Partially Performative Prediction

Jaewook Lee
Stanford University
jwl99@stanford.edu

Tijana Zrnic
Stanford University
tijana.zrnic@stanford.edu

Abstract

Performative prediction studies feedback loops that arise when predictive models are deployed in consequential domains. In these settings, deploying a model can change the population whose patterns the model aims to predict, inducing a distribution shift that is *endogenous* to the learning system. This perspective departs from classical treatments of distribution shift, where shifts are typically modeled as *exogenous* changes in the data-generating process. Yet, in practice, distribution shift is rarely one or the other. Predictive models may influence future data through the decisions they support, while the world itself continues to drift for reasons beyond the learner's control. We study *partially performative prediction*, a framework that captures both endogenous and exogenous sources of distribution shift. The framework generalizes performative prediction by allowing the data distribution to evolve both in response to the deployed model and according to an external, time-varying process. We extend the central notions of performative stability and performative optimality to this setting by defining their online analogues that track the evolving partially performative environment. We analyze practical learning heuristics, including repeated retraining, and characterize when they successfully adapt to partially performative environments.

1 Introduction

When prediction informs decision-making, predictive models do not merely describe a population; they can also shape it. Disease forecasts guide interventions that slow disease spread; recommendation systems influence user preferences; stock and home price forecasts affect the prices that ultimately materialize. Performative prediction [23] formalizes this phenomenon by studying learning problems in which deploying a model changes the distribution of future observations.

The core idea is that deploying a model θ induces a model-dependent distribution over observations, denoted by $\mathcal{D}(\theta)$. The learner's objective is to minimize the expected loss under the induced distribution, called the *performative risk*:

$$\mathrm{PR}(\theta) = \mathbb{E}_{Z\sim\mathcal{D}(\theta)}[\ell(Z,\theta)], \tag{1}$$

where Z denotes an observation, typically a feature–label pair $Z = (X, Y)$. The central object $\mathcal{D}(\theta)$ is called the *distribution map*. Thus, the goal is not simply to choose a model that performs well on the population as currently observed, but one that performs well after the model's impact on the population has been realized. Existing work studies two central solution concepts: *performative stability* and *performative optimality*, and analyzes when retraining procedures converge to such equilibria. In particular, these procedures generate a sequence of model updates $\theta_1, \theta_2, \ldots$, whose limit is either a performatively stable or a performatively optimal point.

Performativity, therefore, surfaces as distribution shift: each model update may induce a new data distribution. What distinguishes performative distribution shift is that it is *endogenous*: it arises from the learner's own actions. This contrasts with classical models of distribution shift in machine

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

learning, which typically treat the shift as driven by *exogenous* factors. For example, online learning studies algorithms for adapting a sequence of actions θ_t to an arbitrary sequence of distributions P_t.

The starting point of this work is the observation that, in realistic performative environments, distribution shift is rarely purely endogenous. Deployed predictive systems are embedded in changing populations, markets, and platforms. Their predictions may influence future data, but so too may seasonal effects, policy changes, market forces, and strategic behavior outside the learner's control.

To capture this interplay, we introduce *partially performative prediction*, an online model that combines endogenous and exogenous sources of distribution shift. At round t, the data distribution is

$$\mathcal{D}_t(\theta) = (1 - \alpha_t)\mathcal{D}(\theta) + \alpha_t P_t,$$

where $\mathcal{D}(\theta)$ captures the performative response to the deployed model θ, P_t captures an exogenous distribution shift, and $\alpha_t \in [0, 1]$ controls their relative strength. This framework strictly generalizes the classical performative setting, recovered when $\alpha_t = 0$, and approaches a purely exogenous-shift setting when $\alpha_t = 1$. Note also that this model allows for a fixed non-performative component, obtained when P_t is constant across t.

This partially performative model requires a reconsideration of the learner's objective. Since the environment is fundamentally non-stationary, the goal can no longer simply be convergence to a fixed model, as in performative stability or performative optimality. Instead, the learner must contend with a moving target whose moves are only partly induced by the learner's own decisions. At the same time, the model highlights several new questions about the interaction between endogenous and exogenous shifts. Is a smaller value of α_t always more favorable to the learner, or can a large exogenous component help? How does the rate of change of P_t affect the performance of learning algorithms? And how do the conditions under which retraining heuristics converge, or fail to converge, change relative to the classical performative setting ($\alpha_t = 0$)?

1.1 Our contributions

We study all the above questions and more in the partially performative setting. Our first contribution is to extend the classical solution concepts of performative prediction to non-stationary environments. In the standard setting, performative stability and performative optimality are fixed models. Under partial performativity, the relevant stable and optimal points may change over time. We therefore introduce regret-based analogues, which we call *performative stability regret* and *performative optimality regret*. These quantities compare the learner's sequence of decisions to the time-varying sequence of performatively stable points $\theta_1^{\mathrm{PS}}, \theta_2^{\mathrm{PS}}, \dots$ and performatively optimal points $\theta_1^{\mathrm{PO}}, \theta_2^{\mathrm{PO}}, \dots$, respectively. We then analyze several natural learning algorithms, including repeated risk minimization [23] and gradient-based methods, under these regret notions.

A key quantity in our bounds is the *path length*: the cumulative movement of the comparator sequence. For performative stability and performative optimality, we denote the one-step changes by

$$\Delta_t^{\mathrm{PS}} = \|\theta_t^{\mathrm{PS}} - \theta_{t+1}^{\mathrm{PS}}\| \quad \text{and} \quad \Delta_t^{\mathrm{PO}} = \|\theta_t^{\mathrm{PO}} - \theta_{t+1}^{\mathrm{PO}}\|.$$

The corresponding path lengths $\sum_t \Delta_t^{(\cdot)}$ measure the degree of non-stationarity in the partially performative problem. Slowly varying exogenous distributions P_t lead to small path length, while rapidly changing P_t can lead to large path length.

Our first main result bounds the stability regret of repeated risk minimization (RRM), an idealized form of retraining given by

$$\theta_{t+1} = \arg\min_{\theta} \mathbb{E}_{Z \sim \mathcal{D}_t(\theta_t)}[\ell(Z, \theta)].$$

Let $\mathcal{A}_{\mathrm{RRM}}$ denote the RRM algorithm.

Theorem 1.1 (Informal)**.** *Suppose that the loss $\ell(z, \theta)$ is β_z-smooth in z and μ-strongly convex in θ, and that the distribution map $\mathcal{D}$ is ϵ-Lispchitz with $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$, where $\alpha_{\min} = \min_{t \in [T]} \alpha_t$. Then, the performative stability regret of repeated risk minimization after T steps satisfies*

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RRM}}) = \mathcal{O}\left(1 + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PS}}\right). \tag{2}$$

This result makes explicit a tradeoff between endogenous and exogenous shifts. When the shifts are mostly performative ($\alpha_t \approx 0, \forall\, t$), the regret has no dependence on T because the path length vanishes. The condition on performativity reduces to the standard condition $\epsilon < \frac{\mu}{\beta_z}$ from [23]. Conversely, when the shifts are mostly exogenous ($\alpha_t \approx 1, \forall\, t$), there is no condition on the strength of performativity, but the regret depends on the non-stationarity of P_t. If P_t is near-constant, the regret is constant too; if P_t fluctuates rapidly, the regret could be large. Thus, retraining on the most recent data is effective under mild distribution shift, but ineffective when shifts are unpredictable.

We prove analogous stability regret bounds for gradient-based algorithms, including those with "greedy" and "lazy" deployment schedules [17].

Our second main contribution concerns performative optimality. We study algorithms that explicitly anticipate performativity by using zeroth-order gradient estimation techniques. Such methods perturb deployed models to estimate how changes in the model affect the induced distribution, and prior work has shown that related procedures can converge to performatively optimal points in the classical setting [18, 9, 16, 15, 3]. We show that, in partially performative environments, these ideas yield regret guarantees relative to the time-varying sequence of performatively optimal points.

Theorem 1.2 (Informal)**.** *Suppose that the performative risk at step t is convex for all t. Then, there exists an algorithm $\mathcal{A}$ that after T steps achieves performative optimality regret of*

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}) = \mathcal{O}\left(d\sqrt{T} \cdot \left(1 + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PO}}\right)\right),$$

where d is the dimension of θ.

Again, we see path length as a key quantity in the bound. Like in the previous result, we can again observe a tradeoff between having a large endogenous and a large exogenous component. A large exogenous component might make the path length large. At the same time, a larger endogenous component can make the performative risk nonconvex, even when the loss $\ell(z, \theta)$ is convex. If $\alpha_t \approx 1$, the performative component is weak; consequently, convexity of the loss directly implies convexity of the performative risk. Thus, exogenous variation can make the target harder to track, while endogenous feedback can make the target harder to optimize.

1.2 Related work

Our work contributes to the growing literature on performative prediction [23], which formalizes settings in which deploying a model changes the future data distribution. A large body of work studies algorithms for finding performatively stable points [17, 4, 19, 27, 21, 24, 5] and performatively optimal points [18, 9, 13, 21, 28, 15], under progressively weaker assumptions.

Most relevant to our work are extensions with stateful observation models [2, 25, 10, 14, 16], which allow prior deployments to influence the current data distribution. These models capture memory, delayed responses, and other forms of endogenous dynamics. However, they still pursue the classical notions of performative stability and performative optimality; the underlying environment is fully determined by the history of deployments. In contrast, our model allows exogenous components of the data-generating process to evolve over time, so the relevant stable and optimal points may drift.

Another closely related thread studies performative regret minimization [12, 3, 29, 22], where the learner deploys models sequentially and seeks regret bounds on the accumulated performative risk. In these works, the performative risk is static: the distribution map $\mathcal{D}(\theta)$ is fixed, and regret is measured against a fixed baseline. In our setting, the distribution map $\mathcal{D}_t(\theta)$ may vary with time. Accordingly, we study dynamic notions of performative regret, benchmarking the learner against time-varying sequences of performatively stable or performatively optimal models.

Our analysis builds on tools from online learning, in particular online convex optimization [34, 7] and zeroth-order optimization [6, 1]. We study dynamic regret, which measures performance relative to a changing sequence of optima rather than a single best baseline in hindsight. Dynamic regret originated in tracking changing experts and comparators [8, 34] and was later sharpened through path and loss variation [11, 20, 30, 32]. Some of our results rely on zeroth-order optimization, where path-length-dependent dynamic regret guarantees are known [33], though existing analyses do not accommodate performative feedback.

2 Preliminaries

We begin by reviewing the formal setup of performative prediction, with a focus on the partially performative setting studied in this work. The learner chooses model parameters $\theta \in \Theta$, where Θ is a convex, compact subset of $\mathbb{R}^d$ of diameter D_Θ. We write Π_Θ for the Euclidean projection onto Θ. Performance is measured by a loss function $\ell(z, \theta)$, where z denotes a data point; thus, $\ell(z, \theta)$ is the loss incurred by model θ on instance z.

In partially performative prediction, the data distribution after deploying model θ at time t is

$$\mathcal{D}_t(\theta) = (1 - \alpha_t)\mathcal{D}(\theta) + \alpha_t P_t, \tag{3}$$

where $\alpha_t \in [0, 1]$ is a fixed sequence of mixture coefficients and P_t is an exogenous distribution that is not affected by the deployed model. The component $\mathcal{D}(\theta)$ captures the performative response to model θ, while P_t captures time-varying external shift. Setting $\alpha_t = 0$ recovers the standard performative prediction model [23]. When the deployed model at time t is θ_t, we sometimes use the shorthand $\mathcal{D}_t := \mathcal{D}_t(\theta_t)$.

The partially performative setting admits natural analogues of the central concepts from performative prediction. We define the *performative risk* and *decoupled performative risk* at time t as

$$\mathrm{PR}_t(\theta) := \mathbb{E}_{Z \sim \mathcal{D}_t(\theta)}[\ell(Z, \theta)], \quad \mathrm{DPR}_t(\theta, \theta') := \mathbb{E}_{Z \sim \mathcal{D}_t(\theta')}[\ell(Z, \theta)].$$

The performative risk evaluates a model on the distribution that it induces. The decoupled performative risk separates the model being evaluated, θ, from the model that induces the distribution, θ'.

The *performatively stable point* at time t is defined as the point $\theta_t^{\mathrm{PS}} \in \Theta$ satisfying

$$\theta_t^{\mathrm{PS}} := \arg\min_{\theta \in \Theta} \mathrm{DPR}_t(\theta, \theta_t^{\mathrm{PS}}).$$

Thus, once θ_t^{PS} is deployed and induces the distribution $\mathcal{D}_t(\theta_t^{\mathrm{PS}})$, it minimizes the expected loss on that induced distribution. The *performatively optimal point* at time t is the point

$$\theta_t^{\mathrm{PO}} := \arg\min_{\theta \in \Theta} \mathrm{PR}_t(\theta).$$

Performative optimality jointly accounts for the loss incurred by predicting with θ and the distributional response induced by deploying θ. In general, stable and optimal points do not coincide [23].

Because the stable and optimal points may vary with time, a natural objective is to achieve low regret relative to the sequences $\{\theta_t^{\mathrm{PS}}\}_{t \in [T]}$ and $\{\theta_t^{\mathrm{PO}}\}_{t \in [T]}$. This leads to our definitions of *performative stability regret* and *performative optimality regret*. Let $\mathcal{A}$ be an algorithm that outputs a possibly stochastic sequence of iterates $\{\theta_t\}_{t \in [T]}$. We define the performative stability regret and the performative optimality regret of $\mathcal{A}$ as

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}) := \sum_{t=1}^{T} \left(\mathbb{E}[\mathrm{PR}_t(\theta_t)] - \mathrm{PR}_t(\theta_t^{\mathrm{PS}}) \right), \quad \mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}) := \sum_{t=1}^{T} \left(\mathbb{E}[\mathrm{PR}_t(\theta_t)] - \mathrm{PR}_t(\theta_t^{\mathrm{PO}}) \right).$$

The expectation is taken over any randomness in the algorithm.

Several previous works have studied regret guarantees in performative prediction [12, 29, 3, 5]. Their guarantees are given in terms of *static* regret, which compares the learner's realized loss to that of a single fixed $\theta \in \Theta$. In the standard performative prediction model, this is natural, since the stable and optimal points do not change over time. In the partially performative setting, however, the exogenous component can make these points time-varying. We therefore study *dynamic* regret, which compares the learner's performance to a sequence of models. This distinction parallels the classical difference between static and dynamic regret in online learning; see, for example, [7].

Finally, we invoke several common assumptions from the performative prediction literature, including (strong) convexity and smoothness conditions on the loss $\ell(z, \theta)$. Since these assumptions are standard in optimization and statistical learning theory, we defer their formal statements to Section A. Here we state one assumption specific to performative settings: sensitivity of the distribution map. We say that $\mathcal{D}$ is ϵ-sensitive if

$$\mathcal{W}_1\big(\mathcal{D}(\theta), \mathcal{D}(\theta')\big) \le \epsilon \|\theta - \theta'\|, \qquad \forall\, \theta, \theta' \in \Theta, \tag{4}$$

where $\mathcal{W}_1$ denotes the Wasserstein-1 distance. This condition requires the performative response to vary smoothly with the deployed model. Intuitively, ϵ measures the strength of performative effects.

3 Repeated risk minimization under partial performativity

We begin by studying repeated risk minimization (RRM) [23]. Given an initial $\theta_1 \in \Theta$, RRM sets

$$\theta_{t+1} = \arg\min_{\theta \in \Theta} \mathbb{E}_{Z \sim \mathcal{D}_t(\theta_t)}[\ell(Z, \theta)]. \tag{RRM}$$

We use $\mathcal{A}_{\mathrm{RRM}}$ to denote this algorithm.

We start from a key technical lemma that characterizes one-step RRM updates.

Lemma 3.1. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4) and $\ell(z, \theta)$ is μ-strongly convex in θ (10) and β_z-smooth in z (12). Then, the iterates of $\mathcal{A}_{\mathrm{RRM}}$ satisfy*

$$\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\| \leq \frac{(1-\alpha_t)\epsilon\beta_z}{\mu} \|\theta_t - \theta_t^{\mathrm{PS}}\|.$$

We note that Lemma 3.1 is a strict generalization of Theorem 3.5 by Perdomo et al. [23], recovered when $\alpha_t = 0$ and $\theta_t^{\mathrm{PS}} \equiv \theta^{\mathrm{PS}}$ for all t. Building on this lemma, we obtain the following stability regret for RRM.

Theorem 3.2. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4), $\ell(z, \theta)$ is μ-strongly convex in θ (10) and β_z-smooth in z (12), and $\mathrm{PR}_t(\theta)$ is L-Lipschitz (13) in θ. Then, if $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$, where $\alpha_{\min} = \min_{t \in [T]} \alpha_t$, the iterates of $\mathcal{A}_{\mathrm{RRM}}$ satisfy*

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RRM}}) \leq \frac{L}{1-\gamma} \cdot \left(\|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PS}} \right), \tag{5}$$

for $\gamma = \frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu}$.

We defer the proof of Theorem 3.2 to Section C.2.

Remark 3.3. Although we state all (strong) convexity and smoothness conditions directly in terms of the loss ℓ for simplicity, all results in the paper also hold under the corresponding conditions on the *expected* loss; see Remark A.1 for details. Also, we show in Lemma B.5 that ℓ being Lipschitz in both z and θ implies that $\mathrm{PR}_t(\theta)$ is Lipschitz.

In the standard performative setting with $\alpha_t \equiv 0$, convergence of RRM to stability required that the contraction factor $\gamma = \frac{\epsilon\beta_z}{\mu}$ be strictly less than one. Under the partially performative setting with possible $\alpha_t > 0$, the contraction factor γ gets smaller by a factor of $(1 - \alpha_t)$, and thus we only need a weaker sensitivity condition $\gamma = \frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu} < 1$, or equivalently $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$. At the same time, large exogenous shifts may lead to large gaps between subsequent stable points Δ_t^{PS}. To provide more intuition for Δ_t^{PS}, we state an upper bound on Δ_t^{PS} as a function of the mixture coefficient α_t, how fast the distribution P_t varies, and the properties of the loss function.

Proposition 3.4. *Fix $t \in \mathbb{N}$. Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4) and $\ell(z, \theta)$ is μ-strongly convex in θ (10) and β_z-smooth in z (12). Further, suppose that $\mathcal{W}_1(\mathcal{D}(\theta_{t+1}^{\mathrm{PS}}), P_{t+1}) \leq C$ for some $C < \infty$. Then, if $\epsilon < \frac{\mu}{(1-\alpha_t)\beta_z}$, we have*

$$\Delta_t^{\mathrm{PS}} \leq \frac{\beta_z}{\mu - \beta_z(1-\alpha_t)\epsilon} \left(\alpha_t \mathcal{W}_1(P_t, P_{t+1}) + |\alpha_t - \alpha_{t+1}| C \right).$$

We defer the proof of Proposition 3.4 to Section C.3. We observe that Δ_t^{PS} is driven by (i) how fast the exogenous component drifts, $\mathcal{W}_1(P_t, P_{t+1})$; (ii) the contribution of the exogenous component α_t; and (iii) the local fluctuations in the mixture coefficient, $|\alpha_t - \alpha_{t+1}|$. In particular, Δ_t^{PS} is small whenever α_t is consistently small or the distributions P_t do not vary significantly.

As a special case, suppose that $\alpha_t \leq ct^{-b}$, for some $c, b > 0$, and that $\mathcal{W}_1(P_t, P_{t+1})$ is uniformly bounded for all t. Then, if $\epsilon < \mu/\beta_z$, we can conclude

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RRM}}) = \mathcal{O}\left(1 + \sum_{t=1}^{T-1} \alpha_t\right).$$

For the case $b \in (0, 1)$, we have $\sum_{t=1}^{T-1} \alpha_t = \mathcal{O}(T^{1-b})$, and therefore $\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RRM}}) = \mathcal{O}(T^{1-b})$. Similarly, for $b = 1$ we have $\sum_{t=1}^{T-1} \alpha_t = \mathcal{O}(\log T)$, and for $b > 1$ we have $\sum_{t=1}^{T-1} \alpha_t = \mathcal{O}(1)$, each implying a regret upper bound of the same rate, respectively.

4 Gradient descent under partial performativity

Next, we consider several variants of gradient descent and prove their stability regret bounds. We begin by considering repeated gradient descent (RGD) with access to the full expected gradient, after which we move on to stochastic gradient descent. In the standard performative prediction setting, the former was studied by [23, 17] and the latter was studied by [17, 4], among others.

4.1 Repeated gradient descent

First, we consider repeated gradient descent (RGD),

$$\theta_{t+1} = \Pi_\Theta\big(\theta_t - \eta_t \mathbb{E}_{Z\sim\mathcal{D}_t(\theta_t)}[\nabla \ell(Z, \theta_t)]\big), \tag{RGD}$$

where $\eta_t > 0$ is a step size sequence. We denote this algorithm by $\mathcal{A}_{\mathrm{RGD}}$.

We obtain a similar rate for the stability regret bound as in the case of RRM.

Theorem 4.1. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4), $\ell(z,\theta)$ is μ-strongly convex (10) and β_θ-smooth (11) in θ and β_z-smooth (12) in z, and $\mathrm{PR}_t(\theta)$ is L-Lipschitz (13) in θ. Then, if $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$, where $\alpha_{\min} = \min_{t\in[T]} \alpha_t$, the iterates of $\mathcal{A}_{\mathrm{RGD}}$ with step size*

$$\eta_t = \frac{\mu - (1-\alpha_t)\epsilon\beta_z}{2(\beta_\theta^2 + (1-\alpha_t)^2\epsilon^2\beta_z^2)}$$

satisfy

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RGD}}) \le \frac{L}{1-\overline{\gamma}} \cdot \left(\|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PS}} \right), \tag{6}$$

for $\overline{\gamma} = 1 - \frac{(\mu-(1-\alpha_{\min})\epsilon\beta_z)^2}{4(\beta_\theta^2+(1-\alpha_{\min})^2\epsilon^2\beta_z^2)}$.

We defer the proof of Theorem 4.1 to Section D.2. As in the case of RRM, we observe a regret upper bound of $\mathcal{O}(1 + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PS}})$. The dependence on the path length mirrors existing results for online gradient descent [20].

