Improved Differentially Private and Lazy Online Convex Optimization: Lower Regret without Smoothness Requirements

Naman Agarwal [1] **Satyen Kale** [2] **Karan Singh** [3] **Abhradeep Thakurta** [1]

Abstract

We design differentially private regret-minimizing algorithms in the online convex optimization (OCO) framework. Unlike recent results, our algorithms and analyses do not require smoothness, thus yielding the first private regret bounds with an optimal leading-order term for non-smooth loss functions. Additionally, even for smooth losses, the resulting regret guarantees improve upon previous results in terms their dependence of dimension. Our results provide the best known rates for DP-OCO in all practical regimes of the privacy parameter, barring when it is exceptionally small. The principal innovation in our algorithm design is the use of sampling from *strongly* log-concave densities which satisfy the Log-Sobolev Inequality. The resulting concentration of measure allows us to obtain a better trade-off for the dimension factors than prior work, leading to improved results. Following previous works on DP-OCO, the proposed algorithm explicitly limits the number of switches via rejection sampling. Thus, independently of privacy constraints, the algorithm also provides improved results for online convex optimization with a switching budget.

1. Introduction

The framework of online convex optimization (OCO) provides a unified treatment of computational and statistical aspects of decision-making and learning under uncertainty. In it, in each round $t = 1, 2, \ldots, T$, a learner chooses an element x_t from a compact convex set $\mathcal{K} \in \mathbb{R}^d$, after which an adversarially chosen Lipschitz convex loss function $l_t : \mathcal{K} \to \mathbb{R}$ in revealed. Thus the learner suffers loss $l_t(x_t)$ in round t. The learner's *regret* is defined as $\sum_{t=1}^{T} l_t(x_t) - \min_{x\in\mathcal{K}} \sum_{t=1}^{T} l_t(x)$. The learner's performance may be assessed via her expected regret, averaging over the randomness in her and the adversary's choices. In this paper, we restrict our discussion to *obliviously* chosen loss functions, i.e., they are independent of the iterates x_t.

Differentially Private OCO (DP-OCO). Online learning algorithms often operate over sensitive data, e.g. user-level data for personalization-oriented services. Hence, in addition to minimizing regret, limiting privacy leakage is desirable, and has been pursued in Jain et al. (2012); Smith & Thakurta (2013); Agarwal & Singh (2017); Kairouz et al. (2021); Asi et al. (2023); Agarwal et al. (2023). The promise of privacy in DP-OCO dictates that if a loss function l_t in any round t were changed to a different function l'_t, then distribution over the *entire* iterate sequence produced by the algorithm is not altered in a quantitative (ε-dependent) distributional sense. Kairouz et al. (2021) established a regret upper bound of $\widetilde{\mathcal{O}}\left(\frac{d^{1/4}\sqrt{T}}{\sqrt{\varepsilon}}\right)$[1]. This was improved in a series of works (Asi et al., 2023; Agarwal et al., 2023), for moderate ranges of ε, with Agarwal et al. (2023) providing the known best bound of $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{dT^{1/3}}{\varepsilon}\right)$; notably, the first term here matches the optimal non-private regret. There are two shortcomings of the result in Agarwal et al. (2023). Firstly, it only applies to smooth loss functions. Moreover, application of artificial smoothing techniques, e.g., Moreau-Yoshida smoothing, to arrive at a result for non-smooth losses yields bounds that are uniformly worse than in Kairouz et al. (2021), and is therefore fruitless. Secondly, the dependence on the dimension in the second term is suboptimal; while specifically for the class of generalized linear model (GLM) functions, the authors proved an improved bound of $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{\sqrt{d}T^{1/3}}{\varepsilon}\right)$, improving the second term by a factor of $\sqrt{d}$, obtaining the above bound for the general class of Lipschitz convex functions was left open. In this paper, we resolve both these problems, and offer an unconditional improvement. Concretely, we provide a DP-OCO algorithm for potentially non-smooth convex Lip-

*Equal contribution [1]Google DeepMind [2]Google Research [3]Tepper School of Business, Carnegie Mellon University. Correspondence to: Naman Agarwal <namanagarwal@google.com>, Satyen Kale <satyenkale@google.com>, Karan Singh <karansingh@cmu.edu>, Abhradeep Thakurta <athakurta@google.com>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024.

[1] $\widetilde{\mathcal{O}}(\cdot)$ hides polylog factors in $1/\delta$ and T.

schitz losses with $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{\sqrt{d}T^{1/3}}{\varepsilon}\right)$ regret.[2] This now provides the known best regret bound for DP-OCO, for all practical regimes of the privacy parameters. We provide a detailed comparison in Table 1.

Lazy OCO. Regret bounds for online leaners subject to limited switching has seen thorough investigation for prediction with expert advice (Merhav et al., 2002; Kalai & Vempala, 2005; Geulen et al., 2010; Altschuler & Talwar, 2021) and more generally for OCO (Anava et al., 2015; Sherman & Koren, 2021). For smooth losses, the best results known so far for OCO appear in Agarwal et al. (2023); Sherman & Koren (2023) who show a regret bound of $\tilde{O}(\sqrt{T} + \frac{dT}{S})$ while switching at most S times in expectation. Additionally, Agarwal et al. (2023) gave an improved $\tilde{O}(\sqrt{T} + \frac{\sqrt{d}T}{S})$ bound under additional restriction to smooth GLM losses. We improve these results by giving an OCO algorithm that has regret at most $\tilde{O}(\sqrt{T} + \frac{\sqrt{d}T}{S})$ for Lipschitz losses, in absence of any smoothness or GLM restrictions. This also improves the known best switching-limited regret bounds for non-smooth losses in Anava et al. (2015); Chen et al. (2020) which scale as $\tilde{O}(\sqrt{dT} + \frac{dT}{S})$ and $O(\frac{T}{\sqrt{S}})$, respectively. In contrast to our result, neither attains an optimal $\sqrt{T}$ leading-order term for a sub-linear switching budget.

Overview of our techniques. The starting point of our algorithm is the Private Shrinking Dartboard algorithm of Asi et al. (2023) which proposes to play a point sampled per step from the distribution with density at x proportional to $\exp\left(-\beta\left(\sum_{\tau=1}^{t} l_\tau(x)\right)\right)$. In this case the regret as well as the differential privacy of the algorithm is governed by the parameter β. As can be easily seen from the regret analyses for such continuous multiplicative weights algorithms, the choice of β scales with the dimension (corresponding to the volume of the domain) leading to an overall regret of $\sqrt{dT}$.

In contrast we propose to sample the played point every round from the distribution with density at x proportional to $\exp\left(-\beta\left(\sum_{\tau=1}^{t} l_\tau(x) + \frac{\lambda}{2}\|x\|^2\right)\right)$. The primary advantage of adding the ℓ_2-norm term is that the stability of the algorithm which affects the regret as well as the differential privacy can now be governed directly by the λ parameter. This is a consequence of that fact that the above density satisfies the Log-Sobolev inequality and thus we are able to leverage the geometry of the underlying domain using the resulting concentration of measure, as opposed to a crude bound on its volume. In particular as we show in Lemma 3.5, the stability of the algorithm which can be measured via the ratio between the densities of distributions at two consecutive timesteps scales as $O\left(\sqrt{\beta/\lambda}\right)$ as opposed to $O(\beta)$ in the Private Shrinking Dartboard algorithm. Tuning the λ parameter appropriately now allows for a significantly better trade-off with respect to the dimension. The idea of using a strongly log-concave distribution has been used previously in the context differentially-private stochastic optimization in the work of Gopi et al. (2022) and Ganesh et al. (2023).

This paper builds upon an earlier paper by the same authors (Agarwal et al., 2023). To maintain continuity and ease understanding for readers who wish to read both papers, the present paper is structured very similarly to Agarwal et al. (2023). In particular, several lemmas are adaptations of analogs in Agarwal et al. (2023); we chose to include them for completeness. Given the similarity between the two papers, we highlight a few key differences and points of technical novelty in Section 5.

2. Preliminaries

Notation. We use $\|\cdot\|$ to denote the standard ℓ_2 norm on $\mathbb{R}^d$. For distributions p and q on the same outcome space, we use $\|p - q\|_{\mathrm{TV}}$ to denote their total variation distance. For a distribution μ on $\mathbb{R}^d$, we use $\mu(A)$ to denote the measure of a measurable set $A \subseteq \mathbb{R}^d$. With some abuse of notation, we also $\mu(x)$ to denote the density of μ at $x \in \mathbb{R}^d$, if it exists.

Problem Setting. We have a convex compact set $\mathcal{K} \subset \mathbb{R}^d$ with diameter $D \triangleq \max_{x,y\in\mathcal{K}} \|x-y\|$. At each step $t \in [T]$, the learner $\mathcal{A}$ chooses a point $x_t \in \mathcal{K}$, after which she sees a loss function $l_t : \mathcal{K} \to \mathbb{R}$, and suffers a loss of $l_t(x_t)$. For any t-indexed sequence of objects, e.g. the loss function l_t, let $l_{1:T} = (l_1, \dots l_T)$ be the concatenated sequence. We consider *oblivious adversaries* in that we assume the loss function sequence $l_{1:T}$ is chosen independently of the iterates x_t picked by the learner.[3] A function $l : \mathcal{K} \to \mathbb{R}$ is said to be G-Lipschitz if $|l(x) - l(y)| \le G\|x - y\|$ for any pair $x, y \in \mathcal{K}$. Without loss of generality, we assume that $\mathcal{K}$ is full-dimensional and contains the origin.

Assumption 2.1. *The loss functions $l_{1:T} \in \mathcal{L}^T$ are chosen* obliviously *from the class $\mathcal{L}$ of G-Lipschitz convex functions.*

The expected regret assigned to the learner is the expected excess aggregate loss of the learner in comparison to the best fixed point in $\mathcal{K}$ chosen with the benefit of hindsight.

$$\mathcal{R}_T(\mathcal{A}, l_{1:T}) \triangleq \mathop{\mathbb{E}}_{\mathcal{A}}\left[\sum_{t=1}^{T} l_t(x_t) - \min_{x^*\in\mathcal{K}} \sum_{t=1}^{T} l_t(x^*)\right]$$

[2] For a small regime of $\varepsilon \in [T^{-1/6}, d^{2/3}T^{-1/6}]$ we get an additional $\frac{\sqrt{d}T^{3/8}}{\varepsilon^{3/4}}$ term; this also arises in Agarwal et al. (2023) for GLM losses. Even with this included, the bound unconditionally improves Agarwal et al. (2023) as shown in Table 1 on page 3.

[3] As noted in Asi et al. (2023), the privacy guarantee is not reliant on obliviousness, but the regret bounds are. The necessity of obliviousness is due to our use of the Shrinking Dartboard algorithm. For adaptive adversaries, Asi et al. (2023) show that no algorithm can achieve sublinear regret when $\varepsilon \le 1/\sqrt{T}$.

Privacy Parameter ε	Previous Best		Theorem 4.4
	With Smoothness	No Smoothness	(No Smoothness)
$\varepsilon \geq dT^{-1/6}$	$\sqrt{T}$ (Agarwal et al., 2023)	$\sqrt{dT}$ (Asi et al., 2023)	$\sqrt{T}$
$\varepsilon \in [d^{2/3}T^{-1/6}, dT^{-1/6}]$	$d \cdot T^{1/3} \cdot \varepsilon^{-1}$ (Agarwal et al., 2023)		$\sqrt{T}$
$\varepsilon \in [\sqrt{d}T^{-1/6}, d^{2/3}T^{-1/6}]$			$\sqrt{d} \cdot T^{3/8} \cdot \varepsilon^{-3/4}$
$\varepsilon \in [T^{-1/6}, \sqrt{d}T^{-1/6}]$		$d \cdot T^{1/3} \cdot \varepsilon^{-1}$ (Asi et al., 2023)	
$\varepsilon \in [d^{3/2}T^{-1/3}, T^{-1/6}]$			$\sqrt{d} \cdot T^{1/3} \cdot \varepsilon^{-1}$
$\varepsilon \in [dT^{-1/3}, d^{3/2}T^{-1/3}]$	$d^{1/4} \cdot T^{1/2} \cdot \varepsilon^{-1/2}$ (Kairouz et al., 2021)		
$\varepsilon \leq dT^{-1/3}$	$d^{1/4} \cdot T^{1/2} \cdot \varepsilon^{-1/2}$ (Kairouz et al., 2021) (Current Best)		$\sqrt{d} \cdot T^{1/3} \cdot \varepsilon^{-1}$

Table 1. Landscape of the best known results for DP-OCO across different regimes of the privacy parameter. In red, we highlight our results where they are *strictly better* than the known best result for *both* smooth and non-smooth losses, and in blue are highlighted improved results for *non-smooth* losses alone. Notice that our algorithm strictly improves the best known results *without assuming smoothness* for all $\varepsilon \geq dT^{-1/3}$. While we focus on factors of d, for the asymptotics we assume $T \gg d$.

