Differentially Private and Lazy Online Convex Optimization

Naman Agarwal NAMANAGARWAL@GOOGLE.COM
Google DeepMind

Satyen Kale SATYENKALE@GOOGLE.COM
Google Research

Karan Singh KARANSINGH@CMU.EDU
Tepper School of Business, Carnegie Mellon University

Abhradeep Thakurta ATHAKURTA@GOOGLE.COM
Google DeepMind

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

We study the task of differentially private online convex optimization (OCO). In the online setting, the release of each distinct decision or iterate carries with it the potential for privacy loss. To limit such privacy leakage, we design an optimization-based OCO algorithm that explicitly limits the number of switches via objective perturbation and rejection sampling. This improves over known results in multiple aspects: an optimal leading-order regret term, in being efficiently implementable without requiring log-concave sampling subroutines, and in matching the non-private regret bound for sub-constant regimes of privacy parameters. Leveraging the fact that the algorithm is designed to explicitly minimize the number of switches of decisions, we show that the algorithm also obtains optimal regret bounds in the lazy OCO setting, where the learner is constrained to perform a limited number of switches. In addition, for one- and two-dimensional decision sets, we present a novel approach for differentially private online Lipschitz learning, where the loss functions are Lipschitz but not necessarily convex, that achieves the optimal regret bound matching known lower bounds.

Keywords: online convex optimization, differential privacy, low switching, regret minimization

1. Introduction

In online convex optimization (OCO), in each round $t = 1, 2, \ldots, T$, a learner is required to choose a point x_t in a compact convex set $\mathcal{K} \in \mathbb{R}^d$, and is provided an adversarially chosen Lipschitz convex loss function $l_t : \mathcal{K} \to \mathbb{R}$ in response. The learner suffers loss $l_t(x_t)$ in round t. The learner's goal is to minimize her *regret* defined as $\sum_{t=1}^{T} l_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} l_t(x)$. We assume that the adversary chooses the loss functions *obliviously*, i.e., independently of the points x_t. When the points x_t are chosen randomly, the corresponding performance metric is the *expected* regret.

Differentially Private OCO (DP-OCO). The goal in DP-OCO is to design an online learning algorithm for this problem that guarantees that if one of the loss functions l_t in an arbitrary round t were changed to another function l'_t, then the *entire* output sequence of the algorithm doesn't change much in a certain precise manner depending on privacy parameters (ε, δ) that we formalize later. DP-OCO has been studied for over a decade (Jain et al., 2012; Smith and Thakurta, 2013; Agarwal and Singh, 2017; Kairouz et al., 2021; Asi et al., 2022). Until the very recent work of Asi et al. (2022), the best known upper bound for the problem was $\widetilde{\mathcal{O}}\left(\frac{d^{1/4}\sqrt{T}}{\sqrt{\varepsilon}}\right)$[1] (Kairouz et al., 2021). The

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

work of Asi et al. (2022), for moderate ranges of ε, improves the bound to $\widetilde{\mathcal{O}}\left(\sqrt{dT} + \frac{dT^{1/3}}{\varepsilon}\right)$. This is a significant improvement over prior work, and in fact doesn't even need convexity of the loss functions, simply Lipschitzness. However, there are a few drawbacks of this result. First, the regret bound is always $\Omega(\sqrt{dT})$ even with low (or even no!) privacy requirements. While this scaling in terms of d is unavoidable for non-convex Lipschitz losses, for convex losses, this is much worse than the optimal $O(\sqrt{T})$ regret achievable for OCO in the non-private setting. Second, the algorithm provided by Asi et al. (2022), while implementable in polynomial time for convex losses, is not very practical as it (occasionally) needs to sample from log-concave distributions. In this paper, we improve the state-of-the-art for DP-OCO on both these fronts. Our contributions are:

1. We give a DP-OCO algorithm for smooth Lipschitz losses with $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{dT^{1/3}}{\varepsilon}\right)$ regret. The key improvement in our regret bound over (Asi et al., 2022) is that we remove the additional $\sqrt{d}$ factor on the leading term, leading to the first algorithm with matches the optimal $O(\sqrt{T})$ regret in the non-private setting for $\varepsilon = \Omega(dT^{-1/6})$. We provide a detailed comparison of our regret vs the previously best known algorithms in Table 1 on page 3. Furthermore, crucially our algorithm improves over (Asi et al., 2022) in terms of computational efficiency. While the (Asi et al., 2022) algorithm is efficient in requiring function evaluations in most rounds, for roughly $\widetilde{\mathcal{O}}\left(T^{2/3}\right)$ rounds it needs to sample from logconcave probability densities, for which the best known algorithms (Chewi, 2023) are not yet practical. On the other hand, our algorithm needs the evaluation of one gradient and one Hessian in most rounds, and in roughly $\widetilde{\mathcal{O}}\left(T^{2/3}\right)$ rounds, needs to solve a convex optimization problem, for which several practical algorithms exist (Boyd and Vandenberghe, 2004). This reduction in the computational burden of a regret minimizer from sampling to optimization has long been a source of motivation for many works in online learning (e.g., (Hazan et al., 2007)).

2. In 1 or 2 dimensions for Lipschitz (but not necessarily convex) losses, we give a differentially private algorithm with a regret of $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{1}{\varepsilon}\right)$. This matches the bound obtained by Agarwal and Singh (2017) for the much weaker class of linear functions in 1 or 2 dimensions (their precise bound is $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{\sqrt{d}}{\varepsilon}\right)$ in dimension d). As a direct corollary, we show via online-to-batch conversion (Cesa-Bianchi et al., 2001), that for 1 or 2 dimensions our one pass algorithm is sufficient to obtain (near) optimal excess population risk of $\widetilde{\mathcal{O}}\left(\frac{1}{\sqrt{T}} + \frac{1}{\varepsilon T}\right)$ in differentially private Stochastic Convex Optimization (DP-SCO) (Bassily et al., 2019; Feldman et al., 2020; Bassily et al., 2020; Kulkarni et al., 2021; Gopi et al., 2022; Asi et al., 2021) with T samples.

3. We consider the case when the loss functions are chosen from the class of GLMs (Generalized Linear Models), i.e., functions that exhibit the structure $l_t(x) = f(v_t^\top x)$ for a fixed known smooth Lipschitz function f, with v_t chosen by the adversary. For such models we show an improved regret bound[2] of $\widetilde{\mathcal{O}}\left(\sqrt{T} + \frac{\sqrt{d}T^{1/3}}{\varepsilon}\right)$. The improvement in the regret bound for GLMs originates from an observation communicated to us by Kifer et al. (2023).

2. Previous results by Kifer et al. (2012) on objective perturbation would seem to suggest that this latter bound is what should be expected for our approach even in the general convex case, however we found an error in the analysis of Kifer et al. (2012) that was subsequently acknowledged by the authors in personal communication. The error is at top of page 25.21: the claim that "Γ is independent of the noise vector" is not justified since $\Gamma = b(\alpha; \mathcal{D}) - b(\alpha; \mathcal{D}')$, leading to a worse dependence on d than claimed in their paper.

$\varepsilon = T^{-\alpha}$	Previous Best	Our Algorithm	
		General Convex	GLM
$\alpha = 0$	$d^{1/4}\sqrt{T}$ (Kairouz et al., 2021)	$\sqrt{T}$	$\sqrt{T}$
$\alpha \in (0, 1/6)$	$\sqrt{dT}$ (Asi et al., 2022)	$\sqrt{T}$	$\sqrt{T}$
$\alpha \in [1/6, 1/3)$	$d \cdot T^{1/3+\alpha}$ (Asi et al., 2022)	$d \cdot T^{1/3+\alpha}$	$\sqrt{d} \cdot T^{1/3+\alpha}$
$\alpha \geq 1/3$	$d^{1/4} \cdot T^{1/2+\alpha/2}$ (Kairouz et al., 2021)	$d \cdot T^{1/3+\alpha}$	$\sqrt{d} \cdot T^{1/3+\alpha}$

Table 1: Comparison between our results and the known best results previously for DP-OCO in different regimes for ε. Entries in red are the known best result in a row. For the asymptotics we assume $T \gg d$.

Lazy OCO. Lazy OCO is the problem of developing OCO algorithms with a limit on the number of switches between the points chosen by the learner. This setting is motivated by real-world applications where changes in the learner's decision are costly. For example, this cost manifests as the need for verifying the safety of the newly proposed controllers in robotics, transaction costs associated with rebalancing portfolios in portfolio optimization, and as the burden of reimplementation in public or organizational policy decisions. Online learning with limited switching has been extensively studied in the context of prediction with expert advice (Merhav et al., 2002; Kalai and Vempala, 2005; Geulen et al., 2010; Altschuler and Talwar, 2021) and OCO (Anava et al., 2015; Sherman and Koren, 2021). For the OCO problem the best results known so far were provided in Anava et al. (2015) who showed that there exists a log-concave sampling based algorithm that achieved regret $\tilde{O}(\sqrt{dT} + \frac{dT}{S})$ while switching at most S times in expectation. Recently, Sherman and Koren (2021) claimed an improved guarantee of $\tilde{O}(\sqrt{T} + \frac{\sqrt{d}T}{S})$ via a much more practical algorithm based on optimization similar to the algorithm presented in this paper. However, we found an error (see Section A) in the paper that was acknowledged by the authors in personal communication. Our contributions for Lazy OCO are the following (summarized in Table 2).

1. Observing that our DP-OCO algorithm is designed to perform a limited number of switches in order to minimize privacy loss, we show that the very same algorithm also achieves $\tilde{O}(\sqrt{T} + \frac{dT}{S})$ regret while switching at most S in expectation for any given S. The algorithm is naturally significantly more efficient than the one proposed in Anava et al. (2015) since the latter algorithm uses log-concave sampling. To highlight the significance of our result, note that due to the additional d factor in the leading term of the regret bound by Anava et al. (2015), prior to our result it was not known whether optimal $O(\sqrt{T})$ regret for OCO could be achieved for any $S = o(T)$. Our result demonstrates that this is indeed possible for any $S = \Omega(d\sqrt{T})$. Furthermore, for strongly-convex losses, we can improve the bound to $\tilde{O}(1 + \frac{d^2T}{S^2})$ by leveraging the increasing noise technique introduced by Sherman and Koren (2021). However, we remark that it is unclear if a scaling of d in the lower-order regret terms is necessary.
2. We show that our bounds can be improved when the losses are GLMs to $\tilde{O}(\sqrt{T} + \frac{\sqrt{d}T}{S})$.

2. Preliminaries

Notation. We use $\|\cdot\|$ to denote the standard ℓ_2 norm in $\mathbb{R}^d$. For distributions p and q on the same outcome space, we use $\|p - q\|_{\text{TV}}$ to denote their total variation distance. For a distribution μ on

$S = T^\alpha$	Previous Best	Our Algorithm	
		General Convex	GLM
$\alpha = 1$	$\sqrt{T}$ (Zinkevich, 2003)	$\sqrt{T}$	$\sqrt{T}$
$\alpha \in (1/2, 1)$	$\sqrt{dT}$ (Anava et al., 2015)	$\sqrt{T}$	$\sqrt{T}$
$\alpha \leq 1/2$	$d \cdot T^{1-\alpha}$ (Anava et al., 2015)	$d \cdot T^{1-\alpha}$	$\sqrt{d} \cdot T^{1-\alpha}$

Table 2: Comparison between our results and the known best results previously for Lazy OCO in different regimes for the switching budget S. Entries in red are the known best results in a row. For the asymptotics we assume $T \gg d$.

$\mathbb{R}^d$, 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 are given a convex compact set $\mathcal{K} \in \mathbb{R}^d$ with diameter D (i.e. $D = \max_{x,y \in \mathcal{K}} \|x - y\|$). In OCO, at the start of each round $t \in [T]$, the learner $\mathcal{A}$ chooses a point $x_t \in \mathcal{K}$ from some compact convex decision set $\mathcal{K} \subset \mathbb{R}^d$, and upon making this choice it observes the 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 restrict our attention to the case of *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] Recall that a function $l : \mathcal{K} \to \mathbb{R}$ is said to be G-Lipschitz if $|l(x) - l(y)| \leq G\|x - y\|$ for any pair $x, y \in \mathcal{K}$ and β-smooth if l is differentiable on $\mathcal{K}$ with $\|\nabla l(x) - \nabla l(y)\| \leq \beta \|x - y\|$ for any pair $x, y \in \mathcal{K}$.