4.2 Stochastic gradient descent: greedy and lazy deploy

We consider two variants of stochastic gradient descent. Both sequentially collect samples from the currently deployed model and take a single gradient step after each. They differ in their data collection schedule: *greedy deploy* collects a single data point between two deployments, while *lazy deploy* collects multiple data points and takes multiple offline steps between two deployments.

Before formally stating the algorithms and their guarantees, we introduce a standard assumption in stochastic optimization. We assume that the stochastic gradients have a quadratic-bounded variance: for all possible $\mathcal{D}_t$, for some C_V we have

$$\mathbb{E}_{Z\sim\mathcal{D}_t}[\|\nabla_\theta \ell(Z,\theta)\|^2] \le \sigma^2 + C_V^2 \|\theta - \theta^\star_{\mathcal{D}_t}\|^2, \tag{7}$$

where $\theta^\star_{\mathcal{D}_t} = \arg\min_\theta \mathbb{E}_{Z\sim\mathcal{D}_t}[\ell(Z,\theta)]$.

Greedy deploy. Consider the *greedy deploy* [17] variant of stochastic gradient descent (SGD):

$$\theta_{t+1} = \Pi_\Theta(\theta_t - \eta_t \nabla_\theta \ell(Z_t, \theta_t)), \quad \text{where } Z_t \sim \mathcal{D}_t, \tag{greedy SGD}$$

which uses a single sample per each deployment. We denote the algorithm by $\mathcal{A}_{\text{SGD-greedy}}$.

Proposition 4.2. *Let the same assumptions as in Theorem 4.1 and the quadratic variance bound (7) hold. Then, the iterates of $\mathcal{A}_{\text{SGD-greedy}}$ with step size $\eta_t \propto \frac{1}{t+t_0}$, for some constant $t_0 > 0$, satisfy*

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\text{SGD-greedy}}) \le \mathcal{O}\left(T^{1/2} + T^{1/4} \left(\sum_{t=1}^{T-1} (t + t_0 + 1)^{5/2} (\Delta_t^{\mathrm{PS}})^2 \right)^{1/2} \right). \tag{8}$$

We defer the proof of Proposition 4.2 to Section E.3. We can observe from (8) that $\mathcal{A}_{\text{SGD-greedy}}$ enjoys sublinear regret whenever we have $\Delta_t^{\text{PS}} \leq Mt^{-b}$ for $b > 1$,

$$\text{Reg}_T^{\text{PS}}(\mathcal{A}_{\text{SGD-greedy}}) = \begin{cases} \mathcal{O}(T^{2-b}), & \text{if } b \in (1, \frac{3}{2}), \\ \mathcal{O}(T^{1/2}), & \text{if } b \geq \frac{3}{2}. \end{cases}$$

As discussed in Section 3, we can have $\Delta_t^{\text{PS}} \leq Mt^{-b}$ when $\alpha_t \leq ct^{-b}$ for some constant $c > 0$.

Remark 4.3. Unlike the other results in Section 3 and Section 4, the upper bound in Proposition 4.2 does not guarantee sublinear regret of $\mathcal{A}_{\text{SGD-greedy}}$ when $\Delta_t^{\text{PS}} \leq Mt^{-b}$ decays with a rate of $b \in (0, 1)$. We leave the search for a tighter upper bound or a matching lower bound to future work.

Lazy deploy. Consider the *lazy deploy* [17] variant of stochastic gradient descent (SGD):

$$\varphi_{t,j+1} = \Pi_\Theta(\varphi_{t,j} - \eta_{t,j} \nabla_\theta \ell(Z_{t,j}, \varphi_{t,j})), \tag{lazy SGD}$$

where $Z_{t,j} \overset{\text{i.i.d.}}{\sim} \mathcal{D}_t$ for $j \in [n(t)]$, with $\varphi_{t,1} = \theta_t$ and $\varphi_{t,n(t)+1} = \theta_{t+1}$. In other words, $\varphi_{t,j}$ are internal iterates computed between two deployments θ_t and θ_{t+1}. Within each deployment, we collect $n(t)$ samples from $\mathcal{D}_t$, and we assume $n(t) \geq n_0 t^r$ for some $r > 0$. We denote this algorithm by $\mathcal{A}_{\text{SGD-lazy}}$.

Proposition 4.4. *Suppose that the same assumptions as in Proposition 4.2 hold. Then, the iterates of $\mathcal{A}_{\text{SGD-lazy}}$ with step sizes $\eta_{t,j} \propto \frac{1}{j+t_0}$, for some constant $t_0 > 0$, and $n(t) \geq n_0 t^r$ with a large enough n_0 satisfy*

$$\text{Reg}_T^{\text{PS}}(\mathcal{A}_{\text{SGD-lazy}}) \leq \mathcal{O}\left(\sum_{t=1}^{T-1} t^{-r/2} + \sum_{t=1}^{T-1} \Delta_t^{\text{PS}}\right). \tag{9}$$

We defer the proof of Proposition 4.4 to Section F.3. Summing up over t, we see that the first term in the regret is upper bounded by $\mathcal{O}(1)$ for $r > 2$, $\mathcal{O}(\log T)$ for $r = 2$, and $\mathcal{O}(T^{1-r/2})$ for $r \in (0, 2)$.

5 Achieving optimality under partial performativity

We now turn our attention to studying performative optimality regret. Performative optimality is harder to achieve than stability because computing the gradients of $\text{PR}_t(\theta)$ is rarely feasible, due to the (typically complex) dependence of $\text{PR}_t(\theta)$ on θ through $\mathcal{D}_t(\theta)$. We thus turn to zeroth-order algorithms, which only require querying $\text{PR}_t(\theta)$, not its gradient.

Throughout we make a couple of mild assumptions: that $|\ell(z, \theta)| \leq F$ for some finite F, for all z, θ, and that Θ contains a ball of radius r centered at 0, for some $r > 0$.

To prove that the regret vanishes, we will need $\text{PR}_t(\theta)$ to be convex or strongly convex. This is not guaranteed by the convexity of $\ell(z, \theta)$ alone. Miller et al. [18] identify conditions that ensure convexity of the performative risk, the most important one being a so-called mixture dominance condition. Miller et al. show that this condition is satisfied by location-scale families. In the partially performative setting, convexity is easier to ensure, since the α_t-fraction of the objective corresponding to the exogenous component is convex whenever $\ell(z, \theta)$ is convex. We formalize this below.

Proposition 5.1 (Theorem 3.1 of [18])**.** *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4), and $\ell(z, \theta)$ is μ-strongly convex in θ (10) and β_z-smooth in z (12). Assume also that the mixture dominance condition (see (14) of Section A) holds. Then, PR_t is $(\mu - 2(1 - \alpha_t)\epsilon\beta_z)$-strongly convex.*

If $\mu - 2(1 - \alpha_t)\epsilon\beta_z$ is negative, then $\text{PR}_t(\theta)$ is *weakly* convex by that amount, meaning that adding a quadratic regularizer of that magnitude makes $\text{PR}_t(\theta)$ convex.

Consider two-point zeroth-order gradient descent (ZGD2) [1]: starting from an initial $\theta_1 \in \Theta$, we set

$$\theta_{t+1} = \Pi_{(1-\rho)\Theta}\left(\theta_t - \eta_t \cdot \tilde{g}_t\right), \tag{ZGD2}$$

$$\tilde{g}_t = \big(\text{PR}_t(\theta_t + \delta u_t) - \text{PR}_t(\theta_t - \delta u_t)\big) u_t \frac{d}{2\delta}, \text{ where } u_t \sim \text{Unif}(\mathbb{S}^{d-1}).$$

Here, $\eta_t > 0$ is a step size, $\delta > 0$ and $\rho \in [0, 1]$ are hyperparameters, and u_t is sampled uniformly on the d-dimensional unit sphere $\mathbb{S}^{d-1}$. In other words, we deploy two models: $\phi_t^+ = \theta_t + \delta u_t$

and $\phi_t^- = \theta_t - \delta u_t$, observe their corresponding risks $\mathrm{PR}_t(\phi_t^+)$ and $\mathrm{PR}_t(\phi_t^-)$, and use those values to perform an approximate gradient update on θ_t. Practically, one can think of the procedure as partitioning the population into two halves and deploying two different treatments ϕ_t^+ and ϕ_t^- within each. It is known that $\mathbb{E}[\tilde{g}_t|\theta_t] = \nabla\widehat{\mathrm{PR}}_t(\theta_t)$ for $\widehat{\mathrm{PR}}_t(\theta) \coloneqq \mathbb{E}_v[\mathrm{PR}_t(\theta + \delta v)]$ with v sampled uniformly within the d-dimensional unit ball, $v \sim \mathrm{Unif}(\mathbb{B}^d)$. We use projections to $(1-\rho)\Theta$ with $\rho \geq \frac{\delta}{r}$ to ensure that the perturbed iterates $\{\phi_t^+, \phi_t^-\}$ are in the domain, i.e., $\theta_t \pm \delta u_t \in \Theta$. We denote the algorithm by $\mathcal{A}_{\mathrm{ZGD2}}$.

Note that here we have applied a slight abuse of notation, since θ_t is not the actual sequence *deployed* by $\mathcal{A}_{\mathrm{ZGD2}}$; θ_t is the sequence used internally for optimization. To accurately capture optimality regret, we measure it with respect to the deployed sequence $\{\phi_t^+, \phi_t^-\}$:

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) = \sum_{t=1}^{T}\left(\max_{\phi_t \in \{\phi_t^+, \phi_t^-\}} \mathbb{E}[\mathrm{PR}_t(\phi_t)] - \mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right).$$

Theorem 5.2. *Suppose that* $\mathrm{PR}_t(\theta)$ *is* L*-Lipschitz* (13)*. If* PR_t *is* μ'*-strongly convex, then* $\mathcal{A}_{\mathrm{ZGD2}}$ *with* $\eta_t = \frac{1}{\mu' t}$, $\delta = \frac{d^2 L}{6\mu' T}$, *and* $\rho = \frac{\delta}{r}$ *satisfies for* $T \geq \frac{d^2 L}{6\mu' r}$,

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) = \mathcal{O}\left(d^2 \log T + \sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right).$$

If PR_t *is convex, then* $\mathcal{A}_{\mathrm{ZGD2}}$ *with* $\eta_t \equiv \eta = \frac{D_\Theta}{dL\sqrt{T}}$, $\delta = \frac{dD_\Theta}{6\sqrt{T}}$, *and* $\rho = \frac{\delta}{r}$ *satisfies for* $T \geq \frac{d^2 D_\Theta^2}{36 r^2}$,

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) = \mathcal{O}\left(d\sqrt{T} \cdot \left(1 + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PO}}\right)\right).$$

We defer the proof of Theorem 5.2 to Section G.1.

If we only have access to stochastic feedback, like in Section 4.2, then two-point zeroth-order gradient descent is not appropriate, since we cannot subject any one individual Z_t to two different treatments ϕ_t^+ and ϕ_t^-. To address this challenge, in Section H.1 we analyze a one-point version of zeroth-order gradient descent and bound its optimality regret. Given the weaker feedback signal, the regret bounds are naturally larger.

Remark 5.3. If we are not explicitly given the Lipschitz constant L, and thus cannot set the algorithm parameters accordingly, the assumptions we have on the upper bound F of $|\ell|$ and that $r\mathbb{B}^d \subseteq \Theta$ allows us to use the effective Lipschitz constant $L_{\mathrm{eff}} = \frac{2F}{\rho r}$ as a proxy, as suggested in [6]. We defer related discussions and prove regret upper bounds of single and two-deployment ZGD under this setting to Section H.2.

Remark 5.4. Chen et al. [3] similarly study how to use two deployed models to construct a zeroth-order estimate of the performative risk gradient in a fully performative setting. Our result gives a new analysis of the *dynamic* optimality regret of the two-deployment algorithm under partial performativity.

6 Experiments

We conduct experiments evaluating the stability regret of the retraining algorithms considered in the paper. We include several main results in this section and defer further empirical results and additional experimental detail to Appendix I.

We use the credit scoring dataset from [31], in particular the processed version from [26]. The dataset contains demographic and socioeconomic features $x_i \in \mathbb{R}^{16}$, including age, education, and credit history, and labels y_i indicating whether a person defaulted on their credit card payment, for $30,000$ individuals. The learner aims to predict the individuals' creditworthiness by minimizing a regularized logistic loss over an ℓ_2 ball in $d = 17$ dimensions (one dimension corresponding to the intercept), after standardizing the features. The performative feedback is modeled as coming from strategic responses: an individual i responds to a deployed model θ by optimizing the score $\theta^\top x_i$ (equivalently, maximizing the prediction), with an additional quadratic penalty for deviating too far

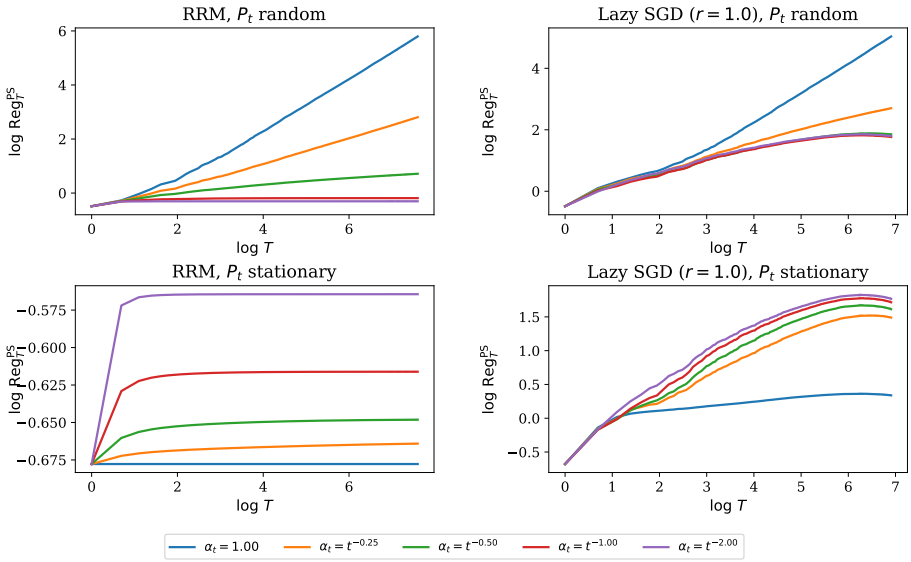

Figure 1: Performative stability regret of $\mathcal{A}_{\mathrm{RRM}}$ (left) and $\mathcal{A}_{\text{SGD-lazy}}$ (right) in the credit scoring experiment, for varying α_t. We consider a randomly varying (top) and a stationary (bottom) exogenous component P_t.

from their initial features $x_{i,0}$. Only certain features are modifiable (for example, age is not), and thus the individuals only manipulate those modifiable features. We model the initial features x_0 in the dataset as being drawn from a normal distribution $\mathcal{N}(m, \Sigma)$, and we fit m and Σ from 1500 data points. The individuals' best response has a closed-form expression in this case: $x \leftarrow x_0 + A\theta$, where $A = a \cdot \mathrm{diag}(p_{\mathrm{mod}})$, a is a constant that captures the individuals' utility–cost tradeoff, and p_{mod} indicates the modifiable coordinates. Putting everything together yields a distribution map $\mathcal{D}(\theta) = \mathcal{N}(A\theta + m, \Sigma)$. We set $a = 0.1$.

We introduce an exogenous component $P_t = \mathcal{N}(m_t, \sigma^2 I)$. To capture "hard" and "easy" exogenous shifts, we consider two cases: in the first, m_t is sampled independently from a bounded ball at each round; in the second, m_t is sampled once at the beginning and held fixed across t. We set $\sigma = 0.3$. We study the effect of different values of α_t; we compare polynomial schedules $\alpha_t = t^{-b}$ with $b \in \{0.25, 0.5, 1.0, 2.0\}$ and the constant schedule $\alpha_t = 1$ (or, equivalently, $b = 0$).

In Figure 1 we plot the stability regret of repeated risk minimization (RRM) and SGD with a lazy deployment schedule, for random and stationary P_t. The figure shows that the effect of the mixing schedule α_t depends on whether the exogenous component drifts. With random shifts, faster decay of α_t suppresses the moving exogenous term and produces smaller stability regret for the displayed stability algorithms. With fixed P_t, the ordering is reversed: keeping α_t larger for longer makes the observed distribution closer to a fixed, nonperformative component and yields smaller regret, $\alpha_t = 1$ being the most desirable scenario. This behavior is consistent with the theoretical results, as well as the path length bound from Proposition 3.4. In the random P_t case, the term $\sum_t \Delta_t^{\mathrm{PS}}$, which is driven by $\alpha_t \mathcal{W}_1(P_t, P_{t+1})$, is the dominant term in the regret bound. Therefore, larger α_t leads to a larger regret. In the fixed P_t case, the path length is zero, and thus the regret bound is dominated by the performative regret component, which decreases as α_t grows.

7 Conclusion

We introduce partially performative prediction—a novel model for learning under distribution shifts that are both induced by performative feedback and exogenous variations. In this setting, we measure performance via regret relative to a sequence of performatively stable and performatively optimal points, rather than a single stable or optimal point, as in the standard performative prediction model. We prove dynamic regret bounds for several standard algorithms, including repeated risk minimization and variants of stochastic gradient descent. Several questions remain open, such as matching regret lower bounds or tighter upper bounds, such as in the case of greedy SGD. It would also be interesting to study partially performative environments beyond linear mixtures.

Acknowledgments and Disclosure of Funding

This work was supported by a Dieter Schwarz Early Career Catalyst Award.

References

[1] Alekh Agarwal, Ofer Dekel, and Lin Xiao. Optimal algorithms for online convex optimization with multi-point bandit feedback. In *Conference on Learning Theory (COLT)*, 2010.

[2] Gavin Brown, Shlomi Hod, and Iden Kalemaj. Performative prediction in a stateful world. In *International conference on artificial intelligence and statistics*, pages 6045–6061. PMLR, 2022.

[3] Yatong Chen, Wei Tang, Chien-Ju Ho, and Yang Liu. Performative prediction with bandit feedback: Learning through reparameterization. In *International Conference on Machine Learning*, pages 7298–7324. PMLR, 2024.

[4] Dmitriy Drusvyatskiy and Lin Xiao. Stochastic optimization with decision-dependent distributions. *Mathematics of Operations Research*, 48(2):954–998, 2023.

[5] Gabriele Farina and Juan Carlos Perdomo. The stability of online algorithms in performative prediction. *arXiv preprint arXiv:2602.24207*, 2026.

[6] Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2005.

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

[8] Mark Herbster and Manfred K. Warmuth. Tracking the best expert. *Machine Learning*, 32(2):151–178, 1998.

[9] Zachary Izzo, Lexing Ying, and James Zou. How to learn when data reacts to your model: performative gradient descent. In *International Conference on Machine Learning*, pages 4641–4650. PMLR, 2021.

[10] Zachary Izzo, James Zou, and Lexing Ying. How to learn when data gradually reacts to your model. In *International Conference on Artificial Intelligence and Statistics*, pages 3998–4035. PMLR, 2022.

[11] Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *International Conference on Artificial Intelligence and Statistics*, pages 398–406. PMLR, 2015.

[12] Meena Jagadeesan, Tijana Zrnic, and Celestine Mendler-Dünner. Regret minimization with performative feedback. In *International Conference on Machine Learning*, pages 9760–9785. PMLR, 2022.

[13] Michael P Kim and Juan C Perdomo. Making decisions under outcome performativity. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251, page 79, 2023.

[14] Qiang Li and Hoi-To Wai. State dependent performative prediction with stochastic approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 3164–3186. PMLR, 2022.

[15] Licong Lin and Tijana Zrnic. Plug-in performative optimization. In *International Conference on Machine Learning*, pages 30546–30565, 2024.

[16] Haitong Liu, Qiang Li, and Hoi To Wai. Two-timescale derivative free optimization for performative prediction with markovian data. In *International Conference on Machine Learning*, pages 31425–31450. PMLR, 2024.

[17] Celestine Mendler-Dünner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic optimization for performative prediction. *Advances in Neural Information Processing Systems*, 33:4929–4939, 2020.

[18] John P Miller, Juan C Perdomo, and Tijana Zrnic. Outside the echo chamber: Optimizing the performative risk. In *International Conference on Machine Learning*, pages 7710–7720. PMLR, 2021.

[19] Mehrnaz Mofakhami, Ioannis Mitliagkas, and Gauthier Gidel. Performative prediction with neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 11079–11093. PMLR, 2023.

[20] Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *IEEE Conference on Decision and Control (CDC)*, 2016.

[21] Adhyyan Narang, Evan Faulkner, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. Multiplayer performative prediction: Learning in decision-dependent games. *Journal of Machine Learning Research*, 24(202):1–56, 2023.

[22] Sungwoo Park, Junyeop Kwon, Byeongnoh Kim, Suhyun Chae, Jeeyong Lee, and Dabeen Lee. Parameter-free algorithms for performative regret minimization under decision-dependent distributions. *arXiv preprint arXiv:2402.15188*, 2024.

[23] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, pages 7599–7609. PMLR, 2020.

[24] Juan Carlos Perdomo. Revisiting the predictability of performative, social events. In *International Conference on Machine Learning*, pages 48948–48961. PMLR, 2025.

[25] Mitas Ray, Lillian J Ratliff, Dmitriy Drusvyatskiy, and Maryam Fazel. Decision-dependent risk minimization in geometrically decaying dynamic environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8081–8088, 2022.

[26] Berk Ustun, Alexander Spangher, and Yang Liu. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, page 10–19. Association for Computing Machinery, 2019.

[27] Xiaolu Wang, Chung-Yiu Yau, and Hoi To Wai. Network effects in performative prediction games. In *International Conference on Machine Learning*, pages 36514–36540. PMLR, 2023.

[28] Songkai Xue and Yuekai Sun. Distributionally robust performative prediction. In *Advances in Neural Information Processing Systems*, volume 37, pages 55030–55052, 2024.

[29] Wenjing Yan and Xuanyu Cao. Zero-regret performative prediction under inequality constraints. *Advances in Neural Information Processing Systems*, 36:1298–1308, 2023.

[30] Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng Yi. Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient. In *International Conference on Machine Learning*, pages 449–457. PMLR, 2016.

[31] I-Cheng Yeh and Che-hui Lien. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert systems with applications*, 36(2):2473–2480, 2009.

[32] Lijun Zhang, Tianbao Yangt, Jinfeng Yi, Rong Jin, and Zhi-Hua Zhou. Improved dynamic regret for non-degenerate functions. In *International Conference on Neural Information Processing Systems*, page 732–741. Curran Associates Inc., 2017.