Number of Switches S	Previous Best		Theorem 4.5
	With Smoothness	No Smoothness	(No Smoothness)
$d\sqrt{T} \leq S \ll T$	$\sqrt{T}$ (Agarwal et al., 2023) (Sherman & Koren, 2023)	$\sqrt{dT}$ (Anava et al., 2015)	$\sqrt{T}$
$\sqrt{dT} \leq S \leq d\sqrt{T}$	$d \cdot T/S$ (Agarwal et al., 2023) (Sherman & Koren, 2023)		$\sqrt{T}$
$d^2 \leq S \leq \sqrt{d} \cdot T$	$d \cdot T/S$ (Anava et al., 2015)		$\sqrt{d} \cdot T/S$
$d \leq S \leq d^2$	$T/\sqrt{S}$ (Chen et al., 2020)		
$S \leq d$	$T/\sqrt{S}$ (Chen et al., 2020) (Current Best)		$\sqrt{d} \cdot T/S$

Table 2. Comparison between our results and the known best results for Lazy OCO in different regimes for the switching budget S. In red, we highlight our results where they are *strictly better* than the known best result for *both* smooth and non-smooth losses, and in blue are highlighted improved results for *non-smooth* losses alone. While we focus on factors of d, for asymptotics we assume $T \gg d$.

Without making any distributional assumptions, we will bound the *worst-case* regret, $\mathcal{R}_T(\mathcal{A}) \triangleq \max_{l_{1:T} \in \mathcal{L}^T} \mathcal{R}_T(\mathcal{A}, l_{1:T})$, taken over all loss sequences.

The expected number of discrete switches for a given learner can be calculated as

$$\mathcal{S}_T(\mathcal{A}, l_{1:T}) \triangleq \mathbb{E}_{\mathcal{A}}\left[\sum_{t=2}^{T} \mathbb{I}_{x_t \neq x_{t-1}}\right].$$

For brevity, henceforth we will simply use $\mathcal{R}_T$ and $\mathcal{S}_T$ to refer to $\mathcal{R}_T(\mathcal{A}, l_{1:T})$ and $\mathcal{S}_T(\mathcal{A}, l_{1:T})$ respectively.

Finally, an online learning algorithm $\mathcal{A}$ is said to (ε, δ)-differentially private if for any loss function sequence pair $l_{1:T}, l'_{1:T} \in \mathcal{L}^T$ such that $l_t = l'_t$ for all but possibly one $t \in [T]$, we have for any Lebesgue measurable $O \subset \mathcal{K}^T$:

$$\Pr_{\mathcal{A}}(x_{1:T} \in O | l_{1:T}) \leq e^{\varepsilon} \Pr_{\mathcal{A}}(x_{1:T} \in O | l'_{1:T}) + \delta.$$

To finish up, we recall the adaptive strong composition lemma for differentially-private mechanisms.

Lemma 2.2 (Whitehouse et al. (2022))**.** *Let $\mathcal{A}_t : \mathcal{L}^{t-1} \times \mathcal{K}^{t-1} \to \mathcal{K}$ be a t-indexed family of $(\varepsilon_t, \delta_t)$-differentially private algorithms, i.e. for every t, for any pair of sequences of loss functions $l_{1:t-1}, l'_{1:t-1} \in \mathcal{L}^{t-1}$ differing in at most one index in $[t-1]$, and any $x_{1:t-1} \in \mathcal{K}^{t-1}$:*

$$P_{\mathcal{A}_t}(x_t | l_{1:t-1}, x_{1:t-1}) \leq e^{\varepsilon} P_{\mathcal{A}_t}(x_t | l'_{1:t-1}, x_{1:t-1}) + \delta.$$

Define a new t-indexed family $\mathcal{B}_t : \mathcal{L}^{t-1} \to \mathcal{K}^t$ recursively starting with $\mathcal{B}_1 = \mathcal{A}_1$ as

$$\mathcal{B}_t(l_{1:t-1}) = \mathcal{B}_{t-1}(l_{1:t-2}) \circ \mathcal{A}_t(l_{1:t-1}, \mathcal{B}_{t-1}(l_{1:t-2})).$$

Then for any $\delta'' > 0$, $\mathcal{B}_T$ is (ε', δ')-differentially private, where

$$\varepsilon' = \frac{3}{2}\sum_{t=1}^{T} \varepsilon_t^2 + \sqrt{6\sum_{t=1}^{T} \varepsilon_t^2 \log\frac{1}{\delta''}}, \qquad \delta' = \delta'' + \sum_{t=1}^{T} \delta_t.$$

Lastly, we state a lemma routinely used in online learning to decompose regret into incremental stability terms.

Lemma 2.3 (FTL-BTL (Hazan, 2016)). *For any loss function sequence $l_{0:T}$ over any set $\mathcal{B}$, define*

$$y_t = \operatorname*{argmin}_{x \in \mathcal{B}} \left\{ \sum_{i=0}^{t-1} l_i(x) \right\}.$$

Then, for any $x \in \mathcal{B}$, we have

$$\sum_{t=0}^{T} l_t(y_{t+1}) \leq \sum_{t=0}^{T} l_t(x).$$

3. Preliminary results for Gibbs measures

In this paper we consider a class of Gibbs distributions over the set $\mathcal{K}$. Given any function $f : \mathcal{K} \in \mathbb{R}$, a temperature constant $\beta \geq 0$ and a regularization parameter $\lambda \geq 0$ we define $\mu(f, \beta, \lambda) : \mathcal{K} \to \mathbb{R}_+$ to be a measure function defined as

$$\mu(f, \beta, \lambda)(x) = \exp\left(-\beta \cdot \left(f(x) + \frac{\lambda}{2}\|x\|^2\right)\right). \quad (3.1)$$

We further define $Z(f, \beta, \lambda)$ to be normalization constant of the above function defined as

$$Z(f, \beta, \lambda) = \int_{x \in \mathcal{K}} \exp\left(-\beta \cdot \left(f(x) + \frac{\lambda}{2}\|x\|^2\right)\right) dx. \quad (3.2)$$

Using the above we can define a probability density $\bar{\mu}(f, \beta, \lambda)(x)$ over $\mathcal{K}$ as follows

$$\bar{\mu}(f, \beta, \lambda)(x) \triangleq \frac{\mu(f, \beta, \lambda)(x)}{Z(f, \beta, \lambda)}. \quad (3.3)$$

We will interchangeably use the notation $\bar{\mu}$ for the probability density function as well as the distribution itself. We will suppress β, λ from the above definitions when they will be clear from the context. In the following we collect some useful definitions and results pertaining to concentration of measure resulting from the Log-Sobolev Inequality.

Definition 3.1. *A distribution P satisfies the Log-Sobolev Inequality (LSI) with constant c if for all smooth functions $g : \mathbb{R}^d \to \mathbb{R}$ with $\mathbb{E}_{x\sim P}[g(x)^2] < \infty$:*

$$\mathbb{E}_{x\sim P}[g(x)^2 \log(g(x)^2)] - \mathbb{E}_{x\sim P}[g(x)^2]\mathbb{E}_{x\sim P}[\log(g(x)^2)] \leq \frac{2}{c}\mathbb{E}_{x\sim P}[\|\nabla g(x)\|^2]$$

Lemma 3.2 (Proposition 3 and Corollaire 2 in Bakry & Émery (2006)). *Given a Λ-strongly convex function l, let Q be the distribution supported over $\mathcal{K}$ with density $\mu(x)$ proportional to* $\exp(-\beta \cdot l(x))$. *Then Q satisfies LSI (Definition 3.1) with constant $c = \beta\Lambda$.*

Lemma 3.3 (Concentration of Measure; follows from Herbst's argument presented in Section 2.3 of Ledoux (1999)). *Let F be a L-Lipschitz function and let Q be a distribution satisfying LSI with a constant c then*

$$\Pr_{X\sim Q}(|F(X) - \mathbb{E}[F(X)]| \geq r) \leq 2\exp\left(-\frac{c \cdot r^2}{2L^2}\right)$$

The following definition defines a notion of closeness for two Gibbs-measures:

Definition 3.4. *Two Gibbs distributions $\bar{\mu}$, $\bar{\mu}'$ on $\mathcal{K}$ are said to be (Φ, δ)-close if the following inequalities hold:*

$$\Pr_{X\sim\bar{\mu}}\left[\frac{1}{\Phi} \leq \frac{\bar{\mu}(X)}{\bar{\mu}'(X)} \leq \Phi\right] \geq 1 - \delta$$
$$\Pr_{X\sim\bar{\mu}'}\left[\frac{1}{\Phi} \leq \frac{\bar{\mu}(X)}{\bar{\mu}'(X)} \leq \Phi\right] \geq 1 - \delta$$

One of the core components of our analysis is to show that the Gibbs-measures are *smooth* under changes of the underlying functions. A similar result was also proved by Gopi et al. (2022, Theorem 4) and a slightly looser bound by Ganesh et al. (2023). Missing proofs in this section can be found in Appendix A.

Lemma 3.5 (Density ratio). *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be convex functions such that $l - l'$ is G-Lipschitz. Further let $\beta, \lambda \geq 0$ be parameters and define the Gibbs-distributions $\bar{\mu} = \bar{\mu}(l, \beta, \lambda)$ and $\bar{\mu}' = \bar{\mu}(l', \beta, \lambda)$ (as defined in (3.1)). Then for any $\delta \in (0, 1]$, we have that $\bar{\mu}$ and $\bar{\mu}'$ are (Φ, δ) close where*

$$\Phi = \exp\left(\frac{2\beta G^2}{\lambda} + \sqrt{\frac{8\beta G^2 \log(2/\delta)}{\lambda}}\right)$$

The proof of the lemma crucially uses the following bound on the Wasserstein-distance of the Gibbs-distributions and other machinery developed by Ganesh et al. (2023).

Lemma 3.6 (Wasserstein Distance). *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be convex functions such that $l - l'$ is G-Lipschitz. Further let $\beta, \lambda \geq 0$ be parameters and define the Gibbs-distributions $\bar{\mu} = \bar{\mu}(l, \beta, \lambda)$ and $\bar{\mu}' = \mu(l', \beta, \lambda)$ (as defined in (3.3)). Then we have that ∞-Wasserstein distance between $\bar{\mu}$ and $\bar{\mu}'$ over the ℓ_2 metric is bounded as*

$$W_\infty(\bar{\mu}, \bar{\mu}') \leq \frac{G}{\lambda}.$$

4. Algorithm and main result

Our proposed algorithm Private Continuous Online Multiplicative Weights with Euclidean Regularization (POMER)

Algorithm 1: Private Online Continuous Multiplicative Weights with Euclidean Regularization (POMER)

Inputs: A temperature parameter β, a regularization parameter $\lambda > 0$, switching rate parameter $p \in [0, 1]$, switching budget $B \geq 0$, a scale parameter $\Phi > 0$.
Let x_1 be a random variable sampled uniformly from $\mathcal{K}$.
for $t = 1$ *to* T **do**
  Play $x_t \in \mathcal{K}$.
  Observe $l_t : \mathcal{K} \to \mathbb{R}$ and suffer a loss of $l_t(x_t)$.
  Define the measure function
  $\mu_{t+1}(x) \triangleq \mu(f, \beta, \lambda)(x) \triangleq \exp\left(-\beta\left(\sum_{\tau=1}^{t} l_\tau(x) + \lambda\frac{\|x\|^2}{2}\right)\right)$
  Accordingly denote $\bar{\mu}_{t+1}(x)$ the probability density resulting from the measure μ_{t+1}. (cf. (3.3))
  Sample
  $S_t \sim \text{Ber}\left(\min\left\{1, \max\left\{\frac{1}{\Phi^2}, \frac{\bar{\mu}_{t+1}(x_t)}{\Phi\cdot\bar{\mu}_t(x_t)}\right\}\right\}\right)$ and $S'_t \sim \text{Ber}(1-p)$.
  if $b_t < B$ *and* ($S'_t = 0$ *or* $S_t = 0$) **then**
    Update $b_{t+1} = b_t + 1$ and draw an independent sample $x_{t+1} \sim \bar{\mu}_{t+1}$.
  end
  else
    Set $b_{t+1} = b_t$ and $x_{t+1} = x_t$.
  end
end

(Algorithm 1) builds upon the Private Shrinking Dartboard algorithm proposed by Asi et al. (2023) (also see Agarwal et al. (2023)). At a high level at every step the algorithm ensures that at every iteration it samples x_t from the distribution $\bar{\mu}_t$ over $\mathcal{K}$ corresponding to the measure function $\mu_t(x)$ defined as

$$\mu_t(x) = \mu\left(\sum_{\tau=1}^{t} l_\tau, \beta, \lambda\right) = \exp\left(-\beta\left(\sum_{\tau=1}^{t-1} l_\tau(x) + \lambda\frac{\|x\|^2}{2}\right)\right).$$

The distribution $\bar{\mu}_t$ is the same distribution as Online Continuous Multiplicative Weights (as used in Asi et al. (2023)) with an added strong-convexity term governed by λ. This additional strong-convexity term is key to the improvements provided in this paper as it provides a better trade-off between switching and regret.