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

As for the domain $\mathcal{K}$, we assume that (a) it is full-dimensional and (b) $0 \in \mathcal{K}$. We define the Minkowski set $\mathcal{K}^\circ = \{(1 - \frac{1}{T})x : x \in \mathcal{K}\}$. This is a convex set and in particular we have that for any $x \in \mathcal{K}$ there is a point $x' \in \mathcal{K}^\circ$ such that $\|x - x'\| \leq \frac{D}{T}$. We also assume we have a twice-differentiable and convex barrier function for $\mathcal{K}$, i.e. a function $\mathcal{B} : \mathbb{R}^d \to \mathbb{R}_+ \cup \{+\infty\}$ such that for $x \in \text{int}(\mathcal{K})$, $\mathcal{B}(x) \in \mathbb{R}$ and for any point $x \notin \text{int}(\mathcal{K})$, $\mathcal{B}(x) = +\infty$. Such barriers are easy to construct for several convex sets of interest: e.g., for the d-dimensional unit sphere, $x \mapsto -\log(1 - \|x\|^2)$ is a barrier function of the type described above. Additionally, via appropriate scaling, we may assume that $\mathcal{B}(x) \leq GD$ for any $x \in \mathcal{K}^\circ$. For example, for the unit sphere, the function $x \mapsto -\frac{GD}{\log(T/2)} \log(1 - \|x\|^2)$ satisfies this assumption.

The possibly random learner's performance through such mode of interaction as outlined above may be assessed via the regret it incurs; this, as defined below, measures the expected excess aggregate loss the learner is subject to in comparison to the best fixed point in $\mathcal{K}$ determined with the benefit of hindsight.

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

3. As remarked in Asi et al. (2022), and as is true for most of the literature on private OCO, our privacy bounds hold in the absence of this assumption – obliviousness – due to the use of adaptive strong composition. The utility or regret bounds are strongly reliant on this assumption on loss functions, however.

Algorithm 1: Couple-The-Regularized-Leader (CTRL)

Inputs: A distribution ν on $\mathbb{R}^d$, a regularization parameter $\eta > 0$, a barrier $\mathcal{B}(x)$, switching rate parameter $p \in [0,1]$, switching budget $B \geq 0$, a scale parameter $\Phi > 0$.

Set $b_1 = 0$, sample $Z_0 \sim \nu$, choose $x_1 = x^\star(0, Z_0)$.

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 $l_t(x_t)$.

  Sample $S_t \sim \text{Ber}\left(\min\left\{1, \max\left\{\frac{1}{\Phi^2}, \frac{\mu_{t+1}(x_t)}{\Phi \cdot \mu_t(x_t)}\right\}\right\}\right)$ and $S'_t \sim \text{Ber}(1-p)$.

  // $\mu_t(\cdot), \mu_{t+1}(\cdot)$ can be computed via Lemma 2.

  if $b_t < B$ *and* $(S'_t = 0$ *or* $S_t = 0)$ **then**

    Update $b_{t+1} = b_t + 1$, sample $Z_t \sim \nu$ and choose $x_{t+1} = x^\star(l_{1:t}, Z_t)$.

  end

  else

    Set $b_{t+1} = b_t$ and $x_{t+1} = x_t$.

  end

end

Later on, since we do not make any distributional assumptions on the loss sequence, the primary quantity of interest will be the *worst-case* regret, i.e. $\mathcal{R}_T(\mathcal{A}) \triangleq \max_{l_{1:T} \in \mathcal{L}^T} \mathcal{R}_T(\mathcal{A}, l_{1:T})$.

Another characteristic of the learner that is relevant to the discussion below is the number of discrete switches the learner makes. To this end, we define the number of switches the learner makes 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.

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 $O \subset \mathcal{K}^T$ that

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

3. Algorithm and main result

We now present our algorithm, dubbed Couple-The-Regularized-Leader (CTRL); see Algorithm 1. At a high level, the algorithm is an instance of Follow-The-Regularized-Leader, i.e., in round t, it plays the point

$$x_t = \arg\min_{x \in \mathcal{K}} \left\{\sum_{\tau=1}^{t-1} l_\tau(x) + \text{Reg}(x)\right\},$$

where $\text{Reg} : \mathcal{K} \to \mathbb{R}$ is a strongly-convex regularizer. Such schemes are known to have low regret; see Hazan (2016). A few factors go into the design of Reg. First, in order to preserve privacy, the regularizer contains a *random* linear function, $x \mapsto Z^\top x$ for some random vector $Z \in \mathbb{R}^d$ drawn

from some distribution ν (we will use Gaussians). Second, for technical reasons that we list shortly, we require the algorithm to choose $x_t \in \text{int}(\mathcal{K})$. To ensure this, we add a barrier function $\mathcal{B}$ in Reg. Finally, to ensure strong convexity, we add $x \mapsto \frac{\|x\|^2}{2\eta}$ to Reg for some learning rate parameter $\eta > 0$ which we specify later.

The above scheme already yields low regret, but leak a lot of private information since the algorithm can potentially alter its decisions in each round. To guard against this, we use a rejection sampling procedure, drawing 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 $\frac{\mu_{t+1}(x_t)}{\Phi \mu_t(x_t)}$, where for any t, μ_t is the distribution of $\arg\min_{x \in \mathcal{K}} \left\{ \sum_{\tau=1}^{t-1} l_\tau(x) + \text{Reg}(x) \right\}$ induced by the random vector Z, and Φ is a scaling factor. With the remaining probability, we sample x_{t+1} from μ_{t+1} (we call this a "switch"). This rejection sampling technique ensures that the distribution of x_{t+1} is indeed μ_{t+1}.

One final issue remains: the (random, but adaptively determined) decision to switch itself may leak private information (specifically, it is possible that a change in one loss function makes the probability of switching 0 whereas it is non-zero prior to this change). To guard against this, we employ a technique developed by Asi et al. (2022): we force switching in each round at a certain base switching rate p. Finally, it is necessary to put a hard cap on the number of switches, so we introduce a switching budget B, again inspired by (Asi et al., 2022), and no longer switch once the budget is exhausted. We also need to scale the density ratio $\frac{\mu_{t+1}(x_t)}{\Phi \mu_t(x_t)}$ appropriately to make sure it is at most unit sized.

To complete the technical description of the algorithm, we need explicit formulas for μ_t. We now provide those formulas. First, given any loss function $l : \mathcal{K} \to \mathbb{R}$, we define the following quantities.

$$\mathcal{J}(l, x) = l(x) + \frac{\|x\|^2}{2\eta} + \mathcal{B}(x), \qquad x^\star(l, Z) = \operatorname*{argmin}_{x \in \mathcal{K}} \mathcal{J}(l, x) + Z^\top x \tag{3.1}$$

With some abuse of notation we use $\mathcal{J}(l_{1:t-1}, \cdot)$ to denote $\mathcal{J}(\sum_{\tau=1}^{t-1} l_\tau, \cdot)$, and similarly $x^\star(l_{1:t-1}, \cdot)$ to denote $x^\star(\sum_{\tau=1}^{t-1} l_\tau, \cdot)$. Note that μ_t is defined to be the distribution of $x^\star(l_{1:t-1}, Z)$ when $Z \sim \nu$. The following lemma provides the necessary explicit formula for the density of μ_t:

Lemma 2 *For a loss function $l : \mathcal{K} \to \mathbb{R}$, let μ be the distribution of $x^\star(l, Z)$ when $Z \sim \nu$. Then we have*

$$\mu(x) = \nu(-\nabla \mathcal{J}(l, x)) |\det(-\nabla^2 \mathcal{J}(l, x))| = \nu(-\nabla \mathcal{J}(l, x)) \det(\nabla^2 \mathcal{J}(l, x)),$$

where the gradient and Hessian above are taken with respect to the x argument of $\mathcal{J}(l, x)$.

Proof Due to the presence of the barrier function $\mathcal{B}$ in Equation (3.1), we have that $x^\star(l, Z) \in \text{int}(\mathcal{K})$ for any Z. Note that $\mathcal{J}(l, x)$ is strongly convex and is differentiable at any point $x \in \text{int}(\mathcal{K})$. Hence, the Fenchel conjugate, $\mathcal{J}^*$, of $\mathcal{J}(l, \cdot)$, is differentiable at any $Z \in \mathbb{R}^d$, and it follows that $x^\star(l, Z) = \nabla \mathcal{J}^*(-Z)$ and $Z = -\nabla \mathcal{J}(l, x^\star(l, Z))$, which further implies that $Z \mapsto \nabla \mathcal{J}^*(-Z)$ is one-to-one, with the inverse map given by $x \mapsto -\nabla \mathcal{J}(l, x)$. The claimed formula for the density of μ then follows by the change-of-variable formula; see, e.g., (Bogachev and Ruas, 2007). ■

We now turn to the regret analysis for Algorithm 1. We state the regret bound below in a somewhat general fashion in order to easily yield results in various specific cases. The analysis hinges on the following definition:

Definition 3 *Probability distributions μ, μ' on $\mathcal{K}$ are said to be (Φ, δ)-close if*

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

We have the following regret bound for Algorithm 1.

Theorem 4 (Regret bound for CTRL) *In Algorithm 1, fix any $\eta, \sigma > 0$, any $\delta \in [0, 1/2]$, any $p \in [0,1]$, set $\nu = \mathcal{N}(0, \sigma^2 I)$ and choose Φ such that for all t the distributions μ_t, μ_{t+1} are (Φ, δ)-close. For any sequence of obliviously chosen G-Lipschitz, β-smooth convex loss functions $l_{1:T}$, the following hold:*

- *If $B = \infty$,*
$$\mathcal{R}_T \le \frac{D^2}{2\eta} + 2G^2\eta T + \sigma\sqrt{d}D + 6GD\delta T^2 + 2GD.$$
- *Let $\tilde{p} = p + 1 - \Phi^{-2}$. If $B = 3\tilde{p}T$,*
$$\mathcal{R}_T \le \frac{D^2}{2\eta} + 2G^2\eta T + \sigma\sqrt{d}D + 2GDT(e^{-\tilde{p}T} + 3\delta T) + 2GD.$$

Proof (Sketch; detailed proof in Appendix C.) The proof of Theorem 4 is based on the standard FTRL analysis, and a sketch follows. While ideally x_t would be sampled from μ_t, due to the budget constrant B and the trimming of the density ratio $\frac{\mu_{t+1}(x_t)}{\Phi\cdot\mu_t(x_t)}$ to the interval $[\frac{1}{\Phi^2}, 1]$, the actual distribution of x_t deviates somewhat from μ_t. We can place bounds on this deviation in terms of δ and p. Now if x_t were indeed sampled exactly from μ_t, then the FTRL analysis can be applied. The main idea there is that in expectation, the regret would be the same if the noise vectors Z_t were all set to be equal to one single noise vector $Z \sim \nu$. So we analyze the algorithm where this is done, treating $Z^\top x$ as part of the regularizer. Standard FTRL analysis then shows that the regret is bounded by $\frac{D^2}{2\eta} + 2G^2\eta T$ plus some excess regret due to the regularization. This excess regret primarily arises from the $Z^\top x$ part of the regularizer, which can be bounded by $\sigma\sqrt{d}D$ in expectation. The scaling of the barrier $\mathcal{B}$ ensures that it contributes no more than $2GD$. Finally, the excess regret of Algorithm 1 over that of the idealized algorithm can be bounded in terms of the deviation of the distribution of x_t from μ_t and yields the other terms in the stated regret bound. ■

The following lemma, proved in Appendix C, 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:

Lemma 5 (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,*

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

Finally, we turn to the privacy guarantee for Algorithm 1.

Theorem 6 (Privacy) *Given $\sigma > 0$ and $\delta \in (0, 1/2]$, for any $T \geq 12\log(1/\delta)$, let $\delta' = \frac{\delta T^{-2}}{60}$, $G' = 3G$, and $\beta' = 2\beta$. Suppose there exists $\Phi' > 0$ such that for all convex functions l, l' where $l - l'$ is G'-Lipschitz and β'-smooth, we have that, the distributions of $x^\star(l, Z)$ and $x^\star(l', Z)$ respectively when $Z \sim \mathcal{N}(0, \sigma^2 I)$, are (Φ', δ')-close. Define*

$$\varepsilon' = 7\log^2(\Phi)T^{2/3} + 2\log^3(\Phi)T + (2G^2/\sigma^2 + 2\eta^2\beta^2 d)^2 T^{5/3}.$$

Then for any sequence of G-Lipschitz, β-smooth convex functions, Algorithm 1 when run with $\Phi = \Phi'^2$, $p = T^{-1/3}$ and $B = 3\tilde{p}T$ is $(\varepsilon, \delta + 3Te^{-(1-\Phi^{-2})T})$-differentially private where

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

Proof (Sketch; detailed proof in Appendix D) At a high level, our proof of the above theorem relies on two main ideas. Firstly at any round due to the perturbation we get privacy proportional to $1/\log(\Phi)$. The argument to establish this follows a similar line of reasoning as known objective perturbation results (Kifer et al., 2012). The probability of switching at any round is a ratio of probabilities of successive distributions. Using a similar argument as in establishing per-round privacy we obtain that the probability of switching scales as $1/\log(\Phi)$. Using adaptive strong composition it can be seen that the total privacy loss scales as $\sqrt{(\#\text{Switches}) \cdot (\text{PerStepPrivacyLoss})^2} \sim \sqrt{T\log(\Phi)^{-3}}$. As we show later $\log(\Phi)$ scales as σ^{-1} and setting σ appropriately gives us the required bound.