[33] Peng Zhao, Guanghui Wang, Lijun Zhang, and Zhi-Hua Zhou. Bandit convex optimization in non-stationary environments. In *International conference on artificial intelligence and statistics*, 2020.

[34] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning*, 2003.

A Definitions and assumptions

We say that a continuously differentiable function $\ell : \mathcal{Z} \times \Theta \to \mathbb{R}$ is μ-strongly convex in θ if

$$\ell(z, \theta') \geq \ell(z, \theta) + \langle \nabla_\theta \ell(z, \theta), \theta' - \theta \rangle + \frac{\mu}{2} \|\theta' - \theta\|^2, \qquad \forall \theta, \theta' \in \Theta,\, z \in \mathcal{Z}, \tag{10}$$

and convex in θ if the above holds with $\mu = 0$.

We say that ℓ is β_θ-smooth in θ if

$$\|\nabla_\theta \ell(z, \theta) - \nabla_\theta \ell(z, \theta')\| \leq \beta_\theta \|\theta - \theta'\|, \qquad \forall\, \theta, \theta' \in \Theta\,, z \in \mathcal{Z}, \tag{11}$$

and β_z-smooth in z if[1]

$$\|\nabla_\theta \ell(z, \theta) - \nabla_\theta \ell(z', \theta)\| \leq \beta_z \|z - z'\|, \qquad \forall\, z, z' \in \mathcal{Z},\, \theta \in \Theta. \tag{12}$$

We say that $\mathrm{PR}_t(\theta)$ is L-Lipschitz if

$$|\mathrm{PR}_t(\theta) - \mathrm{PR}_t(\theta')| \leq L \|\theta - \theta'\|, \qquad \forall\, \theta, \theta' \in \Theta. \tag{13}$$

Remark A.1. When we make assumptions on the loss as a function of z, we require the assumptions to hold for all θ; for example, $z \mapsto \ell(z, \theta)$ is β_z-smooth for all θ. For the assumptions on the loss as a function of θ, however, our results hold under a weaker assumption that the *expected* risk $\mathrm{R}_D(\theta) \coloneqq \mathbb{E}_{Z \sim D}[\ell(Z, \theta)]$ is μ-strongly convex, β_θ-smooth, etc, for all admissible distributions that could be generated by $\mathcal{D}_t(\theta_t)$. For simplicity we state the assumptions for all $z \in \mathcal{Z}$, as above, however all the proofs go through if we consider the expected risk.

Making assumptions on the expected risk is strictly weaker, especially if we have additional structure on the distributions that can be generated by $\mathcal{D}_t(\theta_t)$. For example, consider a simple one-dimensional coin-flip setting with a single feature where the random variable $Z = (X, Y) \in \{-1, +1\}^2$ has a fixed distribution of $X \sim \mathrm{Unif}(\{-1, +1\})$ and only the conditional law of $Y \mid X$ can change due to performative and exogenous factors. Consider

$$\ell(z, \theta) = \frac{1}{2}(\theta - y)^2 - x\theta^2.$$

We have that $\theta \mapsto \ell(z, \theta)$ is nonconvex in θ for $x = 1$. However, if we average the loss over the distribution of (X, Y), we have

$$\mathrm{R}_D(\theta) = \mathbb{E}_{Z \sim D}\left[\frac{1}{2}(\theta - Y)^2 - X\theta^2\right] = \mathbb{E}_{Z \sim D}\left[\frac{1}{2}(\theta - Y)^2\right],$$

which is 1-strongly convex in θ. Similarly, we can also consider

$$\ell(z, \theta) = \frac{1}{2}(\theta - y)^2 - x|\theta|,$$

which is non-smooth at $\theta = 0$, but admits a smooth $\mathrm{R}_D(\theta)$ under the same setup.

We say that the mixture dominance condition holds if for any $\theta, \theta', \theta_0 \in \Theta$ and $\alpha \in (0, 1]$,

$$\mathbb{E}_{Z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')}[\ell(Z, \theta_0)] \leq \mathbb{E}_{Z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')}[\ell(Z, \theta_0)]. \tag{14}$$

As stated in Proposition 5.1 or Theorem 3.1 of [18], the mixture dominance condition is a sufficient condition for the performative risk to be (strongly) convex under ϵ-sensitivity of $\mathcal{D}(\theta)$ and suitable regularity assumptions on the loss.

B Auxiliary lemmas

Lemma B.1. *Suppose that we have a recursive inequality of the form*

$$F_{t+1} \leq \gamma_t F_t + \Delta_t.$$

Then we have

$$\sum_{t=1}^{T} (1 - \gamma_t) F_t \leq F_1 - \gamma_T F_T + \sum_{t=1}^{T-1} \Delta_t. \tag{15}$$

[1] Note that β_z-smoothness is also defined using the gradient with respect to θ.

Proof. Since we have

$$(1-\gamma_t)F_t \le F_t - F_{t+1} + \Delta_t,$$

the summation of the above yields

$$\sum_{t=1}^{T-1}(1-\gamma_t)F_t \le \sum_{t=1}^{T-1}(F_t - F_{t+1}) + \sum_{t=1}^{T-1}\Delta_t = F_1 - F_T + \sum_{t=1}^{T-1}\Delta_t.$$

We add $(1-\gamma_T)F_T$ to both sides to obtain Eq. (15). □

Lemma B.2. *Suppose that we have a recursive inequality of the form*

$$F_{t+1} \le \left(1 - \frac{\alpha}{t+t_0}\right) F_t + H(t+t_0)^{-m-1}$$

for all $t \in \mathbb{N}$, with $t_0 > 0$ and $0 \le m < \alpha \le 1 + t_0$. Let

$$S = \max\left\{F_1(1+t_0)^m,\ \frac{H}{\alpha - m}\right\}.$$

Then we have $F_t \le S(t+t_0)^{-m}$ for all $t \in \mathbb{N}$.

Proof. For $t = 1$ we have $F_1 \le S(1+t_0)^{-m}$ by definition.
Suppose that $F_t \le S(t+t_0)^{-m}$. Then,

$$\begin{aligned}
F_{t+1} &\le \left(1 - \frac{\alpha}{t+t_0}\right) F_t + H(t+t_0)^{-m-1} \\
&\le \left(1 - \frac{\alpha}{t+t_0}\right) S(t+t_0)^{-m} + S(\alpha - m)(t+t_0)^{-1} \cdot (t+t_0)^{-m} \\
&= S\left(1 - \frac{\alpha}{t+t_0} + \frac{\alpha - m}{t+t_0}\right)(t+t_0)^{-m} \\
&= S\left(1 - \frac{m}{t+t_0}\right)(t+t_0)^{-m}.
\end{aligned}$$

Since Bernoulli's inequality (for $m \ge 0$ and $\frac{1}{t+t_0} > -1$) implies

$$1 - \frac{m}{t+t_0} \le \left(1 + \frac{1}{t+t_0}\right)^{-m},$$

we have

$$F_{t+1} \le S\left(1 + \frac{1}{t+t_0}\right)^{-m}(t+t_0)^{-m} = S(t+t_0+1)^{-m}$$

as desired. □

Lemma B.3 (Lemma B.5, [17])**.** *For any $c \in (0,1)$, $r > 0$, and $T \in \mathbb{N}$,*

$$\sum_{t=1}^{T} c^{T-t} t^{-r} \le \frac{c^{T(1-2^{-1/r})}}{1-c} + \frac{2T^{-r}}{1-c}.$$

Proof. (We rewrite the proof from [17] here for completeness.) We write

$$\sum_{t=1}^{T} c^{T-t} t^{-r} = \sum_{t=1}^{M_T} c^{T-t} t^{-r} + \sum_{t=M_T+1}^{T} c^{T-t} t^{-r},$$

where we define $M_T \coloneqq \max\{m \in \mathbb{N} : m^{-r} > 2T^{-r}\}$ (or $M_T = 0$ if the set is empty for small T). Then we have

$$\begin{aligned}\sum_{t=1}^{M_T} c^{T-t}t^{-r} + \sum_{t=M_T+1}^{T} c^{T-t}t^{-r} &\le \sum_{t=1}^{M_T} c^{T-t} + 2T^{-r} \cdot \sum_{t=M_T+1}^{T} c^{T-t} \\ &\le \frac{c^{T-M_T}}{1-c} + \frac{2T^{-r}}{1-c} \\ &\le \frac{c^{T(1-2^{-1/r})}}{1-c} + \frac{2T^{-r}}{1-c},\end{aligned}$$

where we use $M_T < 2^{-\frac{1}{r}} \cdot T$ by its definition. □

Lemma B.4. *Let $a, b \in [0,1]$ and $Q_1, Q_2, R_1, R_2 \in \mathcal{P}(\mathcal{Z})$. Then we have*

$$\begin{aligned}&\mathcal{W}_1\big((1-a)Q_1 + aR_1, (1-b)Q_2 + bR_2\big) \\ &\le (1-a)\mathcal{W}_1(Q_1, Q_2) + a\mathcal{W}_1(R_1, R_2) + |a-b|\,\mathcal{W}_1(Q_2, R_2).\end{aligned}$$

Proof. Since $\mathcal{W}_1$ is convex,

$$\mathcal{W}_1\big((1-a)Q_1 + aR_1, (1-a)Q_2 + aR_2\big) \le (1-a)\mathcal{W}_1(Q_1, Q_2) + a\mathcal{W}_1(R_1, R_2).$$

We also have

$$\mathcal{W}_1\big((1-a)Q_2 + aR_2, (1-b)Q_2 + bR_2\big) \le |b-a|\,\mathcal{W}_1(Q_2, R_2).$$

Assuming $a \le b$ without loss of generality, this holds because the distributions $(1-a)Q_2 + aR_2$ and $(1-b)Q_2 + bR_2$ share a mass of at least $1-b$ from Q_2 and a from R_2. Therefore the optimal transport plan from Q_2 to R_2 also yields a transport plan from $(1-a)Q_2 + aR_2$ to $(1-b)Q_2 + bR_2$ for the remaining mass $b-a = |b-a|$. Finally, applying triangle inequality for $\mathcal{W}_1$ proves the given statement. □

Lemma B.5 shows that sensitivity inherits Lipschitz properties of ℓ to the performative risk.

Lemma B.5. *Suppose that the map $\mathcal{D}(\theta)$ is ϵ-sensitive (4) and the function $\ell(z, \theta)$ is L_θ-Lipschitz in $\theta \in \Theta$ and L_z-Lipschitz in $z \in \mathcal{Z}$. Then, $\mathrm{PR}_t(\theta)$ is L_t-Lipschitz, where $L_t = (1-\alpha_t)\epsilon L_z + L_\theta$.*

Proof. (The proof is similar to that of Lemma 2 of [12].) Observe that $\mathrm{PR}_t(\theta) = \mathrm{R}_{\mathcal{D}_t(\theta)}(\theta)$ if we define $\mathrm{R}_D(\theta) \coloneqq \mathbb{E}_{Z\sim D}[\ell(Z, \theta)]$, and

$$\mathrm{R}_{\mathcal{D}_t(\theta)}(\theta) - \mathrm{R}_{\mathcal{D}_t(\theta')}(\theta') = (1-\alpha_t)\left(\mathrm{R}_{\mathcal{D}(\theta)}(\theta) - \mathrm{R}_{\mathcal{D}(\theta')}(\theta')\right) + \alpha_t\left(\mathrm{R}_{P_t}(\theta) - \mathrm{R}_{P_t}(\theta')\right),$$

for any $\theta, \theta' \in \Theta$. Note that by Lipschitzness and sensitivity of $\mathcal{D}$,

$$\begin{aligned}(1-\alpha_t)\left(\mathrm{R}_{\mathcal{D}(\theta)}(\theta) - \mathrm{R}_{\mathcal{D}(\theta')}(\theta)\right) &\le (1-\alpha_t)L_z \cdot \mathcal{W}_1(\mathcal{D}(\theta), \mathcal{D}(\theta')) \\ &\le (1-\alpha_t)\epsilon L_z\|\theta - \theta'\|,\end{aligned}$$

and also

$$\begin{aligned}&(1-\alpha_t)\left(\mathrm{R}_{\mathcal{D}(\theta')}(\theta) - \mathrm{R}_{\mathcal{D}(\theta')}(\theta')\right) + \alpha_t\big(\mathrm{R}_{P_t}(\theta) - \mathrm{R}_{P_t}(\theta')\big) \\ &= \big((1-\alpha_t)\mathrm{R}_{\mathcal{D}(\theta')}(\theta) + \alpha_t\mathrm{R}_{P_t}(\theta)\big) - \big((1-\alpha_t)\mathrm{R}_{\mathcal{D}(\theta')}(\theta') + \alpha_t\mathrm{R}_{P_t}(\theta')\big) \\ &= \mathrm{R}_{\mathcal{D}_t(\theta')}(\theta) - \mathrm{R}_{\mathcal{D}_t(\theta')}(\theta') \\ &\le L_\theta\|\theta - \theta'\|.\end{aligned}$$

Taking the sum of the two and swapping θ, θ' in the argument concludes the proof. □

C Proofs for repeated risk minimization

We will define and frequently use the simplified notation $\mathcal{D}_t^{\mathrm{PS}} \coloneqq \mathcal{D}_t(\theta_t^{\mathrm{PS}})$ throughout the remaining parts of the appendix.

C.1 Proof of Lemma 3.1

Here we prove Lemma 3.1, restated below for the sake of readability.

Lemma 3.1. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4) and $\ell(z,\theta)$ is μ-strongly convex in θ (10) and β_z-smooth in z (12). Then, the iterates of $\mathcal{A}_{\mathrm{RRM}}$ satisfy*

$$\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\| \le \frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\|\theta_t - \theta_t^{\mathrm{PS}}\|.$$

Proof. (The proof is similar to that of Theorem 3.5 of [23]. For completeness, we provide the full proof adapted to our settings.)

Define $g(\varphi) = \mathrm{R}_{\mathcal{D}_t(\theta)}(\varphi)$ and $g'(\varphi) = \mathrm{R}_{\mathcal{D}_t(\theta')}(\varphi)$ for fixed $\theta, \theta' \in \Theta$. Let $\mathcal{G}(\theta) \coloneqq \arg\min_{\varphi\in\Theta} g(\varphi)$ and $\mathcal{G}(\theta') \coloneqq \arg\min_{\varphi\in\Theta} g'(\varphi)$. Then the first-order optimality conditions yield

$$\nabla g(\mathcal{G}(\theta))^\top \big(\mathcal{G}(\theta') - \mathcal{G}(\theta)\big) \ge 0,$$
$$\nabla g'(\mathcal{G}(\theta'))^\top \big(\mathcal{G}(\theta) - \mathcal{G}(\theta')\big) \ge 0.$$

By strong convexity, we have

$$\begin{aligned} g(\mathcal{G}(\theta)) - g(\mathcal{G}(\theta')) &\ge \nabla g(\mathcal{G}(\theta'))^\top \big(\mathcal{G}(\theta) - \mathcal{G}(\theta')\big) + \frac{\mu}{2}\|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|^2, \\ g(\mathcal{G}(\theta')) - g(\mathcal{G}(\theta)) &\ge \nabla g(\mathcal{G}(\theta))^\top \big(\mathcal{G}(\theta') - \mathcal{G}(\theta)\big) + \frac{\mu}{2}\|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|^2 \\ &\ge \frac{\mu}{2}\|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|^2, \end{aligned}$$

where the last inequality comes from the first-order optimality condition of $\mathcal{G}(\theta) = \arg\min_{\varphi\in\Theta} g(\varphi)$. Therefore we have

$$-\nabla g(\mathcal{G}(\theta'))^\top \big(\mathcal{G}(\theta) - \mathcal{G}(\theta')\big) \ge \mu\|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|^2.$$

Note that

$$\begin{aligned} -\nabla g(\mathcal{G}(\theta'))^\top \big(\mathcal{G}(\theta) - \mathcal{G}(\theta')\big) &\le -\nabla g(\mathcal{G}(\theta'))^\top \big(\mathcal{G}(\theta) - \mathcal{G}(\theta')\big) + \nabla g'(\mathcal{G}(\theta'))^\top \big(\mathcal{G}(\theta) - \mathcal{G}(\theta')\big) \\ &\le \|\nabla g(\mathcal{G}(\theta')) - \nabla g'(\mathcal{G}(\theta'))\| \cdot \|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|, \end{aligned}$$

where the first line comes from the first-order optimality condition of $\mathcal{G}(\theta') = \arg\min_{\varphi\in\Theta} g'(\varphi)$.

Since $\nabla g(\varphi) = \mathbb{E}_{Z\sim\mathcal{D}_t(\theta)}[\nabla_\theta \ell(Z,\varphi)]$ and $\nabla g'(\varphi) = \mathbb{E}_{Z\sim\mathcal{D}_t(\theta')}[\nabla_\theta \ell(Z,\varphi)]$, we can observe that β_z-smoothness and the Kantorovich-Rubinstein duality of the $\mathcal{W}_1$ distance implies that

$$\begin{aligned} \|\nabla g(\mathcal{G}(\theta')) - \nabla g'(\mathcal{G}(\theta'))\| &\le \|\mathbb{E}_{Z\sim\mathcal{D}_t(\theta)}[\nabla_\theta \ell(Z, \mathcal{G}(\theta'))] - \mathbb{E}_{Z\sim\mathcal{D}_t(\theta')}[\nabla_\theta \ell(Z, \mathcal{G}(\theta'))]\| \\ &\le \beta_z \cdot \mathcal{W}_1\big(\mathcal{D}_t(\theta), \mathcal{D}_t(\theta')\big) \end{aligned}$$

and by ϵ-sensitivity of $\mathcal{D}(\theta)$, we have

$$\begin{aligned} \mathcal{W}_1(\mathcal{D}_t(\theta), \mathcal{D}_t(\theta')) &= \mathcal{W}_1\left((1-\alpha_t)\mathcal{D}(\theta) + \alpha_t P_t, (1-\alpha_t)\mathcal{D}(\theta') + \alpha_t P_t\right) \\ &\le (1-\alpha_t)\mathcal{W}_1\left(\mathcal{D}(\theta), \mathcal{D}(\theta')\right) \\ &\le (1-\alpha_t)\epsilon\|\theta - \theta'\|. \end{aligned}$$

Therefore we have

$$\begin{aligned} \mu\|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|^2 &\le -\nabla g(\mathcal{G}(\theta'))^\top \big(\mathcal{G}(\theta) - \mathcal{G}(\theta')\big) \\ &\le (1-\alpha_t)\epsilon\beta_z \cdot \|\theta - \theta'\| \cdot \|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|. \end{aligned}$$

If $\mathcal{G}(\theta) = \mathcal{G}(\theta')$, the desired inequality is trivial. Otherwise, dividing by $\|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|$ gives

$$\|\mathcal{G}(\theta) - \mathcal{G}(\theta')\| \le \frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\|\theta - \theta'\|,$$

for any θ, θ'. This immediately implies

$$\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\| \le \frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\|\theta_t - \theta_t^{\mathrm{PS}}\|,$$

as desired, as $\theta_{t+1} = \mathcal{G}(\theta_t)$ and θ_t^{PS} is defined such that $\theta_t^{\mathrm{PS}} = \mathcal{G}(\theta_t^{\mathrm{PS}})$. □

C.2 Proof of Theorem 3.2

Here we prove Theorem 3.2, restated below for the sake of readability.

Theorem 3.2. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4), $\ell(z,\theta)$ is μ-strongly convex in θ (10) and β_z-smooth in z (12), and $\mathrm{PR}_t(\theta)$ is L-Lipschitz (13) in θ. Then, if $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$, where $\alpha_{\min} = \min_{t\in[T]} \alpha_t$, the iterates of $\mathcal{A}_{\mathrm{RRM}}$ satisfy*

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RRM}}) \le \frac{L}{1-\gamma} \cdot \left(\|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PS}} \right), \tag{5}$$

for $\gamma = \frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu}$.

Proof. Lemma 3.1 immediately implies

$$\begin{aligned}\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\| &\le \|\theta_{t+1} - \theta_t^{\mathrm{PS}}\| + \|\theta_t^{\mathrm{PS}} - \theta_{t+1}^{\mathrm{PS}}\| \\ &\le \frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\|\theta_t - \theta_t^{\mathrm{PS}}\| + \Delta_t^{\mathrm{PS}}.\end{aligned}$$

Therefore we can use Lemma B.1 as

$$\sum_{t=1}^{T}(1-\gamma_t)\|\theta_t - \theta_t^{\mathrm{PS}}\| \le \|\theta_1 - \theta_1^{\mathrm{PS}}\| - \gamma_T\|\theta_T - \theta_T^{\mathrm{PS}}\| + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}},$$

with $\gamma_t = \frac{(1-\alpha_t)\epsilon\beta_z}{\mu}$.

For any choice of nonnegative weights satisfying $w_t \le \frac{1-\gamma_t}{L}$, we can conclude that

$$\begin{aligned}\sum_{t=1}^{T} w_t \left(\mathrm{R}_{\mathcal{D}_t}(\theta_t) - \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}}) \right) &\le \sum_{t=1}^{T} w_t L \|\theta_t - \theta_t^{\mathrm{PS}}\| \\ &\le \sum_{t=1}^{T}(1-\gamma_t)\|\theta_t - \theta_t^{\mathrm{PS}}\| \\ &\le \|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}}.\end{aligned}$$

Finally, we can observe that γ_t is *decreasing* in α_t, and thus the expression $\frac{1-\gamma_t}{L}$ is an *increasing* function of α_t. Therefore, if we let $\alpha_{\min} = \min_{t\in[T]} \alpha_t$ and choose

$$\gamma = \frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu},$$

then choosing a constant $w_t \equiv w = \frac{1-\gamma}{L}$ satisfies $w \le \frac{1-\gamma_t}{L}$ for all t. Then we have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RRM}}) &= \frac{1}{w}\sum_{t=1}^{T} w \left(\mathrm{R}_{\mathcal{D}_t}(\theta_t) - \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}}) \right) \\ &\le \frac{L}{1-\gamma} \cdot \left(\|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PS}} \right),\end{aligned}$$

as in (5). □

C.3 Proof of Proposition 3.4

Here we prove Proposition 3.4, restated below for the sake of readability.