The above scheme was first analyzed by Gopi et al. (2022) and was recently shown to be able to obtain optimal results in the case of stochastic convex optimization (Ganesh et al., 2023). In the online case however a direct application of the above scheme can leak a lot of private information since the algorithm can potentially alter its decisions in each round. To guard against this, as in the work of Asi et al. (2023); Agarwal et al. (2023), we use a rejection sampling procedure which draws inspiration from Geulen et al. (2010). Specifically, for any t, the point x_{t+1} is chosen to be equal to x_t with probability $\frac{\bar{\mu}_{t+1}(x_t)}{\Phi\bar{\mu}_t(x_t)}$, where Φ is a scaling factor.[4] With the remaining probability, we sample x_{t+1} independently from $\bar{\mu}_{t+1}$ (we call this a "switch"). This rejection sampling technique ensures that the distribution of x_{t+1} remains very close to $\bar{\mu}_{t+1}$. We rescale the density ratio $\frac{\bar{\mu}_{t+1}(x_t)}{\Phi\bar{\mu}_t(x_t)}$ appropriately to make sure it is at most unit sized with high probability.

We now turn to the regret analysis for Algorithm 1. The following theorem is proved in Section 4.2:

Theorem 4.1 (Regret bound for POMER). *In Algorithm 1, fix any $\beta, \lambda > 0$, any $\delta \in [0, 1/2]$, any $p \in [0, 1]$, and choose Φ such that for all t the distributions $\bar{\mu}_t, \bar{\mu}_{t+1}$ are (Φ, δ)-close. For any sequence of obliviously chosen G-Lipschitz, convex loss functions $l_{1:T}$, the following hold:*

- *If $B = \infty$,*
$$\mathcal{R}_T \leq \frac{\lambda D^2}{2} + \frac{G^2 T}{\lambda} + \frac{d\log(T)}{\beta} + GD + 6GD\delta T^2.$$

- *Let $\tilde{p} = p + 1 - \Phi^{-2}$. If $B = 3\tilde{p}T$,*
$$\mathcal{R}_T \leq \frac{\lambda D^2}{2} + \frac{G^2 T}{\lambda} + \frac{d\log(T)}{\beta} + 2GDT(e^{-\tilde{p}T} + 3\delta T) + GD.$$

The following lemma (originally proved in (Agarwal et al., 2023)), gives a bound on the number of switches made by the Algorithm 1 and immediately follows by observing that the probability of switching in any round is at most $\tilde{p}$ via a simple Chernoff bound. For completeness we provide a proof in Appendix B.

Lemma 4.2 (Switching bound). *For any $p \in [0, 1]$ and any $\Phi \geq 0$, setting $\tilde{p} = p + 1 - \Phi^{-2}$, we have that the number of switches is bounded in the following manner,*

$$\mathbb{E}\left[\mathcal{S}_T\right] \leq \tilde{p}T, \qquad \Pr\left[\mathcal{S}_T \geq 3\tilde{p}T\right] \leq e^{-\tilde{p}T}.$$

Finally, we turn to the privacy guarantee for Algorithm 1, proved in Appendix C, with a sketch in Section 4.3:

[4] Note one key difference in the definition of the acceptance probability from Asi et al. (2023): we use the ratio of *normalized* instead of unnormalized densities. This is crucial since the ratio of unnormalized densities may introduce unnecessarily high switching probability: consider the case where the two loss sequences which differ only at one time step t_0, with the constant 0 loss function in one sequence, and the constant 1 loss function in the other.

Theorem 4.3 (Privacy). *Given $\beta, \lambda > 0$ and $\delta \in (0, 1/2]$, for any $T \geq 12\log(1/\delta)$, let $\delta' = \frac{\delta T^{-2}}{60}$, $G' = 3G$. Suppose there exists $\Phi' > 0$ such that for all convex functions l, l' where $l - l'$ is G'-Lipschitz, we have that, the distributions $\bar{\mu}(l, \beta, \lambda)$ and $\bar{\mu}(l', \beta, \lambda)$ respectively are (Φ', δ')-close. Then for any sequence of G-Lipschitz convex functions, Algorithm 1 when run with $\Phi = \Phi'^2$, $p = \max\left(T^{-1/3}, \left(\frac{G^4\beta^2}{\lambda^2 \cdot \log^2(\Phi)}\right)^{1/3}\right)$, $\tilde{p} = p + 1 - \Phi^{-2}$ and $B = 3\tilde{p}T$ is $(\varepsilon, \delta + 3Te^{-(1-\Phi^{-2})T})$-differentially private where*

$$\varepsilon = 3\varepsilon'/2 + \sqrt{6\varepsilon'}\sqrt{\log(2/\delta)},$$

with

$$\varepsilon' = 7T^{2/3}\log^2(\Phi) + 12\log^3(\Phi)T + 11\left(\frac{G^4\beta^2}{\lambda^2}\right)^{1/3}\log^{4/3}(\Phi)T.$$

4.1. Bounds for Lipschitz loss functions

In order to apply the above results for OCO with convex G-Lipschitz loss functions, all we need to do is compute Φ. This bound was established by Lemma 3.5. Using Lemma 3.5 and combining Theorem 4.3 and Theorem 4.1, we get the following result via straightforward calculations:

Theorem 4.4 (DP OCO). *For any given $\varepsilon \leq 1, \delta \in (0, 1/2]$ and any $T \geq 12\log(1/\delta)$, set*

$$\lambda = \frac{G}{D}\max\left\{2\sqrt{T}, \frac{10^3T^{1/3}\sqrt{d}\log(T/\delta)}{\varepsilon}, \frac{10^3T^{3/8}\sqrt{d}\log(T/\delta)}{\varepsilon^{3/4}}\right\},$$

$$\beta = \frac{\lambda}{10^5 \cdot G^2\log^2(T/\delta)}\min\left\{\frac{\varepsilon^2}{T^{2/3}}, \frac{\varepsilon^{3/2}}{T^{3/4}}\right\}$$

and other parameters as in Theorem 4.3. Then we get that Algorithm 1 is (ε, δ) differentially private and additionally satisfies

$$\mathcal{R}_T \leq \widetilde{\mathcal{O}}\left(GD\sqrt{T} + GD \cdot \sqrt{d}\left(\frac{T^{1/3}}{\varepsilon} + \frac{T^{3/8}}{\varepsilon^{3/4}}\right)\right).$$

Similarly, for Lazy OCO, using Theorem 4.1, Lemma 4.2 and Lemma 3.5, we get the following result:

Theorem 4.5 (Lazy OCO). *For any $T \geq 3$ and any given bound $S \leq T$ on the number of switches, set $\delta = 2/T^2$, $\lambda = \max\left\{\frac{G\sqrt{2T}}{D}, \frac{\sqrt{512d}G\log(T)}{D} \cdot \frac{T}{S}\right\}$, $\beta = \frac{\lambda}{256G^2\log(T)} \cdot \frac{S^2}{T^2}$, $\Phi = \exp\left(\frac{2\beta G^2}{\lambda} + \sqrt{\frac{8\beta G^2\log(2/\delta)}{\lambda}}\right)$, $p = 0$ and $B = \infty$ in Algorithm 1. Then for any sequence of obliviously chosen G-Lipschitz convex loss functions $l_{1:T}$, where $\mathbb{E}[\mathcal{S}_T] \leq S$, Algorithm 1 satisfies the following:*

$$\mathcal{R}_T \leq GD\sqrt{2T} + 16GD\log(T) \cdot \frac{\sqrt{d} \cdot T}{S} + 13GD.$$

Proof. We begin by first bounding the number of switches using Lemma 4.2. We get that

$$\mathbb{E}[\mathcal{S}_T] \leq \tilde{p}T \leq (1 - \Phi^{-2})T \leq 2\log(\Phi)T \leq 2T\left(\underbrace{\frac{2\beta G^2}{\lambda}}_{=\frac{S^2}{128T^2\log(T)} \leq \frac{S}{128T}} + \underbrace{\sqrt{\frac{8\beta G^2\log(2/\delta)}{\lambda}}}_{\leq \frac{S}{4T}}\right) \leq S$$

To bound the regret note that Lemma 3.5 implies that the distributions μ_t, μ_{t+1} are (Φ, δ)-close and therefore Theorem 4.1 implies

$$\begin{aligned}\mathcal{R}_T &\leq \frac{\lambda D^2}{2} + \frac{G^2T}{\lambda} + \frac{d\log(T)}{\beta} + GD + 6GD\delta T^2 \\ &= \frac{\lambda D^2}{2} + \frac{G^2T}{\lambda} + \frac{256 \cdot d \cdot G^2\log^2(T)}{\lambda} \cdot \frac{T^2}{S^2} + 13GD \\ &\leq GD\sqrt{2T} + 16GD\log(T) \cdot \frac{\sqrt{d} \cdot T}{S} + 13GD.\end{aligned}$$

□

4.2. Regret analysis of Algorithm 1

Here we provide our analysis for the regret of Algorithm 1 given by Theorem 4.1. For notational convenience, define $\Pi : \mathbb{R} \to [\frac{1}{\Phi^2}, 1]$ as $\Pi(x) = \min\{1, \max\{\frac{1}{\Phi^2}, x\}\}$. Also define $\zeta_t \triangleq \mathbb{I}(S_t = 0 \text{ or } S'_t = 0)$. The following lemma adapted from Agarwal et al. (2023) obtains bounds on the actual distribution that x_t is sampled from in terms of $\bar{\mu}_t$:

Lemma 4.6 (Distribution drift). *Given $\delta \in [0, \frac{1}{2}]$ and $\Phi \geq 1$, suppose that for all $t \in [T]$, the Gibbs-measures μ_t, μ_{t+1} are (Φ, δ)-close. If q_t is the marginal distribution induced by Algorithm 1 on its iterates x_t, then we have that*

- *If $B = \infty$, then for all t, $\|q_t - \bar{\mu}_t\|_{\mathrm{TV}} \leq 3\delta(t-1)$.*
- *If $B = 3\tilde{p}T$, then we have*
$$\|q_t - \bar{\mu}_t\|_{TV} \leq e^{-\tilde{p}T} + 3\delta(t-1).$$

We prove this lemma in the Appendix B. Next we prove the main theorem bounding the regret of Algorithm 1, i.e. Theorem 4.1 here:

Proof of Theorem 4.1. Recall that we defined μ_t to be the distribution with density proportional as

$$\mu_t(x) \propto \exp\left(-\beta\left(\sum_{\tau=1}^{t-1} l_\tau(x) + \lambda \cdot \frac{\|x\|^2}{2}\right)\right)$$

Let q_t be the distribution induced by Algorithm 1 on its iterates x_t. Lemma 4.6 establishes that the sequence of iterates x_t played by Algorithm 1 follows μ_t approximately. We define a sequence of random variables $\{y_t\}$ wherein each y_t is sampled from μ_t independently. In the following we only prove the case when $B = 3\tilde{p}T$, the $B = \infty$ can easily be derived by using the bounds from Lemma 4.6 appropriately. We leverage the following lemma,

Lemma 4.7. (Levin & Peres, 2017) *For a pair of probability distributions μ, ν, each supported on $\mathcal{K}$, we have for any function $f : \mathcal{K} \to \mathbb{R}$ that*

$$|\mathbb{E}_{x\sim\mu} f(x) - \mathbb{E}_{x\sim\nu} f(x)| \leq 2\|\mu - \nu\|_{TV} \max_{x\in\mathcal{K}} |f(x)|.$$

We can now apply Lemma 4.7 to pair $x_t \sim q_t$ and $y_t \sim \mu_t$, using Lemma 4.6, and functions $\bar{l}_t(x) = l_t(x) - l_t(\bar{x})$, where $\bar{x} \in \mathcal{K}$ is chosen arbitrarily, to arrive at

$$\begin{aligned}\left|\mathbb{E}\left[\sum_{t=1}^{T}(l_t(x_t) - l_t(y_t))\right]\right| &\leq \sum_{t=1}^{T} |\mathbb{E}\left[l_t(x_t) - l_t(y_t)\right]| \\ &\leq \sum_{t=1}^{T} \left|\mathbb{E}\left[\bar{l}_t(x_t) - \bar{l}_t(y_t)\right]\right| \\ &\leq 2GDT\left(e^{-\tilde{p}T} + 3\delta T\right),\end{aligned} \tag{4.1}$$

where we use that $\max_t \max_{x\in\mathcal{K}} |l_t(x) - l_t(\bar{x})| \leq G\max_{x\in\mathcal{K}} \|x - \bar{x}\| \leq GD$. Therefore hereafter we only focus on showing the expected regret bound for the sequence y_t.

