqrt{d} \log T)$ for the problem of DP-OCO (online convex optimization) with stochastic adversaries, improving over the best existing result $\sqrt{T} d^{1/4} / \varepsilon$ (Kairouz et al., 2021).

Improved rates for DP-OCO (Section 5): Building on our improvements for DP-OPE, we improve the existing rates for DP-OCO where the algorithm picks $x_t \in \mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$ and the adversary picks $\ell_t : \mathcal{X} \rightarrow \mathbb{R}$. Our rates improve over the rates of Kairouz et al. (2021) in certain regimes.

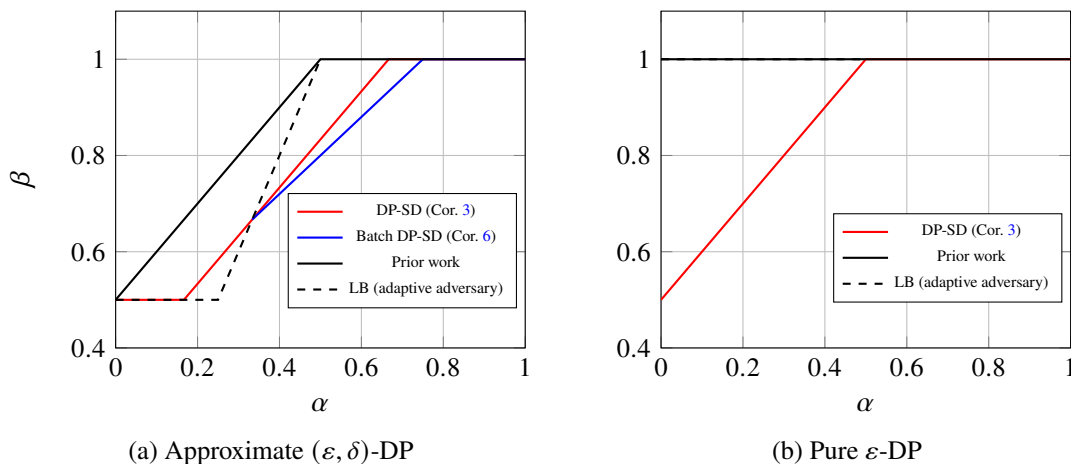


Figure 1: Regret bounds for private prediction from experts against oblivious adversaries for the high-dimensional regime $d \geq T$. We denote the privacy parameter $\varepsilon = T^{-\alpha}$ and regret T^β , and plot β as a function of α (ignoring logarithmic factors).

1.2. Related work

Dwork et al. (2010) were the first to study differentially private learning in the online setting and introduced the binary tree mechanism which is an important building block of many private algorithms in the online setting. In our context of online prediction from experts, there has been several

	Prior work	This work
Stochastic	$\sqrt{T \log d} + \min\left(\frac{\sqrt{T \log d}}{\varepsilon}, \frac{\sqrt{d \log d}}{\varepsilon}\right)$	$\sqrt{T \log d} + \frac{\log d}{\varepsilon}$
Oblivious		$\sqrt{T \log d} + \frac{T^{1/3} \log d}{\varepsilon}$
Adaptive		None

(a) Approximate (ε, δ) -DP.

	Prior work	This work
Stochastic	$\sqrt{T \log d} + \frac{d \log d}{\varepsilon}$	$\sqrt{T \log d} + \frac{\log d}{\varepsilon}$
Oblivious		$\frac{\sqrt{T} \log d}{\varepsilon}$
Adaptive		None

(b) Pure ε -DP.
 Table 1: Upper bounds for DP-OPE with different types of adversaries. For readability, we omit logarithmic factors that depend on T and $1/\delta$.

works that study this problem under the constraint of differential privacy (Jain et al., 2012; Smith and Thakurta, 2013; Jain and Thakurta, 2014; Agarwal and Singh, 2017). The best existing algorithms depend on the dimensionality regime: in the high-dimensional setting, Jain and Thakurta (2014) developed an algorithm based on follow-the-regularized-leader (FTRL) with entropy regularization that achieves regret $O(\varepsilon^{-1} \sqrt{T \log d \log(1/\delta)})$ for (ε, δ) -DP. For low dimensional problems, Agarwal and Singh (2017) developed an improved algorithm that uses the binary tree mechanism to estimate the running sum of the gradients in the FTRL framework. Their algorithm achieves regret $O(\sqrt{T \log d} + \varepsilon^{-1} d \log d \log^2 T)$ for ε -DP. Moreover, extending their algorithm to (ε, δ) -DP results in regret $O(\sqrt{T \log d} + \varepsilon^{-1} \sqrt{d \log(1/\delta)} \log d \log^2 T)$. More recently, Asi et al. (2023) considered DP-OPE in the realizable regime where there is a zero-loss expert, and propose new algorithms that obtain near-optimal regret $(poly)(\log d)/\varepsilon$ up to logarithmic factors.

A slightly related and somewhat easier problem is that of differentially private stochastic convex optimization (DP-SCO) which has been extensively investigated recently (Bassily et al., 2014, 2019; Feldman et al., 2020; Asi et al., 2021b,a). In DP-SCO, we are given n samples from some distribution and we wish to minimize the excess population loss. In contrast to the online setting, here all n samples are given to the algorithm and it is required to produce only one model; this makes the online version harder as the algorithm has to output a model at each time-step. However, we note that our reduction in Section 3.2 shows that DP-SCO is essentially as hard as DP-OCO with stochastic adversaries (up to logarithmic factors). For oblivious and adaptive adversaries, the online setting may be harder as it allows general loss functions that are not necessarily generated from a distribution.

Perhaps most closely related to our problem in the non-private setting is online learning with limited switching (Kalai and Vempala, 2005; Geulen et al., 2010; Altschuler and Talwar, 2018; Chen et al., 2020; Sherman and Koren, 2021). In this setting, the algorithm aims to minimize the regret when having a budget of at most S switches, that is, it can update its decision at most S times. This constraint is (informally) somewhat related to privacy as the less you update the model, the less information is leaked about the data. The ideas developed in this literature turn out to be useful for our DP-OPE setting. Indeed, our algorithms in Section 3.1 build on ideas from Geulen et al. (2010) which developed a lazy version of multiplicative weights to limit the number of switches. Moreover, similarly to our results, the hardness of online learning problems with limited switching can depend on the power of the adversary. For example, for OCO with limited switching, the optimal rate is $\sqrt{T} + T/S$ for oblivious adversaries and $\sqrt{T} + T/\sqrt{S}$ for adaptive adversaries (Chen et al., 2020; Sherman and Koren, 2021). Despite of these similarities, our results do not establish a fundamental connection between privacy and switching constraints and we leave this as an open question for future research.

2. Problem setting and preliminaries

Online prediction from experts (OPE) is an interactive T -round game between an online algorithm \mathcal{A} and adversary Adv . At round t , the algorithm \mathcal{A} chooses an expert $x_t \in [d]$ and the adversary Adv picks a loss function $\ell_t \in \mathcal{L} = \{\ell \mid \ell : [d] \rightarrow [0, 1]\}$ simultaneously. We let $\mathcal{A}_t(\ell_1, \dots, \ell_{t-1}) = x_t$ denote the mapping of algorithm \mathcal{A} at round t . Similarly, we define Adv_t to be the mapping of the adversary at round t (we provide more details on the input of Adv_t below depending on the type of the adversary). The algorithm observes ℓ_t (after choosing x_t) and incurs loss $\ell_t(x_t)$. For a predefined number of rounds T , the regret of the algorithm \mathcal{A} is defined as

$$\text{Reg}_T(\mathcal{A}) = \sum_{t=1}^T \ell_t(x_t) - \min_{x^* \in [d]} \sum_{t=1}^T \ell_t(x^*).$$

We consider three types of adversary Adv for choosing the sequence of loss functions $\{\ell_t\}_{t=1}^T$. To help define privacy, we will diverge from the traditional presentation of these models in online learning literature. The adversary will consist of a sequence of T data points $z_1, \dots, z_T \in \mathcal{Z}$, and an algorithm ℓ that generates the sequence of losses.

For both stochastic and oblivious adversaries the loss function at step t is generated based on data point z_t alone; i.e. $\ell_t(\cdot) = \ell(\cdot; z_t)$, where for all $z \in \mathcal{Z}$, $\ell(\cdot; z)$ is an admissible loss function for the relevant setup. The two models differ in the choice of the sequence z_1, \dots, z_T : for a stochastic adversary, the sequence of z_i 's is chosen i.i.d. from some distribution P (chosen by the adversary). For an oblivious adversary, this sequence itself is adversarially chosen. In other words:

$$\begin{aligned} \text{Reg}_T^{(\text{stochastic})}(\mathcal{A}) &= \sup_{\ell, P} \mathbb{E}_{z_1, \dots, z_T \sim P^T} [\text{Reg}_T(\mathcal{A}) \mid \ell_t(\cdot) = \ell(\cdot; z_t)], \\ \text{Reg}_T^{(\text{oblivious})}(\mathcal{A}) &= \sup_{\ell} \sup_{z_1, \dots, z_T \in \mathcal{Z}^T} [\text{Reg}_T(\mathcal{A}) \mid \ell_t(\cdot) = \ell(\cdot; z_t)]. \end{aligned}$$

In the case of an adaptive adversary, the loss at step t can depend on the algorithm's choices in previous steps. Thus $\ell_t(\cdot) = \ell(\cdot; z_t, x_{1:t-1})$, where as before this loss function is constrained to be

an admissible loss function for all possible values of inputs $z_t, x_{1:t-1}$. The adaptive regret is then the worst case regret over z_1, \dots, z_T and the mapping ℓ :

$$Reg_T^{(\text{adaptive})}(\mathcal{A}) = \sup_{\ell} \sup_{z_1, \dots, z_T \in \mathcal{Z}^T} [Reg_T(\mathcal{A}) | \ell_t(\cdot) = \ell(\cdot; z_t, x_{1:t-1})].$$

Given data $z_{1:T} = (z_1, \dots, z_T) \in \mathcal{Z}^T$, we let $\mathcal{A} \circ \text{Adv}(z_{1:T}) = x_{1:T}$ denote the output of the interaction between the algorithm \mathcal{A} and adversary Adv given inputs $z_{1:T}$.