Proposition 3.4. *Fix $t \in \mathbb{N}$. Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4) and $\ell(z, \theta)$ is μ-strongly convex in θ (10) and β_z-smooth in z (12). Further, suppose that $\mathcal{W}_1(\mathcal{D}(\theta^{\mathrm{PS}}_{t+1}), P_{t+1}) \leq C$ for some $C < \infty$. Then, if $\epsilon < \frac{\mu}{(1-\alpha_t)\beta_z}$, we have*

$$\Delta^{\mathrm{PS}}_t \leq \frac{\beta_z}{\mu - \beta_z(1-\alpha_t)\epsilon} \left(\alpha_t \mathcal{W}_1(P_t, P_{t+1}) + |\alpha_t - \alpha_{t+1}|C\right).$$

Proof. Let us define $\theta \coloneqq \theta^{\mathrm{PS}}_t$, $\theta' \coloneqq \theta^{\mathrm{PS}}_{t+1}$ and $\mathcal{D}^{\mathrm{PS}}_t \coloneqq \mathcal{D}_t(\theta^{\mathrm{PS}}_t)$, $\mathcal{D}^{\mathrm{PS}}_{t+1} \coloneqq \mathcal{D}_{t+1}(\theta^{\mathrm{PS}}_{t+1})$ for simplicity. By the definition of performative stability, θ^{PS}_t and $\theta^{\mathrm{PS}}_{t+1}$ are the constrained minimizers of $\mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}$ and $\mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}$, respectively.

The first-order optimality conditions yield

$$\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta)^\top(\theta' - \theta) \geq 0, \qquad \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta')^\top(\theta - \theta') \geq 0.$$

If $\theta = \theta'$, the desired inequality is trivial. Otherwise, strong monotonicity of $\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}$ and the above inequalities imply

$$\begin{aligned}
\mu\|\theta - \theta'\|^2 &\leq \left(\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta') - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta)\right)^\top (\theta' - \theta) \\
&= \left(\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta') - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta')\right)^\top (\theta' - \theta) \\
&\quad + \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta')^\top(\theta' - \theta) - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta)^\top(\theta' - \theta) \\
&\leq \left(\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta') - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta')\right)^\top (\theta' - \theta) \\
&\leq \|\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta') - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta')\| \cdot \|\theta - \theta'\|.
\end{aligned}$$

Dividing by $\|\theta - \theta'\|$ gives

$$\|\theta - \theta'\| \leq \frac{1}{\mu}\|\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta') - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta')\|,$$

and plugging back in $\theta \coloneqq \theta^{\mathrm{PS}}_t$ and $\theta' \coloneqq \theta^{\mathrm{PS}}_{t+1}$,

$$\Delta^{\mathrm{PS}}_t = \|\theta^{\mathrm{PS}}_t - \theta^{\mathrm{PS}}_{t+1}\| \leq \frac{1}{\mu}\|\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta^{\mathrm{PS}}_{t+1}) - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta^{\mathrm{PS}}_{t+1})\|.$$

By β_z-smoothness, Kantorovich-Rubinstein duality, and Lemma B.4, we have

$$\begin{aligned}
&\frac{1}{\mu}\|\nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_t}(\theta^{\mathrm{PS}}_{t+1}) - \nabla \mathrm{R}_{\mathcal{D}^{\mathrm{PS}}_{t+1}}(\theta^{\mathrm{PS}}_{t+1})\| \leq \frac{\beta_z}{\mu} \cdot \mathcal{W}_1(\mathcal{D}^{\mathrm{PS}}_t, \mathcal{D}^{\mathrm{PS}}_{t+1}) \\
&= \frac{\beta_z}{\mu} \cdot \mathcal{W}_1\left((1-\alpha_t)\mathcal{D}(\theta^{\mathrm{PS}}_t) + \alpha_t P_t, (1-\alpha_{t+1})\mathcal{D}(\theta^{\mathrm{PS}}_{t+1}) + \alpha_{t+1}P_{t+1}\right) \\
&\leq \frac{\beta_z}{\mu}\big((1-\alpha_t)\mathcal{W}_1(\mathcal{D}(\theta^{\mathrm{PS}}_t), \mathcal{D}(\theta^{\mathrm{PS}}_{t+1})) + \alpha_t \mathcal{W}_1(P_t, P_{t+1}) \\
&\qquad\qquad + |\alpha_t - \alpha_{t+1}| \cdot \mathcal{W}_1(\mathcal{D}(\theta^{\mathrm{PS}}_{t+1}), P_{t+1})\big).
\end{aligned}$$

Now, $\mathcal{W}_1(\mathcal{D}(\theta^{\mathrm{PS}}_t), \mathcal{D}(\theta^{\mathrm{PS}}_{t+1})) \leq \epsilon\|\theta^{\mathrm{PS}}_t - \theta^{\mathrm{PS}}_{t+1}\| = \epsilon\Delta^{\mathrm{PS}}_t$ yields

$$\Delta^{\mathrm{PS}}_t \leq \frac{\beta_z}{\mu}\left((1-\alpha_t)\epsilon\Delta^{\mathrm{PS}}_t + \alpha_t \mathcal{W}_1(P_t, P_{t+1}) + |\alpha_t - \alpha_{t+1}| \cdot \mathcal{W}_1(\mathcal{D}(\theta^{\mathrm{PS}}_{t+1}), P_{t+1})\right).$$

Rearranging terms, we have

$$\left(1 - \frac{\beta_z}{\mu} \cdot (1-\alpha_t)\epsilon\right)\Delta^{\mathrm{PS}}_t \leq \frac{\beta_z}{\mu}\left(\alpha_t \mathcal{W}_1(P_t, P_{t+1}) + |\alpha_t - \alpha_{t+1}| \cdot \mathcal{W}_1(\mathcal{D}(\theta^{\mathrm{PS}}_{t+1}), P_{t+1})\right),$$

and $\mathcal{W}_1(\mathcal{D}(\theta^{\mathrm{PS}}_{t+1}), P_{t+1}) \leq C$ yields the given statement. □

D Proofs for repeated gradient descent

D.1 Proof of Lemma D.1

Here we prove Lemma D.1, which provides a one-step inequality for RGD.

Lemma D.1. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive and $\ell(z,\theta)$ is μ-strongly convex (10) and β_θ-smooth (11) in θ and β_z-smooth (12) in z. Then the iterates of $\mathcal{A}_{\mathrm{RGD}}$ satisfy*

$$\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\| \le \gamma_t \|\theta_t - \theta_t^{\mathrm{PS}}\| + \Delta_t^{\mathrm{PS}}, \tag{16}$$

where $\gamma_t \coloneqq 1 - \eta_t(\mu - (1-\alpha_t)\epsilon\beta_z) + \eta_t^2(\beta_\theta^2 + (1-\alpha_t)^2\epsilon^2\beta_z^2)$.

Proof. The proof is similar to that of Proposition 2.6 in [17], but our proof relies only on the first-order optimality conditions in the definition of θ_t^{PS}, rather than assuming $\mathbb{E}\mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}}) = 0$, which in turn requires θ_t^{PS} to lie in the interior of Θ, as in [17].

By definition of θ_t^{PS}, we can observe that

$$\nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})^\top (\vartheta - \theta_t^{\mathrm{PS}}) \ge 0, \qquad \forall \vartheta \in \Theta.$$

For any closed convex Θ, a projection $y = \Pi_\Theta(u)$ of u must satisfy

$$(y-u)^\top (\vartheta - y) \ge 0, \qquad \forall \vartheta \in \Theta.$$

for $y = \theta_t^{\mathrm{PS}}$ and $u = \theta_t^{\mathrm{PS}} - \eta_t \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})$, this is equivalent to

$$\eta_t \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})^\top (\vartheta - \theta_t^{\mathrm{PS}}) \ge 0, \qquad \forall \vartheta \in \Theta,$$

which is true by the definition of θ_t^{PS}. This shows that y is a projection of u onto Θ, i.e.,

$$\theta_t^{\mathrm{PS}} = \Pi_\Theta\left(\theta_t^{\mathrm{PS}} - \eta_t \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)$$

for any $\eta_t > 0$. Therefore, by nonexpansiveness of Π_Θ,

$$\begin{aligned}
\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\| &\le \|\theta_{t+1} - \theta_t^{\mathrm{PS}}\| + \|\theta_t^{\mathrm{PS}} - \theta_{t+1}^{\mathrm{PS}}\| \\
&= \left\|\Pi_\Theta\left(\theta_t - \eta_t \nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t)\right) - \Pi_\Theta\left(\theta_t^{\mathrm{PS}} - \eta_t \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)\right\| + \Delta_t^{\mathrm{PS}} \\
&\le \left\|\theta_t - \theta_t^{\mathrm{PS}} - \eta_t\left(\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)\right\| + \Delta_t^{\mathrm{PS}}.
\end{aligned}$$

Squaring and expanding the first norm gives

$$\begin{aligned}
&\left\|\theta_t - \theta_t^{\mathrm{PS}} - \eta_t\left(\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)\right\|^2 \\
&= \|\theta_t - \theta_t^{\mathrm{PS}}\|^2 - 2\eta_t\left(\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \\
&\quad + \eta_t^2 \left\|\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right\|^2.
\end{aligned} \tag{17}$$

The inner product term of (17) can be decomposed as

$$\begin{aligned}
&\left(\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \\
&= \left(\mathbb{E}_{Z\sim\mathcal{D}_t}[\nabla_\theta \ell(Z,\theta_t)] - \mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\nabla_\theta \ell(Z,\theta_t^{\mathrm{PS}})]\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \\
&= \left(\mathbb{E}_{Z\sim\mathcal{D}_t}[\nabla_\theta \ell(Z,\theta_t)] - \mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\nabla_\theta \ell(Z,\theta_t)]\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \\
&\quad + \left(\mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\nabla_\theta \ell(Z,\theta_t)] - \mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\nabla_\theta \ell(Z,\theta_t^{\mathrm{PS}})]\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}).
\end{aligned} \tag{18}$$

Define the auxiliary function

$$\xi_t(z) \coloneqq \nabla_\theta \ell(z,\theta_t)^\top (\theta_t - \theta_t^{\mathrm{PS}}),$$

which is $\beta_z\|\theta_t - \theta_t^{\mathrm{PS}}\|$-Lipschitz in z. The first term of the RHS of (18) satisfies

$$\begin{aligned}
&\left(\mathbb{E}_{Z\sim\mathcal{D}_t}[\nabla_\theta \ell(Z,\theta_t)] - \mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\nabla_\theta \ell(Z,\theta_t)]\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \\
&= \mathbb{E}_{Z\sim\mathcal{D}_t}[\xi_t(Z)] - \mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\xi_t(Z)] \\
&\geq -\beta_z\|\theta_t - \theta_t^{\mathrm{PS}}\| \cdot \mathcal{W}_1(\mathcal{D}_t, \mathcal{D}_t^{\mathrm{PS}}) \\
&\geq -\beta_z\|\theta_t - \theta_t^{\mathrm{PS}}\| \cdot (1-\alpha_t)\epsilon\|\theta_t - \theta_t^{\mathrm{PS}}\| \\
&= -(1-\alpha_t)\epsilon\beta_z\|\theta_t - \theta_t^{\mathrm{PS}}\|^2,
\end{aligned} \tag{19}$$

where the first inequality follows from the fact that $\xi_t(z)$ is Lipschitz, and the second inequality is by $(1-\alpha_t)\epsilon$-sensitivity of $\mathcal{D}_t(\theta)$. The second term of the RHS of (18) satisfies

$$\left(\mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\nabla_\theta \ell(Z,\theta_t)] - \mathbb{E}_{Z\sim\mathcal{D}_t^{\mathrm{PS}}}[\nabla_\theta \ell(Z,\theta_t^{\mathrm{PS}})]\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \geq \mu\|\theta_t - \theta_t^{\mathrm{PS}}\|^2, \tag{20}$$

by μ-strong monotonicity of the gradients. Plugging (19), (20) into (18), we have

$$\begin{aligned}
&\left(\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \\
&= \left(\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t)\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) + \left(\nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)^\top (\theta_t - \theta_t^{\mathrm{PS}}) \\
&\geq \left(\mu - (1-\alpha_t)\epsilon\beta_z\right)\|\theta_t - \theta_t^{\mathrm{PS}}\|^2.
\end{aligned} \tag{21}$$

The gradient-difference norm term of (17) satisfies

$$\begin{aligned}
&\left\|\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right\|^2 \\
&\leq 2\left\|\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t^{\mathrm{PS}})\right\|^2 + 2\left\|\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t^{\mathrm{PS}}) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right\|^2 \\
&\leq 2\beta_\theta^2\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + 2(1-\alpha_t)^2\epsilon^2\beta_z^2\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 \\
&= 2(\beta_\theta^2 + (1-\alpha_t)^2\epsilon^2\beta_z^2)\|\theta_t - \theta_t^{\mathrm{PS}}\|^2,
\end{aligned} \tag{22}$$

by smoothness and sensitivity.

Plugging (21) and (22) into (17), we have

$$\begin{aligned}
&\left\|\theta_t - \theta_t^{\mathrm{PS}} - \eta_t\left(\nabla \mathrm{R}_{\mathcal{D}_t}(\theta_t) - \nabla \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)\right\|^2 \\
&\leq \left(1 - 2\eta_t(\mu - (1-\alpha_t)\epsilon\beta_z) + 2\eta_t^2(\beta_\theta^2 + (1-\alpha_t)^2\epsilon^2\beta_z^2)\right)\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 \\
&\leq \left(1 - \eta_t(\mu - (1-\alpha_t)\epsilon\beta_z) + \eta_t^2(\beta_\theta^2 + (1-\alpha_t)^2\epsilon^2\beta_z^2)\right)^2\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 \\
&= \gamma_t^2\|\theta_t - \theta_t^{\mathrm{PS}}\|^2,
\end{aligned}$$

where the second inequality uses $1 - 2x \leq (1-x)^2$ for simplification. Therefore we have

$$\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\| \leq \gamma_t\|\theta_t - \theta_t^{\mathrm{PS}}\| + \Delta_t^{\mathrm{PS}},$$

which is equivalent to (16). □

D.2 Proof of Theorem 4.1

Here we prove Theorem 4.1, restated below for the sake of readability.

Theorem 4.1. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive (4), $\ell(z,\theta)$ is μ-strongly convex (10) and β_θ-smooth (11) in θ and β_z-smooth (12) in z, and $\mathrm{PR}_t(\theta)$ is L-Lipschitz (13) in θ. Then, if $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$, where $\alpha_{\min} = \min_{t\in[T]}\alpha_t$, the iterates of $\mathcal{A}_{\mathrm{RGD}}$ with step size*

$$\eta_t = \frac{\mu - (1-\alpha_t)\epsilon\beta_z}{2(\beta_\theta^2 + (1-\alpha_t)^2\epsilon^2\beta_z^2)}$$

satisfy

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RGD}}) \leq \frac{L}{1-\overline{\gamma}} \cdot \left(\|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}}\right), \tag{6}$$

for $\overline{\gamma} = 1 - \frac{(\mu-(1-\alpha_{\min})\epsilon\beta_z)^2}{4(\beta_\theta^2+(1-\alpha_{\min})^2\epsilon^2\beta_z^2)}$.

Proof. By Lemma D.1 (proven in Section D.1), we have

$$\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\| \le \gamma_t \|\theta_t - \theta_t^{\mathrm{PS}}\| + \Delta_t^{\mathrm{PS}},$$

where $\gamma_t \coloneqq 1 - \eta_t(\mu - (1-\alpha_t)\epsilon\beta_z) + \eta_t^2(\beta_\theta^2 + (1-\alpha_t)^2\epsilon^2\beta_z^2)$.

Lemma B.1 yields

$$\sum_{t=1}^{T}(1-\gamma_t)\|\theta_t - \theta_t^{\mathrm{PS}}\| \le \|\theta_1 - \theta_1^{\mathrm{PS}}\| - \gamma_T\|\theta_T - \theta_T^{\mathrm{PS}}\| + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}} \le \|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}},$$

and thus

$$\begin{aligned}\sum_{t=1}^{T} w_t \left(\mathrm{R}_{\mathcal{D}_t}(\theta_t) - \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right) &\le \sum_{t=1}^{T} w_t L\|\theta_t - \theta_t^{\mathrm{PS}}\| \\ &\le \sum_{t=1}^{T}(1-\gamma_t)\|\theta_t - \theta_t^{\mathrm{PS}}\| \\ &\le \|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}},\end{aligned}$$

for any choice of weights satisfying $w_t \le \frac{1-\gamma_t}{L}$. Therefore, if we set

$$\overline{\gamma} = 1 - \frac{(\mu - (1-\alpha_{\min})\epsilon\beta_z)^2}{4(\beta_\theta^2 + (1-\alpha_{\min})^2\epsilon^2\beta_z^2)},$$

then choosing a constant $w_t \equiv w = \frac{1-\overline{\gamma}}{L}$ satisfies $w \le \frac{1-\gamma_t}{L}$ for all t. Then we have

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\mathrm{RGD}}) = \frac{1}{w}\sum_{t=1}^{T} w\left(\mathrm{R}_{\mathcal{D}_t}(\theta_t) - \mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right) \le \frac{1}{w}\left(\|\theta_1 - \theta_1^{\mathrm{PS}}\| + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}}\right),$$

which is equivalent to (6). □

E Proofs for greedy SGD

E.1 Proof of Lemma E.1

Here we prove Lemma E.1, which provides a one-step inequality for greedy SGD.

Lemma E.1. *Suppose that $\mathcal{D}(\theta)$ is ϵ-sensitive and $\ell(z,\theta)$ is μ-strongly convex (10) and β_θ-smooth (11) in θ and β_z-smooth (12) in z. Also, suppose that we have quadratic-bounded variance (Ass. 7). Then, the iterates of $\mathcal{A}_{\text{RSGD-G}}$ satisfy*

$$\mathbb{E}[\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\|^2] \le \gamma_t'\mathbb{E}[\|\theta_t - \theta_t^{\mathrm{PS}}\|^2] + 2\Delta_t^{\mathrm{PS}}\sqrt{\gamma_t'\mathbb{E}[\|\theta_t - \theta_t^{\mathrm{PS}}\|^2]} + 2\eta_t^2\sigma^2 + 2(\Delta_t^{\mathrm{PS}})^2, \quad (23)$$

where $\gamma_t' \coloneqq 1 - 2\eta_t(\mu - (1-\alpha_t)\epsilon\beta_z) + \eta_t^2 C_V^2(1 + \frac{(1-\alpha_t)\epsilon\beta_z}{\mu})^2$ and $\eta_t > 0$ is small enough so that $\gamma_t' \ge 0$.

Proof. We have

$$\begin{aligned}\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\| &\le \|\theta_{t+1} - \theta_t^{\mathrm{PS}}\| + \|\theta_t^{\mathrm{PS}} - \theta_{t+1}^{\mathrm{PS}}\| \\ &= \left\|\Pi_\Theta(\theta_t - \eta_t\nabla_\theta\ell(Z_t,\theta_t)) - \Pi_\Theta(\theta_t^{\mathrm{PS}})\right\| + \|\theta_t^{\mathrm{PS}} - \theta_{t+1}^{\mathrm{PS}}\| \\ &\le \|\theta_t - \theta_t^{\mathrm{PS}} - \eta_t\nabla_\theta\ell(Z_t,\theta_t)\| + \Delta_t^{\mathrm{PS}},\end{aligned} \quad (24)$$

since $\theta_t^{\mathrm{PS}} = \Pi_\Theta(\theta_t^{\mathrm{PS}})$. Also, under Eq. (7), we have (conditioned on θ_t)

$$\begin{aligned}\mathbb{E}_{Z_t\sim\mathcal{D}_t}[\|\nabla_\theta\ell(Z_t,\theta_t)\|^2] &\le \sigma^2 + C_V^2\|\theta_t - \theta_{\mathcal{D}_t}^\star\|^2 \\ &\le \sigma^2 + C_V^2\left(\|\theta_t - \theta_t^{\mathrm{PS}}\| + \|\theta_t^{\mathrm{PS}} - \theta_{\mathcal{D}_t}^\star\|\right)^2 \\ &\le \sigma^2 + C_V^2\left(1 + \frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\right)^2\|\theta_t - \theta_t^{\mathrm{PS}}\|^2\end{aligned} \quad (25)$$

where the last inequality follows from Lemma 3.1. We can observe that

$$\begin{aligned}
&\mathbb{E}_{Z_t\sim\mathcal{D}_t}\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2\\
&\le\|\theta_t-\theta_t^{\mathrm{PS}}\|^2-2\eta_t\mathbb{E}_{Z_t\sim\mathcal{D}_t}[\nabla_\theta\ell(Z_t,\theta_t)^\top(\theta_t-\theta_t^{\mathrm{PS}})]+\eta_t^2\mathbb{E}_{Z_t\sim\mathcal{D}_t}[\|\nabla_\theta\ell(Z_t,\theta_t)\|^2]\\
&=\|\theta_t-\theta_t^{\mathrm{PS}}\|^2-2\eta_t\nabla\mathrm{R}_{\mathcal{D}_t}(\theta_t)^\top(\theta_t-\theta_t^{\mathrm{PS}})+\eta_t^2\mathbb{E}_{Z_t\sim\mathcal{D}_t}[\|\nabla_\theta\ell(Z_t,\theta_t)\|^2].
\end{aligned}$$