We take a distributional approach to the regret bound by defining the function $l_t^{\Delta} : \Delta(\mathcal{K}) \to \mathbb{R}$ as $l_t^{\Delta}(\mu) \triangleq \mathbb{E}_{x\sim\mu} l_t(x)$. We can now redefine the regret in terms of the distributions as follows

$$\mathrm{Regret}(\mu) = \sum_{t=1}^{T} l_t^{\Delta}(\mu_t) - \sum_{t=1}^{T} l_t^{\Delta}(\mu).$$

Let $x^* \triangleq \arg\min_{x\in\mathcal{K}} \sum_{t=1}^{T} l_t(x)$. Note that $\arg\min_{\mu\in\Delta(\mathcal{K})} \sum_{t=1}^{T} l_t^{\Delta}(\mu)$ is the Dirac-delta distribution at x^*, and that $\min_{\mu\in\Delta(\mathcal{K})} \sum_{t=1}^{T} l_t^{\Delta}(\mu) = \sum_{t=1}^{T} l_t(x^*)$. For a given value $\varepsilon \in [0,1]$ define the set $\mathcal{K}_\varepsilon : \{\varepsilon x + (1-\varepsilon)x^* | x \in \mathcal{K}\}$. Let μ_ε^* to be uniform distribution over the set $\mathcal{K}_\varepsilon$. It is now easy to see using the Lipschitzness of l_t,

$$\sum_{t=1}^{T} l_t^{\Delta}(\mu_\varepsilon^*) - \min_{\mu\in\Delta(\mathcal{K})} \sum_{t=1}^{T} l_t^{\Delta}(\mu) \leq GDT\varepsilon. \tag{4.2}$$

Further we define a proxy loss function $l_0(x) = \frac{\lambda}{2}\|x\|^2$ and correspondingly, l_0^{Δ}. Finally define μ_0 as the uniform distribution over the set $\mathcal{K}$. The following lemma establishes an equivalence between sampling from μ_t and a Follow-the-regularized-leader strategy in the space of distributions.

Lemma 4.8. *Consider an arbitrary distribution μ_0 on $\mathcal{K}$ (referred to as the prior) and f be an arbitrary bounded function on $\mathcal{K}$. Define the distribution μ over $\mathcal{K}$ with density $\mu(x) \propto \mu_0(x)e^{-(x)}$. Then we have that*

$$\mu = \arg\min_{\mu'\in\Delta(\mathcal{K})} \left(\mathbb{E}_{x\sim\mu'}[f(x)] + \mathrm{KL}(\mu'\|\mu_0)\right).$$

The lemma follows from the Gibbs variational principle and a proof is included in Appendix B. Using the above lemma, we have that at every step $t \geq 1$,

$$\mu_t = \arg\min_{\mu\in\Delta(\mathcal{K})} \left(\sum_{\tau=0}^{t-1} \beta \cdot l_\tau^{\Delta}(\mu) + \mathrm{KL}(\mu\|\mu_0)\right).$$

Using the above and the FTL-BTL Lemma (Lemma 2.3) we get the following

$$\begin{aligned}&\beta \cdot \left(\sum_{t=1}^{T}\left(l_t^{\Delta}(\mu_t) - l_t^{\Delta}(\mu_\varepsilon^*)\right)\right) \\ &\leq \beta \cdot \left(\sum_{t=1}^{T}\left(l_t^{\Delta}(\mu_t) - l_t^{\Delta}(\mu_{t+1})\right)\right) \\ &\quad + \beta \cdot (l_0^{\Delta}(\mu_\varepsilon^*) - l_0^{\Delta}(\mu_1)) + \mathrm{KL}(\mu_\varepsilon^*\|\mu_0) - \mathrm{KL}(\mu_0\|\mu_0) \\ &\leq \beta \cdot \left(\sum_{t=1}^{T}\left(\mathbb{E}_{x\sim\mu_t}[\beta \cdot l_t(x)] - \mathbb{E}_{x\sim\mu_{t+1}}[\beta \cdot l_t(x)]\right)\right) \\ &\quad + \beta \cdot l_0^{\Delta}(\mu_\varepsilon^*) + \mathrm{KL}(\mu_\varepsilon^*\|\mu_0)\end{aligned}$$

Now using Lemma 3.6, there is a coupling γ between μ_t and μ_{t+1} such that $\sup_{(x,x')\sim\gamma} \|x - x'\| \leq \frac{G}{\lambda}$. Using this coupling we get that,

$$\begin{aligned}&\sum_{t=1}^{T}\left(\mathbb{E}_{x\sim\mu_t}[l_t(x)] - \mathbb{E}_{x\sim\mu_{t+1}}[l_t(x)]\right) \\ &= \sum_{t=1}^{T} \mathbb{E}_{(x,x')\sim\gamma}[l_t(x) - l_t(x')] \\ &\leq \sum_{t=1}^{T} \mathbb{E}_{(x,x')\sim\gamma} G\|x - x'\| \leq \sum_{t=1}^{T} G^2/\lambda \leq \frac{G^2 T}{\lambda}\end{aligned}$$

Combining the above two displays one gets the following

$$\begin{aligned}\mathrm{Regret}(\mu_\varepsilon^*) &= \sum_{t=1}^{T} l_t^{\Delta}(\mu_t) - \sum_{t=1}^{T} l_t^{\Delta}(\mu_\varepsilon^*) \\ &\leq l_0^{\Delta}(\mu_\varepsilon^*) + \frac{G^2 T}{\lambda} + \frac{\mathrm{KL}(\mu_\varepsilon^*\|\mu_0)}{\beta} \\ &\leq \frac{\lambda D^2}{2} + \frac{G^2 T}{\lambda} + \frac{d}{\beta}\log(1/\varepsilon)\end{aligned}$$

where we use that $\mathrm{KL}(\mu_\varepsilon^*\|\mu_0) = d\log(1/\varepsilon)$, since μ_ε^* is the uniform distribution over $\mathcal{K}_\varepsilon \subseteq \mathcal{K}$ and $\frac{\mathrm{Vol}(\mathcal{K}_\varepsilon)}{\mathrm{Vol}(\mathcal{K})} = \varepsilon^d$. Setting $\varepsilon = 1/T$ and using (4.2) we get that for any μ,

$$\mathrm{Regret}(\mu) \leq \frac{\lambda D^2}{2} + \frac{G^2 T}{\lambda} + \frac{d\log(T)}{\beta} + GD.$$

Combining the above with (4.1) finishes the proof. □

4.3. Sketch of the privacy analysis of Algorithm 1

In this section we provide a sketch of the proof of privacy for Algorithm 1, i.e. Theorem 4.3. The full proof is provided in the appendix. Note that while our algorithm is quite similar to the one proposed by Asi et al. (2023), the privacy analysis is complicated by the fact that the switching probabilities depend on the entire sequence of loss functions and not just the latest one due to our use of the ratio of normalized densities to define the switching probabilities, unlike Asi et al. (2023). The analysis sketch presented below leverages techniques from Agarwal et al. (2023), to together with new ideas, to match this challenge.

For brevity of notation, we say two random variables X, Y supported on some set Ω are (ε, δ)-indistinguishable if for any outcome set $O \subseteq \Omega$, we have that

$$\Pr(X \in O) \leq e^\varepsilon \Pr(Y \in O) + \delta.$$

Consider any two t-indexed loss sequences $l_{1:T}, l'_{1:T} \in \mathcal{L}^T$ that differ at not more than one index $t_0 \in [T]$, i.e. it is the case that $l_t(x) = l'_t(x)$ holds for all $x \in \mathcal{K}$ and $t \in T - \{t_0\}$. For ease of argumentation we will show differential privacy for the outputs x_t of the algorithm along with the internal variables ζ_t which constitutes the decision to switch, defined for any t in the algorithm as

$$\zeta_t \triangleq \mathbb{I}\{S'_t = 0 \text{ or } S_t = 0\}.$$

We now provide the claim that is the core of the privacy proof. The claim analyses the privacy loss at three types of timesteps, viz. before the switch of the loss function happened, at the switch and after the switch. To establish definitions, let $\{(x_t, \zeta_t)\}_{t=1}^T$ and $\{(x'_t, \zeta'_t)\}_{t=1}^T$ be the instantiations of the random variables determined by Algorithm 1 upon execution on $l_{1:T}$ and $l'_{1:T}$, respectively. For brevity of notation, we will denote by Σ_t the random variable $\{x_\tau, \zeta_\tau\}_{\tau=1}^t$. We denote by $\mathbf{\Sigma_t}$ all possible values Σ_t can take. We make the following claim,

Claim 4.9. *Let $\delta' \geq 0$ and Φ be as defined in Theorem 4.3. Then for any $t \in [T]$ the random variable pairs (x_t, ζ_t) and (x'_t, ζ'_t) are $(\varepsilon_t, \delta_t)$-indistinguishable when conditioned on Σ_{t-1}, i.e. when conditioned on identical values of random choices made by the algorithm before (but not including) round t, where $\delta_t = 4\delta' + 9\delta' T + 3e^{-\tilde{p}T}$ and*

$$\varepsilon_t = \begin{cases} 0, & t < t_0 \\ \mathbb{I}_{\sum_{s=1}^{t-1}\zeta_s < B} \cdot 2\log(\Phi)/p, & t = t_0 \\ \mathbb{I}_{\sum_{s=1}^{t-1}\zeta_s < B}\left(\zeta_{t-1}\log(\Phi) + \frac{2G^2\beta/\lambda}{p}\right) & t > t_0 \end{cases} \tag{4.3}$$

Next, we sketch the proof of the claim and Theorem 4.3 follows from applying the strong composition Lemma 2.2 to the above claim. In particular note that the privacy loss for all time steps $t > t_0$ depends on whether the switch was made at the step or not via the ζ_{t-1}. Therefore the total privacy loss depends on the number of overall switches which is why the algorithm needs a bounded switching mechanism. The above claim divides the privacy loss into three cases depending upon the time step t.

Case 1: $t < t_0$ At any such time the privacy loss is naturally 0 as the algorithm has seen the same sequence thus far.

Case 2: $t = t_0$ Since x_{t_0} depends only on the past there is no privacy loss for x_{t_0}. Now consider the random variable ζ_{t_0} which depends on the random variables S_{t_0}, S'_{t_0}. If we use only S_t to decide on switching it can be seen that the privacy loss at this step can be infinite (essentially a single loss function change can cause a determinisitic non-switch decision to a deterministic switch). As a result following the idea introduced in Asi et al. (2023) we add a small probability p of forced switching at every step. Via a simple calculation it can be seen that the privacy loss is now bounded by the density ratio between two consecutive time steps is bounded by $\log(\Phi)/p$.

Case 3: $t > t_0$: As remarked before at any such step we see that conditioned on the history thus far being equal privacy loss occurs only if a switch is performed, i.e. $\zeta_{t-1} = 0$. If a switch is indeed performed the privacy loss through x_t is bounded by the ratio of the densities which is at most $\log(\Phi)$. We now consider the privacy loss through the variables ζ_t which depends on the log ratio of consecutive probabilities of success in the Bernoulli trials. Via an argument that utilizes convexity of a certain log-partition function, and the Wasserstein distance bound for Gibbs measures (Lemma 3.6), we can show that the log ratio of probabilities scales like $O(\frac{\beta}{\lambda p})$ (see Appendix C for details). Accounting for these two privacy losses if a switch happens gives the overall privacy loss in this case.

Overall putting these arguments together finishes the proof of the claim and an application of strong composition (Lemma 2.2) implies Theorem 4.3.

5. Comparisons to Agarwal et al. (2023)

Algorithm. The algorithm in this paper differs significantly from Agarwal et al. (2023) in our use of correlated sampling on top of continuous multiplicative weights, instead of Follow the Perturbed Leader (FTPL). The FTPL approach crucially uses smoothness in bounding the number of switches and it is unclear how this might extend to non-smooth settings. In this regard, our algorithm is similar to Asi et al. (2023) with the vital difference being the addition of a ℓ_2 regularization term in the sampling log-density.

Regret analysis. While *prima facie* adding a regularization term to the sampling density might appear to be a minor change it is a vital idea towards obtaining our results. A key observation is that upon adding the regularization term the stability term in the regret analysis becomes independent of the temperature β in exponential density. To gain intuition, consider the two extremes: if $\beta = 0$ we sample from uniform distribution and are stable by definition; if $\beta \to \infty$ our algorithm algorithm essentially reduces to Follow The Regularized Leader for which the stability comes from the regularization term.

Note that without the regularization term in the log-density, β cannot be very large without degrading stability, and overall that leads to a bias which adds a square root of dimension factor to regret. To the best of our knowledge even in the non-private case this simple modification to the sampling density leading to optimal regret is not known in the literature. The main lemma that proves the stability is Lemma 3.6 which is a Wasserstein bound on the successive densities. We believe this observation is useful more broadly.

Privacy analysis. The significant deviation in the privacy analysis from the one in Agarwal et al. (2023) is in the case $t > t_0$; t_0 is the point of change in the loss function sequence. Herein we need to explicitly account for the privacy loss incurred due to the switching decisions S_t which depends on the logarithm of the ratio of consecutive probabilities of success in the Bernoulli trials. Since the distribution used by us is entirely different from the FTPL-based one in Agarwal et al. (2023), we develop a new, substantially different, argument which applies to our case. This argument that utilizes the convexity of a certain log-partition function and the Wasserstein distance bound for Gibbs measures. In particular, we develop and prove Lemma C.4 and the analysis following this lemma in the appendix, which marks a complete departure from the analysis presented in Agarwal et al. (2023).