This rough sketch outlined above is unfortunately incomplete since we also need to consider the privacy of switching decision itself which happens at every round. To overcome this issue we use the forced switching technique developed by Asi et al. (2022) which allows for bounded privacy loss at the point where the loss sequence changes. However unlike in the case of Asi et al. (2022), since the distributions we sample from cannot be broken down as product distributions depending on individual loss functions we need to provide additional analysis that the switching test does not leak privacy at other rounds. To this end, we perform a second-order analysis of the privacy loss incurred by the switching decision. The benefit of the second-order analysis is that it leads to the total privacy loss incurred by switching decisions to be of smaller order than the overall privacy loss thereby not affecting the overall privacy loss bound. This lower order penalty is the source of the $(2G^2/\sigma^2 + 2\eta^2\beta^2 d)T^{5/6}$ in the above theorem. ■

3.1. Bounds for Lipschitz and Smooth loss functions

In order to apply the above results for OCO with G-Lipschitz and β-smooth loss functions, all we need to do is compute Φ. This bound is given by the following lemma, proved in Appendix C:

Lemma 7 (Density ratio) *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be convex twice-differentiable functions such that $l - l'$ is G-Lipschitz and β-smooth. Let μ, μ' be the probability distributions of $x^\star(l, Z)$ and $x^\star(l', Z)$ respectively when $Z \sim \mathcal{N}(0, \sigma^2 I)$. Then for any $\delta \in (0, 1]$, we have that μ and μ' are (Φ, δ) close where*

$$\Phi = \exp\left(\eta\beta d + (G^2 + 2G\sigma\sqrt{2d\log(2/\delta)})/2\sigma^2\right).$$

This lemma follows directly as a corollary of a more general statement proved in Lemma 21 that is needed for the strongly-convex OCO case. Via this bound, and combining Theorem 6 and Theorem 4, we get the following result via straightforward calculations:

Theorem 8 (DP OCO) *For any given $\varepsilon \leq 1, \delta \in (0, 1/2]$ and any $T \geq 12\log(1/\delta)$, set*

$$\eta = \min\left(\frac{D}{2G\sqrt{T}}, \frac{\sqrt{\varepsilon}}{10^3 T^{5/12}\beta\sqrt{d}\sqrt{\log(T/\delta)}}, \frac{\varepsilon}{10^3 T^{1/3}\beta d\sqrt{\log(T/\delta)}}\right),$$

$$\sigma = \max\left(\frac{G\sqrt{T}}{\sqrt{d}}, \frac{10^3 G T^{1/3}\sqrt{d\log(120T/\delta)}}{\varepsilon}\right)$$

and other parameters as in Theorem 6. Then we get that Algorithm 1 is (ε, δ) differentially private and additionally satisfies

$$\mathcal{R}_T \leq \widetilde{\mathcal{O}}\left(GD\sqrt{T} + \frac{(GD + \beta D^2(\max(\beta D/G, 1)))dT^{1/3}\log(T/\delta)}{\varepsilon}\right).$$

Similarly, for Lazy OCO, using Theorem 4 and Lemma 5, we get the following result:

Theorem 9 (Lazy OCO) *For any $T \geq 3$ and any given bound $S \leq T$ on the number of switches, set $\delta = 2/T^2$, $\sigma = \frac{12GT\sqrt{d\log(T)}}{S}$, $\nu = \mathcal{N}(0, \sigma^2 I)$, $\Phi = \exp\left(\eta\beta d + \frac{G^2 + 4G\sigma\sqrt{d\log(T)}}{2\sigma^2}\right)$, $p = 0$, $\eta = \min\left(\frac{D}{2G\sqrt{T}}, \frac{S}{6\beta dT}\right)$, and $B = \infty$ in Algorithm 1. Then for any sequence of obliviously chosen G-Lipschitz β-smooth convex loss functions $l_{1:T}$, Algorithm 1 satisfies the following:*

$$\mathcal{R}_T \leq 2GD\sqrt{T} + \frac{dT}{S}\left(3\beta D^2 + 12GD\sqrt{\log(T)}\right) + 14GD \text{ and } \mathbb{E}[\mathcal{S}_T] \leq S.$$

Proof We begin by first bounding the number of switches using Lemma 5. We get that

$$\mathbb{E}[\mathcal{S}_T] \leq \tilde{p}T \leq (1 - \Phi^{-2})T \leq 2\log(\Phi)T \leq 2T\left(\underbrace{\eta\beta d}_{\leq \frac{S}{6T}} + \underbrace{\frac{G^2}{2\sigma^2}}_{= \frac{S^2}{72T^2\sqrt{d}\log^2(T)} \leq \frac{S}{72T}} + \underbrace{\frac{4G\sigma\sqrt{d\log(T)}}{2\sigma^2}}_{\leq \frac{S}{6T}}\right) \leq S$$

The regret calculation is straightforward. ■

3.2. Bounds for Generalized Linear Models

In this section, we provide improved bounds in terms of the dependence on dimension for Generalized Linear Models (GLMs). The specific setting we consider is the following: the loss function $l_t(x) = f(\langle v_t, x\rangle)$ where $f : [-D, D] \to \mathbb{R}$ is a G-Lipschitz and β-smooth convex function and v_t a vector satisfying $\|v_t\| \leq 1$ that is chosen by an (oblivious) adversary. Since $\nabla l_t(x) = f'(\langle v_t, x\rangle)v_t$, we conclude that l_t is G-Lipschitz, as in the rest of this paper. Similarly, since $\nabla^2 l_t(x) = f'(\langle v_t, x\rangle)v_t v_t^\top$, we conclude that l_t is β-smooth. In order to apply the general results to GLMs, we need to do is compute Φ. This bound is given by the following lemma, which is a direct corollary of Lemma 24 proved in Appendix C:

Lemma 10 *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be convex twice-differentiable functions such that for all $x \in \mathcal{K}$, $l(x) - l'(x) = sf(\langle v, x\rangle)$ for some $\|v\| \le 1$, $s \in \{-1, 1\}$ and $f : [-D, D] \to \mathbb{R}$ is a G-Lipschitz and β-smooth convex function. Given $\sigma > 0$, let μ, μ' be the probability density functions of $x^\star(l, Z)$ and $x^\star(l', Z)$ respectively when $Z \sim \mathcal{N}(0, \sigma^2 I)$. Then for any $\delta \in (0, 1]$, we have that μ and μ' are (Φ, δ) close where*

$$\Phi = \exp\left(\eta\beta + (G^2 + 2G\sigma\sqrt{2\log(2/\delta)})/2\sigma^2\right).$$

The main improvement we have in the above bound compared to Lemma 7 is an improvement of $\sqrt{d}$ (suggested to us by Kifer et al. (2023)) in the terms scaling with $1/\sigma$ above which lead to the following improved results which we achieve via straightforward calculations using the above bound, and combining Theorem 6 and Theorem 4.

Theorem 11 (DP OCO - GLM) *For any given $\varepsilon \le 1, \delta \in (0, 1/2]$ and any $T \ge 12\log(1/\delta)$, set*

$$\eta = \min\left(\frac{D}{2G\sqrt{T}}, \frac{\sqrt{\varepsilon}}{10^3 T^{5/12}\beta\sqrt{\log(T/\delta)}}, \frac{\varepsilon}{10^3 T^{1/3}\beta\sqrt{\log(T/\delta)}}\right),$$

$$\sigma = \max\left(G\sqrt{T}, \frac{10^3 G T^{1/3}\sqrt{\log(120T/\delta)}}{\varepsilon}\right)$$

and other parameters as in Theorem 6. Then we get that Algorithm 1, when run on GLM losses, is (ε, δ) differentially private and additionally satisfies

$$\mathcal{R}_T \le \widetilde{\mathcal{O}}\left(GD\sqrt{T} + \frac{(GD\sqrt{d} + \beta D^2 \max(\beta D/G, 1))T^{1/3}\log(T/\delta)}{\varepsilon}\right).$$

For Lazy OCO, via similar calculations as in the proof of Theorem 9 using Theorem 4 and Lemma 5, we immediately get the following result:

Theorem 12 (Lazy OCO - GLM) *For any $T \ge 3$ and any given bound $S \le T$ on the number of switches, set $\delta = 2/T^2$, $\sigma = \frac{12GT\sqrt{\log(T)}}{S}$, $\nu = \mathcal{N}(0, \sigma^2 I)$, $\Phi = \exp\left(\eta\beta + \frac{G^2 + 4G\sigma\sqrt{\log(T)}}{2\sigma^2}\right)$, $p = 0$, $\eta = \min\left(\frac{D}{2G\sqrt{T}}, \frac{S}{6\beta T}\right)$, and $B = \infty$ in Algorithm 1.Then for any sequence of obliviously chosen G-Lipschitz β-smooth GLM functions $l_{1:T}$, Algorithm 1 satisfies the following:*

$$\mathcal{R}_T \le 2DG\sqrt{T} + \frac{T}{S}\left(3\beta D^2 + 12GD\sqrt{\log(T)}\right) + 14GD \text{ and } \mathbb{E}[\mathcal{S}_T] \le S.$$

3.3. Lazy OCO Bounds for Lipschitz, Smooth and Strongly-Convex loss functions

In this section, we provide improved Lazy OCO bounds for strongly-convex loss functions. Specifically, we assume that the loss functions l_t are G-Lipschitz, β-smooth, and λ-strongly-convex for all $t \in [T]$. Recall that a function $l : \mathcal{K} \to \mathbb{R}$ is λ-strongly-convex if for all $x, y \in \mathcal{K}$, we have $l(y) \ge l(x) + [\nabla l(x)]^\top (y - x) + \frac{\lambda}{2}\|y - x\|^2$.

To obtain the bound, we employ a technique due to Sherman and Koren (2021) where the sampling distributions for the noise are changed in each round t. The exact algorithm is given in the appendix in Algorithm 2. The changing noise distributions, denoted ν_t in round t, necessitates changing values of Φ_t in the different rounds as well. We have the following bound (proved in Appendix E) for Lazy OCO with strongly-convex loss funtions:

Theorem 13 (Lazy OCO - Strongly-Convex) *For any $T \geq 3$ and any given bound on the number of switches $S \geq \left(\frac{4\beta d}{\lambda} + 16d\log(T)\right)\log(T)$, set $\delta = 2/T^2$ and $\sigma = \frac{16G\sqrt{dT\log(T)}}{S}$. For all $t \geq 0$, set $\sigma_t = \sigma\sqrt{t}$, $\nu_t = \mathcal{N}(0, \sigma_t^2 I)$,*

$$\Phi_t = \frac{\beta d}{\lambda t} + \frac{4d\log(T)}{t} + d\log\left(\frac{\sqrt{t+1}}{\sqrt{t}}\right) + \frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t},$$

and $\eta = \infty$ in Algorithm 2. Then for any sequence of obliviously chosen G-Lipschitz β-smooth λ-strongly convex functions $l_{1:T}$, Algorithm 2 satisfies the following:

$$\mathcal{R}_T \leq \frac{2G^2(1+\log(T))}{\lambda} + \frac{512d^2\log(T)(1+\log(T))}{\lambda}\cdot\frac{T}{S^2} + 14GD \text{ and } \mathbb{E}[\mathcal{S}_T] \leq S.$$

4. Differentially Private Online Lipschitz Optimization in 1 or 2 dimensions

We consider Differentially Private Online Lipschitz Optimization in 1 or 2 dimensions in this section. In this problem, the domain $\mathcal{K}$ is a closed convex subset of $\mathbb{R}^d$ with $d \in \{1, 2\}$, with diameter at most D. The loss functions $l : \mathcal{K} \to \mathbb{R}$ are assumed to be G-Lipschitz (but not necessarily convex). The main result in this section is the following:

Theorem 14 *There is an (ε, δ)-differentially private algorithm for the problem of online Lipschitz optimization in $d = 1$ or $d = 2$ dimensions whose expected regret is bounded by*

$$O\left(GD\sqrt{T\log(T)} + \frac{GD\log^2(T)\sqrt{C_d\log(1/\delta)}}{\varepsilon}\right),$$

where $C_1 = 1$ and $C_2 = O(\log(T))$.

The main idea behind this result is to discretize the loss function on a carefully chosen net, and then maintain a running sum of function evaluations at each of these net points using the tree-aggregation mechanism (Chan et al., 2011; Dwork et al., 2010). Essentially, we maintain a running sum of the functional approximation of the original loss function. The net contains $O(T^d)$ points and so is not very efficient (even for $d = 1$ or 2); however the technique works with non-convex (but Lipschitz) losses. Unfortunately, the idea does not extend to higher dimensions greater than 2 — the constructure here bound suffers from the curse of dimensionality and the regret scales exponentially in d. Technically, the result follows from a reduction of the problem to an online learning with experts problem for which the algorithm of Agarwal and Singh (2017) can be applied. The reduction can be abstracted to a more general setting which may be useful in other contexts, and we describe it here.