Since we have $\nabla\mathrm{R}_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})^\top(\theta_t-\theta_t^{\mathrm{PS}})\ge 0$ by the first-order optimality condition of θ_t^{PS}, we have

$$\begin{aligned}
&\nabla\mathrm{R}_{\mathcal{D}_t}(\theta_t)^\top(\theta_t-\theta_t^{\mathrm{PS}})\\
&=\left(\nabla\mathrm{R}_{\mathcal{D}_t}(\theta_t)-\nabla R_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)^\top(\theta_t-\theta_t^{\mathrm{PS}})+\nabla R_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})^\top(\theta_t-\theta_t^{\mathrm{PS}})\\
&\ge\left(\nabla\mathrm{R}_{\mathcal{D}_t}(\theta_t)-\nabla R_{\mathcal{D}_t^{\mathrm{PS}}}(\theta_t^{\mathrm{PS}})\right)^\top(\theta_t-\theta_t^{\mathrm{PS}})\\
&\ge(\mu-(1-\alpha_t)\epsilon\beta_z)\,\|\theta_t-\theta_t^{\mathrm{PS}}\|^2,
\end{aligned}$$

by (21) from the proof of Lemma D.1. We can also use (25) to obtain

$$\begin{aligned}
&\|\theta_t-\theta_t^{\mathrm{PS}}\|^2-2\eta_t\nabla\mathrm{R}_{\mathcal{D}_t}(\theta_t)^\top(\theta_t-\theta_t^{\mathrm{PS}})+\eta_t^2\mathbb{E}_{Z_t\sim\mathcal{D}_t}[\|\nabla_\theta\ell(Z_t,\theta_t)\|^2]\\
&\le\|\theta_t-\theta_t^{\mathrm{PS}}\|^2-2\eta_t\,(\mu-(1-\alpha_t)\epsilon\beta_z)\,\|\theta_t-\theta_t^{\mathrm{PS}}\|^2\\
&\quad+\eta_t^2\left(\sigma^2+C_V^2\left(1+\frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\right)^2\|\theta_t-\theta_t^{\mathrm{PS}}\|^2\right)\\
&=\left(1-2\eta_t\,(\mu-(1-\alpha_t)\epsilon\beta_z)+\eta_t^2C_V^2\left(1+\frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\right)^2\right)\|\theta_t-\theta_t^{\mathrm{PS}}\|^2+\eta_t^2\sigma^2.
\end{aligned}$$

Therefore, choosing

$$\gamma_t'\coloneqq 1-2\eta_t\,(\mu-(1-\alpha_t)\epsilon\beta_z)+\eta_t^2C_V^2\left(1+\frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\right)^2$$

to simplify as

$$\mathbb{E}_{Z_t\sim\mathcal{D}_t}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2]\le\gamma_t'\|\theta_t-\theta_t^{\mathrm{PS}}\|^2+\eta_t^2\sigma^2,$$

conditioned on θ_t, or

$$\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2]\le\gamma_t'\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}\|^2]+\eta_t^2\sigma^2,\tag{26}$$

where the expectations are now taken on the entire history. Therefore, starting from (24), we have

$$\|\theta_{t+1}-\theta_{t+1}^{\mathrm{PS}}\|^2\le\left(\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla\ell(Z_t,\theta_t)\|+\Delta_t^{\mathrm{PS}}\right)^2,$$

and therefore

$$\begin{aligned}
&\mathbb{E}[\|\theta_{t+1}-\theta_{t+1}^{\mathrm{PS}}\|^2]\\
&\le\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2]+2\Delta_t^{\mathrm{PS}}\cdot\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|]+(\Delta_t^{\mathrm{PS}})^2\\
&\le\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2]+2\Delta_t^{\mathrm{PS}}\cdot\sqrt{\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2]}+(\Delta_t^{\mathrm{PS}})^2.
\end{aligned}$$

Finally, we can use (26), the inequality $\sqrt{a+b}\le\sqrt{a}+\sqrt{b}$, and AM-GM to obtain

$$\begin{aligned}
&\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2]+2\Delta_t^{\mathrm{PS}}\cdot\sqrt{\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}-\eta_t\nabla_\theta\ell(Z_t,\theta_t)\|^2]}+(\Delta_t^{\mathrm{PS}})^2\\
&\le\gamma_t'\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}\|^2]+\eta_t^2\sigma^2+2\Delta_t^{\mathrm{PS}}\sqrt{\gamma_t'\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}\|^2]+\eta_t^2\sigma^2}+(\Delta_t^{\mathrm{PS}})^2\\
&\le\gamma_t'\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}\|^2]+\eta_t^2\sigma^2+2\Delta_t^{\mathrm{PS}}\sqrt{\gamma_t'\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}\|^2]}+2\eta_t\sigma\Delta_t^{\mathrm{PS}}+(\Delta_t^{\mathrm{PS}})^2\\
&\le\gamma_t'\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}\|^2]+2\Delta_t^{\mathrm{PS}}\sqrt{\gamma_t'\mathbb{E}[\|\theta_t-\theta_t^{\mathrm{PS}}\|^2]}+2\eta_t^2\sigma^2+2(\Delta_t^{\mathrm{PS}})^2,
\end{aligned}$$

which proves the given statement. □

E.2 Proof of Lemma E.2

Here we prove Lemma E.2, which provides upper bounds of $\mathbb{E}\|\theta_T - \theta_T^{\mathrm{PS}}\|^2$ for a fixed T.

Lemma E.2. *Suppose that $\mathcal{D}\colon \Theta \to \mathcal{P}(\mathcal{Z})$ is ϵ-sensitive and $\ell(z,\theta)$ is μ-strongly convex (10) and β_θ-smooth (11) in θ and β_z-smooth (12) in z Also, suppose that $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$ and we have quadratic-bounded variance (Ass. 7). Then, $\mathcal{A}_{\text{RSGD-G}}$ with step sizes*

$$\eta_t = \frac{4}{A(t+t_0)}, \quad t_0 = \max\left\{1, \left\lceil \frac{8B}{A^2} \right\rceil, \left\lceil \frac{8\mu}{A} \right\rceil \right\} \tag{27}$$

for some constants $A, B > 0$ defined as

$$A \coloneqq 2\left(\mu - (1-\alpha_{\min})\epsilon\beta_z\right), \qquad B \coloneqq C_V^2\left(1 + \frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu}\right)^2$$

satisfies the following inequality for $T \geq 0$,

$$\mathbb{E}\|\theta_{T+1} - \theta_{T+1}^{\mathrm{PS}}\|^2$$
$$\leq \left(\frac{t_0+1}{T+t_0+1}\right)^{3/2}\|\theta_1 - \theta_1^{\mathrm{PS}}\|^2 + \frac{64\sigma^2}{A^2(T+t_0)} + \frac{2}{(T+t_0+1)^{3/2}}\sum_{t=1}^{T}(t+t_0+1)^{5/2}(\Delta_t^{\mathrm{PS}})^2.$$

Proof. First, note that the case $T = 0$ is trivial, and that we have $A > 0$ since $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$.

Let us define $e_t = \theta_t - \theta_t^{\mathrm{PS}}$ for simplicity. We also simplify γ_t' using

$$\begin{aligned}\gamma_t' &= 1 - 2\eta_t\left(\mu - (1-\alpha_t)\epsilon\beta_z\right) + \eta_t^2 C_V^2\left(1 + \frac{(1-\alpha_t)\epsilon\beta_z}{\mu}\right)^2 \\ &\leq 1 - \eta_t \cdot \underbrace{2\left(\mu - (1-\alpha_{\min})\epsilon\beta_z\right)}_{\triangleq A} + \eta_t^2 \cdot \underbrace{C_V^2\left(1 + \frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu}\right)^2}_{\triangleq B} =: \overline{\gamma}_t', \end{aligned} \tag{28}$$

where $\alpha_{\min} = \min_{t\in[T]}\alpha_t$. Recall that we use step sizes

$$\eta_t = \frac{4}{A(t+t_0)}, \quad t_0 = \max\left\{1, \left\lceil \frac{8B}{A^2} \right\rceil, \left\lceil \frac{8\mu}{A} \right\rceil \right\},$$

for time-independent constants A, B defined in (28). Since $1 + t_0 > \frac{8\mu}{A}$, we have $\frac{4}{A(1+t_0)} < \frac{1}{2\mu}$ and therefore $\eta_t = \frac{4}{A(t+t_0)} < \frac{1}{2\mu}$ for all $t \in [T]$, which ensures that $\overline{\gamma}_t' \geq \gamma_t' \geq 0$ for all t.

By (23) of Lemma E.1, we have

$$\mathbb{E}\|e_{t+1}\|^2 \leq \overline{\gamma}_t'\mathbb{E}\|e_t\|^2 + 2\Delta_t^{\mathrm{PS}}\sqrt{\overline{\gamma}_t'\mathbb{E}\|e_t\|^2} + 2\eta_t^2\sigma^2 + 2(\Delta_t^{\mathrm{PS}})^2, \quad \text{where } \overline{\gamma}_t' = 1 - \eta_t A + \eta_t^2 B.$$

Note that

$$\begin{aligned}\overline{\gamma}_t' = 1 - \eta_t A + \eta_t^2 B &= 1 - \frac{4}{t+t_0} + \frac{16B}{A^2(t+t_0)^2} \\ &\leq 1 - \frac{4}{t+t_0} + \frac{8B}{A^2}\cdot\frac{2}{t_0(t+t_0)} \leq 1 - \frac{4}{t+t_0} + \frac{2}{t+t_0} = 1 - \frac{2}{t+t_0}.\end{aligned}$$

Then we have

$$\begin{aligned}\mathbb{E}\|e_{t+1}\|^2 &\leq \overline{\gamma}_t'\mathbb{E}\|e_t\|^2 + 2\Delta_t^{\mathrm{PS}}\sqrt{\overline{\gamma}_t'\mathbb{E}\|e_t\|^2} + 2\eta_t^2\sigma^2 + 2(\Delta_t^{\mathrm{PS}})^2 \\ &\leq \overline{\gamma}_t'\mathbb{E}\|e_t\|^2 + 2\Delta_t^{\mathrm{PS}}\sqrt{\mathbb{E}\|e_t\|^2} + 2\eta_t^2\sigma^2 + 2(\Delta_t^{\mathrm{PS}})^2 \\ &\leq \left(1 - \frac{2}{t+t_0}\right)\mathbb{E}\|e_t\|^2 + \frac{32\sigma^2}{A^2(t+t_0)^2} + 2(\Delta_t^{\mathrm{PS}})^2 \\ &\quad + \left(\frac{1}{2(t+t_0)}\mathbb{E}\|e_t\|^2 + 2(t+t_0)(\Delta_t^{\mathrm{PS}})^2\right) \\ &= \left(1 - \frac{3/2}{t+t_0}\right)\mathbb{E}\|e_t\|^2 + \frac{32\sigma^2}{A^2(t+t_0)^2} + 2(t+t_0+1)(\Delta_t^{\mathrm{PS}})^2.\end{aligned}$$

For simplicity, let us define

$$F_t \coloneqq \mathbb{E}\|e_t\|^2, \quad a_t \coloneqq \left(1 - \frac{3/2}{t+t_0}\right), \quad b_t \coloneqq \frac{32\sigma^2}{A^2(t+t_0)^2}, \quad d_t \coloneqq 2(t+t_0+1)(\Delta_t^{\mathrm{PS}})^2$$

and write

$$F_{t+1} \leq a_t F_t + b_t + d_t.$$

Since $(1 - \frac{3/2}{s+t_0}) \leq \exp(-\frac{3/2}{s+t_0})$, we have

$$\begin{aligned}
\prod_{s=t+1}^{T} a_s &\leq \prod_{s=t+1}^{T} \exp\left(-\frac{3/2}{s+t_0}\right) \\
&= \exp\left(-\sum_{s=t+1}^{T} \frac{3/2}{s+t_0}\right) \\
&\leq \exp\left(-\frac{3}{2} \cdot \log \frac{T+t_0+1}{t+t_0+1}\right) \\
&= \left(\frac{t+t_0+1}{T+t_0+1}\right)^{3/2}.
\end{aligned}$$

Also, we can observe that

$$\begin{aligned}
\frac{64\sigma^2}{A^2(t+t_0)} - \left(1 - \frac{3/2}{t+t_0}\right)\frac{64\sigma^2}{A^2(t+t_0-1)} &= \frac{32\sigma^2}{A^2(t+t_0)(t+t_0-1)} \\
&\geq \frac{32\sigma^2}{A^2(t+t_0)^2},
\end{aligned}$$

and thus $N_{t+1} \geq a_t N_t + b_t$ for $N_t = \frac{64\sigma^2}{A^2(t+t_0-1)}$. This implies that

$$F_{t+1} - N_{t+1} \leq a_t(F_t - N_t) + d_t,$$

and therefore

$$\begin{aligned}
F_{T+1} &\leq N_{T+1} + \left(\prod_{t=1}^{T} a_t\right)(F_1 - N_1) + \sum_{t=1}^{T} d_t \left(\prod_{s=t+1}^{T} a_s\right) \\
&\leq \frac{64\sigma^2}{A^2(T+t_0)} + \left(\frac{t_0+1}{T+t_0+1}\right)^{3/2} \cdot F_1 + 2\sum_{t=1}^{T}(t+t_0+1)(\Delta_t^{\mathrm{PS}})^2 \cdot \left(\frac{t+t_0+1}{T+t_0+1}\right)^{3/2},
\end{aligned}$$

which is equivalent to the given statement. □

E.3 Proof of Proposition 4.2

Here we prove Proposition 4.2, restated below for the sake of readability.

Proposition 4.2. *Let the same assumptions as in Theorem 4.1 and the quadratic variance bound (7) hold. Then, the iterates of $\mathcal{A}_{\text{SGD-greedy}}$ with step size $\eta_t \propto \frac{1}{t+t_0}$, for some constant $t_0 > 0$, satisfy*

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\text{SGD-greedy}}) \leq \mathcal{O}\left(T^{1/2} + T^{1/4}\left(\sum_{t=1}^{T-1}(t+t_0+1)^{5/2}(\Delta_t^{\mathrm{PS}})^2\right)^{1/2}\right). \tag{8}$$

Proof. Using the same choice of learning rates η_t and the same auxiliary definitions (e.g., e_t, F_t) as in Lemma E.2, we start from the result of Lemma E.2 with $T \leftarrow t-1$,

$$\mathbb{E}\|e_t\|^2 \leq \frac{64\sigma^2}{A^2(t+t_0-1)} + \left(\frac{t_0+1}{t+t_0}\right)^{3/2} \cdot \|\theta_1 - \theta_1^{\mathrm{PS}}\|^2 + \frac{2}{(t+t_0)^{3/2}}\sum_{s=1}^{t-1}(s+t_0+1)^{5/2}(\Delta_s^{\mathrm{PS}})^2,$$

which implies

$$\sqrt{\mathbb{E}\|e_t\|^2} \le \frac{8\sigma}{A(t+t_0-1)^{1/2}} + \left(\frac{t_0+1}{t+t_0}\right)^{3/4} \cdot \|\theta_1 - \theta_1^{\mathrm{PS}}\| + \frac{2}{(t+t_0)^{3/4}} \left(\sum_{s=1}^{t-1}(s+t_0+1)^{5/2}(\Delta_s^{\mathrm{PS}})^2\right)^{1/2},$$

since $\sqrt{a+b+c} \le \sqrt{a}+\sqrt{b}+\sqrt{c}$ for $a,b,c \ge 0$. We can observe that

$$\sum_{t=1}^{T} \frac{8\sigma}{A(t+t_0-1)^{1/2}} \le \frac{16\sigma}{A}(T+t_0)^{1/2},$$

$$\sum_{t=1}^{T} \left(\frac{t_0+1}{t+t_0}\right)^{3/4} \cdot \|\theta_1 - \theta_1^{\mathrm{PS}}\| \le 4(t_0+1)^{3/4}\|\theta_1 - \theta_1^{\mathrm{PS}}\| \cdot (T+t_0)^{1/4},$$

and we can substitute the last term with an upper bound

$$\sum_{s=1}^{t-1}(s+t_0+1)^{5/2}(\Delta_s^{\mathrm{PS}})^2 \le \sum_{s=1}^{T-1}(s+t_0+1)^{5/2}(\Delta_s^{\mathrm{PS}})^2,$$

so that

$$\sum_{t=1}^{T} \frac{2}{(t+t_0)^{3/4}} \left(\sum_{s=1}^{t-1}(s+t_0+1)^{5/2}(\Delta_s^{\mathrm{PS}})^2\right)^{1/2} \le 8(T+t_0)^{1/4} \left(\sum_{s=1}^{T-1}(s+t_0+1)^{5/2}(\Delta_s^{\mathrm{PS}})^2\right)^{1/2}$$

Finally, since PR_t is L-Lipschitz, we have

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\text{SGD-greedy}}) \le L\sum_{t=1}^{T}\mathbb{E}\|e_t\| \le L\sum_{t=1}^{T}\sqrt{\mathbb{E}\|e_t\|^2}$$

$$\le 8L(T+t_0)^{1/4}\left((t_0+1)^{3/4}\|\theta_1 - \theta_1^{\mathrm{PS}}\| + \left(\sum_{s=1}^{T-1}(s+t_0+1)^{5/2}(\Delta_s^{\mathrm{PS}})^2\right)^{1/2}\right) + \frac{16\sigma L}{A}(T+t_0)^{1/2},$$

and therefore

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\text{SGD-greedy}}) \le \mathcal{O}\left(T^{1/2} + T^{1/4}\left(\sum_{t=1}^{T-1}(t+t_0+1)^{5/2}(\Delta_t^{\mathrm{PS}})^2\right)^{1/2}\right),$$

as desired. □

F Proofs for lazy SGD

F.1 Proof of Lemma F.1

Here we prove Lemma F.1, which provides a one-step inequality for the iterations of lazy SGD within step t, based on the standard SGD analysis under a fixed distribution $\mathcal{D}_t$.

Lemma F.1. *Suppose that $\ell(z,\theta)$ is μ-strongly convex in θ and the stochastic gradient oracle has quadratic-bounded variance (7). Then, one iteration of $\mathcal{A}_{\text{SGD-lazy}}$ satisfies*

$$\mathbb{E}\|\varphi_{t,j+1} - \theta_{\mathcal{D}_t}\|^2 \le \left(1 - 2\eta_{t,j}\mu + \eta_{t,j}^2 C_V^2\right)\mathbb{E}\|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2 + \eta_{t,j}^2\sigma^2, \tag{29}$$

where $\theta_{\mathcal{D}_t} \coloneqq \arg\min_\theta \mathrm{R}_{\mathcal{D}_t}(\theta)$.

Proof. Similarly as in the previous proofs, we have

$$\begin{aligned}
&\mathbb{E}\|\varphi_{t,j+1} - \theta_{\mathcal{D}_t}\|^2 \\
&\leq \mathbb{E}\|\varphi_{t,j} - \theta_{\mathcal{D}_t} - \eta_{t,j}\nabla_\theta \ell(Z_{t,j}, \varphi_{t,j})\|^2 \\
&= \|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2 - 2\eta_{t,j}\mathbb{E}[\nabla_\theta \ell(Z_{t,j}, \varphi_{t,j})^\top(\varphi_{t,j} - \theta_{\mathcal{D}_t})] + \eta_{t,j}^2\mathbb{E}[\|\nabla_\theta \ell(Z_{t,j}, \varphi_{t,j})\|^2] \\
&\leq \|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2 - 2\eta_{t,j}\nabla \mathrm{R}_{\mathcal{D}_t}(\varphi_{t,j})^\top(\varphi_{t,j} - \theta_{\mathcal{D}_t}) + \eta_{t,j}^2\mathbb{E}[\|\nabla_\theta \ell(Z_{t,j}, \varphi_{t,j})\|^2]
\end{aligned}$$

where the expectation here is over $Z_{t,j} \sim \mathcal{D}_t$ conditioned on $\varphi_{t,j}$.

Since $\theta_{\mathcal{D}_t} = \arg\min_\theta \mathrm{R}_{\mathcal{D}_t}(\theta)$, the first-order optimality condition of $\theta_{\mathcal{D}_t}$ and μ-strong monotonicity of $\nabla \mathrm{R}_{\mathcal{D}_t}(\theta)$ yields

$$\begin{aligned}
\nabla \mathrm{R}_{\mathcal{D}_t}(\varphi_{t,j})^\top(\varphi_{t,j} - \theta_{\mathcal{D}_t}) &\geq \left(\nabla \mathrm{R}_{\mathcal{D}_t}(\varphi_{t,j}) - \nabla \mathrm{R}_{\mathcal{D}_t}(\theta_{\mathcal{D}_t})\right)^\top(\varphi_{t,j} - \theta_{\mathcal{D}_t}) \\
&\geq \mu\|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2,
\end{aligned}$$

and by Eq. (7), we have

$$\mathbb{E}[\|\nabla_\theta \ell(Z_{t,j}, \varphi_{t,j})\|^2] \leq \sigma^2 + C_V^2\|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2.$$

Therefore we have

$$\mathbb{E}\|\varphi_{t,j+1} - \theta_{\mathcal{D}_t}\|^2 \leq \left(1 - 2\eta_{t,j}\mu + \eta_{t,j}^2 C_V^2\right)\|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2 + \eta_{t,j}^2\sigma^2,$$

conditioned on $\varphi_{t,j}$, and thus (29) by taking expectations over the entire history. □

F.2 Proof of Lemma F.2

Here we prove Lemma F.2, which provides upper bounds of $\mathbb{E}\|\theta_T - \theta_T^{\mathrm{PS}}\|^2$ for a fixed T.

Lemma F.2. *Suppose that $\mathcal{D}\colon \Theta \to \mathcal{P}(\mathcal{Z})$ is ϵ-sensitive, $z \mapsto \ell(z, \theta)$ is β_z-smooth for all $\theta \in \Theta$, and $\mathrm{R}_D(\theta)$ is μ-strongly convex and β_θ-smooth for all $D \in \mathcal{C}$. Also, suppose that $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$ and we have quadratic-bounded variance (Ass. 7). Then $\mathcal{A}_{\text{SGD-lazy}}$ with step sizes*

$$\eta_{t,j} = \frac{4}{A(j + t_0)}, \qquad t_0 = \max\left\{1, \left\lceil\frac{8B}{A^2}\right\rceil, \left\lceil\frac{8\mu}{A}\right\rceil\right\}, \tag{30}$$

for constants $A = 2\mu$ and $B = C_V^2$, and $n(t) \geq n_0 t^r$ samples per deployment with large enough n_0 satisfies

$$\mathbb{E}\|\theta_{T+1} - \theta_{T+1}^{\mathrm{PS}}\|^2 \leq \tilde{\gamma}^T\|\theta_1 - \theta_1^{\mathrm{PS}}\|^2 + J\cdot\left(\tilde{\gamma}^{T(1-2^{-\frac{1}{r}})} + 2T^{-r}\right) + (1+\tau)\sum_{t=1}^{T}\tilde{\gamma}^{T-t}(\Delta_t^{\mathrm{PS}})^2,$$

where $\tau, J > 0$ and $0 < \tilde{\gamma} < 1$ are chosen constants independent of T.

Proof. First, note that the case $T = 0$ is trivial.