6. Conclusion

We studied the task of differentially private online convex optimization, and presented an algorithm Private Continuous Online Multiplicative Weights with Euclidean Regularization (POMER) that is (ε, δ)-differentially private and has a regret of at most $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{\sqrt{d}T^{1/3}}{\varepsilon}\right)$ for Lipschitz loss functions. This improves the known best bound for smooth loss functions by a factor of $\sqrt{d}$. Furthermore, for non-smooth loss functions, it gives the first regret bound with an optimal leading-order regret term. While the addition of strongly-convex terms in general does not yield improved regret bounds for the unregularized objective in OCO, our improvement leverages LSI properties of log-strongly-convex measures induced by additional regularization. To the best of our knowledge this is the best rate known for DP-OCO.

A central open question that remains is whether this rate can be further improved. In particular for linear functions a rate of $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{\sqrt{d}}{\varepsilon}\right)$ was shown by Agarwal & Singh (2017) and it remains an open question to show such a rate of convex functions. This rate would close the existing gap between the online and the one-pass stochastic setting. An intermediate goal would be to show a slightly improved bound of $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{\sqrt{d}T^{1/3}}{\varepsilon^{2/3}}\right)$ which has a more appropriate scaling of T, ε than our result. On the other hand showing any separation between the online and the stochastic setting in terms of regret is also open.

A second open question concerns lazy OCO in the *strongly-convex* setting. A straightforward application of the techniques in this paper unfortunately do not seem to yield improvements in this setting, and new ideas may be needed.

Impact statement

This paper presents work that advances the state of the art for differentially private learning in the OCO framework. By lowering the cost of privacy to data-efficient learning, we hope such a foundational advances leads to greater adoption of privacy preserving measures in digital interactions by platforms, and greater willingness for data sharing by participants to enable social goods.

References

Agarwal, N. and Singh, K. The price of differential privacy for online learning. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pp. 32–40. PMLR, 2017. URL `http://proceedings.mlr.press/v70/agarwal17a.html`.

Agarwal, N., Kale, S., Singh, K., and Thakurta, A. Differentially private and lazy online convex optimization. In Neu, G. and Rosasco, L. (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 4599–4632. PMLR, 12–15 Jul 2023. URL `https://proceedings.mlr.press/v195/agarwal23d.html`.

Altschuler, J. M. and Talwar, K. Online learning over a finite action set with limited switching. *Math. Oper. Res.*, 46(1):179–203, 2021. doi: 10.1287/moor.2020.1052. URL https://doi.org/10.1287/moor.2020.1052.

Anava, O., Hazan, E., and Mannor, S. Online learning for adversaries with memory: price of past mistakes. In *Advances in Neural Information Processing Systems*, pp. 784–792, 2015.

Asi, H., Feldman, V., Koren, T., and Talwar, K. Private online prediction from experts: Separations and faster rates. In *The Thirty Sixth Annual Conference on Learning Theory*, pp. 674–699. PMLR, 2023.

Bakry, D. and Émery, M. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84: Proceedings*, pp. 177–206. Springer, 2006.

Chen, L., Yu, Q., Lawrence, H., and Karbasi, A. Minimax regret of switching-constrained online convex optimization: No phase transition. *Advances in Neural Information Processing Systems*, 33:3477–3486, 2020.

Donsker, M. D. and Varadhan, S. S. Asymptotic evaluation of certain markov process expectations for large time, i. *Communications on Pure and Applied Mathematics*, 28 (1):1–47, 1975.

Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Ganesh, A., Thakurta, A., and Upadhyay, J. Universality of langevin diffusion for private optimization, with applications to sampling from rashomon sets. In Neu, G. and Rosasco, L. (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 1730–1773. PMLR, 12–15 Jul 2023. URL https://proceedings.mlr.press/v195/ganesh23a.html.

Geulen, S., Vöcking, B., and Winkler, M. Regret minimization for online buffering problems using the weighted majority algorithm. In Kalai, A. T. and Mohri, M. (eds.), *COLT*, pp. 132–143. Omnipress, 2010. URL http://colt2010.haifa.il.ibm.com/papers/COLT2010proceedings.pdf#page=140.

Gopi, S., Lee, Y. T., and Liu, D. Private convex optimization via exponential mechanism. In Loh, P.-L. and Raginsky, M. (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 1948–1989. PMLR, 02–05 Jul 2022. URL https://proceedings.mlr.press/v178/gopi22a.html.

Hazan, E. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.

Jain, P., Kothari, P., and Thakurta, A. Differentially private online learning. In *Proc. of the 25th Annual Conf. on Learning Theory (COLT)*, volume 23, pp. 24.1–24.34, June 2012.

Kairouz, P., McMahan, B., Song, S., Thakkar, O., Thakurta, A., and Xu, Z. Practical and private (deep) learning without sampling or shuffling. In *ICML*, 2021.

Kalai, A. and Vempala, S. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.

Ledoux, M. Concentration of measure and logarithmic sobolev inequalities. In *Seminaire de probabilites XXXIII*, pp. 120–216. Springer, 1999.

Levin, D. A. and Peres, Y. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

Merhav, N., Ordentlich, E., Seroussi, G., and Weinberger, M. J. On sequential strategies for loss functions with memory. *IEEE Trans. Inf. Theory*, 48(7):1947–1958, 2002. doi: 10.1109/TIT.2002.1013135. URL https://doi.org/10.1109/TIT.2002.1013135.

Sherman, U. and Koren, T. Lazy oco: Online convex optimization on a switching budget. In *Conference on Learning Theory*, pp. 3972–3988. PMLR, 2021.

Sherman, U. and Koren, T. Lazy oco: Online convex optimization on a switching budget. *arXiv preprint arXiv:2102.03803 version 7 (see also version 5)*, 2023. URL https://arxiv.org/abs/2102.03803.

Smith, A. and Thakurta, A. (nearly) optimal algorithms for private online learning in full-information and bandit settings. In *Advances in Neural Information Processing Systems*, pp. 2733–2741, 2013.

Whitehouse, J., Ramdas, A., Rogers, R., and Wu, Z. S. Fully adaptive composition in differential privacy. *arXiv preprint arXiv:2203.05481*, 2022.

A. Proofs of smoothness of Gibbs measures

In this section we prove the Lemmas concerning the smoothness of Gibbs measures, i.e. Lemmas 3.6 and 3.5. We begin by restating and proving Lemma 3.6.

Lemma A.1 (Wasserstein Distance)**.** *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be convex functions such that $l - l'$ is G-Lipschitz. Further let $\beta, \lambda \geq 0$ be parameters and define the Gibbs-distributions $\bar{\mu} = \bar{\mu}(l, \beta, \lambda)$ and $\bar{\mu}' = \mu(l', \beta, \lambda)$ (as defined in (3.3)). Then we have that ∞-Wasserstein distance between $\bar{\mu}$ and $\bar{\mu}'$ over the ℓ_2 metric is bounded as*

$$W_\infty(\bar{\mu}, \bar{\mu}') \leq \frac{G}{\lambda}.$$

Proof of Lemma 3.6. By definition $W_\infty(\bar{\mu}, \bar{\mu}') = \inf_{\gamma \in \Gamma(\bar{\mu}, \bar{\mu}')} \sup_{(X, X') \sim \gamma} \|X - X'\|$, where the notation $\sup_{(X,X')\sim\gamma}$ is shorthand for all (X, X') in the support of γ. To bound W_∞ we consider the following coupling between $\bar{\mu}, \bar{\mu}'$. Define the functions $L(x) = l(x) + \frac{\lambda}{2}\|x\|^2$, $L'(x) = l'(x) + \frac{\lambda}{2}\|x\|^2$ and consider the following "Projected" Langevin diffusions given by the following SDEs (see (Ganesh et al., 2023) for details):

$$dX_{t+1} = -\beta \nabla L(X_t) + \sqrt{2} dW_t - \nu_t \zeta(d_t)$$

$$dX'_{t+1} = -\beta \nabla L'(X_t) + \sqrt{2} dW_t - \nu'_t \zeta'(d_t)$$

where ζ and ζ' are measures supported on $\{t : X_t \in \partial K\}$ and $\{t : X'_t \in \partial K\}$ respectively, and ν_t and ν'_t are outer unit normal vectors at X_t and X'_t respectively. It is known that $\lim_{t\to\infty} X_t$ converges in distribution to $\bar{\mu}$ and similarly $\lim_{t\to\infty} X'_t$ converges in distribution to $\bar{\mu}'$. Our desired coupling γ is defined by sampling a Brownian motion sequence $\{W_t\}_{t=1}^\infty$ and the output sample is set to $\lim_{t\to\infty} X_t$ and $\lim_{t\to\infty} X'_t$ with the same $\{W_t\}_{t=1}^\infty$ sequence. For a fixed Brownian motion sequence $\{W_t\}_{t=1}^\infty$, we get the following calculations (by defining $\Delta_t = \|X_t - X'_t\|$):

$$\begin{aligned}
\frac{1}{2}\frac{d\Delta_t^2}{dt} = \frac{1}{2}\frac{d\|X_t - X'_t\|^2}{dt} &= \left\langle \frac{dX_t}{dt} - \frac{dX'_t}{dt}, X_t - X'_t \right\rangle \\
&= -\beta\langle \nabla l(X_t) - \nabla l'(X'_t), X_t - X'_t\rangle - \langle \nu_t, X_t - X'_t\rangle \frac{\zeta(d_t)}{d_t} + \langle \nu'_t, X_t - X'_t\rangle \frac{\zeta'(d_t)}{d_t} \\
&\leq -\beta\langle \nabla l(X_t) - \nabla l'(X'_t), X_t - X'_t\rangle \\
&(\because \langle \nu_t, X'_t - X_t\rangle \leq 0 \text{ and } \langle \nu'_t, X_t - X'_t\rangle \leq 0 \text{ since } K \text{ is convex}) \\
&= -\beta\langle \nabla l(X_t) - \nabla l(X'_t), X_t - X'_t\rangle + \beta\langle \nabla l'(X'_t) - \nabla l(X'_t), X_t - X'_t\rangle \\
&\leq \beta\left(-\lambda\|X_t - X'_t\|^2 + G\|X_t - X'_t\|\right) = \beta\left(-\lambda\Delta_t^2 + G\Delta_t\right) \\
&\leq \beta\left(-\lambda\Delta_t^2 + \frac{\lambda}{2}\Delta_t^2 + \frac{G^2}{2\lambda}\right) \\
&\leq \beta\left(-\frac{\lambda}{2}\Delta_t^2 + \frac{G^2}{2\lambda}\right)
\end{aligned}$$

Defining $F_t = \Delta_t^2 - \frac{G^2}{\lambda^2}$, the above implies that $\frac{dF_t}{dt} \leq -\beta\lambda F_t$ which implies, via Grönwall's inequality, that $F_t \leq F_0 \exp(-\beta\lambda t)$. Therefore we have that $\lim_{t\to\infty} F_t \to 0$ which implies that $\lim_{t\to\infty} \Delta_t \to \frac{G}{\lambda}$.

Therefore we get that under the above coupling γ we have that $\sup_{(x,y)\sim\gamma} \|x - y\| \leq \frac{G}{\lambda}$ which finishes the proof. □

Using the above we restate and prove Lemma 3.5 below.