Under this setting, the goal is to design private algorithms that minimize the appropriate notion of regret. To this end, we extend the standard definition of (ϵ, δ) -differential privacy (Dwork et al., 2006b,a) to the online setting. Like most previous works, we study a stronger notion of privacy that holds even against adaptive adversaries.¹

Definition 1 (Adaptive DP). A randomized algorithm \mathcal{A} is (ϵ, δ) -differentially private against adaptive adversaries ((ϵ, δ) -DP) if, for all sequences $\mathcal{S} = (z_1, \dots, z_T) \in \mathcal{Z}^T$ and $\mathcal{S}' = (z'_1, \dots, z'_T) \in \mathcal{Z}^T$ that differ in a single element, for any ℓ defining an adaptive adversary Adv , and for all events O in the output space of $\mathcal{A} \circ \text{Adv}$, we have

$$\Pr[\mathcal{A} \circ \text{Adv}(\mathcal{S}) \in O] \leq e^\epsilon \Pr[\mathcal{A} \circ \text{Adv}(\mathcal{S}') \in O] + \delta.$$

As remarked earlier, all our algorithms will be differentially private against adaptive adversaries, and the other adversary models are considered only from the the point of view of utility. This is consistent with a long line of work on private learning algorithms, where privacy is proven for worst-case inputs while utility bounds often make distributional or other assumptions on the data.

3. Private algorithms for prediction from experts

We begin our algorithmic contribution by developing new algorithms for oblivious (Section 3.1) and stochastic adversaries (Section 3.2). The main idea is to save the privacy budget by limiting the number of model updates. Though, the way in which this is done varies significantly depending on the adversary: for stochastic adversaries, we allow for $\log(T)$ updates at fixed time-steps $t = 2^i$, while for oblivious adversaries we employ a more adaptive time-step strategy for updates based on the underlying data.

3.1. Oblivious adversaries using shrinking dartboard

In this section, we present our main algorithms for DP-OPE with oblivious adversaries. We build on the shrinking dartboard algorithm (Geulen et al., 2010) to develop a private algorithm that improves over the best existing results for both pure and approximate DP. For ϵ -DP, our algorithm obtains regret $\epsilon^{-1} \sqrt{T} \log d$ which nearly matches the non-private regret for constant ϵ ; this is the first algorithm for pure DP that achieves sub-linear regret for oblivious adversaries. For approximate DP, our algorithm obtains regret $\sqrt{T \log d} + \epsilon^{-1} (T \log(1/\delta))^{1/3} \log d$: our algorithm achieves negligible privacy cost compared to the non-private regret in the high-dimensional regime. Previous algorithms obtain regret roughly $\epsilon^{-1} \sqrt{T \log d \log(1/\delta)}$ for $d \geq T$ which is $\epsilon^{-1} \sqrt{\log(1/\delta)}$ larger than the non-private regret.

1. Jain et al. (2021) recently formalized DP against adaptive adversaries for a different online learning problem. Their notion is equivalent to ours, but our presentation may be easier to work with.

Key to our improvements in this section (similarly to limited switching OPE; [Geulen et al., 2010](#)) is the observation that the distribution of the multiplicative weights (MW) algorithm does not change significantly between consecutive time-steps. Thus, by using a correlated sampling procedure, we can preserve the same marginal distribution as MW while updating the model with small probability. As this limits the number of model updates, this will allow us to assign higher privacy budget for each model update, hence improving its utility.

Another obstacle that arises from this approach is that the switching probability (to preserve the same marginal distribution) depends on the underlying data $p_{\text{switch}} = 1 - (1 - \eta)^{\ell_{t-1}(x_{t-1})}$. This is clearly not private as the probability of switching is zero whenever $\ell_{t-1}(x_{t-1}) = 0$. To tackle this issue, we guarantee privacy by adding another switching step that forces the algorithm to switch with probability p regardless of the underlying data.

Algorithm 1 Private Shrinking Dartboard

Require: Step size $\eta > 0$, switching probability $p \in [0, 1]$, switching budget K

- 1: Set $w_i^1 = 1$, $p_i^1 = \frac{1}{d}$ for all $i \in [d]$
 - 2: Choose expert x_1 from the distribution $P^1 = (p_1^1, \dots, p_d^1)$
 - 3: Set $k = 1$
 - 4: **for** $t = 2$ to T **do**
 - 5: Set $w_i^t = w_i^{t-1}(1 - \eta)^{\ell_{t-1}(i)}$ for all $i \in [d]$
 - 6: Set $p_i^t = \frac{w_i^t}{\sum_{i'=1}^d w_{i'}^t}$ for all $i \in [d]$
 - 7: Sample $z_t \sim \text{Ber}(1 - p)$
 - 8: **if** $z_t = 1$ **then**
 - 9: Sample $z_t \sim \text{Ber}(w_{x_{t-1}}^t / w_{x_{t-1}}^{t-1})$
 - 10: **if** $z_t = 1$ **then**
 - 11: $x_t = x_{t-1}$
 - 12: **else if** $k < K$ **then**
 - 13: $k = k + 1$
 - 14: Sample x_t from P^t
 - 15: Receive $\ell_t : [d] \rightarrow [0, 1]$
 - 16: Pay cost $\ell_t(x_t)$
-

The following theorem states the privacy guarantees and the upper bounds on the regret of Algorithm 1. We defer the proof to Appendix A.1.

Theorem 2. Let $\ell_1, \dots, \ell_T \in [0, 1]^d$ be chosen by an oblivious adversary. Algorithm 1 with $p < 1/2$, $\eta < 1/2$, and $K = 4Tp$ has regret

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq \eta T + \frac{\ln d}{\eta} + 2Te^{-Tp/3}.$$

Moreover, Algorithm 1 is (ε, δ) -DP where

$$\varepsilon = \begin{cases} \frac{5\eta}{p} + 100Tp\eta^2 + 20\eta\sqrt{Tp \log(1/\delta)} & \text{if } \delta > 0; \\ \frac{\eta}{p} + 16Tp\eta & \text{if } \delta = 0. \end{cases}$$

Before proving Theorem 2, we illustrate the implication of this result when optimizing the parameters in the algorithm. We have the following regret for approximate DP (see Appendix A.2 for the proof).

Corollary 3. *Let $\ell_1, \dots, \ell_T \in [0, 1]^d$ be chosen by an oblivious adversary. Let $\varepsilon \leq 1$ and set $\varepsilon_0 = \min(\varepsilon/2, \log^{1/3}(1/\delta)T^{-1/6}\sqrt{\ln d})$ and $0 < \delta \leq 1$ such that $T \geq \Omega(\log(1/\delta))$. Setting $p = 1/(T \log(1/\delta))^{1/3}$ and $\eta = p\varepsilon_0/20$, Algorithm 1 is (ε, δ) -DP and has regret*

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq O \left(\sqrt{T \ln d} + \frac{T^{1/3} \log^{1/3}(1/\delta) \ln d}{\varepsilon} \right).$$

Moreover, we have the following regret for pure DP (proof in Appendix A.3).

Corollary 4. *Let $\ell_1, \dots, \ell_T \in [0, 1]^d$ be chosen by an oblivious adversary. Let $\varepsilon \leq 1$. Setting $p = 1/\sqrt{T}$ and $\eta = p\varepsilon/20$, Algorithm 1 is ε -DP and has regret*

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq O \left(\frac{\sqrt{T \ln d}}{\varepsilon} \right).$$

3.1.1. PRIVATE BATCHED SHRINKING DARTBOARD

For smaller values of ε , we present a batch version of Algorithm 1 that groups losses in a batch of size B and applies the same update. For a batch size B , we define the grouped loss

$$\tilde{\ell}_t = \frac{1}{B} \sum_{i=Bt}^{B(t+1)} \ell_i.$$

The batch version of Algorithm 1 then runs Algorithm 1 on the grouped loss $\tilde{\ell}_t$ for $\tilde{T} = \lceil T/B \rceil$ iterations. The following theorem state the regret of this algorithm. We prove the theorem in Appendix A.5

Theorem 5. *Let $\ell_1, \dots, \ell_T \in [0, 1]^d$ be chosen by an oblivious adversary. Algorithm 1 with batch size $1 \leq B \leq T$, $p < 1/2$, $\eta < 1/2$, and $K = 4Tp/B$ has regret*

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq \eta T + \frac{B \ln d}{\eta} + 2Te^{-Tp/3B}.$$

Moreover, for $\delta > 0$, Algorithm 1 is (ε, δ) -DP where

$$\varepsilon = \frac{5\eta}{Bp} + 100Tp\eta^2/B^3 + \frac{20\eta}{B} \sqrt{12Tp/B \log(1/\delta)}.$$

Optimizing the batch size B , we obtain the following regret. Fig. 1 illustrates that this algorithm offers improvements over the original version (without batches) in the high-privacy regime. We defer the proof to Appendix A.4.