Detour into Differentially Private Online Learning with Experts. Consider an online learning with experts problem with a set of experts $\mathcal{K}$. The loss functions $l : \mathcal{K} \to \mathbb{R}$ are drawn from a class of loss functions $\mathcal{L}$. With some abuse of notation, also use $l \in \mathbb{R}^{|\mathcal{K}|}$ to denote the vector of losses indexed by $\mathcal{K}$, given by the function $l : \mathcal{K} \to \mathbb{R}$. For this problem, Agarwal and Singh (2017) give an (ε, δ)-differentially private online learning with experts algorithm using the standard entropy

regularizer. Theorem 3.1 in (Agarwal and Singh, 2017) gives the following result when specialized to the online learning with experts problem:[4]

Theorem 15 (Theorem 3.1 in Agarwal and Singh (2017), restated) *There is an (ε, δ)-differentially private online learning algorithm with experts in $\mathcal{K}$ and losses in $\mathcal{L}$ with expected regret bounded by*

$$O\left(\sqrt{T\log(|\mathcal{K}|)} + \sigma\sqrt{\log(|\mathcal{K}|)}\right), \text{ where } \sigma = \max_{l\in\mathcal{L}}\{\|l\|\} \cdot \frac{\log(T)\sqrt{\log(1/\delta)}}{\varepsilon}.$$

In the above result, the $\sigma\sqrt{\log(|\mathcal{K}|)}$ term is the only one that arises due to noise added in the algorithm to preserve privacy; the first term appears in standard regret bounds without privacy requirements. In certain cases, such as the problem of differentially private online Lipschitz optimization in 1 or 2 dimensions, the following trick to reduce σ becomes useful. Call a matrix $M \in \mathbb{R}^{d\times d}$ invertible on $\mathcal{L}$ if there is a matrix $M^\dagger$ such that $M^\dagger M l = l$ for all $l \in \mathcal{L}$; we will call $M^\dagger$ a pseudoinverse of M. The following lemma (proved in Appendix F) describes the trick:

Lemma 16 *Let $M \in \mathbb{R}^{d\times d}$ be any matrix that is invertible on $\mathcal{L}$ with pseudoinverse $M^\dagger$. In Theorem 15, σ can be replaced by*

$$\sigma' = \|M^\dagger\|_{2\to\infty} \cdot \max_{l\in\mathcal{L}}\{\|Ml\|\} \cdot \frac{\log(T)\sqrt{\log(1/\delta)}}{\varepsilon}.$$

Thus, for particular applications of interest, one may be able to get better bounds by searching for a suitable matrix M that is invertible on $\mathcal{L}$.

Proof (Of Theorem 14) We make a few normalizations that do not affect the problem. First, by shifting $\mathcal{K}$, we may assume that $\mathcal{K} \subseteq [0, D]^d$. Next, we extend l to all of $[0, D]^d$ by setting $l(x) = l_t(\Pi(x))$, where $\Pi(x)$ is the projection of x on $\mathcal{K}$. This extension remains G-Lipschitz: for any $x, x' \in [0, D]^d$, we have

$$|l(x) - l(x')| = |l(\Pi(x)) - l(\Pi(x'))| \le G\|\Pi(x) - \Pi(x')\| \le G\|x - x'\|.$$

Thus, we may now assume that the learner can play points in $[0, D]^d$ instead of $\mathcal{K}$: whenever the learner wants to play x, we can instead play $\Pi(x)$ for the exact same loss. Hence, we may simply assume $\mathcal{K} = [0, D]^d$ henceforth. Finally, by shifting l appropriately, we may assume that $l(0) = 0$ without affecting regret.

Let $k := \lceil \log_2(T) \rceil$. Let $\mathcal{G} \subset \mathcal{K}$ be the set of points in $\mathcal{K}$ whose coordinates can be written as $\frac{t}{2^k}D$ for some $t \in \{0, 1, \ldots, 2^k\}$. We now reduce the differentially-private online Lipschitz optimization (DP-OLO) problem to a differentially-private online learning with experts problem over $|\mathcal{G}| \le (2T)^d$ experts identified with the points in $\mathcal{G}$, where the loss of expert $x \in \mathcal{G}$ in round t equals $l_t(x)$. It is easy to see that the total loss of the best expert in the experts problem, over all T rounds, is at most $O(GD)$ higher than the loss of the best point in K in the DP-OLO problem; hence the regret of any algorithm for the experts problem is at most $O(GD)$ less than the regret of the same algorithm for the DP-OLO problem.

4. While the result in Agarwal and Singh (2017) is stated for pure ε-differential privacy, the same analysis yields the statement given in Theorem 15 using Gaussian instead of Laplacian noise.

We can now apply the Agarwal and Singh (2017) algorithm along with the trick from Lemma 16. Let $\mathcal{L}$ denote the loss vectors arising from G-Lipschitz functions $l : \mathcal{K} \to \mathbb{R}$ with $l(0) = 0$. We need to define an appropriate linear transformation of the loss vectors in $\mathcal{L}$. For this, we first define a mapping $p : \mathcal{G} \to \mathcal{G}$ as follows. For $d = 1$, we first define $p(0) = 0$. Then, for any $x \in \mathcal{G}$ that is non-zero, let $x = \frac{a}{2^b} D$ be the unique representation of x where a is an odd integer and $b \geq 0$ is an integer. Then $p(x) := \frac{a-1}{2^b} D$. This construction implies that $\sum_x |x - p(x)|^2 \leq \sum_{b=0}^{k} 2^{b-1} \cdot \frac{D^2}{4^b} \leq D^2$; this fact will be useful later on. For $d = 2$, we again define $p(0) = 0$. Then, for any $x \in \mathcal{G}$ that is non-zero, x can be uniquely written as $x = (\frac{a_1}{2^b} D, \frac{a_2}{2^b} D)$ where at least one of a_1 and a_2 is an odd integer, and b is a non-negative integer. Then define $p(x) = (\frac{a_1 - \mathbb{I}(a_1 \text{ is odd})}{2^b} D, \frac{a_2 - \mathbb{I}(a_2 \text{ is odd})}{2^b} D)$. This construction implies that $\sum_x \|x - p(x)\|^2 \leq \sum_{b=0}^{k} 2^b \cdot 2^b \cdot \frac{2D^2}{4^b} \leq 2(k+1)D^2$. Thus, we have $\sum_x \|x - p(x)\|^2 \leq C_d D^2$ where $C_1 = 1$ and $C_2 = 2(k+1)$. It is also easy to see that for either $d = 1$ or $d = 2$, for any $x \in \mathcal{G}$, we have $p^k(x) = 0$, where $p^k(\cdot)$ denotes the k-fold application of p.

Now we are ready to define the linear transformation M of the loss vectors. For convenience, we index the loss vectors with $\mathcal{G}$ and interpret loss vectors as functions mapping $\mathcal{G}$ to $\mathbb{R}$. Then M maps the loss function $l : \mathcal{G} \to \mathbb{R}$ to the function $M(l) : \mathcal{G} \to \mathbb{R}$ defined as $M(l)(x) := l(x) - l(p(x))$ for all $x \in \mathcal{G}$. With some abuse of notation, we also use M to denote the matrix corresponding to the linear transformation M. We note a few desirable properties:

1. M is invertible on $\mathcal{L}$: to reconstruct $l(x)$ from $l' = M(l)$, the fact that $p^k(x) = 0$ and $l(0) = 0$ implies that
$$l(x) = l'(x) + l'(p(x)) + l'(p^2(x)) + \ldots + l'(p^{k-1}(x)). \tag{4.1}$$

2. We have
$$\|M(l)\| = \sqrt{\sum_{x \in G} (l(x) - l(p(x)))^2} \leq \sqrt{\sum_{x \in G} G^2 \|x - p(x)\|^2} \leq \sqrt{C_d} G D, \tag{4.2}$$
where the penultimate inequality follows from the G-Lipschitzness of l.

We can now apply the algorithm from Theorem 15 for online learning with experts using the entropy regularization. To apply Lemma 16, we need to compute a bound on σ'. Let $M^\dagger$ be the pseudoinverse for M defined by the (4.1). From (4.1), it is evident that
$$\|M^\dagger\|_{2 \to \infty} = \sqrt{k} = O(\sqrt{\log(T)}). \tag{4.3}$$

Using (4.2) and (4.3), we conclude that
$$\sigma' \leq GD\sqrt{C_d \log(T)} \cdot \frac{\log(T)\sqrt{\log(1/\delta)}}{\varepsilon}.$$

Using Lemma 16 and the fact that $\log(|\mathcal{G}|) = O(\log(T))$ in Theorem 15, we obtained the claimed bound. ∎

Acknowledgments

We thank Daniel Kifer and Adam Smith for discussions regarding their 2012 result on objective perturbation (joint with the last author of the present paper), and for pointing us to an improved analysis for GLMs. We thank Uri Sherman and Tomer Koren for discussing their work on lazy OCO with us. Finally, we express our indebtedness to Elad Hazan for early explorations into the feasibility of the present work. We also appreciate various careful comments from anonymous reviewers that aided us in making the writing clear and precise.

References

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

Jason M. Altschuler and Kunal Talwar. 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.

Oren Anava, Elad Hazan, and Shie Mannor. Online learning for adversaries with memory: price of past mistakes. In *Advances in Neural Information Processing Systems*, pages 784–792, 2015.

Hilal Asi, Daniel Asher Nathan Levy, and John Duchi. Adapting to function difficulty and growth conditions in private optimization. In *Advances in Neural Information Processing Systems*, 2021.

Hilal Asi, Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private online prediction from experts: Separations and faster rates. *arXiv preprint arXiv:2210.13537*, 2022.

Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Thakurta. Private stochastic convex optimization with optimal rates. In *Advances in Neural Information Processing Systems*, pages 11279–11288, 2019.

Raef Bassily, Vitaly Feldman, Cristóbal Guzmán, and Kunal Talwar. Stability of stochastic gradient descent on nonsmooth convex losses. *arXiv preprint arXiv:2006.06914*, 2020.

Vladimir Igorevich Bogachev and Maria Aparecida Soares Ruas. *Measure theory*, volume 1. Springer, 2007.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

Nicolò Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. In Thomas G. Dietterich, Suzanna Becker, and Zoubin Ghahramani, editors, *NeurIPS*, pages 359–366. MIT Press, 2001. URL https://proceedings.neurips.cc/paper/2001/hash/01931a6925d3de09e5f87419d9d55055-Abstract.html.

T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *ACM Trans. on Information Systems Security*, 14(3):26:1–26:24, November 2011.

Sinho Chewi. Log-concave sampling, 2023. URL https://chewisinho.github.io/main.pdf.

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, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In *Proc. of the Forty-Second ACM Symp. on Theory of Computing (STOC'10)*, pages 715–724, 2010.

Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: Optimal rates in linear time. In *Proc. of the Fifty-Second ACM Symp. on Theory of Computing (STOC'20)*, 2020.

Sascha Geulen, Berthold Vöcking, and Melanie Winkler. Regret minimization for online buffering problems using the weighted majority algorithm. In Adam Tauman Kalai and Mehryar Mohri, editors, *COLT*, pages 132–143. Omnipress, 2010. URL http://colt2010.haifa.il.ibm.com/papers/COLT2010proceedings.pdf#page=140.

Sivakanth Gopi, Yin Tat Lee, and Daogao Liu. Private convex optimization via exponential mechanism. *arXiv preprint arXiv:2203.00263*, 2022.

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

Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

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

Peter Kairouz, Brendan McMahan, Shuang Song, Om Thakkar, Abhradeep Thakurta, and Zheng Xu. Practical and private (deep) learning without sampling or shuffling. In *ICML*, 2021.

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

Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *Conference on Learning Theory*, pages 25–1, 2012.

Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Personal communication, 2023.

Janardhan Kulkarni, Yin Tat Lee, and Daogao Liu. Private non-smooth erm and sco in subquadratic steps. *Advances in Neural Information Processing Systems*, 34, 2021.

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

Neri Merhav, Erik Ordentlich, Gadiel Seroussi, and Marcelo J. Weinberger. 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.

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

Uri Sherman and Tomer Koren. Lazy oco: Online convex optimization on a switching budget. *arXiv preprint arXiv:2102.03803 version 5*, 2023. URL https://arxiv.org/pdf/2102.03803v5.pdf.

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*, pages 2733–2741, 2013.

Justin Whitehouse, Aaditya Ramdas, Ryan Rogers, and Zhiwei Steven Wu. Fully adaptive composition in differential privacy. *arXiv preprint arXiv:2203.05481*, 2022.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, pages 928–936, 2003.