Suppose that $T > 0$. We start with Lemma F.1 (proved in Section F.1), which shows

$$\mathbb{E}\|\varphi_{t,j+1} - \theta_{\mathcal{D}_t}\|^2 \leq \left(1 - 2\eta_{t,j}\mu + \eta_{t,j}^2 C_V^2\right)\mathbb{E}\|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2 + \eta_{t,j}^2\sigma^2.$$

Let $F_{t,j} = \mathbb{E}\|\varphi_{t,j} - \theta_{\mathcal{D}_t}\|^2$. Then we can rewrite as

$$F_{t,j+1} \leq (1 - \eta_{t,j}A + \eta_{t,j}^2 B)F_{t,j} + \eta_{t,j}^2\sigma^2,$$

with $A = 2\mu$ and $B = C_V^2$. Similarly to the proof of Lemma E.2 in Section E.2, the choice of

$$\eta_{t,j} = \frac{4}{A(j + t_0)}, \qquad t_0 = \max\left\{1, \left\lceil\frac{8B}{A^2}\right\rceil, \left\lceil\frac{8\mu}{A}\right\rceil\right\}$$

yields

$$F_{t,j+1} \leq \left(1 - \frac{2}{j + t_0}\right)F_{t,j} + \frac{4\sigma^2}{\mu^2(j + t_0)^2}.$$

By Lemma B.2 with $m = 1$ and $\alpha = 2$, we have

$$F_{t,j+1} \leq S(j + t_0)^{-1},$$

for $t_0 = \max\{1, \lceil \frac{8B}{A^2} \rceil, \lceil \frac{8\mu}{A} \rceil\} = \max\{1, \lceil \frac{2C_V^2}{\mu^2} \rceil, \lceil \frac{8\mu}{A} \rceil\}$ and

$$\begin{aligned} S &= \max\left\{F_1(1 + t_0), \frac{H}{\alpha - m}\right\} \\ &= \max\left\{(1 + t_0)\mathbb{E}\|\theta_t - \theta_{\mathcal{D}_t}\|^2, \frac{4\sigma^2}{\mu^2}\right\} \\ &\leq (1 + t_0)\mathbb{E}\|\theta_t - \theta_{\mathcal{D}_t}\|^2 + \frac{4\sigma^2}{\mu^2}. \end{aligned}$$

In particular, for $j = n(t)$ we have

$$\mathbb{E}\|\theta_{t+1} - \theta_{\mathcal{D}_t}\|^2 \leq \frac{1}{n(t) + t_0} \cdot \left((1 + t_0)\mathbb{E}\|\theta_t - \theta_{\mathcal{D}_t}\|^2 + \frac{4\sigma^2}{\mu^2}\right). \tag{31}$$

Observe that

$$\|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\| \leq \frac{(1 - \alpha_t)\epsilon\beta_z}{\mu}\|\theta_t - \theta_t^{\mathrm{PS}}\|,$$

and also,

$$\|\theta_t - \theta_{\mathcal{D}_t}\| \leq \|\theta_t - \theta_t^{\mathrm{PS}}\| + \|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\| \leq \left(1 + \frac{(1 - \alpha_t)\epsilon\beta_z}{\mu}\right)\|\theta_t - \theta_t^{\mathrm{PS}}\|.$$

Therefore, from each inequality we can deduce that

$$\mathbb{E}\|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\|^2 \leq \left(\frac{(1 - \alpha_t)\epsilon\beta_z}{\mu}\right)^2 \mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 \tag{32}$$

and

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta_{\mathcal{D}_t}\|^2 &\leq \frac{1}{n(t) + t_0} \cdot \left((1 + t_0)\mathbb{E}\|\theta_t - \theta_{\mathcal{D}_t}\|^2 + \frac{4\sigma^2}{\mu^2}\right) \\ &\leq \frac{1}{n(t) + t_0} \cdot \left((1 + t_0)\left(1 + \frac{(1 - \alpha_t)\epsilon\beta_z}{\mu}\right)^2 \mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{4\sigma^2}{\mu^2}\right) \\ &\leq \frac{1}{n(t) + t_0} \cdot \left(4(1 + t_0)\mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{4\sigma^2}{\mu^2}\right). \end{aligned} \tag{33}$$

We can compute

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\|^2 &= \mathbb{E}\|\theta_{t+1} - \theta_{\mathcal{D}_t}\|^2 + 2\mathbb{E}\left[(\theta_{t+1} - \theta_{\mathcal{D}_t})^\top(\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}})\right] + \mathbb{E}\|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\|^2 \\ &\leq \mathbb{E}\|\theta_{t+1} - \theta_{\mathcal{D}_t}\|^2 + 2\mathbb{E}\left[\|\theta_{t+1} - \theta_{\mathcal{D}_t}\| \cdot \|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\|\right] + \mathbb{E}\|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\|^2 \\ &\leq \mathbb{E}\|\theta_{t+1} - \theta_{\mathcal{D}_t}\|^2 + 2\sqrt{\mathbb{E}\|\theta_{t+1} - \theta_{\mathcal{D}_t}\|^2 \cdot \mathbb{E}\|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\|^2} + \mathbb{E}\|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\|^2 \\ &\leq (1 + \zeta)\mathbb{E}\|\theta_{t+1} - \theta_{\mathcal{D}_t}\|^2 + \left(1 + \frac{1}{\zeta}\right)\mathbb{E}\|\theta_{\mathcal{D}_t} - \theta_t^{\mathrm{PS}}\|^2, \end{aligned}$$

for any $\zeta > 0$, which we choose later.

By Eqs. (31) to (33), we have

$$\begin{aligned} &\mathbb{E}\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\|^2 \\ &\leq \left(\frac{4(1 + t_0)(1 + \zeta)}{n(t) + t_0} + \left(1 + \frac{1}{\zeta}\right)\left(\frac{(1 - \alpha_t)\epsilon\beta_z}{\mu}\right)^2\right)\mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{1 + \zeta}{n(t) + t_0} \cdot \frac{4\sigma^2}{\mu^2} \\ &\leq \left(\frac{4(1 + t_0)(1 + \zeta)}{n_0} + \left(1 + \frac{1}{\zeta}\right)\left(\frac{(1 - \alpha_{\min})\epsilon\beta_z}{\mu}\right)^2\right)\mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{1 + \zeta}{n(t) + t_0} \cdot \frac{4\sigma^2}{\mu^2}, \end{aligned}$$

where the last inequality is valid by $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$. Substituting

$$c = \frac{4(1+t_0)(1+\zeta)}{n_0} + \left(1+\frac{1}{\zeta}\right)\left(\frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu}\right)^2,$$

we can write

$$\mathbb{E}\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\|^2 \le c\mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{1+\zeta}{n(t)+t_0}\cdot\frac{4\sigma^2}{\mu^2}.$$

Recall that $\Delta_t^{\mathrm{PS}} = \|\theta_t^{\mathrm{PS}} - \theta_{t+1}^{\mathrm{PS}}\|$. By the AM-GM inequality, we have

$$\begin{aligned}
\mathbb{E}\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\|^2 &\le \mathbb{E}\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\|^2 + 2\Delta_t^{\mathrm{PS}}\sqrt{\mathbb{E}\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\|^2} + (\Delta_t^{\mathrm{PS}})^2 \\
&\le \left(1+\frac{1}{\tau}\right)\mathbb{E}\|\theta_{t+1} - \theta_t^{\mathrm{PS}}\|^2 + (1+\tau)(\Delta_t^{\mathrm{PS}})^2 \\
&\le c\left(1+\frac{1}{\tau}\right)\mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{(1+\frac{1}{\tau})(1+\zeta)}{n(t)+t_0}\cdot\frac{4\sigma^2}{\mu^2} + (1+\tau)(\Delta_t^{\mathrm{PS}})^2
\end{aligned}$$

for any $\tau > 0$, which we choose later.

Let $\tilde{\gamma} = c(1+\frac{1}{\tau})$. Now we choose constants $\zeta, \tau > 0$ that satisfy

$$\left(1+\frac{1}{\tau}\right)\left(1+\frac{1}{\zeta}\right)\left(\frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu}\right)^2 < 1,$$

and n_0 as a large enough constant so that

$$\tilde{\gamma} = \left(1+\frac{1}{\tau}\right)\left(\frac{4(1+t_0)(1+\zeta)}{n_0} + \left(1+\frac{1}{\zeta}\right)\left(\frac{(1-\alpha_{\min})\epsilon\beta_z}{\mu}\right)^2\right) < 1,$$

which is possible as we assume $\epsilon < \frac{\mu}{(1-\alpha_{\min})\beta_z}$. Then, we have

$$\begin{aligned}
\mathbb{E}\|\theta_{t+1} - \theta_{t+1}^{\mathrm{PS}}\|^2 &\le \tilde{\gamma}\cdot\mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{(1+\frac{1}{\tau})(1+\zeta)}{n(t)+t_0}\cdot\frac{4\sigma^2}{\mu^2} + (1+\tau)(\Delta_t^{\mathrm{PS}})^2 \\
&\le \tilde{\gamma}\cdot\mathbb{E}\|\theta_t - \theta_t^{\mathrm{PS}}\|^2 + \frac{4\sigma^2(1+\frac{1}{\tau})(1+\zeta)}{n_0\mu^2}t^{-r} + (1+\tau)(\Delta_t^{\mathrm{PS}})^2,
\end{aligned}$$

since $n_0 t^r \le n(t) \le n(t) + t_0$.

Finally, by Lemma B.3, we have

$$\begin{aligned}
&\mathbb{E}\|\theta_{T+1} - \theta_{T+1}^{\mathrm{PS}}\|^2 \\
&\le \tilde{\gamma}^T\|\theta_1 - \theta_1^{\mathrm{PS}}\|^2 + \frac{4\sigma^2(1+\frac{1}{\tau})(1+\zeta)}{n_0\mu^2}\sum_{t=1}^{T}\tilde{\gamma}^{T-t}t^{-r} + (1+\tau)\sum_{t=1}^{T}\tilde{\gamma}^{T-t}(\Delta_t^{\mathrm{PS}})^2 \\
&\le \tilde{\gamma}^T\|\theta_1 - \theta_1^{\mathrm{PS}}\|^2 + \frac{4\sigma^2(1+\frac{1}{\tau})(1+\zeta)}{n_0\mu^2(1-\tilde{\gamma})}\left(\tilde{\gamma}^{T(1-2^{-\frac{1}{r}})} + 2T^{-r}\right) + (1+\tau)\sum_{t=1}^{T}\tilde{\gamma}^{T-t}(\Delta_t^{\mathrm{PS}})^2,
\end{aligned}$$

which proves the given statement for $J = \frac{4\sigma^2(1+\frac{1}{\tau})(1+\zeta)}{n_0\mu^2(1-\tilde{\gamma})}$. □

F.3 Proof of Proposition 4.4

Here we prove Proposition 4.4, restated below for the sake of readability.

Proposition 4.4. *Suppose that the same assumptions as in Proposition 4.2 hold. Then, the iterates of $\mathcal{A}_{\text{SGD-lazy}}$ with step sizes $\eta_{t,j} \propto \frac{1}{j+t_0}$, for some constant $t_0 > 0$, and $n(t) \ge n_0 t^r$ with a large enough n_0 satisfy*

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\text{SGD-lazy}}) \le \mathcal{O}\left(\sum_{t=1}^{T-1} t^{-r/2} + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}}\right). \tag{9}$$

Proof. Using the same step sizes $\eta_{t,j}$ as in Lemma F.2, we start from Lemma F.2 with $T \leftarrow t-1$,

$$\mathbb{E}\|e_t\|^2 \le \tilde{\gamma}^{t-1}\|e_1\|^2 + J\cdot\left(\tilde{\gamma}^{\beta(t-1)} + 2(t-1)^{-r}\right) + (1+\tau)\sum_{s=1}^{t-1}\tilde{\gamma}^{t-s-1}(\Delta_s^{\mathrm{PS}})^2,$$

where we define $e_t = \theta_t - \theta_t^{\mathrm{PS}}$ and $\beta = 1-2^{-1/r}$ for simplicity. This implies

$$\sqrt{\mathbb{E}\|e_t\|^2} \le \tilde{\gamma}^{\frac{t-1}{2}}\|e_1\| + \sqrt{J}\cdot\tilde{\gamma}^{\frac{\beta(t-1)}{2}} + \sqrt{2J}\cdot(t-1)^{-r/2} + \sqrt{1+\tau}\left(\sum_{s=1}^{t-1}\tilde{\gamma}^{t-s-1}(\Delta_s^{\mathrm{PS}})^2\right)^{1/2},$$

since $\sqrt{a+b+c+d} \le \sqrt{a}+\sqrt{b}+\sqrt{c}+\sqrt{d}$ for $a,b,c,d \ge 0$. We can immediately observe that

$$\sum_{t=1}^{T}\tilde{\gamma}^{\frac{t-1}{2}}\|e_1\| \le \frac{\|e_1\|}{1-\tilde{\gamma}^{1/2}},$$
$$\sum_{t=1}^{T}\sqrt{J}\cdot\tilde{\gamma}^{\frac{\beta(t-1)}{2}} \le \frac{\sqrt{J}}{1-\tilde{\gamma}^{\beta/2}},$$

and the last term can be upper-bounded using the inequality

$$\sum_{s=1}^{t-1}\tilde{\gamma}^{t-s-1}(\Delta_s^{\mathrm{PS}})^2 \le \left(\sum_{s=1}^{t-1}\tilde{\gamma}^{\frac{t-s-1}{2}}\Delta_s^{\mathrm{PS}}\right)^2,$$

which implies

$$\begin{aligned}\sum_{t=1}^{T}\left(\sum_{s=1}^{t-1}\tilde{\gamma}^{t-s-1}(\Delta_s^{\mathrm{PS}})^2\right)^{1/2} &\le \sum_{t=1}^{T}\sum_{s=1}^{t-1}\tilde{\gamma}^{\frac{t-s-1}{2}}\Delta_s^{\mathrm{PS}} \\ &= \sum_{s=1}^{T-1}\Delta_s^{\mathrm{PS}}\left(\sum_{t=s+1}^{T}\tilde{\gamma}^{\frac{t-s-1}{2}}\right) \\ &\le \frac{1}{1-\tilde{\gamma}^{1/2}}\cdot\sum_{s=1}^{T-1}\Delta_s^{\mathrm{PS}}.\end{aligned}$$

Therefore (recalling that $e_1 = \theta_1 - \theta_1^{\mathrm{PS}}$) we have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\text{SGD-lazy}}) &\le L\sum_{t=1}^{T}\mathbb{E}\|e_t\| \le L\sum_{t=1}^{T}\sqrt{\mathbb{E}\|e_t\|^2} \\ &\le L\left(\frac{\|\theta_1-\theta_1^{\mathrm{PS}}\|}{1-\tilde{\gamma}^{1/2}} + \frac{\sqrt{J}}{1-\tilde{\gamma}^{\beta/2}} + \sqrt{2J}\sum_{t=1}^{T-1}t^{-r/2} + \frac{\sqrt{1+\tau}}{1-\tilde{\gamma}^{1/2}}\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}}\right).\end{aligned}$$

Eliminating all $\mathcal{O}(1)$ terms, we have

$$\mathrm{Reg}_T^{\mathrm{PS}}(\mathcal{A}_{\text{SGD-lazy}}) \le \mathcal{O}\left(\sum_{t=1}^{T-1}t^{-r/2} + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PS}}\right),$$

as desired. □

G Proofs for zeroth-order gradient descent

As we state in Section 5, we make the following mild assumptions throughout Sections G and H: $|\ell(z,\theta)| \le F$ for all z,θ for some $F < \infty$, and Θ contains a ball of radius $r > 0$ centered at 0.

G.1 Proof of Theorem 5.2

Here we prove Theorem 5.2, restated below for the sake of readability.

Theorem 5.2. *Suppose that* $\mathrm{PR}_t(\theta)$ *is* L*-Lipschitz* (13)*. If* PR_t *is* μ'*-strongly convex, then* $\mathcal{A}_{\mathrm{ZGD2}}$ *with* $\eta_t = \frac{1}{\mu' t}$*,* $\delta = \frac{d^2 L}{6\mu' T}$*, and* $\rho = \frac{\delta}{r}$ *satisfies for* $T \geq \frac{d^2 L}{6\mu' r}$*,*

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) = \mathcal{O}\left(d^2 \log T + \sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right).$$

If PR_t *is convex, then* $\mathcal{A}_{\mathrm{ZGD2}}$ *with* $\eta_t \equiv \eta = \frac{D_\Theta}{dL\sqrt{T}}$*,* $\delta = \frac{dD_\Theta}{6\sqrt{T}}$*, and* $\rho = \frac{\delta}{r}$ *satisfies for* $T \geq \frac{d^2 D_\Theta^2}{36 r^2}$*,*

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) = \mathcal{O}\left(d\sqrt{T} \cdot \left(1 + \sum_{t=1}^{T-1} \Delta_t^{\mathrm{PO}}\right)\right).$$

Proof. Since we assume Θ is a bounded (convex) set, we denote the (finite) diameter of Θ by D_Θ, i.e., $D_\Theta \coloneqq \sup\{\|\theta - \theta'\| : \theta, \theta' \in \Theta\}$.

We first prove the strongly convex case, and then prove the convex case.

Strongly convex case. Recall that for two-point ZGD, we use

$$\tilde{g}_t = \big(\mathrm{PR}_t(\theta_t + \delta u_t) - \mathrm{PR}_t(\theta_t - \delta u_t)\big) u_t \cdot \frac{d}{2\delta}.$$

Then, we have $\mathbb{E}[\tilde{g}_t] = \nabla \widehat{\mathrm{PR}}_t(\theta_t)$ and

$$\mathbb{E}[\|\tilde{g}_t\|^2] \leq d^2 L^2, \tag{34}$$

since PR_t and $\widehat{\mathrm{PR}}_t$ are both L-Lipschitz.

Then, μ'-strong convexity of $\widehat{\mathrm{PR}}_t(\theta)$ implies

$$\begin{aligned}
\mathbb{E}[\|\theta_{t+1} - \theta_t^u\|^2] &= \mathbb{E}[\|\Pi_{(1-\rho)\Theta}(\theta_t - \eta_t \tilde{g}_t) - \theta_t^u\|^2] \\
&\leq \mathbb{E}[\|\theta_t - \theta_t^u - \eta_t \tilde{g}_t\|^2] \\
&= \mathbb{E}[\|\theta_t - \theta_t^u\|^2 - 2\eta_t \tilde{g}_t^\top (\theta_t - \theta_t^u) + \eta_t^2 \|\tilde{g}_t\|^2] \\
&= \|\theta_t - \theta_t^u\|^2 - 2\eta_t \nabla \widehat{\mathrm{PR}}_t(\theta_t)^\top (\theta_t - \theta_t^u) + \eta_t^2 \mathbb{E}[\|\tilde{g}_t\|^2],
\end{aligned} \tag{35}$$

for any comparator $\theta_t^u \in (1-\rho)\Theta$ which we choose later, since $\mathbb{E}[\tilde{g}_t] = \nabla \widehat{\mathrm{PR}}_t(\theta_t)$. (The expectations here are conditioned on θ_t.)

By μ'-strong convexity of $\widehat{\mathrm{PR}}_t$, (34), and (35), we have

$$\begin{aligned}
\widehat{\mathrm{PR}}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t^u) &\leq \nabla \widehat{\mathrm{PR}}_t(\theta_t)^\top (\theta_t - \theta_t^u) - \frac{\mu'}{2}\|\theta_t - \theta_t^u\|^2 \\
&\leq \left(\frac{1}{2\eta_t} - \frac{\mu'}{2}\right)\|\theta_t - \theta_t^u\|^2 - \frac{1}{2\eta_t}\mathbb{E}[\|\theta_{t+1} - \theta_t^u\|^2] + \eta_t \frac{d^2 L^2}{2},
\end{aligned} \tag{36}$$

for any choice of comparators $\theta_t^u, \theta_{t+1}^u \in (1-\rho)\Theta$. By triangle inequality, we have

$$\|\theta_{t+1} - \theta_t^u\| \geq \|\theta_{t+1} - \theta_{t+1}^u\| - \|\theta_t^u - \theta_{t+1}^u\|,$$

and since $\|\theta_{t+1} - \theta_{t+1}^u\| \leq D_\Theta$, we can use

$$\begin{aligned}
\|\theta_{t+1} - \theta_t^u\|^2 &\geq \|\theta_{t+1} - \theta_{t+1}^u\|^2 - 2(\theta_{t+1} - \theta_{t+1}^u)^\top (\theta_t^u - \theta_{t+1}^u) + \|\theta_t^u - \theta_{t+1}^u\|^2 \\
&\geq \|\theta_{t+1} - \theta_{t+1}^u\|^2 - 2D_\Theta \|\theta_t^u - \theta_{t+1}^u\|.
\end{aligned}$$

Plugging in $\eta_t = \frac{1}{\mu' t}$ and noting that $\frac{1}{2\eta_1} - \frac{\mu'}{2} = 0$, we have

$$\begin{aligned}
&\widehat{\mathrm{PR}}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t^u) \\
&\leq \frac{\mu'(t-1)}{2}\|\theta_t - \theta_t^u\|^2 - \frac{\mu' t}{2}\mathbb{E}[\|\theta_{t+1} - \theta_{t+1}^u\|^2] + \mu' D_\Theta t \|\theta_t^u - \theta_{t+1}^u\| + \eta_t \frac{d^2 L^2}{2},
\end{aligned}$$

and taking the sum and expectations over the entire history,

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T}\left(\widehat{\mathrm{PR}}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t^u)\right)\right] &\le \frac{d^2L^2}{2\mu'}\sum_{t=1}^{T}\frac{1}{t}+\mu' D_\Theta\sum_{t=1}^{T} t\|\theta_t^u-\theta_{t+1}^u\| \\ &\le \frac{d^2L^2}{2\mu'}(1+\log T)+\mu' D_\Theta\sum_{t=1}^{T} t\|\theta_t^u-\theta_{t+1}^u\|.\end{aligned} \tag{37}$$

Recall that $\phi_t^+=\theta_t+\delta u_t$ and $\widehat{\mathrm{PR}}_t(\theta)=\mathbb{E}_{v\sim\mathrm{Unif}(\mathbb{B}^d)}[\mathrm{PR}_t(\theta+\delta v)]$. Since PR_t and $\widehat{\mathrm{PR}}_t$ are both L-Lipschitz, we have

$$\mathrm{PR}_t(\phi_t^+)-\mathrm{PR}_t(\theta_t)\le\delta L,\quad \mathrm{PR}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t)\le\delta L,\quad -\mathrm{PR}_t(\theta_t^u)+\widehat{\mathrm{PR}}_t(\theta_t^u)\le\delta L.$$

The same arguments hold when we replace ϕ_t^+ with ϕ_t^-. Therefore, for any choice of $\phi_t\in\{\phi_t^+,\phi_t^-\}$, we have

$$\mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t)-\mathrm{PR}_t(\theta_t^u)\right)\right]\le\frac{d^2L^2}{2\mu'}(1+\log T)+3\delta LT+\mu' D_\Theta\sum_{t=1}^{T} t\|\theta_t^u-\theta_{t+1}^u\|.$$

Now choose $\theta_t^u=(1-\rho)\theta_t^{\mathrm{PO}}$ for all $t\in[T]$, where we can check $\theta_t^u\in(1-\rho)\Theta$, and $\theta_{T+1}^u=\theta_T^u$. Then, by convexity of PR_t and $0\in\Theta$, we have

$$\begin{aligned}\mathrm{PR}_t(\theta_t^u) &\le (1-\rho)\mathrm{PR}_t(\theta_t^{\mathrm{PO}})+\rho\mathrm{PR}_t(0) \\ &= \mathrm{PR}_t(\theta_t^{\mathrm{PO}})+\rho\big(\mathrm{PR}_t(0)-\mathrm{PR}_t(\theta_t^{\mathrm{PO}})\big) \\ &\le \mathrm{PR}_t(\theta_t^{\mathrm{PO}})+2\rho F.\end{aligned} \tag{38}$$

Therefore, taking the maximum over all possible choices of $\phi_t\in\{\phi_t^+,\phi_t^-\}$, we have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) &= \sum_{t=1}^{T}\left(\max_{\phi_t\in\{\phi_t^+,\phi_t^-\}}\mathbb{E}[\mathrm{PR}_t(\phi_t)]-\mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right) \\ &\le \frac{d^2L^2}{2\mu'}(1+\log T)+3\delta LT+2\rho FT+\mu' D_\Theta(1-\rho)\sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}},\end{aligned}$$

where the $t=T$ term is eliminated by choosing $\theta_{T+1}^u=\theta_T^u$.