Corollary 6. *Let $\frac{\log^{2/3}(1/\delta) \log(d)}{T} \leq \varepsilon \leq \frac{\log^{2/3}(1/\delta) \log(d)}{T^{1/3}}$ and $\delta \leq 1$. Setting $B = \frac{\log^{2/5}(1/\delta) \log^{3/5}(d)}{T^{1/5} \varepsilon^{3/5}}$, $p = (\frac{B}{T \log(1/\delta)})^{1/3}$ and $\eta = Bp\varepsilon/40$, Algorithm 1 is (ε, δ) -DP and has regret*

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq O \left(\frac{T^{2/5} \log^{1/5}(1/\delta) \log^{4/5}(d)}{\varepsilon^{4/5}} + 2Te^{-Tp/3B} \right).$$

3.2. Stochastic adversaries

In this section, we consider stochastic adversaries and present a reduction from private online learning problems—including OPE and online convex optimization (OCO) in general—to the (offline) differentially private stochastic convex optimization (DP-SCO) (Bassily et al., 2019). This reduction demonstrates that the (offline) DP-OCO problem is not harder than private online learning for stochastic adversaries: any algorithm for DP-SCO can be transformed into an online algorithm with stochastic adversaries with nearly the same rate (up to logarithmic factors). Using existing algorithms for DP-SCO (Asi et al., 2021b), this reduction then results in regret $\sqrt{T \log d} + \varepsilon^{-1} \log d \log T$ for private prediction from experts (Corollary 8) and regret $\sqrt{T} + \varepsilon^{-1} \sqrt{d} \log T$ for general DP online convex optimization in ℓ_2 -geometry (Corollary 13).

Our reduction builds on algorithms for DP-SCO. In this problem, the loss functions are sampled i.i.d. from some distribution $\ell_i \stackrel{\text{iid}}{\sim} P$ where $\ell_i : \mathcal{X} \rightarrow \mathbb{R}$ and the goal is to minimize the population loss $L(x) = \mathbb{E}_{\ell \sim P}[\ell(x)]$ given n samples $\ell_1, \dots, \ell_n \stackrel{\text{iid}}{\sim} P$. The performance of an algorithm \mathcal{A} given n samples is measured by its excess population loss, that is, $\Delta_n(\mathcal{A}) = \mathbb{E}[L(\mathcal{A}(\ell_1, \dots, \ell_n)) - \inf_{x \in \mathcal{X}} L(x)]$.

Given an algorithm \mathcal{A} for DP-SCO, we design an online algorithm that updates the model only a logarithmic number of times—during the time-steps $t = 1, 2, 4, 8, \dots, T$. For each such time-step t , we run \mathcal{A} on the past $t/2$ samples to release the next model. As the loss functions are from the same distribution, the previous model generated by \mathcal{A} should perform well for future loss functions. We present the full details in Algorithm 2.

Algorithm 2 Limited Updates for Online Optimization with Stochastic Adversaries

Require: Parameter space \mathcal{X} , DP-SCO algorithm \mathcal{A}

- 1: Set $x_0 \in \mathcal{X}$
 - 2: **for** $t = 1$ to T **do**
 - 3: **if** $t = 2^\ell$ for some integer $\ell \geq 1$ **then**
 - 4: Run an optimal (ε, δ) -DP-SCO algorithm \mathcal{A} over \mathcal{X} with samples $\ell_{t/2}, \dots, \ell_{t-1}$.
 - 5: Let x_t denote the output of the private algorithm
 - 6: **else**
 - 7: Let $x_t = x_{t-1}$
 - 8: Receive $\ell_t : \mathcal{X} \rightarrow \mathbb{R}$.
 - 9: Pay cost $\ell_t(x_t)$
-

We have the following regret for Algorithm 2. We defer the proof to Appendix B.1.

Theorem 7. *Let $\ell_1, \dots, \ell_T : \mathcal{X} \rightarrow \mathbb{R}$ be convex functions chosen by a stochastic adversary, $\ell_i \stackrel{\text{iid}}{\sim} P$. Let \mathcal{A} be a (ε, δ) -DP algorithm for DP-SCO. Then Algorithm 2 is (ε, δ) -DP and has regret*

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^T \ell_t(x) \right] \leq \sum_{i=1}^{\log T} 2^i \Delta_{2^i}(\mathcal{A}).$$

Now we present the implications of this reduction for DP-OPE with stochastic adversaries. We use the algorithm for DP-SCO in ℓ_1 -geometry from Asi et al. (2021b). This algorithm, denoted \mathcal{A}_{ℓ_1} , has excess population loss $\Delta_n(\mathcal{A}_{\ell_1}) = O(\sqrt{\log(d)/n} + \log(d)/n\varepsilon)$. We have the following result which we prove in Appendix B.2.

Corollary 8. Let $\ell_1, \dots, \ell_T : [d] \rightarrow [0, 1]$ be chosen by a stochastic adversary, $\ell_i \stackrel{\text{iid}}{\sim} P$. Then Algorithm 2 using \mathcal{A}_{ℓ_i} is ε -DP and has regret

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq O \left(\sqrt{T \log d} + \frac{\log d \log T}{\varepsilon} \right).$$

4. Lower bounds

In this section, we prove new lower bounds for adaptive adversaries which show a separation from the non-adaptive case. In particular, we show that for $\varepsilon = 1/\sqrt{T}$, an adaptive adversary can force any (ε, δ) -DP algorithm to incur a linear regret $\Omega(T)$. Similarly, any ε -DP algorithm with $\varepsilon \leq 1/10$ must incur linear regret against adaptive adversaries. On the other hand, the results of Section 3 show that a sub-linear regret is possible for both privacy regimes with oblivious adversaries.

Our lower bound is based on finger-printing lower bound constructions (Bun et al., 2018) which are the basic technique for proving lower bounds in the offline setting. The idea is to design a reduction that uses a DP-OPE algorithm for estimating the signs of the mean of high-dimensional inputs; a problem which is known to be hard under differential privacy. The following theorem summarizes our main lower bound for (ε, δ) -DP. We defer the proof to Appendix C.1.

Theorem 9. Let T be sufficiently large and $d \geq 2T$. Let $\varepsilon \leq 1$ and $\delta \leq 1/T^3$. If \mathcal{A} is (ε, δ) -DP then there is an adaptive adversary such that

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \geq \Omega \left(\min \left(T, \frac{1}{(\varepsilon \log T)^2} \right) \right).$$

Proof. (sketch) We illustrate the main ideas of the lower bound for $\varepsilon \leq 1/(\sqrt{T} \log T)$, where we have to prove that the regret has to be linear in this case. We will reduce the problem of private sign estimation to DP-OPE with adaptive adversaries and use existing lower bounds for private sign estimation (see Theorem 15). In this problem, we are given an input matrix $X \in \{-1, +1\}^{n \times p}$ (each row X_i is a user) and the goal is to estimate the sign of the columns for each consensus columns (e.g. constant column).

To this end, given an algorithm \mathcal{A} for DP-OPE and an input $X \in \{-1, +1\}^{n \times p}$, we have the following procedure for estimating the signs of the columns of X . We design an online experts problem that has $d = 2p$ experts where column $j \in [p]$ in X will have two corresponding experts $2j$ and $2j + 1$ (corresponding to the sign of column j). We design the loss functions in a way that will allow to estimate the signs of the columns based on the experts chosen by the algorithm. In particular, at round $t \in [T]$, we pick a random user $i \in [n]$ and define a loss function based on the user X_i . We will make $\ell_t(2j + 1) = 0$ and $\ell_t(2j + 2) = 1$ if the sign of the j 'th column is $X_{ij} = -1$, and will set $\ell_t(2j + 1) = 1$ and $\ell_t(2j + 2) = 0$ if the sign is $X_{ij} = +1$. This implies that the algorithm will either estimate the correct sign of some column $j \in [p]$ or suffer +1 in regret. However, the algorithm can keep estimating the sign of the same column and still suffer low regret. In order to prevent this, once the algorithm has estimated the sign of a column j , we increase the loss of all experts corresponding to that column to 1, hence preventing the algorithm from iteratively estimating the sign of the same column. Overall, this implies that a low-regret algorithm will be able to estimate the signs of most of the columns, which is a contradiction to lower bounds for private sign estimation. We formalize these ideas and complete the proof in Appendix C.1.

■

We also have the following lower bound for pure differential privacy. It shows that pure DP algorithms cannot learn against adaptive adversaries, that is, they must suffer linear regret for constant ε .

Theorem 10. *Let $\varepsilon \leq 1/10$ and $d \geq 2T$. If \mathcal{A} is ε -DP then there is an adaptive adversary such that*

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \geq \Omega(T).$$

We can also extend the previous lower bounds to larger values of ε (see proof in Appendix C.3).