Appendix A. Error in Sherman and Koren (2021)

In this section we provide a quick description of the error we discovered in Sherman and Koren (2021) which provides the previously best known bounds for the Lazy OCO problem. The core algorithm (Algorithm 2) proposed by Sherman and Koren (2021) couples the distribution of the noise employed at the next step via the LazySample procedure they describe. In particular in Line 6 of Algorithm 2 the next perturbation is chosen as

$$p_{t+1} = \text{LazySample}(p_t - \nabla f_t(w_t), \mathcal{N}(-\nabla f(w_t), \sigma_t^2), \mathcal{N}(0, \sigma_{t+1}^2))).$$

One of the core requirements of the guarantee(Lemma 1) provided for the LazySample procedure (Algorithm 1, (Sherman and Koren, 2021)) is that the first argument ($p_t - \nabla f_t(w_t)$) above should be drawn from the distribution supplied in the second argument ($\mathcal{N}(-\nabla f(w_t), \sigma_t^2)$). However as set in the paper w_t actually depends on p_t (in a potentially complicated manner) as

$$w_t = \text{argmin}(\sum_{i=1}^{t-1} f_i(w) + p_t^\top w + R(w)).$$

This correlation however breaks the requirement of LazySample as it is no longer clear that $p_t - \nabla f_t(w_t)$ is distributed as $\mathcal{N}(-\nabla f(w_t, \sigma_t^2))$. We note that this error was acknowledged by the authors in personal communication (and subsequently noted in a revision Sherman and Koren (2023)), and they acknowledged that they do not know of simple fix to the algorithm to make it correct.

Appendix B. Useful Results

In this section, we recall some standard results in differential privacy and online learning. The first of these standard results is the adaptive strong composition lemma for differentially private mechanisms.

Lemma 17 (e.g. 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 $x_{1:t-1} \in \mathcal{K}^{t-1}$, it holds that*

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

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

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

Next, we state the follow-the-leader be-the-leader lemma that is helpful in bounding the regret of an online learner as a sum of stability-related and regularization-related terms.

Lemma 18 (FTL-BTL Hazan (2016)) *For any loss function sequence $l_{0:T}$, define*

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

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

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

Appendix C. Analysis of Algorithm 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 $\zeta_t \triangleq \mathbb{I}(S_t = 0 \text{ or } S'_t = 0)$.

We restate and prove Lemma 5 first:

Lemma 19 (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,*

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

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

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

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, $\mathbb{E}[\mathcal{S}_T] \leq \tilde{p}T$ and the Chernoff bound[5] implies

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

■

The following key lemma obtains bounds on the actual distribution that x_t is sampled from in terms of μ_t:

Lemma 20 (Distribution drift) *Given $\delta \in [0, \frac{1}{2}]$ and $\Phi \geq 1$, suppose that for all $t \in [T]$, the distributions 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, $\|q_t - \mu_t\|_{TV} \leq 3\delta(t-1)$.*
- *If $B = 3\tilde{p}T$, then we have*

$$\|q_t - \mu_t\|_{TV} \leq e^{-\tilde{p}T} + 3\delta(t-1).$$

Proof We first consider the $B = \infty$ case. We prove that $\|q_t - \mu_t\|_{\text{TV}} \leq 3\delta(t-1)$ 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 $M = \{x \in \mathcal{K} \mid \Phi^{-1} \leq \frac{\mu_{t+1}(x)}{\mu_t(x)} \leq \Phi\}$. Then by Definition 3, we have $\mu_t(M) \geq 1 - \delta$ and $\mu_{t+1}(M) \geq 1 - \delta$. Next, let $\tilde{\mu}_t$ be the distribution of $X \sim \mu_t$ conditioned on the event $X \in M$. Since $\mu_t(M) \geq 1 - \delta$, it is easy to see that $\|\mu_t - \tilde{\mu}_t\|_{\text{TV}} \leq \delta$. 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 $\Pi(\frac{\mu_{t+1}(x)}{\Phi\mu_t(x)}) = \frac{\mu_{t+1}(x)}{\Phi\mu_t(x)}$, and that $\tilde{\mu}_t(x) = \frac{\mu_t(x)}{\mu_t(M)}$, we have

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

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 $\Pr[\sum_{t=1}^{T} X_t \geq (1+\delta)\mu] \leq e^{-\delta^2\mu/(2+\delta)}$.

Thus,

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

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

$$\|\tilde{q}_{t+1} - \mu_{t+1}\|_{\text{TV}} \le 2\delta.$$

Furthermore, we have

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

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

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

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 5 implies that both $q_{\text{all}}(\mathcal{E}), q'_{\text{all}}(\mathcal{E}) \le e^{-\tilde{p}T}$. Therefore we have that,

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

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

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

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

Next, we prove a general result on (Φ, δ)-closeness of distributions when the sampled noise is Gaussian:

Lemma 21 (Density ratio) *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be convex twice-differentiable functions such that $l - l'$ is G-Lipschitz and β-smooth. Let l, l' be Λ-strongly convex for some $\Lambda > 0$. Given any 2 positive numbers σ and σ', let μ, μ' be the probability density functions of $x(l, Z) \triangleq \operatorname{argmin}_{x \in \mathcal{K}} l(x) + \mathcal{B}(x) + Z^\top x$ and $x(l', Z') \triangleq \operatorname{argmin}_{x \in \mathcal{K}} l'(x) + \mathcal{B}(x) + (Z')^\top x$ respectively when $Z \sim \mathcal{N}(0, \sigma^2 I)$ and $Z' \sim \mathcal{N}(0, (\sigma')^2 I)$. Then for any $\delta \in (0, 1]$, we have that μ and μ' are (Φ, δ) close where*

$$\Phi = \exp\left(\frac{\beta d}{\Lambda} + \max\left\{\frac{G^2 + 2G\sigma\sqrt{2d\log(2/\delta)}}{2(\sigma')^2} + 2d\log(2/\delta)\left(\frac{|\sigma^2 - (\sigma')^2|}{(\sigma')^2}\right) + d\log(\sigma'/\sigma),\right.\right.$$
$$\left.\left.\frac{G^2 + 2G\sigma'\sqrt{2d\log(2/\delta)}}{2(\sigma)^2} + 2d\log(2/\delta)\left(\frac{|\sigma^2 - (\sigma')^2|}{(\sigma)^2}\right) + d\log(\sigma/\sigma')\right\}\right)$$

Proof We begin first by proving the direction

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

and reverse direction follows easily by switching the roles of μ, μ' through the analysis.

For the purpose of this proof, with some abuse of notation, we define $\mathcal{J} = l + \mathcal{B}$ and $\mathcal{J}' = l' + \mathcal{B}$. By Lemma 2, we have

$$\mu(x) = \nu(-\nabla\mathcal{J}(x)) \cdot |\det(-\nabla^2\mathcal{J}(x))| \quad \text{and} \quad \mu'(x) = \nu'(-\nabla\mathcal{J}'(x)) \cdot |\det(-\nabla^2\mathcal{J}'(x))|.$$

Therefore we have

$$\frac{\mu(x)}{\mu'(x)} = \underbrace{\frac{\nu(-\nabla\mathcal{J}(x))}{\nu'(-\nabla\mathcal{J}'(x))}}_{A} \cdot \underbrace{\frac{|\det(-\nabla^2\mathcal{J}(x))|}{|\det(-\nabla^2\mathcal{J}'(x))|}}_{B}.$$

The lemma is proved by bounding the A and B terms above separately. First, we bound B: Note that $\nabla^2\mathcal{J}'(x) = \nabla^2 l'(x) + \nabla^2\mathcal{B}(x) \succeq \Lambda I$, and $\nabla^2\mathcal{J}(x) - \nabla^2\mathcal{J}'(x) = \nabla^2(l - l')(x)$, and since $l - l'$ is β-smooth, we conclude that $\|\nabla^2\mathcal{J}(x) - \nabla^2\mathcal{J}'(x)\| \le \beta$, where the norm is the spectral norm. Thus, if $\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_d$ are the eigenvalues of $\nabla^2\mathcal{J}(x)$ arranged in non-increasing order, $\lambda'_1 \ge \lambda'_2 \ge \cdots \ge \lambda'_d$ are the eigenvalues of $\nabla^2\mathcal{J}'(x)$,then we have that $|\sum_i \lambda_i - \sum_i \lambda'_i| = |\mathrm{Tr}(\nabla^2\mathcal{J}(x)) - \mathrm{Tr}(\nabla^2\mathcal{J}'(x))| \le \beta d$ and $\lambda'_i \ge \lambda$ for all $i \in [d]$. Thus, we have

$$B = \prod_{i=1}^{d}\frac{\lambda_i}{\lambda'_i} = \prod_{i=1}^{d}(1 + (\lambda_i - \lambda'_i)/\lambda'_i) \le \prod_{i=1}^{d}(1 + (\lambda_i - \lambda'_i)/\Lambda) \le \exp((\textstyle\sum_i \lambda_i - \sum_i \lambda'_i)/\Lambda) \le \exp(\beta d/\Lambda).$$

The same argument above, applied to $\frac{1}{B}$, implies that $B \ge \exp(-\beta d/\Lambda)$. We now turn to term A. In this case, to sample $X \sim \mu$, we first sample $Z \sim \mathcal{N}(0, \sigma^2 I)$ and set $X = x(l, Z)$. Note that $X = x(l, Z)$ implies that $Z = -\nabla\mathcal{J}(X)$. This further implies that $-\nabla\mathcal{J}'(X) = Z + \nabla l(X) - \nabla l'(X)$. We define v to be $\nabla l(X) - \nabla l'(X)$ for the rest of the proof. Via standard tail bounds on norm of a Gaussian random vector, $\Pr[\|Z\|_2 \le \sigma\sqrt{2d\log(2/\delta)}] \ge 1 - \delta$ (Boucheron et al., 2013). We condition on this event for the rest of the proof.

$$A = \left(\frac{\sigma'}{\sigma}\right)^d \cdot \frac{\exp(-\|Z\|^2/2\sigma^2)}{\exp(-\|Z + v\|^2/2(\sigma')^2)} = \left(\frac{\sigma'}{\sigma}\right)^d \cdot \exp\left(\frac{\|v\|^2 + 2v^\top Z}{2(\sigma')^2} + \|Z\|^2\left(\frac{\sigma^2 - (\sigma')^2}{(\sigma')^2\sigma^2}\right)\right)$$
$$\le \left(\frac{\sigma'}{\sigma}\right)^d \cdot \exp\left(\frac{G^2 + 2G\sigma\sqrt{2d\log(2/\delta)}}{2(\sigma')^2} + 2d\log(2/\delta)\left(\frac{|\sigma^2 - (\sigma')^2|}{(\sigma')^2}\right)\right).$$

One can similarly show that

$$A \geq \left(\frac{\sigma'}{\sigma}\right)^d \cdot \exp\left(-\frac{2G\sigma\sqrt{2d\log(2/\delta)}}{2(\sigma')^2} - 2d\log(2/\delta)\left(\frac{|\sigma^2-(\sigma')^2|}{(\sigma')^2}\right)\right).$$

Putting all these bounds together, the proof of the lemma is complete. ■

Finally, we restate and prove Theorem 4 here:

Theorem 4 (Regret bound for CTRL) *In Algorithm 1, fix any $\eta, \sigma > 0$, any $\delta \in [0, 1/2]$, any $p \in [0, 1]$, set $\nu = \mathcal{N}(0, \sigma^2 I)$ and choose Φ such that for all t the distributions μ_t, μ_{t+1} are (Φ, δ)-close. For any sequence of obliviously chosen G-Lipschitz, β-smooth convex loss functions $l_{1:T}$, the following hold:*

- *If $B = \infty$,*
$$\mathcal{R}_T \leq \frac{D^2}{2\eta} + 2G^2\eta T + \sigma\sqrt{d}D + 6GD\delta T^2 + 2GD.$$
- *Let $\tilde{p} = p + 1 - \Phi^{-2}$. If $B = 3\tilde{p}T$,*
$$\mathcal{R}_T \leq \frac{D^2}{2\eta} + 2G^2\eta T + \sigma\sqrt{d}D + 2GDT(e^{-\tilde{p}T} + 3\delta T) + 2GD.$$

Proof Let $Z \sim \nu$ be an independently chosen random variable. Given Z, define a sequence of points $y_{1:T}$ as $y_t \triangleq x^\star(l_{1:t-1}, Z)$. Recall that we defined μ_t to be the distribution of $x^\star(l_{1:t-1}, Z)$. Let q_t be the distribution induced by Algorithm 1 on its iterates x_t. Lemma 20 establishes that the sequence of iterates x_t played by Algorithm 1 follows μ_t approximately. 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 20 appropriately. We leverage the following lemma,