Setting $\rho=\frac{\delta}{r}$ and $\delta=\frac{d^2L}{6\mu' T}$, we have[2]

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) &\le \frac{d^2L^2}{2\mu'}\left(2+\log T+\frac{2F}{3Lr}\right)+\mu' D_\Theta(1-\rho)\sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}} \\ &= \mathcal{O}\left(d^2\log T+\sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right),\end{aligned}$$

and $T\ge\frac{d^2L}{6\mu' r}$ ensures that $\rho=\frac{\delta}{r}\le 1$, which proves the given statement.

Convex case. Suppose that PR_t and thus $\widehat{\mathrm{PR}}_t$ are convex and Lipschitz. We use $\eta_t\equiv\eta$, which we choose later. Since this case corresponds to $\mu'=0$, from (36) and the triangle inequality argument, we have

$$\widehat{\mathrm{PR}}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t^u)\le\frac{1}{2\eta}\|\theta_t-\theta_t^u\|^2-\frac{1}{2\eta}\mathbb{E}[\|\theta_{t+1}-\theta_{t+1}^u\|^2]+\frac{D_\Theta}{\eta}\|\theta_t^u-\theta_{t+1}^u\|+\eta\frac{d^2L^2}{2},$$

[2]We can improve the leading constant a bit by choosing $\delta=\frac{d^2L^2}{2\mu'}(3L+2\frac{F}{r})^{-1}\cdot\frac{1+\log T}{T}$, but we choose the above δ for the sake of simplicity.

and therefore

$$\mathbb{E}\left[\sum_{t=1}^{T}\left(\widehat{\mathrm{PR}}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t^u)\right)\right] \le \frac{1}{2\eta}\|\theta_1-\theta_1^u\|^2+\frac{\eta d^2L^2T}{2}+\frac{D_\Theta}{\eta}\sum_{t=1}^{T}\|\theta_t^u-\theta_{t+1}^u\|$$
$$\le \frac{D_\Theta^2}{2\eta}+\frac{\eta d^2L^2T}{2}+\frac{D_\Theta}{\eta}\sum_{t=1}^{T}\|\theta_t^u-\theta_{t+1}^u\|.$$

Following the same steps as in the strongly convex case, we have

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) = \sum_{t=1}^{T}\left(\max_{\phi_t\in\{\phi_t^+,\phi_t^-\}}\mathbb{E}[\mathrm{PR}_t(\phi_t)]-\mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right)$$
$$\le \frac{D_\Theta^2}{2\eta}+\frac{\eta d^2L^2T}{2}+3\delta LT+2\rho FT+\frac{D_\Theta(1-\rho)}{\eta}\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}}.$$

If we choose $\rho=\frac{\delta}{r}$, for which we have

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) \le \frac{D_\Theta^2}{2\eta}+\frac{\eta d^2L^2T}{2}+\left(3L+2\frac{F}{r}\right)\delta T+\frac{D_\Theta}{\eta}\left(1-\frac{\delta}{r}\right)\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}},$$

and then choose[3]

$$\eta=\frac{D_\Theta}{dL\sqrt{T}},\qquad \delta=\frac{dD_\Theta}{6\sqrt{T}},$$

then we have

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) \le \frac{3dD_\Theta L\sqrt{T}}{2}+\frac{dD_\Theta F\sqrt{T}}{3r}+dL\sqrt{T}\left(1-\frac{dD_\Theta}{6r\sqrt{T}}\right)\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}}$$
$$= \mathcal{O}\left(d\sqrt{T}\cdot\left(1+\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}}\right)\right),$$

and $T\ge\frac{d^2D_\Theta^2}{36r^2}$ ensures that $\rho=\frac{\delta}{r}\le 1$, which proves the given statement. □

H Additional results for optimality regret

H.1 Single-deployment algorithm

Here we consider one-point zeroth-order gradient descent (ZGD) [6], which we denote by $\mathcal{A}_{\mathrm{ZGD}}$,

$$u_t\sim\mathrm{Unif}(\mathbb{S}^{d-1})$$
$$\phi_t=\theta_t+\delta u_t,\qquad Z_t\sim\mathcal{D}_t(\phi_t)$$
$$\theta_{t+1}=\Pi_{(1-\rho)\Theta}(\theta_t-\nu_t\ell(Z_t,\phi_t)u_t),$$

where $(1-\rho)\Theta\coloneqq\{(1-\rho)\theta:\theta\in\Theta\}$. Note that for one-point ZGD, we can use $\ell(Z_t,\phi_t)$ with $Z_t\sim\mathcal{D}_t(\phi_t)$ instead of computing the whole population risk $\mathrm{PR}(\phi_t)$ from stochastic feedback.

Let us define

$$g_t\coloneqq\left(\frac{d}{\delta}\cdot\ell(Z_t,\theta_t+\delta u_t)\right)u_t.$$

Similar to $\mathcal{A}_{\mathrm{ZGD2}}$, it is known that $\mathbb{E}[g_t]=\frac{d}{\delta}\mathbb{E}_{u_t}[\mathrm{PR}_t(\theta_t+\delta u_t)u_t]=\nabla\widehat{\mathrm{PR}}_t(\theta_t)$ for $\widehat{\mathrm{PR}}_t(\theta)\coloneqq\mathbb{E}_v[\mathrm{PR}_t(\theta+\delta v)]$ with $v\sim\mathrm{Unif}(\mathbb{B}^d)$, and then we can rewrite the iterates of $\mathcal{A}_{\mathrm{ZGD}}$ as

$$\theta_{t+1}=\Pi_{(1-\rho)\Theta}\left(\theta_t-\nu_t\ell(Z_t,\phi_t)u_t\right)=\Pi_{(1-\rho)\Theta}\left(\theta_t-\nu_t\frac{\delta}{d}\cdot g_t\right),$$

[3] We can improve the leading constant a bit by choosing $(3L+2\frac{F}{r})^{-1}\frac{dD_\Theta L}{2\sqrt{T}}$, but we choose the above δ for the sake of simplicity.

which is equivalent to projected gradient descent with learning rate $\eta_t \coloneqq \nu_t \frac{\delta}{d}$.

We will provide upper bounds of the optimality regret in terms of the sequence $\{\phi_t\}_{t\in[T]}$,

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) = \mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t) - \mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right)\right].$$

We focus on the regret in terms of ϕ_t because $\phi_t \in \Theta$, while $\theta_t \in (1-\rho)\Theta$.

Theorem H.1. *Suppose that PR_t is μ'-strongly convex and L-Lipschitz. Then, $\mathcal{A}_{\mathrm{ZGD}}$ with $\eta_t = \frac{1}{\mu' t}$, $\delta \propto d^{2/3}(\frac{1+\log T}{T})^{1/3}$, and $\rho = \frac{\delta}{r}$ satisfies the following optimality regret upper bound for $T = \tilde{\Omega}(d^2)$:*

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) = \mathcal{O}\left(d^{\frac{2}{3}}T^{\frac{2}{3}}(\log T)^{\frac{1}{3}} + \sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right).$$

If PR_t is convex, then $\mathcal{A}_{\mathrm{ZGD}}$ with $\eta_t \equiv \eta \propto d^{-1/2}T^{-3/4}$, $\delta \propto d^{1/2}T^{-1/4}$, and $\rho = \frac{\delta}{r}$ satisfies the following optimality regret upper bound for $T = \tilde{\Omega}(d^2)$:

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) = \mathcal{O}\left(d^{\frac{1}{2}}T^{\frac{3}{4}} \cdot \left(1 + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}}\right)\right).$$

Proof. We first prove the strongly convex case, and then prove the convex case.

Strongly convex case. For any comparator $\theta_t^u \in (1-\rho)\Theta$ (which we choose later), μ'-strong convexity of $\widehat{\mathrm{PR}}_t(\theta)$ implies

$$\begin{aligned}
\mathbb{E}[\|\theta_{t+1} - \theta_t^u\|^2] &= \mathbb{E}[\|\Pi_{(1-\rho)\Theta}(\theta_t - \eta_t g_t) - \theta_t^u\|^2] \\
&\le \mathbb{E}[\|\theta_t - \theta_t^u - \eta_t g_t\|^2] \\
&= \mathbb{E}[\|\theta_t - \theta_t^u\|^2 - 2\eta_t g_t^\top(\theta_t - \theta_t^u) + \eta_t^2\|g_t\|^2] \\
&= \|\theta_t - \theta_t^u\|^2 - 2\eta_t \nabla\widehat{\mathrm{PR}}_t(\theta_t)^\top(\theta_t - \theta_t^u) + \eta_t^2\mathbb{E}[\|g_t\|^2],
\end{aligned} \tag{39}$$

since $\mathbb{E}[g_t] = \nabla\widehat{\mathrm{PR}}_t(\theta_t)$. Also, since $g_t = \frac{d}{\delta}\cdot\ell(Z_t, \theta_t + \delta u_t)u_t$ with $|\ell(z,\theta)| \le F$ and $\|u_t\| = 1$,

$$\mathbb{E}[\|g_t\|^2] \le \frac{d^2F^2}{\delta^2}.$$

By μ'-strong convexity of $\widehat{\mathrm{PR}}_t$ and (39), we have

$$\begin{aligned}
\widehat{\mathrm{PR}}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t^u) &\le \nabla\widehat{\mathrm{PR}}_t(\theta_t)^\top(\theta_t - \theta_t^u) - \frac{\mu'}{2}\|\theta_t - \theta_t^u\|^2 \\
&\le \left(\frac{1}{2\eta_t} - \frac{\mu'}{2}\right)\|\theta_t - \theta_t^u\|^2 - \frac{1}{2\eta_t}\mathbb{E}[\|\theta_{t+1} - \theta_t^u\|^2] + \eta_t\frac{d^2F^2}{2\delta^2},
\end{aligned} \tag{40}$$

for any choice of comparators $\theta_t^u, \theta_{t+1}^u \in (1-\rho)\Theta$. By triangle inequality, we have

$$\|\theta_{t+1} - \theta_t^u\| \ge \|\theta_{t+1} - \theta_{t+1}^u\| - \|\theta_t^u - \theta_{t+1}^u\|,$$

and since $\|\theta_{t+1} - \theta_{t+1}^u\| \le D_\Theta$, we can use

$$\begin{aligned}
\|\theta_{t+1} - \theta_t^u\|^2 &\ge \|\theta_{t+1} - \theta_{t+1}^u\|^2 - 2(\theta_{t+1} - \theta_{t+1}^u)^\top(\theta_t^u - \theta_{t+1}^u) + \|\theta_t^u - \theta_{t+1}^u\|^2 \\
&\ge \|\theta_{t+1} - \theta_{t+1}^u\|^2 - 2D_\Theta\|\theta_t^u - \theta_{t+1}^u\|.
\end{aligned}$$

Plugging in $\eta_t = \frac{1}{\mu' t}$ and noting that $\frac{1}{2\eta_1} - \frac{\mu'}{2} = 0$, we have

$$\begin{aligned}
&\widehat{\mathrm{PR}}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t^u) \\
&\le \frac{\mu'(t-1)}{2}\|\theta_t - \theta_t^u\|^2 - \frac{\mu' t}{2}\mathbb{E}[\|\theta_{t+1} - \theta_{t+1}^u\|^2] + \mu' D_\Theta t\|\theta_t^u - \theta_{t+1}^u\| + \eta_t\frac{d^2F^2}{2\delta^2},
\end{aligned}$$

and taking the sum and expectations over the entire history,

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T}\left(\widehat{\mathrm{PR}}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t^u)\right)\right] &\le \frac{d^2F^2}{2\delta^2\mu'}\sum_{t=1}^{T}\frac{1}{t}+\mu' D_\Theta\sum_{t=1}^{T}t\|\theta_t^u-\theta_{t+1}^u\| \\ &\le \frac{d^2F^2}{2\delta^2\mu'}(1+\log T)+\mu' D_\Theta\sum_{t=1}^{T}t\|\theta_t^u-\theta_{t+1}^u\|.\end{aligned} \tag{41}$$

Recall that $\phi_t=\theta_t+\delta u_t$ and $\widehat{\mathrm{PR}}_t(\theta)=\mathbb{E}_{v\sim\mathrm{Unif}(\mathbb{B}^d)}[\mathrm{PR}_t(\theta+\delta v)]$. Since PR_t and $\widehat{\mathrm{PR}}_t$ are both L-Lipschitz, we have

$$\mathrm{PR}_t(\phi_t)-\mathrm{PR}_t(\theta_t)\le\delta L,\quad \mathrm{PR}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t)\le\delta L,\quad -\mathrm{PR}_t(\theta_t^u)+\widehat{\mathrm{PR}}_t(\theta_t^u)\le\delta L,$$

and therefore

$$\mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t)-\mathrm{PR}_t(\theta_t^u)\right)\right]\le\frac{d^2F^2}{2\delta^2\mu'}(1+\log T)+3\delta LT+\mu' D_\Theta\sum_{t=1}^{T}t\|\theta_t^u-\theta_{t+1}^u\|.$$

Now we choose $\theta_t^u=(1-\rho)\theta_t^{\mathrm{PO}}$ for all $t\in[T]$, where we can check $\theta_t^u\in(1-\rho)\Theta$, and $\theta_{T+1}^u=\theta_T^u$. Then, by convexity of PR_t and $0\in\Theta$, we have

$$\begin{aligned}\mathrm{PR}_t(\theta_t^u) &\le (1-\rho)\mathrm{PR}_t(\theta_t^{\mathrm{PO}})+\rho\mathrm{PR}_t(0) \\ &= \mathrm{PR}_t(\theta_t^{\mathrm{PO}})+\rho\big(\mathrm{PR}_t(0)-\mathrm{PR}_t(\theta_t^{\mathrm{PO}})\big) \\ &\le \mathrm{PR}_t(\theta_t^{\mathrm{PO}})+2\rho F.\end{aligned} \tag{42}$$

Therefore we have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) &= \mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t)-\mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right)\right] \\ &\le \frac{d^2F^2}{2\delta^2\mu'}(1+\log T)+3\delta LT+2\rho FT+\mu' D_\Theta(1-\rho)\sum_{t=1}^{T-1}t\Delta_t^{\mathrm{PO}}.\end{aligned}$$

If we choose $\rho=\frac{\delta}{r}$ and

$$\delta=\left(\frac{d^2F^2}{2\mu'}(1+\log T)\right)^{\frac{1}{3}}\cdot\left(\left(3L+\frac{2F}{r}\right)T\right)^{-\frac{1}{3}},$$

we have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) &\le 2\left(\frac{d^2F^2}{2\mu'}(1+\log T)\right)^{\frac{1}{3}}\cdot\left(\left(3L+\frac{2F}{r}\right)T\right)^{\frac{2}{3}}+\mu' D_\Theta\sum_{t=1}^{T-1}t\Delta_t^{\mathrm{PO}} \\ &= \mathcal{O}\left(d^{\frac{2}{3}}T^{\frac{2}{3}}(\log T)^{\frac{1}{3}}+\sum_{t=1}^{T-1}t\Delta_t^{\mathrm{PO}}\right),\end{aligned}$$

and $\frac{T}{1+\log T}\ge\frac{d^2F^2}{2\mu' r^3\left(3L+\frac{2F}{r}\right)}$ ensures that $\rho=\frac{\delta}{r}\le 1$, which proves the given statement.

Convex case. For the case when $\widehat{\mathrm{PR}}_t$ is convex, we use constant $\eta_t\equiv\eta$ which we choose later. Since this case corresponds to $\mu'=0$, from (40) and the triangle inequality argument, we have

$$\widehat{\mathrm{PR}}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t^u)\le\frac{1}{2\eta}\|\theta_t-\theta_t^u\|^2-\frac{1}{2\eta}\mathbb{E}[\|\theta_{t+1}-\theta_{t+1}^u\|^2]+\frac{D_\Theta}{\eta}\|\theta_t^u-\theta_{t+1}^u\|+\eta\frac{d^2F^2}{2\delta^2},$$

and therefore

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T}\left(\widehat{\mathrm{PR}}_t(\theta_t)-\widehat{\mathrm{PR}}_t(\theta_t^u)\right)\right] &\le \frac{1}{2\eta}\|\theta_1-\theta_1^u\|^2+\frac{\eta d^2F^2T}{2\delta^2}+\frac{D_\Theta}{\eta}\sum_{t=1}^{T}\|\theta_t^u-\theta_{t+1}^u\| \\ &\le \frac{D_\Theta^2}{2\eta}+\frac{\eta d^2F^2T}{2\delta^2}+\frac{D_\Theta}{\eta}\sum_{t=1}^{T}\|\theta_t^u-\theta_{t+1}^u\|.\end{aligned}$$

Following the same steps as in the strongly convex case, we have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) &= \mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t) - \mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right)\right] \\ &\le \frac{D_\Theta^2}{2\eta} + \frac{\eta d^2F^2T}{2\delta^2} + 3\delta LT + 2\rho FT + \frac{D_\Theta(1-\rho)}{\eta}\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}}.\end{aligned}$$

If we choose $\rho = \frac{\delta}{r}$, for which we have

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) \le \frac{D_\Theta^2}{2\eta} + \frac{\eta d^2F^2T}{2\delta^2} + \left(3L + 2\frac{F}{r}\right)\delta T + \frac{D_\Theta}{\eta}\left(1 - \frac{\delta}{r}\right)\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}},$$

and then choose

$$\eta = \sqrt{\frac{D_\Theta^3}{dF}\cdot\frac{1}{3L + \frac{2F}{r}}}\,T^{-3/4}, \qquad \delta = \sqrt{\frac{D_\Theta dF}{3L + \frac{2F}{r}}}\,T^{-1/4},$$

so that

$$\frac{D_\Theta^2}{2\eta} = \frac{\eta d^2F^2T}{2\delta^2} = \frac{1}{2}\sqrt{D_\Theta dF\left(3L + \frac{2F}{r}\right)}\,T^{3/4},$$

$$\left(3L + 2\frac{F}{r}\right)\delta T = \sqrt{D_\Theta dF\left(3L + \frac{2F}{r}\right)}\,T^{3/4},$$

then we have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) &\le 2\sqrt{D_\Theta dF\left(3L + \frac{2F}{r}\right)}\,T^{3/4} + \sqrt{\frac{dF}{D_\Theta}\left(3L + \frac{2F}{r}\right)}\,T^{3/4}\left(1 - \frac{\delta}{r}\right)\sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}} \\ &= \mathcal{O}\left(d^{\frac{1}{2}}T^{\frac{3}{4}}\cdot\left(1 + \sum_{t=1}^{T-1}\Delta_t^{\mathrm{PO}}\right)\right),\end{aligned}$$

and $T \ge \left(\frac{D_\Theta dF}{r^2\left(3L + \frac{2F}{r}\right)}\right)^2$ ensures that $\rho = \frac{\delta}{r} \le 1$, which proves the given statement. □

H.2 Optimality regret upper bounds with unknown Lipschitz constants

As we have discussed in Remark 5.3, if we are not explicitly given the Lipschitz constant L of PR_t, the assumptions we have on the upper bound F of $|\ell|$ and that $r\mathbb{B}^d \subseteq \Theta$ allow us to use the effective Lipschitz constant $L_{\mathrm{eff}} = \frac{2F}{\rho r}$ as a proxy, as suggested in [6]. For completeness, we include the proof by [6] as well.

Lemma H.2 (Observation 3, [6])**.** *Suppose that $r\mathbb{B}^d \subseteq \Theta$ and $\rho \in (0, 1]$. Then, any convex function f with domain Θ that has bounded function value $|f(\theta)| \le F$ satisfies the inequality $|f(\theta) - f(\theta')| \le \frac{2F}{\rho r}\|\theta - \theta'\|$ for any $\theta \in (1-\rho)\Theta$ and $\theta' \in \Theta$.*

Proof. Let $\theta' = \theta + \vartheta$ If $\vartheta = 0$, then the statement trivially holds. Therefore, we may assume $\vartheta \ne 0$. If $\|\vartheta\| > \rho r$, then we immediately have $|f(\theta) - f(\theta')| \le 2F \le \frac{2F}{\rho r}\|\vartheta\|$. Therefore we may assume $\|\vartheta\| \le \rho r$. Define

$$\phi^+ = \theta + \rho r\cdot\frac{\vartheta}{\|\vartheta\|}, \qquad \phi^- = \theta - \rho r\cdot\frac{\vartheta}{\|\vartheta\|}.$$

Then we have $\phi^+, \phi^- \in \Theta$ because $\theta \in (1-\rho)\Theta$ and $r\mathbb{B}^d \subseteq \Theta$.