Theorem 11. *Let $1 \leq k \leq O(p^{1-\rho})$ for $0 < \rho < 1$, $d = 2^k \binom{p}{k} = 2^{\Theta(k \log p)}$, $T = p/k$, and $\varepsilon \leq \frac{\sqrt{k/T}}{200\rho \log(T)}$ where p is sufficiently large. If \mathcal{A} is (ε, δ) -DP with $\delta \leq 1/T^3$, then there is an adaptive adversary such that*

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \geq \Omega \left(\frac{\sqrt{T \log d}}{\varepsilon \log^{3/2} T} \right).$$

Finally, we note that for stochastic adversaries, existing lower bounds for DP-SCO immediately imply lower bounds for the online setting using online-to-batch transformations. As there is a lower bound of $\log(d)/T\varepsilon$ for private selection (Steinke and Ullman, 2017b), this implies a lower bound on the (normalized) excess loss for DP-SCO with linear functions in ℓ_1 -geometry (as one can reduce private selection to this problem; Asi et al., 2021b). This implies a lower bound of $\log(d)/\varepsilon$ for DP-OPE.

5. Implications for DP-OCO in ℓ_2 -geometry

In this section, we derive several implications of our techniques for *differentially private online convex optimization* (DP-OCO) in ℓ_2 -geometry. In this setting, the algorithm chooses $x_t \in \mathcal{X}$ where $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq D\}$ and the adversary responds with loss functions $\ell_t : \mathcal{X} \rightarrow \mathbb{R}^d$ that are convex and L -Lipschitz. Building on our techniques for DP-OPE, we propose new algorithms that improve over the best existing regret bounds for DP-OCO (Kairouz et al., 2021) which achieve $d^{1/4} \sqrt{T}/\varepsilon$ for stochastic and adaptive adversaries. Our algorithms obtain (up to logarithmic factors) near-optimal regret $\sqrt{T} + \sqrt{d}/\varepsilon$ for stochastic adversaries, and $\sqrt{Td} + T^{1/3}d/\varepsilon$ for adaptive adversaries.

5.1. Oblivious adversaries

Using our private shrinking dartboard algorithm, in this section we develop algorithms that improve the regret for oblivious adversaries. Our algorithms construct a covering of the parameter space \mathcal{X} then apply our private shrinking dartboard algorithm where the experts are the elements of the cover. By optimizing the size of the cover to balance the error from the approximation error and the error due to the number of experts, we obtain the following regret for DP-OCO in ℓ_2 -geometry. We defer the proof to Appendix D.1.

Theorem 12. Let $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq D\}$ and $\ell_1, \dots, \ell_T : \mathcal{X} \rightarrow \mathbb{R}$ be convex and L -Lipschitz functions chosen by an oblivious adversary. There is an (ε, δ) -DP that has regret

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^T \ell_t(x) \right] \leq LD \cdot O \left(\sqrt{Td \log T} + \frac{T^{1/3} d \log^{1/3}(1/\delta) \log T}{\varepsilon} \right).$$

In the high-privacy regime, this result can improve over the previous work (Kairouz et al., 2021) which has regret $\sqrt{T}d^{1/4}/\sqrt{\varepsilon}$. For example, if $d = 1$ and $\varepsilon = T^{-1/4}$, then our regret is roughly $T^{7/12}$ while their regret is $T^{5/8}$.

5.2. Stochastic adversaries

For stochastic adversaries, we use the reduction in Section 3.2 (Algorithm 2) with optimal algorithms for DP-SCO in ℓ_2 -geometry to obtain optimal regret bounds. More precisely, we use an optimal (ε, δ) -DP-SCO algorithm from Feldman et al. (2020), which we call \mathcal{A}_{ℓ_2} . As this algorithm has excess loss $\Delta_n = LD \cdot O(1/\sqrt{n} + \sqrt{d}/n\varepsilon)$, the following result follows immediately from Theorem 7. We defer the proof to Appendix D.2. We also note that we can also obtain regret bounds for pure ε -DP using existing (pure) DP algorithms for DP-SCO (Asi et al., 2021c) with our reduction.

Corollary 13 (DP-OCO in ℓ_2 -geometry). Let $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq D\}$ and $\ell_1, \dots, \ell_T : \mathcal{X} \rightarrow \mathbb{R}$ be convex and L -Lipschitz functions chosen by a stochastic adversary, $\ell_i \stackrel{\text{iid}}{\sim} P$. Then Algorithm 2 using \mathcal{A}_{ℓ_2} is (ε, δ) -DP and has regret

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq LD \cdot O \left(\sqrt{T} + \frac{\sqrt{d} \log T}{\varepsilon} \right)$$

This regret is near-optimal up to logarithmic factors since we have the lower bound $\sqrt{T} + \sqrt{d}/\varepsilon$ for the offline version of this problem (DP-SCO in ℓ_2 -geometry) where all of the samples are given in advance (Bassily et al., 2014, 2019).

References

- Naman Agarwal and Karan Singh. The price of differential privacy for online learning. In *Proceedings of the 34th International Conference on Machine Learning*, pages 32–40, 2017.
- Jason Altschuler and Kunal Talwar. Online learning over a finite action set with limited switching. In *Proceedings of the Thirty First Annual Conference on Computational Learning Theory*, 2018.
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- Hilal Asi, John Duchi, Alireza Fallah, Omid Javidi, and Kunal Talwar. Private adaptive gradient methods for convex optimization. In *Proceedings of the 38th International Conference on Machine Learning*, pages 383–392, 2021a.
- Hilal Asi, Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: Optimal rates in ℓ_1 geometry. In *Proceedings of the 38th International Conference on Machine Learning*, 2021b.

- Hilal Asi, Daniel Levy, and John Duchi. Adapting to function difficulty and growth conditions in private optimization. In *Advances in Neural Information Processing Systems 34*, volume 34, pages 19069–19081, 2021c.
- Hilal Asi, Vitaly Feldman, Tomer Koren, and Kunal Talwar. Near-optimal algorithms for private online optimization in the realizable regime. *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *55th Annual Symposium on Foundations of Computer Science*, pages 464–473, 2014.
- Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Thakurta. Private stochastic convex optimization with optimal rates. In *Advances in Neural Information Processing Systems*, volume 32, pages 11282–11291, 2019.
- Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. *SIAM Journal on Computing*, 47(5):1888–1938, 2018. doi: 10.1137/15M1033587. URL <https://doi.org/10.1137/15M1033587>.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Lin Chen, Qian Yu, Hannah Lawrence, and Amin Karbasi. Minimax regret of switching-constrained online convex optimization: No phase transition. *Advances in Neural Information Processing Systems*, 33:3477–3486, 2020.
- John C. Duchi. Information theory and statistics. Lecture Notes for Statistics 311/EE 377, Stanford University, 2019. URL <http://web.stanford.edu/class/stats311/lecture-notes.pdf>. Accessed May 2019.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3 & 4):211–407, 2014.
- Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology (EUROCRYPT 2006)*, 2006a.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*, pages 265–284, 2006b.
- Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the Forty-Second Annual ACM Symposium on the Theory of Computing*, pages 715–724, 2010.
- Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: optimal rates in linear time. In *Proceedings of the 52nd Annual ACM on the Theory of Computing*, pages 439–449, 2020.

- Sascha Geulen, Berthold Vöcking, and Melanie Winkler. Regret minimization for online buffering problems using the weighted majority algorithm. In *Proceedings of the Twenty Third Annual Conference on Computational Learning Theory*, 2010.
- Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam Smith. The price of differential privacy under continual observation. *arXiv:2112.00828 [cs.DS]*, 2021.
- Prateek Jain and Abhradeep Thakurta. (Near) dimension independent risk bounds for differentially private learning. In *Proceedings of the 31st International Conference on Machine Learning*, pages 476–484, 2014.
- Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Proceedings of the Twenty Fifth Annual Conference on Computational Learning Theory*, 2012.
- Peter Kairouz, Brendan McMahan, Shuang Song, Om Thakkar, Abhradeep Thakurta, and Zheng Xu. Practical and private (deep) learning without sampling or shuffling. *arXiv:2103.00039 [cs.CR]*, 2021.
- A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.
- Uri Sherman and Tomer Koren. Lazy oco: Online convex optimization on a switching budget. In *Proceedings of the Thirty Fourth Annual Conference on Computational Learning Theory*, 2021.
- Adam Smith and Abhradeep Thakurta. (Nearly) optimal algorithms for private online learning in full-information and bandit settings. In *Advances in Neural Information Processing Systems 26*, 2013.
- Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2):3–22, 2017a.
- Thomas Steinke and Jonathan Ullman. Tight lower bounds for differentially private selection. In *58th Annual Symposium on Foundations of Computer Science*, pages 552–563. IEEE, 2017b.
- Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Nearly optimal private Lasso. In *Advances in Neural Information Processing Systems*, volume 28, pages 3025–3033, 2015.

Appendix A. Proofs for Section 3.1

A.1. Proof of Theorem 2

We build on the following lemma.