Lemma 23 (Levin and 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 23 to pair $x_t \sim q_t$ and $y_t \sim \mu_t$, using Lemma 20, 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

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

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

Before proceeding to prove the regret bound, note that for an arbitrary point $x \in \mathcal{K}$, there exists a point $x^\circ \in \mathcal{K}^\circ$ (by the definition of the Minkowski set $\mathcal{K}^\circ$) such that $\|x^\circ - x\| \leq \frac{D}{T}$ and therefore

$$l_{1:T}(x^\circ) \leq l_{1:T}(x) + GD. \tag{C.3}$$

As a consequence of the above display, hereafter we will only focus on proving a regret bound against an arbitrary point $x^\circ \in \mathcal{K}^\circ$, with the bargain that we suffer an extra GD term in the true regret. Define $l_0(x) = \frac{\|x\|^2}{2\eta} + \mathcal{B}(x) + Z^\top x$. Now, by Lemma 18, we deterministically have that

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

Therefore, it holds that

$$\begin{aligned}\sum_{t=1}^{T} (l_t(y_t) - l_t(x^\circ)) &= \sum_{t=1}^{T} (l_t(y_t) - l_t(y_{t+1})) + \sum_{t=0}^{T} (l_t(y_{t+1}) - l_t(x^\circ)) + l_0(x^\circ) - l_0(y_1) \\ &\leq \sum_{t=1}^{T} (l_t(y_t) - l_t(y_{t+1})) + l_0(x^\circ) - l_0(y_1)\end{aligned}$$

The second term may be bounded in expectation as

$$\begin{aligned}\mathbb{E}[l_0(x^\circ) - l_0(y_1)] &= \mathbb{E}\left[Z^\top (x^\circ - y_1)\right] + \frac{1}{2\eta}\|x^\circ\|^2 + \mathcal{B}(x^\circ) - \frac{1}{2\eta}\|y_1\|^2 - \mathcal{B}(y_1) \\ &\leq D\mathbb{E}[\|Z\|] + \frac{D^2}{2\eta} + GD \\ &\leq \sigma\sqrt{d}D + \frac{D^2}{2\eta} + GD,\end{aligned}$$

where in the last equality we appeal to that fact that $\mathcal{B}(x^\circ) \leq GD$ since $x^\circ \in \mathcal{K}^\circ$. For bounding the first term, recall the definition $\mathcal{J}_t(l, x) = l(x) + \frac{\|x\|^2}{2\eta} + \mathcal{B}(x)$. For all $t \in [T]$, define $\tilde{\mathcal{J}}_t : \mathcal{K} \to \mathbb{R} \cup \{+\infty\}$ as

$$\tilde{\mathcal{J}}_t(x) = \mathcal{J}(l_{1:t-1}, x) + Z^\top x.$$

Note that since y_t minimizes $\tilde{\mathcal{J}}_t$ over $\mathcal{K}$, for any $x \in \mathcal{K}$ we have $\left\langle \nabla \tilde{\mathcal{J}}_t(y_t), x - y_t \right\rangle \geq 0$. Since $\mathcal{J}_t(x)$ is $\frac{1}{\eta}$-strongly-convex, we have

$$\begin{aligned}\tilde{\mathcal{J}}_{t+1}(y_{t+1}) &\geq \tilde{\mathcal{J}}_{t+1}(y_t) + \langle \nabla \tilde{\mathcal{J}}_{t+1}(y_t), y_{t+1} - y_t \rangle + \frac{1}{2\eta}\|y_t - y_{t+1}\|^2 \\ &\geq \tilde{\mathcal{J}}_{t+1}(y_t) + \underbrace{\langle \nabla \tilde{\mathcal{J}}_t(y_t), y_{t+1} - y_t \rangle}_{\geq 0} + \langle \nabla l_t(y_t), y_{t+1} - y_t \rangle + \frac{1}{2\eta}\|y_t - y_{t+1}\|^2 \\ &\geq \tilde{\mathcal{J}}_{t+1}(y_t) - \|\nabla l_t(y_t)\|\|y_{t+1} - y_t\| + \frac{1}{2\eta}\|y_t - y_{t+1}\|^2,\end{aligned}$$

where the last step follows from the Cauchy-Schwarz inequality. Since y_{t+1} minimizes $\tilde{\mathcal{J}}_{t+1}$ over $\mathcal{K}$, the inequality above implies that

$$\|y_t - y_{t+1}\| \leq 2\eta\|\nabla l_t(y_t)\| \leq 2\eta G.$$

To conclude the claim, it is sufficient to observe that $l_t(y_t) - l_t(y_{t+1}) \leq G\|y_t - y_{t+1}\| \leq 2\eta G^2$. The statements established above along with (C.3) imply that for any $x \in \mathcal{K}$

$$\mathbb{E}\left[\sum_{t=1}^{T} l_t(y_t) - \sum_{t=1}^{T} l_t(x)\right] \leq \frac{D^2}{2\eta} + 2G^2\eta T + \sigma\sqrt{d}D + 2GD. \tag{C.4}$$

Combining the above with (C.2) completes the proof of the theorem for $B = 3\tilde{p}T$. The proof for $B = \infty$ follows by repeating the argument for that case. ■

C.1. Density Ratio for GLMs

The following lemma proves a stronger bound (in terms of d) than Lemma 21 specifically for the GLM case. This lemma immediately implies Lemma 10:

Lemma 24 *Let $l, l' : \mathcal{K} \to \mathbb{R}$ be Λ-strongly convex twice-differentiable functions such that for all $x \in \mathcal{K}$, $l(x) - l'(x) = sf(\langle v, x\rangle)$ for some $\|v\| \le 1$, $s \in \{-1, 1\}$ and $f : [-D, D] \to \mathbb{R}$ is a G-Lipschitz and β-smooth convex function. Given $\sigma > 0$, let μ, μ' be the probability density functions of $x(l, Z) \triangleq \operatorname{argmin}_{x\in\mathcal{K}} l(x) + \mathcal{B}(x) + Z^\top x$ and $x(l', Z) \triangleq \operatorname{argmin}_{x\in\mathcal{K}} l'(x) + \mathcal{B}(x) + Z^\top x$ respectively when $Z \sim \mathcal{N}(0, \sigma^2 I)$. Then for any $\delta \in (0, 1]$, we have that μ and μ' are (Φ, δ) close where*

$$\Phi = \exp\left(\Lambda\beta + (G^2 + 2G\sigma\sqrt{2\log(2/\delta)})/2\sigma^2\right).$$

Proof Following the proof of Lemma 21, we note that for any $x \in \operatorname{int}(\mathcal{K})$, we have

$$\frac{\mu(x)}{\mu'(x)} = \underbrace{\frac{\nu(-\nabla\mathcal{J}(x))}{\nu(-\nabla\mathcal{J}'(x))}}_{A} \cdot \underbrace{\frac{|\det(-\nabla^2\mathcal{J}(x))|}{|\det(-\nabla^2\mathcal{J}'(x))|}}_{B}.$$

First, we bound B, as in Lemma 21. Note that $\nabla^2\mathcal{J}'(x) = \nabla^2 l'(x) \succeq \Lambda I$, and $\nabla^2\mathcal{J}(x) - \nabla^2\mathcal{J}'(x) = \nabla^2(l - l')(x) = sf''(\langle v, x\rangle)vv^\top$, and since $f''(\cdot) \le \beta$ and $\|v\| \le 1$, we conclude that $|\mathrm{Tr}(\nabla^2\mathcal{J}(x) - \mathcal{J}'(x))| \le \beta$. Thus, if $\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_d$ are the eigenvalues of $\nabla^2\mathcal{J}(x)$ arranged in non-increasing order, $\lambda'_1 \ge \lambda'_2 \ge \cdots \ge \lambda'_d$ are the eigenvalues of $\nabla^2\mathcal{J}'(x)$, then we have that $|\sum_i \lambda_i - \sum_i \lambda'_i| \le |\mathrm{Tr}(\nabla^2\mathcal{J}(x)) - \mathrm{Tr}(\nabla^2\mathcal{J}'(x))| \le \beta$ and $\lambda'_i \ge \Lambda$ for all $i \in [d]$. Thus, we have

$$B = \prod_{i=1}^{d} \frac{\lambda_i}{\lambda'_i} = \prod_{i=1}^{d}(1+(\lambda_i-\lambda'_i)/\lambda'_i) \le \prod_{i=1}^{d}(1+(\lambda_i-\lambda'_i)/\Lambda) \le \exp((\textstyle\sum_i \lambda_i - \sum_i \lambda'_i)/\Lambda) \le \exp(\beta/\Lambda).$$

The same argument above, applied to $\frac{1}{B}$, implies that $B \ge \exp(-\beta/\Lambda)$.

We now turn to term A. In this case, to sample $X \sim \mu$, we first sample $Z \sim \mathcal{N}(0, \sigma^2 I)$ and set $X = x(l, Z)$. Note that $X = x(l, Z)$ implies that $Z = -\nabla\mathcal{J}(X)$. This further implies that $-\nabla\mathcal{J}'(X) = Z - \nabla_X\left(sf(\langle v, X\rangle)\right) = Z - sf'(\langle v, X\rangle)v$. We have

$$A = \frac{\exp(-\|Z\|^2/2\sigma^2)}{\exp(-\|Z + sf'(\langle v, X\rangle)v\|^2/2\sigma^2)} = \exp((\|sf'(\langle v, X\rangle)v\|^2 + 2\langle sf'(\langle v, X\rangle)v, Z\rangle)/2\sigma^2).$$

Note that $\|sf'(\langle v, X\rangle)v\| \le G$. Further, we have

$$|\langle sf'(\langle v, X\rangle)v, Z\rangle)| \le |\langle sf'(\langle v, X\rangle)||\langle v, Z\rangle| \le G|\langle v, Z\rangle|.$$

Note that $\langle v, Z\rangle$ is a Gaussian distributed as $\mathcal{N}(0, \|v\|^2\sigma^2 I)$. Hence, with probability at least $1 - \delta$, we have $|\langle v, Z\rangle| \le \|v\|\sigma\sqrt{2\log(2/\delta)} \le \sigma\sqrt{2\log(2/\delta)}$. Putting all these bounds together, the proof of the lemma is complete. ■

Appendix D. 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 6:

Theorem 6 (Privacy) *Given $\sigma > 0$ and $\delta \in (0, 1/2]$, for any $T \geq 12 \log(1/\delta)$, let $\delta' = \frac{\delta T^{-2}}{60}$, $G' = 3G$, and $\beta' = 2\beta$. Suppose there exists $\Phi' > 0$ such that for all convex functions l, l' where $l - l'$ is G'-Lipschitz and β'-smooth, we have that, the distributions of $x^\star(l, Z)$ and $x^\star(l', Z)$ respectively when $Z \sim \mathcal{N}(0, \sigma^2 I)$, are (Φ', δ')-close. Define*

$$\varepsilon' = 7 \log^2(\Phi) T^{2/3} + 2 \log^3(\Phi) T + (2G^2/\sigma^2 + 2\eta^2\beta^2 d)^2 T^{5/3}.$$

Then for any sequence of G-Lipschitz, β-smooth convex functions, Algorithm 1 when run with $\Phi = \Phi'^2$, $p = T^{-1/3}$ and $B = 3\tilde{p}T$ is $(\varepsilon, \delta + 3Te^{-(1-\Phi^{-2})T})$-differentially private where

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

Proof 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 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 show the following claim,

Claim 26 *Let $\delta' \geq 0$ and Φ be as defined in Theorem 6. 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*

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

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

We intend to use adaptive strong composition for differential privacy (Lemma 17) with Claim 26 and to that end consider the following calculations

$$\begin{aligned} \sum_{t=1}^T \varepsilon_t^2 &\leq \frac{4\log^2(\Phi)}{p^2} + B\log^2(\Phi) + \frac{(2G^2/\sigma^2 + 2\eta^2\beta^2 d)^2}{p^2} T \\ &\leq 4\log^2(\Phi) T^{2/3} + (3T^{2/3} + 2\log(\Phi)T)\log^2(\Phi) + (2G^2/\sigma^2 + 2\eta^2\beta^2 d)^2 T^{5/3} \\ &\leq 7\log^2(\Phi) T^{2/3} + 2\log^3(\Phi) T + (2G^2/\sigma^2 + 2\eta^2\beta^2 d)^2 T^{5/3} \end{aligned}$$

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

Using the above calculations and applying Lemma 17 with $\delta' = \delta/2$ (in Lemma 17) concludes the proof. ■

Proof [Of Claim 26] We begin by defining a subset $\mathcal{E}_t \in \mathcal{K}$ for all t as

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

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 27 *Let Φ be as defined in Theorem 6, 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 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 $(\varepsilon_x, \varepsilon_y, \delta_x, \delta_y)$. 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 and 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.

Ante Case: $t \leq t_0$: 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

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

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.