Lemma A.2 (Density ratio)**.** *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be convex functions such that $l - l'$ is G-Lipschitz. Further let $\beta, \lambda \geq 0$ be parameters and define the Gibbs-distributions $\bar{\mu} = \bar{\mu}(l, \beta, \lambda)$ and $\bar{\mu}' = \bar{\mu}(l', \beta, \lambda)$ (as defined in (3.1)). Then for any $\delta \in (0, 1]$, we have that $\bar{\mu}$ and $\bar{\mu}'$ are (Φ, δ) close where*

$$\Phi = \exp\left(\frac{2\beta G^2}{\lambda} + \sqrt{\frac{8\beta G^2 \log(2/\delta)}{\lambda}}\right)$$

Proof of Lemma 3.5. We begin first by proving the direction

$$\Pr_{X\sim\bar\mu}\left[\frac{1}{\Phi} \le \frac{\bar\mu(X)}{\bar\mu'(X)} \le \Phi\right] \ge 1-\delta$$

and reverse direction follows easily by switching the roles of $\bar\mu, \bar\mu'$ through the analysis. To this end define the function $g(X) = \log\left(\frac{\bar\mu(X)}{\bar\mu'(X)}\right)$. Therefore we are required to show that

$$\Pr_{X\sim\bar\mu}(|g(X)| > \log(\Phi)) \le \delta.$$

We will show this by first bounding $\mathbb{E}_{X\sim\bar\mu}[g(X)]$ and then showing that it concentrates around its expectation. We first show that g is a $2\beta G$-Lipschitz function. To see this consider the following

$$|g(X)-g(X')| = \Big|\log\left(\frac{\bar\mu(X)}{\bar\mu(X')}\right) + \log\left(\frac{\bar\mu'(X')}{\bar\mu'(X)}\right)\Big| = |-\beta\left(l(X)-l(X')+l'(X')-l'(X)\right)| \le 2\beta G\|X-X'\|.$$

Using the proof of Lemma 3.6 we get that there is a coupling γ between $\bar\mu, \bar\mu'$ such that $\sup_{(X,X')\sim\gamma}\|X-X'\| \le \frac{G}{\Lambda}$, therefore sampling from the coupling and using the Lipschitzness of g, we get that

$$\mathbb{E}_{(X,X')\sim\bar\mu}[|g(X)-g(X')|] \le 2\beta G\cdot\mathbb{E}_{(X,X')\sim\bar\mu}[\|X-X'\|] \le 2\beta G\cdot\frac{G}{\Lambda},$$

which implies that

$$\mathbb{E}_{X\sim\bar\mu}[g(X)] \le \mathbb{E}_{X'\sim\bar\mu'}[g(X')] + 2\beta G\cdot\frac{G}{\Lambda}$$

Now noticing that $\mathbb{E}_{X\sim\bar\mu'}[g(X)] = -\mathrm{KL}(\bar\mu'\|\bar\mu) \le 0$, we get that

$$\mathbb{E}_{X\sim\bar\mu}[g(X)] \le \frac{2\beta G^2}{\Lambda}.$$

Furthermore, note that $\mathbb{E}_{X\sim\bar\mu}[g(X)] = \mathrm{KL}(\bar\mu\|\bar\mu') \ge 0$. Thus, we have

$$0 \le \mathbb{E}_{X\sim\bar\mu}[g(X)] \le \frac{2\beta G^2}{\Lambda}.$$

Next we give a high probability bound on g. Lemma 3.2 implies that the distribution corresponding to $\bar\mu$ satisfies LSI (Definition 3.1) with constant $\beta\Lambda$. Now by Lemma 3.3, plugging in the LSI constant and Lipschitzness bound for g, we have that

$$\Pr_{X\sim\bar\mu}[|g(X)-\mathbb{E}[g(X)]| \ge r] \le 2\exp\left(-\frac{\Lambda r^2}{8\beta G^2}\right)$$

Thereby setting $r = \sqrt{\frac{8\beta G^2\log(2/\delta)}{\Lambda}}$ we have that

$$\Pr_{X\sim\bar\mu}\left(|g(X)| > \frac{2\beta G^2}{\Lambda} + \sqrt{\frac{8\beta G^2\log(2/\delta)}{\Lambda}}\right) \le \delta.$$

□

B. Analysis of Algorithm 1

As a reminder for the notation, $\Pi:\mathbb{R}\to[\frac{1}{\Phi^2},1]$ as $\Pi(x) = \min\{1, \max\{\frac{1}{\Phi^2}, x)\}\}$. Also, $\zeta_t \triangleq \mathbb{I}(S_t = 0 \text{ or } S'_t = 0)$. We restate and prove Lemma 4.2 first:

Lemma B.1 (Switching bound)**.** *For any $p\in[0,1]$ and any $\Phi \ge 0$, setting $\tilde p = p + 1 - \Phi^{-2}$, we have that the number of switches is bounded in the following manner,*

$$\mathbb{E}\left[\mathcal{S}_T\right] \le \tilde pT, \qquad \Pr\left[\mathcal{S}_T \ge 3\tilde pT\right] \le e^{-\tilde pT}.$$

Proof of Lemma 4.2. Since $S_t \sim \mathrm{Ber}\left(\Pi\left(\frac{\mu_{t+1}(x_t)}{\Phi\mu_t(x_t)}\right)\right)$, we have $\Pr[S_t = 0] \le 1 - \Phi^{-2}$. From the definition of ζ_t, we have

$$\mathbb{E}[\zeta_t] = \Pr(S'_t = 0) + (1 - \Pr(S'_t = 0)) \cdot \Pr(S_t = 0) \le p + (1-p) \cdot (1 - \Phi^{-2}) \le \tilde{p}. \tag{B.1}$$

Thus, the random variable $S_T = \sum_{t=1}^{T} \zeta_t$ is stochastically dominated by the sum of T Bernoulli random variables with parameter $\tilde{p}$. Hence, $\mathbb{E}[\mathcal{S}_T] \le \tilde{p}T$ and the Chernoff bound[5] implies

$$\Pr\left[\mathcal{S}_T \ge 3\tilde{p}T\right] \le e^{-\tilde{p}T}.$$

□

Next we restate and prove Lemma 4.6.

Lemma B.2 (Distribution drift)**.** *Given $\delta \in [0, \frac{1}{2}]$ and $\Phi \ge 1$, suppose that for all $t \in [T]$, the Gibbs-measures μ_t, μ_{t+1} are (Φ, δ)-close. If q_t is the marginal distribution induced by Algorithm 1 on its iterates x_t, then we have that*

- *If $B = \infty$, then for all t, $\|q_t - \bar{\mu}_t\|_{\mathrm{TV}} \le 3\delta(t-1)$.*
- *If $B = 3\tilde{p}T$, then we have*
$$\|q_t - \bar{\mu}_t\|_{TV} \le e^{-\tilde{p}T} + 3\delta(t-1).$$

Proof of Lemma 4.6. We first consider the $B = \infty$ case. We prove that $\|q_t - \bar{\mu}_t\|_{\mathrm{TV}} \le 3\delta(t-1)$ by induction on t. For $t = 1$, the claim is trivially true. So assume it is true for some t and now we prove it for $t+1$. Let $M = \{x \in \mathcal{K} \mid \Phi^{-1} \le \frac{\bar{\mu}_{t+1}(x)}{\bar{\mu}_t(x)} \le \Phi\}$. Then by Definition 3.4, we have $\bar{\mu}_t(M) \ge 1 - \delta$ and $\bar{\mu}_{t+1}(M) \ge 1 - \delta$. Next, let $\tilde{\mu}_t$ be the distribution of $X \sim \bar{\mu}_t$ conditioned on the event $X \in M$. Since $\bar{\mu}_t(M) \ge 1 - \delta$, it is easy to see that $\|\bar{\mu}_t - \tilde{\mu}_t\|_{\mathrm{TV}} \le \delta$. Let $\tilde{q}_{t+1}$ be the distribution of x_{t+1} if x_t were sampled from $\tilde{\mu}_t$ instead of q_t. Let E be any measurable subset of $\mathcal{K}$. Using the facts that for any $x \in M$, we have $\Pi(\frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}) = \frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}$, and that $\tilde{\mu}_t(x) = \frac{\bar{\mu}_t(x)}{\bar{\mu}_t(M)}$, we have

$$\begin{aligned}
\tilde{q}_{t+1}(E) = \int_{x \in E} \Big(&\Pr(S'_t = 0 | x_t = x) \Pr(x_{t+1} \in E | x_t = x, S'_t = 0) \\
&+ \Pr((S'_t = 1 \wedge S_t = 0) | x_t = x) \Pr(x_{t+1} \in E | x_t = x, (S'_t = 1 \wedge S_t = 0)) \\
&+ \Pr((S'_t = 1 \wedge S_t = 1) | x_t = x) \Pr(x_{t+1} \in E | x_t = x, (S'_t = 1 \wedge S_t = 1)) \Big) \tilde{\mu}_t(x) dx \\
= p\bar{\mu}_{t+1}(E) &+ (1-p)\bar{\mu}_{t+1}(E) \int_M \left(1 - \frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}\right) \left(\frac{\bar{\mu}_t(x)}{\bar{\mu}_t(M)}\right) dx \\
+ (1-p) &\int_{E \cap M} \left(\frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}\right) \left(\frac{\bar{\mu}_t(x)}{\bar{\mu}_t(M)}\right) dx \\
= p\bar{\mu}_{t+1}(E) &+ (1-p)\bar{\mu}_{t+1}(E) \left(1 - \frac{\bar{\mu}_{t+1}(M)}{\Phi\bar{\mu}_t(M)}\right) + (1-p) \frac{\bar{\mu}_{t+1}(E \cap M)}{\Phi\bar{\mu}_t(M)}.
\end{aligned}$$

Thus,

$$\begin{aligned}
|\tilde{q}_{t+1}(E) - \bar{\mu}_{t+1}(E)| &= \frac{1-p}{\Phi\bar{\mu}_t(M)} |\bar{\mu}_{t+1}(E)\bar{\mu}_{t+1}(M) - \bar{\mu}_{t+1}(E \cap M)| \\
&= \frac{1-p}{\Phi\bar{\mu}_t(M)} |\bar{\mu}_{t+1}(E \setminus M) - \bar{\mu}_{t+1}(E \cap M)\bar{\mu}_{t+1}(M^c)| \\
&\le \frac{\delta}{1-\delta},
\end{aligned}$$

since $\bar{\mu}_t(M) \ge 1 - \delta$ and $\bar{\mu}_{t+1}(M) \ge 1 - \delta$. Since $\delta \le \frac{1}{2}$, we conclude that

$$\|\tilde{q}_{t+1} - \bar{\mu}_{t+1}\|_{\mathrm{TV}} \le 2\delta.$$

[5]The specific bound used is that for independent Bernoulli random variables $X_1, X_2, \ldots, X_T$, if $\mu = \mathbb{E}[\sum_{t=1}^{T} X_t]$, then for any $\delta > 0$, we have $\Pr[\sum_{t=1}^{T} X_t \ge (1+\delta)\mu] \le e^{-\delta^2\mu/(2+\delta)}$.

Furthermore, we have

$$\|q_{t+1} - \tilde{q}_{t+1}\|_{\mathrm{TV}} \le \|q_t - \tilde{\mu}_t\|_{\mathrm{TV}} \le \|q_t - \bar{\mu}_t\|_{\mathrm{TV}} + \|\bar{\mu}_t - \tilde{\mu}_t\|_{\mathrm{TV}} \le 3\delta(t-1) + \delta,$$

where the first inequality follows by the data-processing inequality for f-divergences like TV-distance (note that q_{t+1} and $\tilde{q}_{t+1}$ are obtained from q_t and $\tilde{\mu}_t$ respectively via the same data-processing channel), and the second inequality is due to the induction hypothesis. Thus, we conclude that

$$\|q_{t+1} - \bar{\mu}_{t+1}\|_{\mathrm{TV}} \le \|q_{t+1} - \tilde{q}_{t+1}\|_{\mathrm{TV}} + \|\tilde{q}_{t+1} - \bar{\mu}_{t+1}\|_{\mathrm{TV}} \le 3\delta(t-1) + \delta + 2\delta = 3\delta t,$$

completing the induction.

We now turn to the $B = 3\tilde{p}T$ case. Let q'_t be the distribution of x_t if $B = \infty$. We now relate q'_t and q_t. We start by defining q_{all} as the probability distributions over all possible random variables, i.e. $S_{1:T}, S'_{1:T}, Z_{1:T}, x_{1:T}$, sampled by Algorithm 1. Similarly, let q'_{all} be the analogue for the infinite switching budget variant. Let $\mathcal{E}$ be the event that $\sum_{t=1}^{T} \zeta_t \ge 3\tilde{p}T$. Note that Lemma 4.2 implies that both $q_{\text{all}}(\mathcal{E}), q'_{\text{all}}(\mathcal{E}) \le e^{-\tilde{p}T}$. Therefore we have that,

$$\begin{aligned}
\|q_{\text{all}} - q'_{\text{all}}\|_{\mathrm{TV}} &= \sup_{\text{measurable } A} \left(q_{\text{all}}(A) - q'_{\text{all}}(A)\right) \\
&= \sup_{\text{measurable } A} \left(q_{\text{all}}(A \cap \mathcal{E}) - q'_{\text{all}}(A \cap \mathcal{E}) + \underbrace{q_{\text{all}}(A \cap \neg\mathcal{E}) - q'_{\text{all}}(A \cap \neg\mathcal{E})}_{=0}\right) \\
&= \sup_{\text{measurable } A} \left(q_{\text{all}}(A \cap \mathcal{E}) - q'_{\text{all}}(A \cap \mathcal{E})\right) \\
&\le e^{-\tilde{p}T}
\end{aligned}$$

Now, for any t, since q_t, q'_t are marginals of $q_{\text{all}}, q'_{\text{all}}$ respectively, we have

$$\|q_t - q'_t\|_{\mathrm{TV}} \le \|q_{\text{all}} - q'_{\text{all}}\|_{\mathrm{TV}} \le e^{-\tilde{p}T}.$$

Since we have $\|\mu_t - q'_t\|_{\mathrm{TV}} \le 3\delta(t-1)$ by the $B = \infty$ analysis, the proof is complete by the triangle inequality. □

We finish this section by repeating and proving Lemma 4.8.