Lemma 14. *Let \hat{P}_t be the marginal distribution of x_t of Algorithm 1. Then*

$$\|\hat{P}_t - P^t\|_{TV} \leq e^{-Tp/3}.$$

Proof. Let k_t be the value of k at iteration t . We roughly show that if $k_t < K$ then $P_t = \hat{P}_t$. As $P(k_t > K)$ is very small, this will prove the claim. Recall that $k_t = \sum_{i \leq t} 1\{z_i = 0\}$. Note that $P(z_t = 0) \leq p + (1-p)\eta \leq 2p$. Therefore, letting $y_t \sim \text{Ber}(p + (1-p)\eta)$ we have

$$\begin{aligned} P(k_t > K) &\leq P(k_T > K) \\ &= P\left(\sum_{i=1}^T 1\{z_i = 0\} > K\right) \\ &\leq P\left(\sum_{i=1}^T 1\{y_i = 0\} > K\right) \\ &\leq e^{-Tp/3}, \end{aligned}$$

where the last inequality follows from a Chernoff bound (Lemma 17).

Now we proceed to show that \hat{P}_t and P_t are close. To this end, we first define Q_t to be the marginal distribution of x_t in Algorithm 1 when $K = T + 1$ (that is, no limit on switching). We prove by induction that $Q_t = P^t$. The base case for $t = 1$ is trivial. Assuming correctness for t , we have that for $x \in [d]$

$$\begin{aligned} Q_t(x) &= pP_x^t + (1-p)\frac{w_x^t}{w_x^{t-1}}Q_{t-1}(x) + (1-p)p_x^t \sum_{x'=1}^d Q_{t-1}(x')\left(1 - \frac{w_{x'}^t}{w_{x'}^{t-1}}\right) \\ &= pP_x^t + (1-p)\frac{w_x^t}{w_x^{t-1}}\frac{w_x^{t-1}}{W^{t-1}} + (1-p)\frac{w_x^t}{W^t} \sum_{x'=1}^d \frac{w_{x'}^{t-1}}{W^{t-1}} \frac{w_{x'}^{t-1} - w_{x'}^t}{w_{x'}^{t-1}} \\ &= pP_x^t + (1-p)\left(\frac{w_x^t}{W^{t-1}} + \frac{w_x^t}{W^t} \frac{W^{t-1} - W^t}{W^{t-1}}\right) \\ &= P_x^t. \end{aligned}$$

Now consider \hat{P} . Let Q_t^0 and Q_t^1 be the conditional distribution of Q_t given $k_t < K$ or $k_t \geq K$, respectively. Moreover, let \hat{P}_t^0 and \hat{P}_t^1 be the conditional distribution of \hat{P}_t given $k_t < K$ or $k_t \geq K$, respectively. Note that $Q_t(x) = P(k_t < K)Q_t^0 + P(k_t \geq K)Q_t^1$ and that $\hat{P}_t(x) = P(k_t < K)\hat{P}_t^0 + P(k_t \geq K)\hat{P}_t^1$. Noting that $P_t^0 = Q_t^0$, we have

$$\begin{aligned} \|\hat{P}_t - P^t\|_{TV} &= \|\hat{P}_t - Q_t\|_{TV} \\ &= \|P(k_t < K)(\hat{P}_t^0 - Q_t^0) + P(k_t \geq K)(\hat{P}_t^1 - Q_t^1)\|_{TV} \\ &\leq P(k_t < K)\|\hat{P}_t^0 - Q_t^0\| + P(k_t \geq K)\|\hat{P}_t^1 - Q_t^1\|_{TV} \\ &\leq P(k_t \geq K). \end{aligned}$$

■

Proof. First, we begin by analyzing the regret. Lemma 14 shows that \hat{P}_t the marginal distribution of x_t is the same as that of the (non-private) shrinking dartboard algorithm P_t , therefore Theorem 3 of Geulen et al. (2010) shows that for $\eta \leq 1/2$

$$\begin{aligned}
 \mathbb{E}_{x_t \sim \hat{P}_t} \left[\sum_{t=1}^T \ell_t(x_t) \right] &= \mathbb{E}_{x_t \sim P^t} \left[\sum_{t=1}^T \ell_t(x_t) \right] + \mathbb{E}_{x_t \sim \hat{P}_t} \left[\sum_{t=1}^T \ell_t(x_t) \right] - \mathbb{E}_{x_t \sim P^t} \left[\sum_{t=1}^T \ell_t(x_t) \right] \\
 &\leq \mathbb{E}_{x_t \sim P^t} \left[\sum_{t=1}^T \ell_t(x_t) \right] + 2T \|\hat{P}_t - P^t\|_{TV} \\
 &\leq (1 + \eta) \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) + \frac{\ln d}{\eta} + 2T e^{-Tp/3} \\
 &\leq \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) + \eta T + \frac{\ln d}{\eta} + 2T e^{-Tp/3}.
 \end{aligned}$$

Let us now analyze privacy. Assume we have two neighboring sequences that differ at time-step t_1 . Let Z_t and X_t denote the random variables for z_t and x_t in the algorithm when run for the first sequence and let $Y_t = 1 - Z_t$. Similarly, let Z'_t , Y'_t , and X'_t denote the same for the neighboring sequence. We consider the pairs $W_t = (X_t, Z_{t+1})$ (where $X_0 = 0$) and prove that W_t given $\{W_\ell\}_{\ell=0}^{t-1}$ and W'_t given $\{W'_\ell\}_{\ell=0}^{t-1}$ are ε_t -indistinguishable where

$$\varepsilon_t = \begin{cases} 0 & \text{if } t < t_1 \\ \eta/p & \text{if } t = t_1 \\ 1\{\sum_{\ell=1}^{t-1} Y_\ell < K\} 4Y_t \eta & \text{if } t > t_1 \end{cases}$$

The result then follows from advanced composition (Dwork and Roth, 2014): note that $Y_t \in \{0, 1\}$ therefore we have that the final privacy parameter is

$$\begin{aligned}
 \varepsilon_f &\leq \frac{3}{2} \sum_{t=1}^T \varepsilon_t^2 + \sqrt{6 \sum_{t=1}^T \varepsilon_t^2 \log(1/\delta)} \\
 &\leq \frac{3}{2} \left(\frac{\eta^2}{p^2} + 16K\eta^2 \right) + \sqrt{6 \left(\frac{\eta^2}{p^2} + 16K\eta^2 \right) \log(1/\delta)} \\
 &\leq \frac{5\eta}{p} + 24K\eta^2 + \eta \sqrt{100K \log(1/\delta)} \\
 &\leq \frac{5\eta}{p} + 100Tp\eta^2 + 20\eta \sqrt{Tp \log(1/\delta)}.
 \end{aligned}$$

Similarly, the result for $\delta = 0$ follows from basic composition. To finish the proof, consider the pair W_t and W'_t . First, note that if $t < t_1$ then clearly W_t and W'_t are 0-indistinguishable as they do not depend on ℓ_{t_1} or ℓ'_{t_1} . For $t = t_1$, note that X_{t_1} and X'_{t_1} has the same distribution. Moreover, the

definition of Z_t implies that Z_t and Z'_t are η/p -indistinguishable since

$$\begin{aligned} \frac{P(Z_t = 1)}{P(Z'_t = 1)} &\leq \frac{(1-p)}{(1-p)(1-\eta)} \\ &= \frac{1}{1-\eta} \\ &= 1 + \frac{\eta}{1-\eta} \\ &\leq 1 + 2\eta \\ &\leq e^{2\eta}. \end{aligned}$$

Moreover, since $\eta \leq p\varepsilon$ we have

$$\begin{aligned} \frac{P(Z_t = 0)}{P(Z'_t = 0)} &\leq \frac{p + (1-p)\eta}{p} \\ &\leq 1 + \frac{\eta}{p} \leq e^{\eta/p}. \end{aligned}$$

Now consider $t > t_1$. If $\sum_{\ell=1}^{t-1} Y_\ell \geq K$ or $Y_t = Y'_t = 0$ then $X_t = X_{t-1}$ and $X'_t = X'_{t-1}$ and thus X_t and X'_t are 0-indistinguishable. If $Y_t = Y'_t = 0$ then X_t and X'_t are 4η -indistinguishable since $w_x^t/w_x'^t \leq 1/(1-\eta) \leq e^{2\eta}$ which implies that $P(x_t = x)/P(x'_t = x) \leq e^{4\eta}$. Overall, X_t and X'_t are $4Y_t\eta$ -indistinguishable. Moreover, since $t > t_1$, we have that Z_{t+1} is a function of X_t and ℓ_t and Z'_{t+1} is a function of X'_t and $\ell'_t = \ell_t$, hence by post-processing we get that Z_{t+1} and Z'_{t+1} are $4Y_t\eta$ -indistinguishable. Overall, we have that W_t and W'_t are $1\{\sum_{\ell=1}^{t-1} Y_\ell < K\} 4Y_t\eta$ -indistinguishable. \blacksquare

A.2. Proof of Corollary 3

For these parameters, Algorithm 1 has privacy

$$\varepsilon_0/4 + Tp^3\varepsilon_0^2/4 + \varepsilon_0\sqrt{Tp^3\log(1/\delta)} \leq 2\varepsilon_0.$$

As $\varepsilon_0 \leq \varepsilon/2$, this proves the claim about privacy. Moreover, its regret is

$$\begin{aligned} \eta T + \frac{\ln d}{\eta} + 2Te^{-Tp/3} &\leq Tp\varepsilon_0/20 + 20\ln d/(p\varepsilon_0) + 2Te^{-Tp/3} \\ &\leq \frac{T^{2/3}\varepsilon_0}{\log^{1/3}(1/\delta)} + \frac{20T^{1/3}\log^{1/3}(1/\delta)\ln d}{\varepsilon_0} + 2Te^{-Tp/3} \\ &\leq \sqrt{T\ln d} + \frac{20T^{1/3}\log^{1/3}(1/\delta)\ln d}{\varepsilon_0} + 2Te^{-Tp/3} \\ &\leq O\left(\sqrt{T\ln d} + \frac{T^{1/3}\log^{1/3}(1/\delta)\ln d}{\varepsilon}\right), \end{aligned}$$

where the last inequality follows as $\varepsilon_0 = \min(\varepsilon/2, \frac{\log^{1/3}(1/\delta)\sqrt{\ln d}}{T^{1/6}})$.