At Case: $t = t_0$**:** For the *at case*, the last display in the *ante* 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 $\Sigma \in \mathbf{\Sigma_{t_0-1}}$,

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

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 $x \in \mathcal{E}_{t_0}$,

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

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 27 we get that conditioned on Σ_{t-1}

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

Post Case: $t > t_0$: 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 $x_t \sim \mu_t$ and $x'_t \sim \mu'_t$ in round t. Once again by applying the condition on Φ as stated in Theorem 6 we have that x_t, x'_t are $(\log\Phi, \delta')$-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. Using Lemma 2, we have

$$\frac{\mu_{t+1}(x)}{\mu_t(x)} = \underbrace{\frac{e^{-\|-\nabla\mathcal{J}(l_{1:t},x)\|^2/2\sigma^2}}{e^{-\|-\nabla\mathcal{J}(l_{1:t-1},x)\|^2/2\sigma^2}}}_{\triangleq A_t(x)} \underbrace{\frac{|\det(-\nabla^2\mathcal{J}(l_{1:t},x)|}{|\det(-\nabla^2\mathcal{J}(l_{1:t-1},x)|}}_{\triangleq B_t(x)}$$

We bound $A_t(x)$ and $B_t(x)$ separately. We now have that

$$\begin{aligned}
2\sigma^2 \log A_t(x) &= \|\nabla \mathcal{J}(l_{1:t-1}, x)\|^2 - \|\nabla \mathcal{J}(l_{1:t}, x)\|^2 \\
&= \|\nabla \mathcal{J}(l_{1:t-1}, x)\|^2 - \|\nabla \mathcal{J}(l_{1:t-1}, x) + \nabla l_t(x)\|^2 \\
&= 2\langle \nabla \mathcal{J}(l_{1:t-1}, x), \nabla l_t(x)\rangle - \|\nabla l_t(x)\|^2 \\
\log B_t(x) &= \log\det(\nabla^2 \mathcal{J}(l_{1:t}, x)) - \log\det(\nabla^2 \mathcal{J}(l_{1:t-1}, x)) \\
&= \log\det(\nabla^2 \mathcal{J}(l_{1:t-1}, x) + \nabla^2 l_t(x)) - \log\det(\nabla^2 \mathcal{J}(l_{1:t-1}, x)) \\
&\leq \langle (\nabla^2 \mathcal{J}(l_{1:t-1}, x))^{-1}, \nabla^2 l_t(x)\rangle \\
\log B_t(x) &\geq \langle (\nabla^2 \mathcal{J}(l_{1:t}, x))^{-1}, \nabla^2 l_t(x)\rangle.
\end{aligned}$$

The last two inequalities follow from the concavity of $\log\det(\cdot)$ and the fact that $\nabla \log\det(X) = X^{-1}$ (Boyd and Vandenberghe, 2004), which implies that for two positive definite symmetric matrices A, B, we have

$$\langle A^{-1}, A - B\rangle \leq \log\det(A) - \log\det(B) \rangle \langle B^{-1}, A - B\rangle,$$

where $\langle A, B\rangle := \mathrm{Tr}(A^\top B)$ denotes the Frobenius inner product. It now follows that

$$\begin{aligned}
\log \frac{\mu_{t+1}(x)}{\mu_t(x)} &\leq \frac{2\langle \nabla \mathcal{J}(l_{1:t-1}, x), \nabla l_t(x)\rangle + \|\nabla l_t(x)\|^2}{2\sigma^2} + \langle (\nabla^2 \mathcal{J}(l_{1:t-1}, x))^{-1}, \nabla^2 l_t(x)\rangle, \\
\log \frac{\mu_{t+1}(x)}{\mu_t(x)} &\geq \frac{2\langle \nabla \mathcal{J}(l_{1:t-1}, x), \nabla l_t(x)\rangle + \|\nabla l_t(x)\|^2}{2\sigma^2} + \langle (\nabla^2 \mathcal{J}(l_{1:t}, x))^{-1}, \nabla^2 l_t(x)\rangle.
\end{aligned}$$

Similarly for μ', one may establish

$$\begin{aligned}
\log \frac{\mu'_{t+1}(x)}{\mu'_t(x)} &\leq \frac{2\langle \nabla \mathcal{J}(l'_{1:t-1}, x), \nabla l'_t(x)\rangle + \|\nabla l'_t(x)\|^2}{2\sigma^2} + \langle (\nabla^2 \mathcal{J}(l'_{1:t-1}, x))^{-1}, \nabla^2 l'_t(x)\rangle, \\
\log \frac{\mu'_{t+1}(x)}{\mu'_t(x)} &\geq \frac{2\langle \nabla \mathcal{J}(l'_{1:t-1}, x), \nabla l'_t(x)\rangle + \|\nabla l'_t(x)\|^2}{2\sigma^2} + \langle (\nabla^2 \mathcal{J}(l'_{1:t}, x))^{-1}, \nabla^2 l'_t(x)\rangle.
\end{aligned}$$

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.

$$\begin{aligned}
&\log \frac{\frac{\mu'_{t+1}(x)}{\Phi\mu'_t(x)}}{\frac{\mu_{t+1}(x)}{\Phi\mu_t(x)}} \\
\leq &\frac{2\langle \nabla \mathcal{J}(l'_{1:t-1})(x) - \nabla \mathcal{J}(l_{1:t-1})(x), \nabla l_t(x)\rangle}{2\sigma^2} \\
&+ \langle (\nabla^2 \mathcal{J}(l'_{1:t-1})(x))^{-1} - (\nabla^2 \mathcal{J}(l_{1:t})(x))^{-1}, \nabla^2 l_t(x)\rangle \\
= &\frac{2\langle \nabla l'_{t_0}(x) - \nabla l_{t_0}(x)), \nabla l_t(x)\rangle}{2\sigma^2} \\
&+ \langle (\nabla^2 \mathcal{J}(l'_{1:t-1})(x))^{-1} - (\nabla^2 \mathcal{J}(l'_{1:t-1})(x) + \nabla^2 l_{t_0}(x) - \nabla^2 l'_{t_0}(x) + \nabla^2 l_t(x))^{-1}, \nabla^2 l_t(x)\rangle \\
\leq &\frac{2\langle \nabla l'_{t_0}(x) - \nabla l_{t_0}(x)), \nabla l_t(x)\rangle}{2\sigma^2} \\
&+ \eta^2 \langle (\nabla^2 l_{t_0}(x) - \nabla^2 l'_{t_0}(x) + \nabla^2 l_t(x)), \nabla^2 l_t(x)\rangle \\
\leq &\frac{2G^2}{\sigma^2} + 2\eta^2\beta^2 d,
\end{aligned}$$

where we use that $l_{1:t} - l'_{1:t-1} = l_t + l_{t_0} - l'_{t_0}$ is 2β-smooth, and additionally that for arbitrary matrices $X \succ \frac{1}{\eta} I$ and $Y, \Delta \succeq 0$, it is true that

$$\langle Y, X^{-1} - (X+\Delta)^{-1} \rangle \leq \langle Y, X^{-1} \Delta X^{-1} \rangle \leq \eta^2 \langle Y, \Delta \rangle.$$

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

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

Now, for the remaining possibility, we have

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

The above displays thereby imply that conditioned on Σ_{t-1} and the event $x_t \in \mathcal{E}_t$, we have that ζ_t and ζ'_t are $(\frac{(2G^2/\sigma^2 + 2\eta^2\beta^2 d)}{p}, 0)$-indistinguishable. Thereby combining with Claim 27 we get that conditioned on Σ_{t-1}

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

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

■

Proof [Of Claim 27] 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 20, we have that $\|\mu_t - q_t\| \leq e^{-\tilde{p}T} + 3T\delta'$ 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,

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

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

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

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

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

The above statements imply the claim. ■

Appendix E. Improved Analysis for Strongly Convex functions

For strongly convex functions, we need to change Algorithm 1 to use changing noise distributions ν_t in different rounds, which also necessitates using changing scaling parameter Φ_t. The pseudocode is given in Algorithm 2.

Algorithm 2: Couple-The-Regularized-Leader (CTRL)

Inputs: A sequence of distributions $\nu_{0:T}$ on $\mathbb{R}^d$, a regularization parameter $\eta > 0$, a barrier $\mathcal{B}(x)$, a sequence of scaling parameters $\Phi_t \ge 0$.

Set $b_1 = 0$, sample $Z_0 \sim \nu_0$, choose $x_1 = x^\star(0, Z_0)$.

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 $l_t(x_t)$.

  Sample $S_t \sim \text{Ber}\left(\min\left\{1, \max\left\{\frac{1}{\Phi_t^2}, \frac{\mu_{t+1}(x_t)}{\Phi_t \cdot \mu_t(x_t)}\right\}\right\}\right)$.

  // $\mu_t(\cdot), \mu_{t+1}(\cdot)$ can be computed incrementally via Lemma 2.

  if $S_t = 0$ **then**

    Choose $x_{t+1} = x^\star(l_{1:t}, Z_t)$.

  end

  else

    Set $x_{t+1} = x_t$.

  end

end

Since we have changing values of Φ_t in the different rounds, in analogy with $\tilde{p}$, we define $\tilde{p}_t := 1 - \Phi_t^{-2}$. In the following analysis, we will use analogs of Lemma 5 and Lemma 20 with p set to 0 and with the $\tilde{p}T$ terms replaced by $\sum_{t=1}^{T} \tilde{p}_t$. The proofs are identical and are hence omitted. We can now present the following tighter regret bound for strongly-convex loss functions.

Theorem 28 (Regret bound for CTRL with Strongly-Convex losses) *In Algorithm 2, fix any $\eta > 0$ and $\delta \in (0, 1/2]$. Choose any $\sigma > 0$, and for all $t \ge 0$ set $\sigma_t = \sigma\sqrt{t}$ and $\nu_t = \mathcal{N}(0, \sigma_t^2 I)$; choose Φ_t such that the distributions μ_t, μ_{t+1} are (Φ_t, δ)-close. Set $\eta = \infty$. Then for any sequence of obliviously chosen G-Lipschitz, β-smooth, λ-strongly convex loss functions $l_{1:T}$, the following hold:*

$$\mathcal{R}_T \le \frac{2(G^2 + d\sigma^2)(1 + \log(T))}{\lambda} + 6GD\delta T^2 + 2GD.$$

Proof For the regret bound we follow the argument laid out in Sherman and Koren (2021) in the strongly convex case. For the purpose of analysis, define an auxiliary sequence of random variables $Z'_{0:T}$ as follows. Set $Z'_0 = 0$ deterministically. Then sample $Z'_1 \sim \nu_1$. Then for all $t \ge 2$ set

$Z'_t = Z'_1\sqrt{t}$. Note that marginally for all $t \geq 0$, Z'_t is distributed as ν_t. Define a sequence of random variables $y_{1:T}$ as $y_t \triangleq x^\star(l_{1:t-1}, Z'_{t-1})$. Recall that since we defined μ_t to be the distribution of $x^\star(l_{1:t-1}, Z)$ where Z is sampled from ν_{t-1}, we have that y_t are also marginally distributed as μ_t. Let q_t be the marginal distribution induced by Algorithm 2 on its iterates x_t. Lemma 20 establishes that the sequence of iterates x_t played by Algorithm 2 follows μ_t approximately. As in the proof of Theorem 4 we apply Lemma 23 to the the distribution pair $x_t \sim q_t$ and $y_t \sim \mu_t$, using Lemma 20, 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

$$\left|\mathbb{E}\left[\sum_{t=1}^{T}(l_t(x_t) - l_t(y_t))\right]\right| \leq \sum_{t=1}^{T}\left|\mathbb{E}\left[l_t(x_t) - l_t(y_t)\right]\right| \leq 6GD\delta T^2, \tag{E.1}$$

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

Before proceeding to prove the regret bound, note that for an arbitrary point $x \in \mathcal{K}$, there exists a point $x^\circ \in \mathcal{K}^\circ$ (by the definition of the Minkowski set) such that $\|x^\circ - x\| \leq \frac{D}{T}$ and therefore

$$l_{1:T}(x^\circ) \leq l_{1:T}(x) + GD. \tag{E.2}$$

As a consequence of the above display, hereafter we will only focus on proving a regret bound against an arbitrary point $x^\circ \in \mathcal{K}^\circ$, with the bargain that we suffer an extra GD term in the true regret. Define auxiliary functions for all $t \geq 1$

$$l'_t(x) \triangleq l_t(x) + (Z'_t - Z'_{t-1})^\top x,$$

with $l'_0(x) \triangleq \mathcal{B}(x)$ and $Z'_0 \triangleq 0$. It is now immediate to see that

$$y_t \triangleq \operatorname*{argmin}_{x\in\mathcal{K}} \sum_{\tau=1}^{t-1} l(x) + (Z'_{t-1})^\top x + \mathcal{B}(x) = \operatorname*{argmin}_{x\in\mathcal{K}} \sum_{\tau=0}^{t-1} l'_\tau(x).$$