Lemma B.3. *Consider an arbitrary distribution μ_0 on $\mathcal{K}$ (referred to as the prior) and f be an arbitrary bounded function on $\mathcal{K}$. Define the distribution μ over $\mathcal{K}$ with density $\mu(x) \propto \mu_0(x)e^{-(x)}$. Then we have that*

$$\mu = \arg\min_{\mu' \in \Delta(\mathcal{K})} \left(\mathbb{E}_{x\sim\mu'}[f(x)] + \mathrm{KL}(\mu'\|\mu_0)\right).$$

Proof of Lemma 4.8. The proof follows from the following lemma appearing as in (Donsker & Varadhan, 1975)

Lemma B.4 (Lemma 2.1 in (Donsker & Varadhan, 1975), rephrased)**.** *Let $\mathcal{U}$ be the set of continuous functions on $\mathcal{K}$ satisfying $u(x) \in [c_1, c_2]$ for all $u \in \mathcal{U}, x \in \mathcal{K}$, for some constants $c_1, c_2 > 0$. Let ν_1 and ν_2 be any distributions on $\mathcal{K}$, then we have that*

$$\mathrm{KL}(\nu_1\|\nu_2) = \sup_{u\in\mathcal{U}} \left(\mathbb{E}_{x\sim\nu_1}[\log(u(x))] - \log(\mathbb{E}_{x\sim\nu_2}[u(x)])\right)$$

Using the above lemma, setting $\nu_1 = \mu$, $\nu_2 = \mu_0$, $u(x) = e^{-f(x)}$, we get that

$$-\log(\mathbb{E}_{x\sim\mu_0}[e^{-f(x)}]) \le \mathbb{E}_{x\sim\mu}[f(x)] + \mathrm{KL}(\mu\|\mu_0).$$

Let $Z = \int_K e^{-f(x)}\mu_0(x)dx$, then we have that

$$\mathbb{E}_{x\sim\mu}[f(x)] + \mathrm{KL}(\mu\|\mu_0) = \mathbb{E}_{x\sim\mu}[f(x)] + \int_K \mu(x)\log(e^{-f(x)}/Z)dx = -\log(Z) = -\log(\mathbb{E}_{x\sim\mu_0}[e^{-f(x)}]).$$

Combining the above two displays finishes the proof. □

C. Privacy Analysis

For brevity of notation, we say two random variables X, Y supported on some set Ω are (ε, δ)-indistinguishable if for any outcome set $O \subseteq \Omega$, we have that

$$\Pr(X \in O) \leq e^{\varepsilon} \Pr(Y \in O) + \delta.$$

We restate and prove Theorem 4.3:

Theorem 4.3 (Privacy)**.** *Given $\beta, \lambda > 0$ and $\delta \in (0, 1/2]$, for any $T \geq 12\log(1/\delta)$, let $\delta' = \frac{\delta T^{-2}}{60}$, $G' = 3G$. Suppose there exists $\Phi' > 0$ such that for all convex functions l, l' where $l - l'$ is G'-Lipschitz, we have that, the distributions $\bar{\mu}(l, \beta, \lambda)$ and $\bar{\mu}(l', \beta, \lambda)$ respectively are (Φ', δ')-close. Then for any sequence of G-Lipschitz convex functions, Algorithm 1 when run with $\Phi = {\Phi'}^2$, $p = \max\left(T^{-1/3}, \left(\frac{G^4\beta^2}{\lambda^2 \cdot \log^2(\Phi)}\right)^{1/3}\right)$, $\tilde{p} = p + 1 - \Phi^{-2}$ and $B = 3\tilde{p}T$ is $(\varepsilon, \delta + 3Te^{-(1-\Phi^{-2})T})$-differentially private where*

$$\varepsilon = 3\varepsilon'/2 + \sqrt{6\varepsilon'}\sqrt{\log(2/\delta)},$$

with

$$\begin{aligned}\varepsilon' = {} & 7T^{2/3}\log^2(\Phi) + 12\log^3(\Phi)T \\ & + 11\left(\frac{G^4\beta^2}{\lambda^2}\right)^{1/3}\log^{4/3}(\Phi)T.\end{aligned}$$

Proof. Consider any two loss sequences $l_{1:T}, l'_{1:T} \in \mathcal{L}^T$ that possibly differ at some index $t_0 \in [T]$, i.e. $l_t(x) = l'_t(x)$ holds for all $x \in \mathcal{K}$ and $t \in T - \{t_0\}$. For ease of argumentation we will show differential privacy for the outputs x_t of the algorithm along with the internal variables ζ_t which are defined for any t in the algorithm as

$$\zeta_t \triangleq \mathbb{I}\{S'_t = 0 \text{ or } S_t = 0\}.$$

To establish privacy, let $\{(x_t, \zeta_t)\}_{t=1}^T$ and $\{(x'_t, \zeta'_t)\}_{t=1}^T$ be the instantiations of the random variables determined by Algorithm 1 upon execution on $l_{1:T}$ and $l'_{1:T}$, respectively. For brevity of notation, we will denote by Σ_t the random variable $\{x_\tau, \zeta_\tau\}_{\tau=1}^t$. We denote by $\mathbf{\Sigma_t}$ all possible values Σ_t can take. We now prove Claim 4.9, which we restate here for convenience:

Claim C.2. *Let $\delta' \geq 0$ and Φ be as defined in Theorem 4.3. Then for any $t \in [T]$ the random variable pairs (x_t, ζ_t) and (x'_t, ζ'_t) are $(\varepsilon_t, \delta_t)$-indistinguishable when conditioned on Σ_{t-1}, i.e. when conditioned on identical values of random choices made by the algorithm before (but not including) round t, where $\delta_t = 4\delta' + 9\delta' T + 3e^{-\tilde{p}T}$ and*

$$\varepsilon_t = \begin{cases} 0, & t < t_0 \\ \mathbb{I}_{\sum_{s=1}^{t-1}\zeta_s < B} \cdot 2\log(\Phi)/p, & t = t_0 \\ \mathbb{I}_{\sum_{s=1}^{t-1}\zeta_s < B}\left(\zeta_{t-1}\log(\Phi) + \frac{2G^2\beta/\lambda}{p}\right) & t > t_0 \end{cases} \tag{C.1}$$

The proof of the above claim appears after the present proof.

We intend to use adaptive strong composition for differential privacy (Lemma 2.2) with Claim 4.9 and to that end consider the following calculations

$$\sum_{t=1}^{T} \varepsilon_t^2 \leq \frac{4\log^2(\Phi)}{p^2} + 2B\log^2(\Phi) + \frac{8G^4\beta^2/\lambda^2}{p^2}T$$
$$\leq \frac{4\log^2(\Phi)}{p^2} + 6pT\log^2(\Phi) + 12\log^3(\Phi)T + \frac{8G^4\beta^2/\lambda^2}{p^2}T$$
$$\text{(Using } B = 3pT + 3(1-\Phi^{-2})T \leq 3pT + 6\log(\Phi)T\text{)}$$
$$= \frac{4\log^2(\Phi)}{p^2} + 3pT\log^2(\Phi) + 12\log^3(\Phi)T + 3pT\log^2(\Phi) + \frac{8G^4\beta^2/\lambda^2}{p^2}T$$
$$\leq 7T^{2/3}\log^2(\Phi) + 12\log^3(\Phi)T + 11\left(\frac{G^4\beta^2}{\lambda^2}\right)^{1/3}\log^{4/3}(\phi)\cdot T$$

and

$$\sum_{t=1}^{T}\delta_t = 4\delta' T + 9T^2\delta' + 3Te^{-\tilde{p}T} \leq \frac{\delta}{6} + 3Te^{-pT} + 3Te^{-(1-\Phi^{-2})T} \leq \frac{\delta}{3} + 3Te^{-(1-\Phi^{-2})T}.$$

Using the above calculations and applying Lemma 2.2 with $\delta' = \delta/2$ (in Lemma 2.2) concludes the proof. □

Proof Of Claim 4.9. We begin by defining a subset $\mathcal{E}_t \in \mathcal{K}$ for all t as

$$\mathcal{E}_t = \left\{ x \in \mathcal{K} \middle| \left(\frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)} \in \left[\frac{1}{\Phi^2}, 1\right]\right) \wedge \left(\frac{\bar{\mu}'_{t+1}(x)}{\Phi\bar{\mu}'_t(x)} \in \left[\frac{1}{\Phi^2}, 1\right]\right)\right\}.$$

The following claim whose proof is presented after the present proof shows that $\mathcal{E}_t$ occurs with high probability conditioned on Σ_{t-1} taking any value Σ in its domain.

Claim C.3. *Let Φ be as defined in Theorem 4.3, then we have that for all $\Sigma \in \mathbf{\Sigma}_t$,*

$$\Pr(x_t \in \mathcal{E}_t | \Sigma_{t-1} = \Sigma) \geq 1 - 3\delta' - 9T\delta' - 3e^{-\tilde{p}T}.$$

The general recipe we will follow in the proof is to show that x_t, x'_t are $(\varepsilon_x, \delta_x)$-indistinguishable conditioned on Σ_{t-1} and the event that $x_t \in \mathcal{E}_t$, for some $(\varepsilon_x, \delta_x)$. We will then show that ζ_t, ζ'_t are $(\varepsilon_\zeta, \delta_\zeta)-$indistinguishable after conditioning on $\Sigma_{t-1}, x_t = x$ (and $x'_t = x$ respectively) for an arbitrary $\mathcal{E}_t$. Then, by standard composition of differential privacy (Dwork & Roth, 2014), it is implied that (x_t, ζ_t), (x'_t, ζ'_t) are $(\varepsilon_x + \varepsilon_\zeta, \delta_x + \delta_\zeta)$ indistinguishable when conditioned on Σ_{t-1} and the event that $x_t \in \mathcal{E}_t$. It then follows that the same pair is $(\varepsilon_x + \varepsilon_\zeta, \delta_x + \delta_\zeta + \Pr(x_t \notin \mathcal{E}_t|\Sigma_{t-1}))$ indistinguishable when conditioned on Σ_{t-1}.

To execute the above strategy, we will examine the three cases – *ante* $t < t_0$, *at* $t = t_0$, and *post* $t > t_0$ – separately. Recall that $l_{1:T}$ and $l'_{1:T}$ are loss function sequences that differ only at the index t_0.

Case 1: $t \leq t_0$: Observe that since $l_{1:t_0-1} = l'_{1:t_0-1}$, having not yet encountered a change (at $t = t_0$) in loss, the algorithm produces identically distributed outputs for the first t_0 rounds upon being fed either loss sequence. Therefore we have that

$$\forall t < t_0, \quad (x_t, \zeta_t) \text{ and } (x'_t, \zeta'_t) \text{ are } (0,0) - \text{indistinguishable} \tag{C.2}$$

For the remaining two cases, we first assume that number of switches so far have not exceeded B, i.e. $\sum_{s=1}^{t-1}\zeta_s = \sum_{s=1}^{t-1}\zeta_s < B$ (conditioned on the same history). If not then both algorithms become deterministic from this point onwards and are $(0,0)$-indistinguishable.

Case 2: $t = t_0$**:** The last display in the previous case also means that x_{t_0} and x'_{t_0} are identically distributed random variables. Therefore, to conclude the claim for t_0, we need to demonstrate that ζ_{t_0} and ζ'_{t_0} are indistinguishable when also additionally conditioned on $x_{t_0} = x'_{t_0}$. We now observe that for any $x \in \mathcal{E}_{t_0}$ and any $\Sigma \in \mathbf{\Sigma_{t_0-1}}$,

$$\begin{aligned}
\frac{\Pr(\zeta'_{t_0} = 1|\Sigma_{t_0-1} = \Sigma, x'_{t_0} = x)}{\Pr(\zeta_{t_0} = 1|\Sigma_{t_0-1} = \Sigma, x_{t_0} = x)} &= \frac{p + (1-p)\left(1 - \frac{\bar{\mu}'_{t_0+1}(x)}{\Phi\bar{\mu}'_{t_0}(x)}\right)}{p + (1-p)\underbrace{\left(1 - \frac{\bar{\mu}_{t_0+1}(x)}{\Phi\bar{\mu}_{t_0}(x)}\right)}_{\geq 0}} \\
&\leq \frac{p + (1-p)\left(\underbrace{1 - \frac{\bar{\mu}'_{t_0+1}(x)}{\Phi\bar{\mu}'_{t_0}(x)}}_{\geq 0}\right)}{p} \\
&\leq 1 + \frac{1}{p}\left(1 - \frac{\bar{\mu}'_{t_0+1}(x)}{\Phi\bar{\mu}'_{t_0}(x)}\right) \leq 1 + \frac{1}{p}\left(1 - \Phi^{-2}\right) \\
&\leq 1 + \frac{1}{p}(1 - e^{-2\log\Phi}) \leq 1 + \frac{2\log(\Phi)}{p} \leq e^{2\log\Phi/p},
\end{aligned}$$

using the definition of the set $\mathcal{E}_{t_0}$ and that for any real x $1 + x \leq e^x$. Similarly, we have for any $x \in \mathcal{E}_{t_0}$,

$$\frac{\Pr(\zeta'_{t_0} = 0|\Sigma_{t_0-1} = \Sigma, x'_{t_0} = x)}{\Pr(\zeta_{t_0} = 0|\Sigma_{t_0-1} = \Sigma, x_{t_0} = x)} = \frac{(1-p)\frac{\bar{\mu}'_{t_0+1}(x)}{\Phi\bar{\mu}'_{t_0}(x)}}{(1-p)\frac{\bar{\mu}_{t_0+1}(x)}{\Phi\bar{\mu}_{t_0}(x)}} = \frac{\bar{\mu}'_{t_0+1}(x)}{\bar{\mu}'_{t_0}(x)}\frac{\bar{\mu}_{t_0}(x)}{\bar{\mu}_{t_0+1}(x)} \leq e^{2\log\Phi}.$$