A.3. Proof of Corollary 4

For these parameters, Algorithm 1 has privacy

$$\varepsilon/20 + 16Tp\eta \leq \varepsilon/10 + 16Tp^2\varepsilon/20 \leq \varepsilon.$$

Moreover, its regret is

$$\begin{aligned} \eta T + \frac{\ln d}{\eta} + 2Te^{-Tp/3} &\leq Tp\varepsilon/20 + 20 \ln d/(p\varepsilon) + 2Te^{-Tp/3} \\ &\leq \sqrt{T} + \frac{20\sqrt{T} \ln d}{\varepsilon} + 2Te^{-Tp/3}, \end{aligned}$$

where the last inequality follows since $\varepsilon \leq 1$.

A.4. Proof of Corollary 6

For these parameters, Algorithm 1 has privacy

$$\frac{5\eta}{Bp} + \frac{100Tp\eta^2}{B^3} + \frac{20\eta}{B^{3/2}} \sqrt{Tp \log(1/\delta)} \leq \varepsilon/8 + \frac{Tp^3\varepsilon^2}{16B} + \frac{\varepsilon}{2} \sqrt{Tp^3 \log(1/\delta)/B} \leq \varepsilon.$$

Moreover, its regret is

$$\begin{aligned} \eta T + \frac{B \ln d}{\eta} + 2Te^{-Tp/3B} &\leq TBp\varepsilon/40 + 40 \ln d/(p\varepsilon) + 2Te^{-Tp/3B} \\ &\leq \frac{T^{2/3} B^{4/3} \varepsilon}{\log^{1/3}(1/\delta)} + \frac{40T^{1/3} \log^{1/3}(1/\delta) \ln d}{B^{1/3} \varepsilon} + 2Te^{-Tp/3B} \\ &\leq O\left(\frac{T^{2/5} \log^{1/5}(1/\delta) \log^{4/5}(d)}{\varepsilon^{4/5}} + 2Te^{-Tp/3B}\right), \end{aligned}$$

where the last inequality follows by choosing $B = \frac{\log^{2/5}(1/\delta) \log^{3/5}(d)}{T^{1/5} \varepsilon^{3/5}}$ (note that $B \geq 1$ for $\varepsilon \leq \frac{\log^{2/3}(1/\delta) \log(d)}{T^{1/3}}$) and noticing that $Tp/B \geq \Omega(T^{2/5})$ for these parameters as we have a lower bound on ε .

A.5. Proof of Theorem 5

The same analysis as in Theorem 2 yields regret

$$\mathbb{E} \left[\sum_{t=1}^{\tilde{T}} \tilde{\ell}_t(\tilde{x}_t) - \min_{x \in [d]} \sum_{t=1}^{\tilde{T}} \tilde{\ell}_t(x) \right] \leq \eta \tilde{T} + \frac{\ln d}{\eta} + 2\tilde{T}e^{-\tilde{T}p/3}.$$

Setting $x_t = \tilde{x}_{\lfloor t/B \rfloor}$ and multiplying both sides by B , we have regret

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] \leq \eta T + \frac{B \ln d}{\eta} + 2Te^{-\tilde{T}p/3}.$$

Let us now analyze privacy. The privacy follows the same steps as in the proof of Theorem 2 with two main differences. First, let $t = t_1$ be the time such that $\tilde{\ell}_{t_1}$ contains the differing loss function and let $\tilde{\ell}_t = \tilde{\ell}_t(x_{t-1})$ and $\tilde{\ell}'_t = \tilde{\ell}'_t(x_{t-1})$. Note that $|\tilde{\ell}_{t_1} - \tilde{\ell}'_{t_1}| \leq 1/B$ thus we have that Z_t and Z'_t are $\eta/(Bp)$ -indistinguishable since

$$\begin{aligned} \frac{P(Z'_t = 1)}{P(Z_t = 1)} &\leq \frac{(1-p)(1-\eta)^{\tilde{\ell}'_{t-1}}}{(1-p)(1-\eta)^{\tilde{\ell}_{t-1}}} \\ &\leq (1-\eta)^{-|\tilde{\ell}_{t-1} - \tilde{\ell}'_{t-1}|} \\ &\leq e^{2\eta/B}. \end{aligned}$$

Moreover, assuming w.l.o.g. that $\tilde{\ell}'_{t-1} \geq \tilde{\ell}_{t-1}$, we have

$$\begin{aligned} \frac{P(Z'_t = 0)}{P(Z_t = 0)} &\leq \frac{p + (1-p)(1 - (1-\eta)^{\tilde{\ell}'_{t-1}})}{p + (1-p)(1 - (1-\eta)^{\tilde{\ell}_{t-1}})} \\ &\leq 1 + \frac{(1-p)|1 - (1-\eta)^{\tilde{\ell}'_{t-1} - \tilde{\ell}_{t-1}}|}{p + (1-p)(1 - (1-\eta)^{\tilde{\ell}_{t-1}})} \\ &\leq 1 + \frac{|1 - (1-\eta)^{\tilde{\ell}'_{t-1} - \tilde{\ell}_{t-1}}|}{p} \\ &\leq 1 + \frac{|\tilde{\ell}'_{t-1} - \tilde{\ell}_{t-1}|}{p} \leq e^{\eta/(Bp)}. \end{aligned}$$

The second difference in the privacy analysis is that the sensitivity of the score of the exponential mechanism is now $1/B$ hence X_t and X'_t are now $4\eta/B$ -DP. This shows that W_t given $\{W_\ell\}_{\ell=0}^{t-1}$ and W'_t given $\{W'_\ell\}_{\ell=0}^{t-1}$ are ε_t -indistinguishable where

$$\varepsilon_t = \begin{cases} 0 & \text{if } t < t_1 \\ \eta/(Bp) & \text{if } t = t_1 \\ 1\{\sum_{\ell=1}^{t-1} Y_\ell < K\} 4Y_t\eta/B & \text{if } t > t_1 \end{cases}$$

The result then follows from advanced composition (Dwork and Roth, 2014): the final privacy parameter is

$$\begin{aligned} \varepsilon_f &\leq \frac{3}{2} \sum_{t=1}^T \varepsilon_t^2 + \sqrt{6 \sum_{t=1}^T \varepsilon_t^2 \log(1/\delta)} \\ &\leq \frac{3}{2} \left(\frac{\eta}{Bp} + 16K\eta^2/B^2 \right) + \sqrt{6 \left(\frac{\eta^2}{B^2 p^2} + 16K\eta^2/B^2 \right) \log(1/\delta)} \\ &\leq \frac{5\eta}{Bp} + 24K\eta^2/B^2 + \frac{10\eta}{B} \sqrt{K \log(1/\delta)} \\ &\leq \frac{5\eta}{Bp} + 100T p \eta^2 / B^3 + \frac{20\eta}{B} \sqrt{12T p / B \log(1/\delta)}. \end{aligned}$$

Appendix B. Proofs for Section 3.2

B.1. Proof of Theorem 7

The privacy claim is immediate as each sample ℓ_i is used only once in running a single (ε, δ) -DP algorithm. Now we prove the claim about utility. Consider time-step $t = 2^i$ where we invoke a DP-SCO algorithm with $t/2 = 2^{i-1}$ samples. Therefore the guarantees of the algorithm imply that at iteration t we have

$$\mathbb{E}_{\ell_t \sim \mathcal{P}} \left[\ell_t(x_t) - \min_{x \in [d]} \ell_t(x) \right] \leq O(\Delta_{2^i}).$$

Therefore at phase i , that is $2^i \leq t \leq 2^{i+1}$, the total regret is at most

$$\mathbb{E} \left[\sum_{t=2^i}^{2^{i+1}} \ell_t(x_t) - \min_{x \in [d]} \sum_{t=2^i}^{2^{i+1}} \ell_t(x) \right] \leq O(2^i \Delta_{2^i}).$$

Summing over i proves the claim.