Therefore, by Lemma 18, we deterministically have that

$$\sum_{t=0}^{T}\left(l'_t(y_{t+1}) - l'_t(x^\circ)\right) \leq 0.$$

Further we have that

$$\sum_{t=0}^{T}\left(l'_t(y_{t+1}) - l'_t(x^\circ)\right) = \sum_{t=1}^{T}(l_t(y_{t+1}) - l_t(x^\circ)) + \sum_{t=1}^{T}(Z'_t - Z'_{t-1})^\top(y_{t+1} - x^\circ) + \mathcal{B}(y_1) - \mathcal{B}(x^\circ)$$

which implies that the following holds deterministically

$$\sum_{t=1}^{T}(l_t(y_{t+1}) - l_t(x^\circ)) \leq \sum_{t=1}^{T}(Z'_t - Z'_{t-1})^\top(x^\circ - y_{t+1}) + \mathcal{B}(x^\circ) \leq \sum_{t=1}^{T}(Z'_t - Z'_{t-1})^\top(x^\circ - y_{t+1}) + GD,$$

since $\mathcal{B}(x^\circ) \leq GD$ as $x^\circ \in \mathcal{K}^\circ$. Further define the deterministic sequence $y^*_t = \operatorname{argmin}_{x\in\mathcal{K}} \sum_{\tau=1}^{t-1} l_\tau(x) + \mathcal{B}(x)$. Using a simple perturbation bound over the minima of strongly convex functions it is easy to

observe that $\|y_t - y_t^*\| \leq \frac{2\|Z'_{t-1}\|}{(t-1)\lambda}$. Using this fact we get that

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T}(Z'_t - Z'_{t-1})^\top(x^\circ - y_{t+1})\right] &= \underbrace{\mathbb{E}\left[\sum_{t=1}^{T}(Z'_t - Z'_{t-1})^\top(x^\circ - y^*_{t+1})\right]}_{=0} \\
&\quad + \mathbb{E}\left[\sum_{t=1}^{T}(Z'_t - Z'_{t-1})^\top(y^*_{t+1} - y_{t+1})\right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^{T}\|(Z'_t - Z'_{t-1})\|\|y^*_{t+1} - y_{t+1}\|\right] \\
&\leq \frac{2\mathbb{E}[\|Z'_1\|^2]}{\lambda}\frac{\sqrt{t} - \sqrt{t-1}}{\sqrt{t}}
\end{aligned}$$

Therefore combining the above equations we get that

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T}(l_t(y_{t+1}) - l_t(x^\circ))\right] &\leq \sum_{t=1}^{T}\left(\frac{2\mathbb{E}[\|Z'_1\|^2]}{\lambda}\frac{\sqrt{t} - \sqrt{t-1}}{\sqrt{t}}\right) + GD \\
&\leq \frac{2d\sigma^2(1+\log(T))}{\lambda} + GD. \qquad \text{(E.3)}
\end{aligned}$$

Now furthermore we have that

$$\sum_{t=1}^{T}(l_t(y_t) - l_t(x^\circ)) = \sum_{t=1}^{T}(l_t(y_t) - l_t(y_{t+1})) + \sum_{t=1}^{T}(l_t(y_{t+1}) - l_t(x^\circ)). \qquad \text{(E.4)}$$

For bounding the first term, consider the definition

$$\mathcal{J}_t(x) = \sum_{i=1}^{t-1} l_i(x) + (Z'_{t-1})^\top x + \mathcal{B}(x),$$

and note that y_t minimizes $\mathcal{J}_t$ over $\mathcal{K}$. Further note that $\mathcal{J}_t(x)$ is $(t-1)\lambda$-strongly convex. Now using the same analysis as in the proof of Theorem 4, we get that $l_t(y_t) - l_t(y_{t+1}) \leq G\|y_t - y_{t+1}\| \leq \frac{2G^2}{t\lambda}$ and therefore we get that

$$\sum_{t=1}^{T} l_t(y_t) - l_t(y_{t+1}) \leq \frac{2G^2(1+\log(T))}{\lambda}$$

Combining the above with (E.2), (E.3) and (E.4) we get that

$$\mathbb{E}\left[\sum_{t=1}^{T} l_t(y_t) - \sum_{t=1}^{T} l_t(x)\right] \leq \frac{2(G^2 + d\sigma^2)(1+\log(T))}{\lambda} + 2GD. \qquad \text{(E.5)}$$

Combining the above with (E.1) completes the proof of the theorem. ■

We restate and prove Theorem 13 now:

Theorem 13 (Lazy OCO - Strongly-Convex) *For any $T \geq 3$ and any given bound on the number of switches $S \geq \left(\frac{4\beta d}{\lambda} + 16d\log(T)\right)\log(T)$, set $\delta = 2/T^2$ and $\sigma = \frac{16G\sqrt{dT\log(T)}}{S}$. For all $t \geq 0$, set $\sigma_t = \sigma\sqrt{t}$, $\nu_t = \mathcal{N}(0, \sigma_t^2 I)$,*

$$\Phi_t = \frac{\beta d}{\lambda t} + \frac{4d\log(T)}{t} + d\log\left(\frac{\sqrt{t+1}}{\sqrt{t}}\right) + \frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t},$$

and $\eta = \infty$ in Algorithm 2. Then for any sequence of obliviously chosen G-Lipschitz β-smooth λ-strongly convex functions $l_{1:T}$, Algorithm 2 satisfies the following:

$$\mathcal{R}_T \leq \frac{2G^2(1+\log(T))}{\lambda} + \frac{512d^2\log(T)(1+\log(T))}{\lambda}\cdot\frac{T}{S^2} + 14GD \text{ and } \mathbb{E}[\mathcal{S}_T] \leq S.$$

Proof For any σ setting $\sigma_t = \sigma\sqrt{T}$ and setting $\delta = 2T^{-2}$ using 21 we can see that the distributions μ_t, μ_{t+1} are $(\Phi_t, 2T^{-2})$-close as long as we have that

$$\begin{aligned}
\Phi_t &\geq \frac{\beta d}{\lambda t} + \frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t} + \frac{4d\log(T)}{t} + d\log\left(\frac{\sqrt{t+1}}{\sqrt{t}}\right) \\
&= \frac{\beta d}{\lambda t} + \frac{4d\log(T)}{t} + d\log\left(\frac{\sqrt{t+1}}{\sqrt{t}}\right) + \frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t}.
\end{aligned}$$

As mentioned earlier, the appropriate analog of Lemma 5 bounds the expected number of switches by $\sum_{t=1}^T \tilde{p}_t$. Using the fact that $1 - \exp(-2x) \leq 2x$ for all $x \in \mathbb{R}$, we have

$$\begin{aligned}
\sum_{t=1}^T \tilde{p}_t &= \sum_{t=1}^T 1 - \Phi_t^{-2} \\
&\leq \sum_{t=1}^T \left(\frac{\beta d}{\lambda t} + \frac{4d\log(T)}{t} + d\log\left(\frac{\sqrt{t+1}}{\sqrt{t}}\right) + \frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t}\right) \\
&\leq \left(\frac{\beta d}{\lambda} + d\log(T)\right)(1+\log(T)) + 4d\log(T) + \frac{G^2}{2\sigma^2 t} + \sum_{t=1}^T \left(\frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t}\right) \\
&\leq \left(\frac{\beta d}{\lambda} + 4d\log(T)\right)(1.25+\log(T)) + \sum_{t=1}^T \left(\frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t}\right) \\
&\leq \left(\frac{\beta d}{\lambda} + 4d\log(T)\right)(1.25+\log(T)) + \sum_{t=1}^T \left(\frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t}\right) \\
&\leq \left(\frac{2\beta d}{\lambda} + 8d\log(T)\right)\log(T) + \sum_{t=1}^T \left(\frac{G^2}{2\sigma^2 t} + \frac{2G\sqrt{d\log(T)}}{\sigma}\frac{\sqrt{t+1}}{t}\right) \\
&\leq \left(\frac{2\beta d}{\lambda} + 8d\log(T)\right)\log(T) + \frac{G^2(1+\log(T))}{2\sigma^2} + \frac{4G\sqrt{d\log(T)}\sqrt{T}}{\sigma}
\end{aligned}$$

Now using the value of σ and the bound on S provided in the theorem we see the following

$$\underbrace{\left(\frac{2\beta d}{\lambda} + 8d\log(T)\right)\log(T)}_{\leq S/2} + \underbrace{\frac{G^2(1+\log(T))}{2\sigma^2}}_{\leq S/4} + \underbrace{\frac{4G\sqrt{d\log(T)}\sqrt{T}}{\sigma}}_{\leq S/4} \leq S.$$

The regret bound follows via a direct substitution in Theorem 28, which immediately yields the stated bound. ■

Appendix F. Analysis for Online Lipschitz Optimization in One or Two Dimensions

We restate and prove Lemma 16:

Lemma 30 *Let $M \in \mathbb{R}^{d\times d}$ be any matrix that is invertible on $\mathcal{L}$ with pseudoinverse $M^\dagger$. In Theorem 15, σ can be replaced by*

$$\sigma' = \|M^\dagger\|_{2\to\infty} \cdot \max_{l\in\mathcal{L}}\{\|Ml\|\} \cdot \frac{\log(T)\sqrt{\log(1/\delta)}}{\varepsilon}.$$

Proof The idea is simple and is based on the algorithm in Agarwal and Singh (2017), which works as follows. It uses tree-based aggregation algorithm with Gaussian noise to maintain the prefix sums $l_1 + l_2 + \ldots + l_t$ for $t \in [T]$ with (ε, δ)-differential privacy. The standard Hedge algorithm, which is FTRL using the entropy regularizer, only needs these prefix sums to operate. The private prefix sums are fed into the Hedge algorithm to compute the predictions in each round, which are private due to post-processing. This noise is drawn from $\mathcal{N}(0, \sigma^2 I_{|\mathcal{K}|})$, and the excess regret due to this noise is bounded by (see the D_Z term in Theorem 3.1 of Agarwal and Singh (2017))

$$2\mathbb{E}_{Z\sim\mathcal{N}(0,\sigma^2 I_{|\mathcal{K}|})}\left[\max_{x\in\mathcal{K}} Z(x)\right],$$

which in turn is bounded by $O(\sigma\sqrt{\log(|\mathcal{K}|)})$.

We use the algorithm as follows. The matrix M defines a linear transform mapping $\mathbb{R}^d$ to $\mathbb{R}^d$. So we can use the tree-based aggregation algorithm to maintain the transformed prefix sums $Ml_1 + Ml_2 + \ldots + Ml_t$ for $t \in [T]$ with (ε, δ)-differential privacy. To maintain these sums privately, the tree-based aggregation algorithm adds Gaussian noise that scales as $\max_{l\in\mathcal{L}}\{\|Ml\|\}$.

Before feeding the private linear-transformed prefix sums to the FTRL algorithm, we "undo" the linear transformation by multiplying by $M^\dagger$. Specifically, the output of the tree-based aggregation algorithm in round t is $Ml_1 + Ml_2 + \ldots + Ml_{t-1} + Z$ where $Z \sim \mathcal{N}(0, \tilde{\sigma}^2 I_{|\mathcal{K}|})$, where $\tilde{\sigma}^2 = \max_{l\in\mathcal{L}}\{\|Ml\|\} \cdot \frac{\log(T)\sqrt{\log(1/\delta)}}{\varepsilon}$. We multiply this vector by $M^\dagger$ before feeding it into the Hedge algorithm. Thus, the vector fed into Hedge is $l_1 + l_2 + \ldots + l_{t-1} + M^\dagger Z$. Note that $M^\dagger Z$ is another Gaussian vector whose covariance matrix is $\tilde{\sigma}^2 M^\dagger {M^\dagger}^\top$. Working through the same analysis as done by Agarwal and Singh (2017), the excess regret due to the noise is bounded by

$$2\mathbb{E}_{Z\sim\mathcal{N}(0,\tilde{\sigma}^2 I_{|\mathcal{K}|})}\left[\max_{x\in\mathcal{K}}(M^\dagger Z)(x)\right] = 2\mathbb{E}_{Z'\sim\mathcal{N}(0,\tilde{\sigma}^2 M^\dagger {M^\dagger}^\top)}\left[\max_{x\in\mathcal{K}} Z'(x)\right].$$

The diagonal entries of $\tilde{\sigma}^2 M^\dagger {M^\dagger}^\top$ are bounded by $\tilde{\sigma}^2 \|M^\dagger\|_{2\to\infty}^2 = {\sigma'}^2$. Thus, via standard bounds (e.g., Theorem 2.5 in Boucheron et al. (2013)) on the expected maximum of Gaussians, the above expression is bounded by $O(\sigma' \sqrt{\log(|\mathcal{K}|)})$. ■