The above displays thereby imply that conditioned on Σ_{t_0-1} and the event $x_t \in \mathcal{E}_{t_0}$, we have that (x_{t_0}, ζ_{t_0}) and (x'_{t_0}, ζ'_{t_0}) are $(2\log(\Phi)/p, 0)$-indistinguishable. Thereby combining with Claim C.3 we get that conditioned on Σ_{t-1}

$$(x_{t_0}, \zeta_{t_0}) \text{ and } (x'_{t_0}, \zeta'_{t_0}) \text{ are } (2\log(\Phi)/p, 3\delta' + 9T\delta' + 3e^{-\tilde{p}T}) - \text{indistinguishable} \tag{C.3}$$

Case 3: $t > t_0$**:** Recall that while claiming indistinguishability of appropriate pair of random variables, we condition on a shared past of Σ_{t-1}. In particular, this means that $x'_{t-1} = x_{t-1}$ and that $\zeta_{t-1} = \zeta'_{t-1}$. Now, if $\zeta_{t-1} = 0$, then $x'_t = x'_{t-1} = x_{t-1} = x_t$. If $\zeta_{t-1} = 1$, the iterates are sampled as $x_t \sim \bar{\mu}_t$ and $x'_t \sim \bar{\mu}'_t$ in round t. Once again by applying the condition on Φ as stated in Theorem 4.3 we have that x_t, x'_t are $(\zeta_{t-1}\log\Phi, \delta')$-indistinguishable.

To conclude the claim and hence the proof, we need to establish the indistinguishability of ζ_t and ζ'_t conditioned additionally on the event $x_t = x'_t$. Unlike for $t = t_0$, the analysis here for ζ's is more involved. To proceed, we first obtain a second-order perturbation result. We have

$$\begin{aligned}
\frac{\bar{\mu}_{t+1}(x)}{\bar{\mu}_t(x)} &= \frac{\exp\left(-\beta\left(l_{1:t}(x) + \frac{\lambda}{2}\|x\|^2\right)\right)}{\exp\left(-\beta\left(l_{1:t-1}(x) + \frac{\lambda}{2}\|x\|^2\right)\right)} \cdot \frac{\int_{x\in\mathcal{K}}\exp\left(-\beta\left(l_{1:t}(x) + \frac{\lambda}{2}\|x\|^2\right)\right)dx}{\int_{x\in\mathcal{K}}\exp\left(-\beta\left(l_{1:t-1}(x) + \frac{\lambda}{2}\|x\|^2\right)\right)dx} \\
&\triangleq \exp(-\beta \cdot l_t(x)) \cdot \frac{Z(l_{1:t-1})}{Z(l_{1:t})}
\end{aligned}$$

where we have defined $Z(l) = \int_{x\in\mathcal{K}}\exp\left(-\beta\left(l(x) + \frac{\lambda}{2}\|x\|^2\right)\right)dx$. Define $B_t = \frac{Z(l_{1:t-1})}{Z(l_{1:t})}$. To bound B_t we define the following scalar function $p(t) : [0,1] \to \mathbb{R}$ as $p(t) = \log(Z(l_{1:t-1} + t \cdot l_t), \beta, \lambda)$. The following lemma shows that $p(t)$ is a convex function and characterizes the derivative of p.

Lemma C.4. *Given two differentiable loss functions f, g, and any number $t \in \mathbb{R}$ define the measure $\mu(t)(x)$ over a convex set $\mathcal{K}$ as $\mu(t) = \exp(-(f(x) + tg(x)))$. Further define the log partition function of $\mu(t)$, $p(t) \triangleq \log\left(\int_{x\in K}\exp(-(f(x) + tg(x))dx\right)$. Define the probability disitrbution $\bar{\mu}(t)(x) = \frac{\mu(t)(x)}{\exp(p(t))}$. We have that $p(t)$ is a convex function of t. Furthermore $p'(t) = \mathbb{E}_{x\sim\bar{\mu}(t)}[-g(x)]$.*

Proof of Lemma C.4. We first derive the expression for the derivative. Consider the following calculation

$$p'(t) = \frac{\int_{x\in K} -g(x) \cdot \exp(-(f(x) + tg(x))dx}{\int_{x\in K}\exp(-(f(x) + tg(x))dx} = \mathbb{E}_{x\sim\bar{\mu}(t)}[-g(x)]$$

To prove convexity we consider $p''(t)$. Once again, we can calculate as follows:

$$\begin{aligned} p''(t) &= \frac{\int_{x\in K} g^2(x)\cdot \exp(-(f(x)+tg(x))dx}{\int_{x\in K} \exp(-(f(x)+tg(x))dx} - \left(\frac{\int_{x\in K} g(x)\cdot \exp(-(f(x)+tg(x))dx}{\int_{x\in K} \exp(-(f(x)+tg(x))dx}\right)^2 \\ &= \mathrm{Var}_{\bar{\mu}(t)}(g(x)) \geq 0. \end{aligned}$$

Since $p''(t) \geq 0$ this proves that the function is convex. □

In particular using the above lemma we get that

$$\log(B_t) = p(0) - p(1) \leq -\frac{\partial p(0)}{\partial t} = \mathbb{E}_{y\sim\bar{\mu}_t}[\beta \cdot l_t(y)]$$

$$\log(B_t) = p(0) - p(1) \geq -\frac{\partial p(1)}{\partial t} = \mathbb{E}_{y\sim\bar{\mu}_{t+1}}[\beta \cdot l_t(y)]$$

It now follows that

$$\begin{aligned} \log\frac{\bar{\mu}_{t+1}(x)}{\bar{\mu}_t(x)} &\leq -\beta\cdot l_t(x) + \mathbb{E}_{y\sim\bar{\mu}_t}[\beta\cdot l_t(y)] \\ \log\frac{\bar{\mu}_{t+1}(x)}{\bar{\mu}_t(x)} &\geq -\beta\cdot l_t(x) + \mathbb{E}_{y\sim\bar{\mu}_{t+1}}[\beta\cdot l_t(y)]. \end{aligned}$$

Similarly for $\bar{\mu}'$, one can establish

$$\begin{aligned} \log\frac{\bar{\mu}'_{t+1}(x)}{\bar{\mu}'_t(x)} &\leq -\beta\cdot l'_t(x) + \mathbb{E}_{y\sim\bar{\mu}'_t}[\beta\cdot l'_t(y)] \\ \log\frac{\bar{\mu}'_{t+1}(x)}{\bar{\mu}'_t(x)} &\geq -\beta\cdot l'_t(x) + \mathbb{E}_{y\sim\bar{\mu}'_{t+1}}[\beta\cdot l'_t(y)]. \end{aligned}$$

At this point, note that since $t > t_0$, $l'_t = l_t$, and that $l_{1:t-1} - l'_{1:t-1} = l_{t_0} - l'_{t_0}$, we can now bound the term of interest for privacy for all x.

$$\log\frac{\frac{\bar{\mu}'_{t+1}(x)}{\Phi\bar{\mu}'_t(x)}}{\frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}} \leq \ \mathbb{E}_{y\sim\bar{\mu}'_t}[\beta\cdot l_t(y)] - \mathbb{E}_{y\sim\bar{\mu}_{t+1}}[\beta\cdot l_t(y)].$$

Now using Lemma 3.6 twice we get that $W_\infty(\bar{\mu}'_t, \bar{\mu}_{t+1}) \leq \frac{2G}{\lambda}$ which implies that there is a coupling γ between $\bar{\mu}'_t$ and $\bar{\mu}'_{t+1}$ such that $sup_{(y,y')\sim\gamma}\|y - y'\| \leq \frac{2G}{\lambda}$. Therefore we have that

$$\mathbb{E}_{y\sim\bar{\mu}'_t}[\beta\cdot l_t(y)] - \mathbb{E}_{y\sim\bar{\mu}_{t+1}}[\beta\cdot l_t(y)] = \beta\cdot\mathbb{E}_{(y,y')\sim\gamma}[l_t(y) - l_t(y')] \leq \beta\cdot G\cdot\mathbb{E}_{(y,y')\sim\gamma}[|y-y'|] \leq \frac{\beta\cdot 2G^2}{\lambda}.$$

The above display immediately gives that for all $\Sigma \in \mathbf{\Sigma_{t-1}}$ and for all $x \in \mathcal{E}_t$,

$$\frac{\Pr(\zeta'_t = 0|\Sigma'_{t-1} = \Sigma, x'_t = x)}{\Pr(\zeta_t = 0|\Sigma_{t-1} = \Sigma, x_t = x)} = \frac{(1-p)\frac{\bar{\mu}'_{t+1}(x)}{\Phi\bar{\mu}'_t(x)}}{(1-p)\frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}} \leq e^{\frac{2G^2\beta}{\lambda}}.$$

Now, for the remaining possibility, we have

$$\begin{aligned}
\frac{\Pr(\zeta_t' = 1|\Sigma_{t-1}' = \Sigma, x_t' = x)}{\Pr(\zeta_t = 1|\Sigma_{t-1} = \Sigma, x_t = x)} &= \frac{p + (1-p)\left(1 - \frac{\bar{\mu}_{t+1}'(x)}{\Phi\bar{\mu}_t'(x)}\right)}{p + (1-p)\left(1 - \frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}\right)} \\
&\leq \frac{p + (1-p)\left(1 - \frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)} e^{-\frac{2G^2\beta}{\lambda}}\right)}{p + (1-p)\left(1 - \frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}\right)} \\
&\leq 1 + \frac{\underbrace{\frac{\bar{\mu}_{t+1}(x)}{\Phi\bar{\mu}_t(x)}}_{\leq 1}\left(1 - e^{-\frac{2G^2\beta}{\lambda}}\right)}{p} \\
&\leq e^{\frac{1}{p}\cdot\frac{2G^2\beta}{\lambda}}.
\end{aligned}$$

The above displays thereby imply that conditioned on Σ_{t-1} and $x_t \in \mathcal{E}_t$ we have that ζ_t and ζ_t' are $(\frac{2G^2\beta/\lambda}{p}, 0)$-indistinguishable. Thereby we get that conditioned on Σ_{t-1}

$$(x_t, \zeta_t) \text{ and } (x_t', \zeta_t') \text{ are } \left(\zeta_{t-1}\log\Phi + \frac{2G^2\beta/\lambda}{p}, 4\delta' - 9T\delta' - 3e^{-\tilde{p}T}\right) - \text{indistinguishable} \tag{C.4}$$

Combining the statements in Equations (C.2), (C.3) and (C.4) finishes the proof.

□

Proof Of Claim C.3. Let q_t be the probability distribution induced on the iterates chosen by Algorithm 1 when run on a loss sequence $l_{1:T}$. Using the conditions in the theorem and by Lemma 4.6, we have that $\|\bar{\mu}_t - q_t\| \leq e^{-\tilde{p}T} + 3T\delta'$ for any $t \in [T]$. From this, noting that $l_{1:t} - l_{1:t-1}$ is G-Lipschitz and β-smooth, we have that for all t,

$$\Pr_{X \sim q_t}\left[\frac{1}{\sqrt{\Phi}} \leq \frac{\bar{\mu}_{t+1}(X)}{\bar{\mu}_t(X)} \leq \sqrt{\Phi}\right] \geq 1 - \delta' - 3T\delta' - e^{-\tilde{p}T}$$

Furthermore noting that $l_{1:t-1} - l_{1:t-1}'$ is $2G$-Lipschitz and 2β-smooth we have that for all t,

$$\Pr_{X \sim q_t}\left[\frac{1}{\sqrt{\Phi}} \leq \frac{\bar{\mu}_t(X)}{\bar{\mu}_t'(X)} \leq \sqrt{\Phi}\right] \geq 1 - \delta' - 3T\delta' - e^{-\tilde{p}T}$$

Similarly noting that $l_{1:t}' - l_{1:t-1}$ is $3G$-Lipschitz and 2β-smooth we can apply the same argument to obtain

$$\Pr_{X \sim q_t}\left[\frac{1}{\sqrt{\Phi}} \leq \frac{\bar{\mu}_{t+1}'(X)}{\bar{\mu}_t(X)} \leq \sqrt{\Phi}\right] \geq 1 - \delta' - 3T\delta' - e^{-\tilde{p}T}$$

The above statements imply the claim. □