B.2. Proof of Corollary 8

The algorithm \mathcal{A}_{ℓ_1} is ε -DP and has excess population loss $\Delta_n = O(\sqrt{\log(d)/n} + \log(d)/n\varepsilon)$ (Asi et al., 2021b, Theorem 6). Thus, Theorem 7 implies that

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] &\leq \sum_{i=1}^{\log T} 2^i \Delta_i \\ &\leq O \left(\sum_{i=1}^{\log T} 2^{i/2} \sqrt{\log(d)} + \log(d)/\varepsilon \right) \\ &\leq O \left(\sqrt{T \log(d)} + \log(d) \log(T)/\varepsilon \right). \end{aligned}$$

Appendix C. Proofs for Section 4

C.1. Proof of Theorem 9

We build on the following property of the padded Tardos code as done in finger-printing lower bounds. Given a matrix $X \in \{-1, +1\}^{(n+1) \times p}$, we say that $j \in [p]$ is a consensus column if the column is equal to the all one vector or its negation. Let $X_{(i)} \in \{-1, +1\}^{n \times p}$ denote the matrix that results from removing the i 'th row in X . Moreover, we let $\bar{X} \in \mathbb{R}^p$ denote the sum of the rows of X , that is, $\bar{X}_j = \sum_{i=1}^{n+1} X_{ij}$. Finally, for $v \in \mathbb{R}^p$ let $\text{sign}(v) \in \{-1, +1\}^p$ denote the signs of the entries of v

Theorem 15 (Talwar et al., 2015, Theorem 3.2). *Let $p = 1000m^2$ and $n = m/\log m$ for sufficiently large m . There exists a matrix $X \in \{-1, +1\}^{(n+1) \times p}$ such that*

- *There are at least $0.999p$ consensus columns in $X_{(i)}$*
- *Any algorithm $\mathcal{A} : \{-1, +1\}^{n \times p} \rightarrow \{-1, +1\}^p$ such that $\|\mathcal{A}(X_{(i)}) - \text{sign}(\bar{X}_{(i)})\|_0 \leq 1/4$ for all $i \in [n+1]$ with probability at least $2/3$ then \mathcal{A} is not $(1, n^{-1.1})$ -DP.*

Building on Theorem 15, we can now prove our main lower bound.

Proof of Theorem 9. First, we prove the lower bound for $\varepsilon \leq 1/(\sqrt{T} \log T)$, that is, we prove the regret has to be linear in this case. We will reduce the problem of private sign estimation to DP-OPE with adaptive adversaries and use the lower bound of Theorem 15. To this end, given an algorithm \mathcal{A} for DP-OPE and an input $X \in \{-1, +1\}^{n \times p}$, we have the following procedure for estimating the signs of the columns of X . We design an online experts problem that has $d = 2p$ experts where column $j \in [p]$ in X will have two corresponding experts $2j$ and $2j + 1$ (corresponding to the sign of column j). We initialize the vector of signs $s_j = 0$ for all $1 \leq j \leq p$. We have $T = 0.9p$ rounds and at round $1 \leq t \leq T$ we sample a user $i_t \sim [n]$ (arbitrarily while enforcing that each $i \in [n]$ appears at most $2T/n$ times) and play a loss function $\ell_t : [d] \rightarrow \{0, 1\}$ such that

$$\ell_t(2j + 1) = \begin{cases} 1 & \text{if } s_j \neq 0 \\ \frac{X_{i_t, j+1}}{2} & \text{otherwise} \end{cases}$$

We also set

$$\ell_t(2j + 2) = \begin{cases} 1 & \text{if } s_j \neq 0 \\ \frac{-X_{i_t, j+1}}{2} & \text{otherwise} \end{cases}$$

The idea of this loss function is that the $2j + 1$ and $2j + 2$ experts will represent the signs of the j 'th column. If the sign of the j 'th column is $+1$, then expert $2j + 2$ will have better loss and hence should be picked by the algorithm. Moreover, whenever the algorithm has estimated the sign of the j 'th column ($s_j \neq 0$), we set the loss to be 1 for both experts $2j + 1$ and $2j + 2$, in order to force the online algorithm to estimate the sign of new columns.

Then, given the output of the algorithm \mathcal{A} at time t , that is $x_t = \mathcal{A}(\ell_1, \dots, \ell_{t-1})$ we set $s_j = -1$ if $x_t = 2j + 1$ and $s_j = 1$ if $x_t = 2j + 2$ and otherwise we keep s_j unchanged. Moreover, there is an expert that achieves optimal loss, that is, for some $x^* \in [d]$ we have

$$\sum_{t=1}^T \ell_t(x^*) = 0.$$

This follows since X has at least $0.999p$ consensus columns hence there is a zero-loss expert after $T = 0.9p$ iterations. Now we show that if an algorithm \mathcal{A} has small regret, then the vector s estimates the sign of at least $0.8p$ columns. To this end, let $j_t = \lfloor x_t/2 \rfloor$ denote the column corresponding to the expert picked by the algorithm at time t , $S = \{j_t : t \in [T]\}$, and $S_{\text{cons}} = \{j \in [p] :$

column j is a consensus column}. Observe that the regret of the algorithm is

$$\begin{aligned}
 \sum_{t=1}^T \ell_t(x_t) &= \sum_{j_t \in S} 1\{s_{j_t} = 1\} \ell_t(2j_t + 2) + 1\{s_{j_t} = -1\} \ell_t(2j_t + 1) \\
 &= \frac{1}{2} \sum_{j_t \in S} 1\{s_{j_t} = 1\} (-X_{i_t, j_t} + 1) + 1\{s_{j_t} = -1\} (X_{i_t, j_t} - 1) \\
 &= \frac{1}{2} \sum_{j_t \in S} 1\{s_{j_t} = 1, X_{i_t, j_t} = -1\} + 1\{s_{j_t} = -1, X_{i_t, j_t} = 1\} \\
 &= \frac{1}{2} \sum_{j_t \in S} 1\{s_{j_t} \neq X_{i_t, j_t}\} \\
 &\geq -0.001p + \frac{1}{2} \sum_{j_t \in S \cap S_{\text{cons}}} 1\{s_{j_t} \neq X_{i_t, j_t}\} \\
 &\geq -0.001p + \frac{1}{2} \sum_{j_t \in S \cap S_{\text{cons}}} 1\{s_{j_t} \neq \text{sign}(\bar{X})_{j_t}\}.
 \end{aligned}$$

Assume towards a contradiction that \mathcal{A} is $(1/200\sqrt{T} \log(T), \delta)$ -DP where $\delta \leq 1/T^3$ and that the expected regret is at most $T/1000$. Markove inequality implies that with probability at least $9/10$ the regret is at most $T/100$. Under this event we have

$$\sum_{j_t \in S \cap S_{\text{cons}}} 1\{s_{j_t} \neq \text{sign}(\bar{X})_{j_t}\} \leq 0.002T.$$

Now note that we can assume that the online algorithm picks x_t such that each j_t appears at most one. Otherwise we can modify the algorithm to satisfy this property while not increasing the regret: whenever the algorithm picks x_t such that j_t appeared before, the loss of this expert is 1, hence we can randomly pick another expert x_t such that j_t has not appeared. This implies that $|S| = T = 0.9p$ and hence $|S \cap S_{\text{cons}}| \geq 0.85p$. Therefore we have that $s_{j_t} = \text{sign}(\bar{X})_{j_t}$ for at least $0.8p$ columns from S_{cons} with probability 0.9. To finish the proof, we need to argue about the final privacy guarantee of the sign vector s ; we will prove that s is $(1, T\delta)$ -DP which will give a contradiction to Theorem 15 and prove the claim. To this end, note that the algorithm \mathcal{A} is $(1/200\sqrt{T} \log(T), \delta)$ -DP. Moreover, recall that each row $i \in [n]$ appears at most $k \leq 2T/n \leq 2p/n \leq 200\sqrt{p} \log(p)$ times, hence group privacy implies the final output s is $(k\varepsilon, k\delta)$ -DP, that is, $(1, 1/T^2)$ -DP.

Now we proceed to prove the lower bound for larger values $\varepsilon \geq 1/(\sqrt{T} \log T)$. Note that if $\varepsilon \geq \log(T)/T^{1/4}$ then the non-private lower bound of $\sqrt{T} \log d$ is sufficient. Otherwise, consider an algorithm \mathcal{A} that is ε -DP and consider an adversary that in the first $T_0 < T$ iterations behaves the same as the above where $\varepsilon = 1/(\sqrt{T_0} \log T_0)$. Then in the last $T - T_0$ iterations it sends $\ell_t(x) = 0$ for all $x \in [d]$. The above lower bound implies that the algorithm has to pay regret $\Omega(T_0) = \Omega(1/(\varepsilon \log T_0)^2)$. The claim follows as $T_0 \leq T$. ■

C.2. Proof for Theorem 10

To prove a lower bound for pure DP, we use the following version of Theorem 15 for this setting.

Theorem 16 (Steinke and Ullman, 2017a, Theorem A.1). *Let $d = 1000n$ and n sufficiently large. Let $\mathcal{X} = \{X \in \{-1, +1\}^{n \times d} : \text{all the columns in } X \text{ are consensus columns}\}$. Let $\mathcal{A} : \{-1, +1\}^{n \times d} \rightarrow \{-1, +1\}^d$ be an algorithm such that for all $X \in \mathcal{X}$,*

$$\mathbb{E}[\|\mathcal{A}(X) - \text{sign}(X)\|_0] \leq 1/4.$$

Then \mathcal{A} is not 1-DP.

Using the bound of Theorem 16 and following the same steps as in the proof of Theorem 15, the lower bound of Theorem 10 now follows.

C.3. Proof of Theorem 11

We use similar ideas to the one in the proof of Theorem 9 where we used a DP-OPE algorithm for sign estimation. Instead of designing two experts for each column, the idea here is to look at subsets of columns of size k and design 2^k experts to represent the sign vector of these k columns.

Given an input $X \in \{-1, +1\}^{n \times p}$ where we assume for simplicity that p/k is an integer, we design an expert problem with $d = 2^k \binom{p}{k}$ experts. Instead of representing the experts as integers $x \in [d]$, we use an equivalent representation where an expert is a pair (S, v) where $S \subset [p]$ is a set of columns of size k and $v \in \{-1, +1\}^k$ represents the signs that this expert assigns for columns in S . We initialize the vector of signs $s_j = 0$ for all $1 \leq j \leq p$.

Here we have $T = 0.9p/k$ rounds and at round $1 \leq t \leq T$ we sample a user $i_t \sim [n]$ (arbitrarily while enforcing that each $i \in [n]$ appears at most $2T/n$ times) and play a loss function ℓ_t such that

$$\ell_t(S, v) = \begin{cases} 1 & \text{if } s_j \neq 0 \text{ for some } j \in S \\ 0 & \text{otherwise if } \text{sign}(\bar{X}_S) = v \\ 1 & \text{otherwise} \end{cases}$$

Now, given the output of the algorithm \mathcal{A} at time t , that is $x_t = (S_t, v_t)$ we set $s_{S_t} = v_t$ (we assume without loss of generality that each $j \in [p]$ will appear in at most a single S_t . Otherwise, similarly to the proof of Theorem 9, we can ensure this property while not increasing the regret). Moreover, at the end of the game, there is a set $S \subset [p]$ of size k that contains only consensus columns which were not estimated earlier ($S \cap S_t = \emptyset$ for all t). This follows from the fact that X has at least $0.999p$ consensus columns hence there is at least $0.05p \geq k$ consensus columns that have not appeared in S_1, \dots, S_T , hence there is an expert (S, v) such that

$$\sum_{t=1}^T \ell_t(S, v) = 0.$$

Now we show that if an algorithm \mathcal{A} has small regret, then the vector s estimates the sign of at least $0.8p$ columns. Observe that the regret of the algorithm is

$$\begin{aligned} \sum_{t=1}^T \ell_t(x_t) &= \sum_{t=1}^T \ell_t(S_t, v_t) \\ &= \sum_{t=1}^T 1\{\text{sign}(\bar{X}_{S_t}) \neq v_t\} \\ &= \sum_{t=1}^T 1\{\text{sign}(\bar{X}_{S_t}) \neq s_{S_t}\}. \end{aligned}$$

Assume towards a contradiction that \mathcal{A} is (ε, δ) -DP where $\varepsilon \leq \frac{\sqrt{k/T}}{200 \log(T)}$ and $\delta \leq 1/T^3$ and that the expected regret is at most $T/1000$. Markov inequality implies that with probability at least $9/10$ the regret is at most $T/100$. Note that $|S| = kT = 0.9p$. Under this event we have

$$\sum_{t=1}^T 1\{\text{sign}(\bar{X}_{S_t}) \neq s_{S_t}\} \leq 0.002T.$$

Hence we have that $\text{sign}(\bar{X}_{S_t}) = s_{S_t}$ for at least $0.9T$ rounds. As each round has k distinct columns, we have $s_j = \text{sign}(\bar{X}_j)$ for at least $0.9kT \geq 0.8p$. As there are at most $0.001p$ non-consensus columns, this means that $s_j = \text{sign}(\bar{X}_j)$ for at least $0.75p$ consensus columns. Now we prove that s is also $(1, 1/T^2)$ -DP which gives a contradiction to Theorem 15. To this end, note that the algorithm \mathcal{A} is (ε, δ) -DP where $\varepsilon \leq \frac{\sqrt{k/T}}{200 \log(T)} \leq \frac{k/\sqrt{p}}{200\rho \log(p)}$. Moreover, recall that each row $i \in [n]$ appears at most $k_i \leq 2T/n \leq 2p/(nk) \leq 200\sqrt{p} \log(p)/k$ times, hence group privacy implies the final output s is $(\max_i k_i \varepsilon, T\delta)$ -DP, that is, $(1, 1/T^2)$ -DP.

Appendix D. Proofs for Section 5

D.1. Proof of Theorem 12

We assume without loss of generality that $L = 1$ (otherwise divide the loss by L). As \mathcal{X} has diameter D , we can construct a cover $C = \{c_1, \dots, c_M\}$ of \mathcal{X} such that $\min_{i \in [M]} \|x - c_i\|_2 \leq \rho$ for all $x \in \mathcal{X}$ where $M \leq 2^{d \log(4/\rho)}$ (Duchi, 2019, Lemma 7.6). Consider the following algorithm: run Algorithm 1 where the experts are the elements of the cover C . Corollary 3 now implies that this algorithm has regret

$$\mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in C} \sum_{t=1}^T \ell_t(x) \right] \leq O \left(\sqrt{T \ln M} + \frac{T^{1/3} \log^{1/3}(1/\delta) \ln M}{\varepsilon} \right).$$

Since ℓ_t is 1-Lipschitz, we now get

$$\begin{aligned}
 \mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^T \ell_t(x) \right] &\leq \mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in \mathcal{C}} \sum_{t=1}^T \ell_t(x) + \min_{x \in \mathcal{C}} \sum_{t=1}^T \ell_t(x) - \min_{x \in \mathcal{X}} \sum_{t=1}^T \ell_t(x) \right] \\
 &= O \left(\sqrt{T \ln M} + \frac{T^{1/3} \log^{1/3}(1/\delta) \ln M}{\varepsilon} + T\rho \right) \\
 &= O \left(\sqrt{Td \log(1/\rho)} + \frac{T^{1/3} \log^{1/3}(1/\delta) d \log(1/\rho)}{\varepsilon} + T\rho \right) \\
 &= O \left(\sqrt{Td \log(T)} + \frac{T^{1/3} d \log^{1/3}(1/\delta) \log(T)}{\varepsilon} \right),
 \end{aligned}$$

where the last inequality follows by setting $\rho = 1/T$.

D.2. Proof of Corollary 13

The algorithm \mathcal{A}_{ℓ_2} is (ε, δ) -DP and has excess loss $\Delta_n = LD \cdot O(1/\sqrt{n} + \sqrt{d}/n\varepsilon)$. Thus, Theorem 7 implies that

$$\begin{aligned}
 \mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \min_{x \in [d]} \sum_{t=1}^T \ell_t(x) \right] &\leq \sum_{i=1}^{\log T} 2^i \Delta_i \\
 &\leq O(LD) \sum_{i=1}^{\log T} 2^{i/2} + \sqrt{d}/\varepsilon \\
 &\leq LD \cdot O(\sqrt{T} + \sqrt{d} \log(T)/\varepsilon).
 \end{aligned}$$

Appendix E. Concentration for sums of geometric variables

In this section, we prove a concentration result for the sum of geometric random variables, which allows us to upper bound the number of switches in the sparse-vector based algorithm. We say that Z is geometric random variable with success probability p if $P(W = k) = (1-p)^{k-1}p$ for $k \in \{1, 2, \dots\}$. To this end, we use the following Chernoff bound.

Lemma 17 (Mitzenmacher and Upfal, 2005, Ch. 4.2.1). *Let $X = \sum_{i=1}^n X_i$ for $X_i \stackrel{\text{iid}}{\sim} \text{Ber}(p)$. Then for $\delta \in [0, 1]$,*

$$\mathbb{P}(X > (1 + \delta)np) \leq e^{-np\delta^2/3} \quad \text{and} \quad \mathbb{P}(X < (1 - \delta)np) \leq e^{-np\delta^2/2}.$$

The following lemma demonstrates that the sum of geometric random variables concentrates around its mean with high probability.

Lemma 18. *Let W_1, \dots, W_n be iid geometric random variables with success probability p . Let $W = \sum_{i=1}^n W_i$. Then for any $k \geq n$*

$$\mathbb{P}(W > 2k/p) \leq \exp(-k/4).$$

Proof. Notice that W is distributed according to the negative binomial distribution where we can think of W as the number of Bernoulli trials until we get n successes. More precisely, let $\{B_i\}$ for $i \geq 1$ be Bernoulli random variables with probability p . Then the event $W > t$ has the same probability as $\sum_{i=1}^t B_i < n$. Thus we have that

$$\mathbb{P}(W > t) \leq \mathbb{P}\left(\sum_{i=1}^t B_i < n\right).$$

We can now use Chernoff inequality (Lemma 17) to get that for $t = 2n/p$:

$$\mathbb{P}\left(\sum_{i=1}^t B_i < n\right) \leq \exp(-tp/8) = \exp(-n/4).$$

This proves that

$$\mathbb{P}(W > 2n/p) \leq \exp(-n/4).$$

The claim now follows by noticing that $\sum_{i=1}^n W_i \leq \sum_{i=1}^k W_i$ for W_i iid geometric random variable when $k \geq n$, thus $\mathbb{P}(\sum_{i=1}^n W_i \geq 2k/p) \leq \mathbb{P}(\sum_{i=1}^k W_i \geq 2k/p) \leq \exp(-k/4)$

■