Observe that

$$\theta' = \left(1 - \frac{\|\vartheta\|}{\rho r}\right)\theta + \frac{\|\vartheta\|}{\rho r}\phi^+,$$

i.e., θ' is a convex combination of θ and ϕ^+. By convexity of f,

$$\begin{aligned} f(\theta') &\le \left(1 - \frac{\|\vartheta\|}{\rho r}\right) f(\theta) + \frac{\|\vartheta\|}{\rho r} f(\phi^+) \\ &= f(\theta) + \frac{\|\vartheta\|}{\rho r}\left(f(\phi^+) - f(\theta)\right) \\ &\le f(\theta) + \frac{2F\|\vartheta\|}{\rho r}, \end{aligned}$$

and therefore

$$f(\theta') - f(\theta) \le \frac{2F}{\rho r}\|\theta - \theta'\|.$$

Similarly, observe that

$$\theta = \frac{\rho r}{\|\vartheta\| + \rho r}\theta' + \frac{\|\vartheta\|}{\|\vartheta\| + \rho r}\phi^-,$$

i.e., θ is a convex combination of θ' and ϕ^-. By convexity of f,

$$\begin{aligned} f(\theta) &\le \frac{\rho r}{\|\vartheta\| + \rho r} f(\theta') + \frac{\|\vartheta\|}{\|\vartheta\| + \rho r} f(\phi^-) \\ &= f(\theta') + \frac{\|\vartheta\|}{\|\vartheta\| + \rho r}\left(f(\phi^-) - f(\theta')\right) \\ &\le f(\theta') + \frac{2F\|\vartheta\|}{\rho r}, \end{aligned}$$

and therefore

$$f(\theta) - f(\theta') \le \frac{2F}{\rho r}\|\theta - \theta'\|,$$

which proves the given statement. □

By Lemma H.2, we have the following result for one-point (single-deployment) ZGD when the Lipschitz constant L is unknown or not explicitly given.

Theorem H.3. *Suppose that PR_t is μ'-strongly convex. Then, $\mathcal{A}_{\mathrm{ZGD}}$ with $\eta_t = \frac{1}{\mu' t}$, $\delta \propto d^{\frac{4}{5}}(\frac{1+\log T}{T})^{\frac{2}{5}}$, and $\rho = \sqrt{3\delta/r}$ satisfies the following optimality regret upper bound for $T = \tilde{\Omega}(d^2)$:*

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) = \mathcal{O}\left(d^{\frac{2}{5}} T^{\frac{4}{5}} (\log T)^{\frac{1}{5}} + \sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right).$$

Proof. Since we use the same η_t and assumptions, we have (41) as in the proof of Theorem H.1,

$$\mathbb{E}\left[\sum_{t=1}^{T}\left(\widehat{\mathrm{PR}}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t^u)\right)\right] \le \frac{d^2F^2}{2\delta^2\mu'}(1 + \log T) + \mu' D_\Theta \sum_{t=1}^{T-1} t\|\theta_t^u - \theta_{t+1}^u\|.$$

Since $\theta_t, \theta_t^u \in (1-\rho)\Theta$, we can make any choice of $\phi_t \in \{\phi_t^+, \phi_t^-\}$ and apply Lemma H.2 to the μ'-strongly convex functions PR_t and $\widehat{\mathrm{PR}}_t$,

$$\mathrm{PR}_t(\phi_t) - \mathrm{PR}_t(\theta_t) \le \delta\frac{2F}{\rho r}, \quad \mathrm{PR}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t) \le \delta\frac{2F}{\rho r}, \quad -\mathrm{PR}_t(\theta_t^u) + \widehat{\mathrm{PR}}_t(\theta_t^u) \le \delta\frac{2F}{\rho r}.$$

Therefore we have

$$\mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t) - \mathrm{PR}_t(\theta_t^u)\right)\right] \le \frac{d^2F^2}{2\delta^2\mu'}(1 + \log T) + \frac{6\delta FT}{\rho r} + \mu' D_\Theta \sum_{t=1}^{T} t\|\theta_t^u - \theta_{t+1}^u\|.$$

Then, we can choose $\theta_t^u = (1-\rho)\theta_t^{\mathrm{PO}}$ for $t \le T$ and $\theta_{T+1}^u = \theta_T^u$ (similar to the previous proofs) and apply (42) as in the proof of Theorem H.1 to obtain

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) &= \mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t) - \mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right)\right] \\ &\le \frac{d^2F^2}{2\delta^2\mu'}(1+\log T) + \frac{6\delta FT}{\rho r} + 2\rho FT + \mu' D_\Theta(1-\rho)\sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}.\end{aligned}$$

If we choose

$$\delta = \left(\frac{d^2F\sqrt{r}}{4\sqrt{3}\mu'}\cdot\frac{1+\log T}{T}\right)^{\frac{2}{5}}, \qquad \rho = \sqrt{\frac{3\delta}{r}}$$

so that we can balance the terms

$$\frac{d^2F^2}{2\delta^2\mu'}(1+\log T) = \frac{6\delta FT}{\rho r} = 2\rho FT = 2FT\sqrt{\frac{3}{r}}\cdot\left(\frac{d^2F\sqrt{r}}{4\sqrt{3}\mu'}\cdot\frac{1+\log T}{T}\right)^{\frac{1}{5}},$$

and we also choose T to satisfy

$$\frac{T}{1+\log T} \ge \frac{d^2F}{108\mu' r^2} \quad\Longleftrightarrow\quad \delta \le 3r \quad\Longleftrightarrow\quad \rho \ge \frac{\delta}{r},$$

and

$$\frac{T}{1+\log T} \ge \frac{9d^2F}{4\mu' r^2} \quad\Longleftrightarrow\quad \rho \le 1,$$

then we finally have

$$\begin{aligned}\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD}}) &\le 6FT\sqrt{\frac{3}{r}}\cdot\left(\frac{d^2F\sqrt{r}}{4\sqrt{3}\mu'}\cdot\frac{1+\log T}{T}\right)^{\frac{1}{5}} + \mu' D_\Theta(1-\rho)\sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}} \\ &= \mathcal{O}\left(d^{\frac{2}{5}}T^{\frac{4}{5}}(\log T)^{\frac{1}{5}} + \sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right),\end{aligned}$$

which proves the given statement. □

We also have an analogous result for two-point (two-deployment) ZGD.

Theorem H.4. *Suppose that* PR_t *is* μ'*-strongly convex. Then,* $\mathcal{A}_{\mathrm{ZGD2}}$ *with* $\eta_t = \frac{1}{\mu' t}$, $\delta \propto d^{\frac{4}{3}}(\frac{1+\log T}{T})^{\frac{2}{3}}$, *and* $\rho = \sqrt{\frac{3\delta}{r}}$ *satisfies the following optimality regret upper bound for* $T = \tilde{\Omega}(d^2)$*:*

$$\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) = \mathcal{O}\left(d^{\frac{2}{3}}T^{\frac{2}{3}}(\log T)^{\frac{1}{3}} + \sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right).$$

Proof. For two-point ZGD we use

$$\tilde{g}_t = \left(\mathrm{PR}_t(\theta_t + \delta u_t) - \mathrm{PR}_t(\theta_t - \delta u_t)\right)u_t \cdot \frac{d}{2\delta}.$$

We can observe that

$$\begin{aligned}\|\tilde{g}_t\| &= \left\|\left(\mathrm{PR}_t(\theta_t + \delta u_t) - \mathrm{PR}_t(\theta_t - \delta u_t)\right)u_t \cdot \frac{d}{2\delta}\right\| \\ &\le \left\|\left(\mathrm{PR}_t(\theta_t + \delta u_t) - \mathrm{PR}_t(\theta_t)\right)u_t \cdot \frac{d}{2\delta}\right\| + \left\|\left(\mathrm{PR}_t(\theta_t + \delta(-u_t)) - \mathrm{PR}_t(\theta_t)\right)(-u_t) \cdot \frac{d}{2\delta}\right\| \\ &= \left(\,|\mathrm{PR}_t(\theta_t + \delta u_t) - \mathrm{PR}_t(\theta_t)| + |\mathrm{PR}_t(\theta_t - \delta u_t) - \mathrm{PR}_t(\theta_t)|\,\right)\cdot\frac{d}{2\delta} \\ &\le 2\delta\cdot\left(\frac{2F}{\rho r}\right)\cdot\frac{d}{2\delta} = \frac{2dF}{\rho r},\end{aligned}$$

where the last inequality follows from Lemma H.2 with $\theta_t \in (1-\rho)\Theta$, and therefore we have

$$\mathbb{E}[\|\tilde{g}_t\|^2] \leq \frac{4d^2F^2}{\rho^2 r^2}. \tag{43}$$

Therefore we can follow the same steps to obtain a modified version of (37) in the proof of Theorem 5.2 with the only difference being the constant from the upper bound of $\mathbb{E}[\|\tilde{g}_t\|^2]$,

$$\mathbb{E}\left[\sum_{t=1}^{T}\left(\widehat{\mathrm{PR}}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t^u)\right)\right] \leq \frac{2d^2F^2}{\mu'\rho^2r^2}(1+\log T) + \mu' D_\Theta \sum_{t=1}^{T} t\|\theta_t^u - \theta_{t+1}^u\|.$$

Since $\theta_t, \theta_t^u \in (1-\rho)\Theta$, we can make any choice of $\phi_t \in \{\phi_t^+, \phi_t^-\}$ and apply Lemma H.2 to the μ'-strongly convex functions PR_t and $\widehat{\mathrm{PR}}_t$ to obtain

$$\mathrm{PR}_t(\phi_t) - \mathrm{PR}_t(\theta_t) \leq \delta\frac{2F}{\rho r}, \quad \mathrm{PR}_t(\theta_t) - \widehat{\mathrm{PR}}_t(\theta_t) \leq \delta\frac{2F}{\rho r}, \quad -\mathrm{PR}_t(\theta_t^u) + \widehat{\mathrm{PR}}_t(\theta_t^u) \leq \delta\frac{2F}{\rho r}.$$

Therefore we have

$$\mathbb{E}\left[\sum_{t=1}^{T}\left(\mathrm{PR}_t(\phi_t) - \mathrm{PR}_t(\theta_t^u)\right)\right] \leq \frac{2d^2F^2}{\mu'\rho^2r^2}(1+\log T) + \frac{6\delta FT}{\rho r} + \mu' D_\Theta \sum_{t=1}^{T} t\|\theta_t^u - \theta_{t+1}^u\|.$$

Then, we can choose $\theta_t^u = (1-\rho)\theta_t^{\mathrm{PO}}$ for $t \leq T$ and $\theta_{T+1}^u = \theta_T^u$ (similar to the previous proofs) and apply (42) as in the proof of Theorem H.1 to obtain

$$\begin{aligned}
\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) &= \sum_{t=1}^{T}\left(\max_{\phi_t \in \{\phi_t^+, \phi_t^-\}} \mathbb{E}[\mathrm{PR}_t(\phi_t)] - \mathrm{PR}_t(\theta_t^{\mathrm{PO}})\right) \\
&\leq \frac{2d^2F^2}{\mu'\rho^2r^2}(1+\log T) + \frac{6\delta FT}{\rho r} + 2\rho FT + \mu' D_\Theta(1-\rho)\sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}.
\end{aligned}$$

If we choose

$$\delta = \frac{1}{3}\left(\frac{d^4F^2}{\mu'^2 r}\right)^{\frac{1}{3}}\left(\frac{1+\log T}{T}\right)^{\frac{2}{3}}, \qquad \rho = \sqrt{\frac{3\delta}{r}},$$

so that we can balance the terms

$$\frac{2d^2F^2}{\mu'\rho^2r^2}(1+\log T) = \frac{6\delta FT}{\rho r} = 2\rho FT = 2\left(\frac{d^2F^4}{\mu' r^2}\right)^{1/3}(1+\log T)^{1/3}T^{2/3}$$

and we also choose T to satisfy

$$\frac{T}{1+\log T} \geq \frac{d^2F}{27\mu' r^2} \quad \Longleftrightarrow \quad \delta \leq 3r \quad \Longleftrightarrow \quad \rho \geq \frac{\delta}{r},$$

and

$$\frac{T}{1+\log T} \geq \frac{d^2F}{\mu' r^2} \quad \Longleftrightarrow \quad \rho \leq 1,$$

then we finally have

$$\begin{aligned}
\mathrm{Reg}_T^{\mathrm{PO}}(\mathcal{A}_{\mathrm{ZGD2}}) &\leq 6\left(\frac{d^2F^4}{\mu' r^2}\right)^{\frac{1}{3}}(1+\log T)^{\frac{1}{3}}T^{\frac{2}{3}} + \mu' D_\Theta(1-\rho)\sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}} \\
&= \mathcal{O}\left(d^{\frac{2}{3}}T^{\frac{2}{3}}(\log T)^{\frac{1}{3}} + \sum_{t=1}^{T-1} t\Delta_t^{\mathrm{PO}}\right),
\end{aligned}$$

which proves the given statement. □

I Experimental details and further results

I.1 Credit scoring experiment

Environment. Here we provide a more detailed description of the credit scoring experiment summarized in Section 6. We minimize the regularized logistic loss,

$$\ell(z,\theta) = \log\left(1 + e^{-\theta^\top z}\right) + \frac{\lambda}{2}\|\theta\|^2,$$

over the Euclidean ball $\Theta = \{\theta : \|\theta\|_2 \le 1\}$ with $\lambda = 1$. The environment is built from the processed credit data used in the actionable-recourse experiments, available as the `credit` data set from `https://github.com/ustunb/actionable-recourse`. We first standardize each 17-dimensional credit feature to zero mean and unit variance. Then, we construct the signed examples by setting $z = y_{\text{signed}}x$, where $y_{\text{signed}} \in \{-1, +1\}$, and we apply the (regularized) logistic loss to these signed examples. The Gaussian base distribution is also fitted to 1500 signed feature vectors randomly sampled without replacement, from which we compute the empirical mean m and covariance Σ. The first six feature coordinates, corresponding to marital-status and age indicators, are treated as non-modifiable. The remaining eleven coordinates, corresponding to education, balance, payment, spending, and overdue-history features, are treated as modifiable. This split determines which coordinates can move under the strategic response.

As we discuss in Section 6, the performative component can be designed as

$$\mathcal{D}(\theta) = \mathcal{N}(A\theta + m, \Sigma),$$

where Σ is the empirical covariance and $A = a\,\mathrm{diag}(p_{\text{mod}})$ with $a = 0.1$. The vector p_{mod} is zero on the non-modifiable coordinates and one on the modifiable coordinates, which implies that $\|A\|_{\text{op}} = a = 0.1$. The exogenous component is

$$P_t = \mathcal{N}(m_t, \Sigma_P), \qquad \Sigma_P = \sigma_P^2 I, \quad \sigma_P = 0.3.$$

For the exogenous sequence, we report both an i.i.d. randomly varying sequence for m_t in the unit Euclidean ball and the stationary sequence $m_t \equiv m_0$, where m_0 is drawn once from the same bounded construction. The experiments use $d = 17$, $T_{\text{stab}} = 2000$ for the algorithms $\mathcal{A}_{\text{RRM}}$, $\mathcal{A}_{\text{RGD}}$, and $\mathcal{A}_{\text{SGD-greedy}}$, and $T_{\text{lazy}} = 1000$ for $\mathcal{A}_{\text{SGD-lazy}}$. The α_t grid consists of $\alpha_t = 1$ and polynomial schedules $\alpha_t = \alpha_0 t^{-b}$ with $\alpha_0 = 1$ and $b \in \{0.25, 0.5, 1.0, 2.0\}$.

By computing the (block) Hessians of ℓ and doing some algebra, we can observe that the logistic loss is λ-strongly convex in θ, and if $\|z\| \le R_z$ and $\|\theta\| \le R_\theta$ for some finite R_z, R_θ, then ℓ is $\lambda_z + \frac{R_z^2}{4}$-smooth in z and $\lambda_z + \frac{R_z R_\theta}{4}$-smooth in θ. Since Gaussian distributions technically have unbounded support, we use slightly conservative constants as a high-probability surrogate for both the smoothness constants above and the Lipschitz constants that are only used in computing the theoretical upper bounds.

We can also compute

$$\mathcal{W}_1\big(\mathcal{D}(\theta), \mathcal{D}(\theta')\big) = \|A(\theta - \theta')\| \le 0.1\|\theta - \theta'\|.$$

Consequently, the stability condition is satisfied.

Computing PS references and stability metrics. For θ_t^{PS}, we solve the fixed-point problem

$$\theta_t^{\text{PS}} = \arg\min_{\theta' \in \Theta} \mathbb{E}_{Z \sim \mathcal{D}_t(\theta_t^{\text{PS}})}[\ell(Z, \theta')]$$

by a bilevel fixed-point iteration heuristic. At each outer step, the inner fixed-distribution risk minimization is approximated by projected gradient descent using common random numbers. The experiment uses 20 outer iterations, 60 inner gradient steps, inner step size 0.2, and 512 Monte Carlo samples for the gradient estimate. Risk evaluations for plotting use 16384 common random samples, and these are shared across algorithms and comparator evaluations for a consistent comparison.

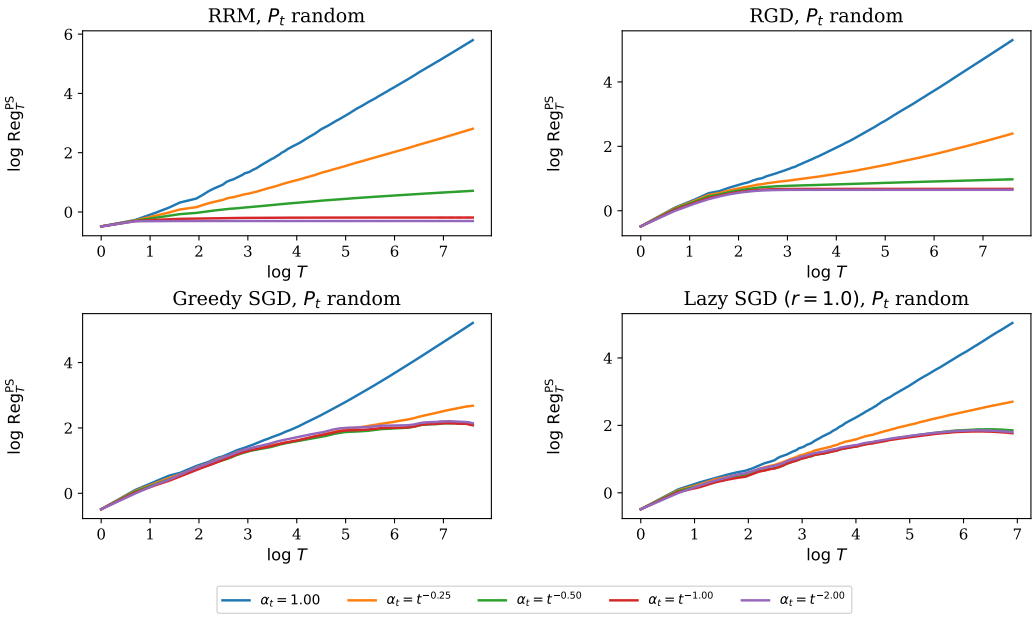

Figure 2: Performative stability regret of $\mathcal{A}_{\mathrm{RRM}}$, $\mathcal{A}_{\mathrm{RGD}}$, $\mathcal{A}_{\text{SGD-greedy}}$, and $\mathcal{A}_{\text{SGD-lazy}}$ in the credit scoring experiment, for varying α_t. We consider a randomly varying exogenous component P_t.

Results. Figures 2 and 3 give the alpha-schedule ablations for different algorithms. Overall, Figs. 2 and 3 show that the results follow the qualitative distinction between moving and stationary exogenous components for most cases.

Figure 2 shows that for the randomly varying case, $\mathcal{A}_{\mathrm{RRM}}$ and $\mathcal{A}_{\mathrm{RGD}}$ have a clear ordering with respect to the alpha schedule, with $\alpha_t \propto t^{-1.0}$ producing smaller regret than $\alpha_t \propto t^{-0.5}$, which matches the trend we can expect from the path-length dependence in Theorems 3.2 and 4.1 and our discussions in Section 6. The stochastic algorithms have less separation among the decaying schedules, yet we can observe that the constant schedule remains visibly worse with a slope near 1 (i.e., linear regret) for $\mathcal{A}_{\text{SGD-greedy}}$ and $\mathcal{A}_{\text{SGD-lazy}}$ as well.

Figure 3 shows that for the stationary case, we have the opposite trend with that of the randomly varying case in general, with the constant schedule performing better than the decaying schedules for $\mathcal{A}_{\mathrm{RRM}}$ and the stochastic methods $\mathcal{A}_{\text{SGD-greedy}}$ and $\mathcal{A}_{\text{SGD-lazy}}$. One interesting observation in Fig. 3 is that only $\mathcal{A}_{\mathrm{RGD}}$ has almost similar regret values for all of the different alpha schedules, which was consistent for multiple runs on different random seeds and initializations. This appears to be specific to the credit scoring experiment, since the same behavior does not appear in the synthetic experiment we later present in Section I.2.

Figure 4 compares the empirical stability regret of $\mathcal{A}_{\mathrm{RRM}}$ and $\mathcal{A}_{\mathrm{RGD}}$ to the finite upper bounds we derive in Theorems 3.2 and 4.1. In the randomly varying case, Fig. 4 again shows that $\alpha_t \propto t^{-1.0}$ produces smaller regret than $\alpha_t \propto t^{-0.5}$ for both algorithms, matching the path-length dependence in Theorems 3.2 and 4.1. We can also observe that, even though Gaussian distributions are unbounded, the finite upper bounds are still qualitatively correct in their ordering and level, which suggests that the conservative high-probability envelope is a reasonable surrogate for the Lipschitz constants in this case. Moreover, for decaying α_t, the slopes of the log-log plots near large T of the empirical regret curves are slightly smaller than the rates predicted by the upper bounds, which leaves room for both an improvement in the path-length dependence of the upper bounds or the existence of a more carefully constructed adversarial lower bound, which we leave for future investigation.

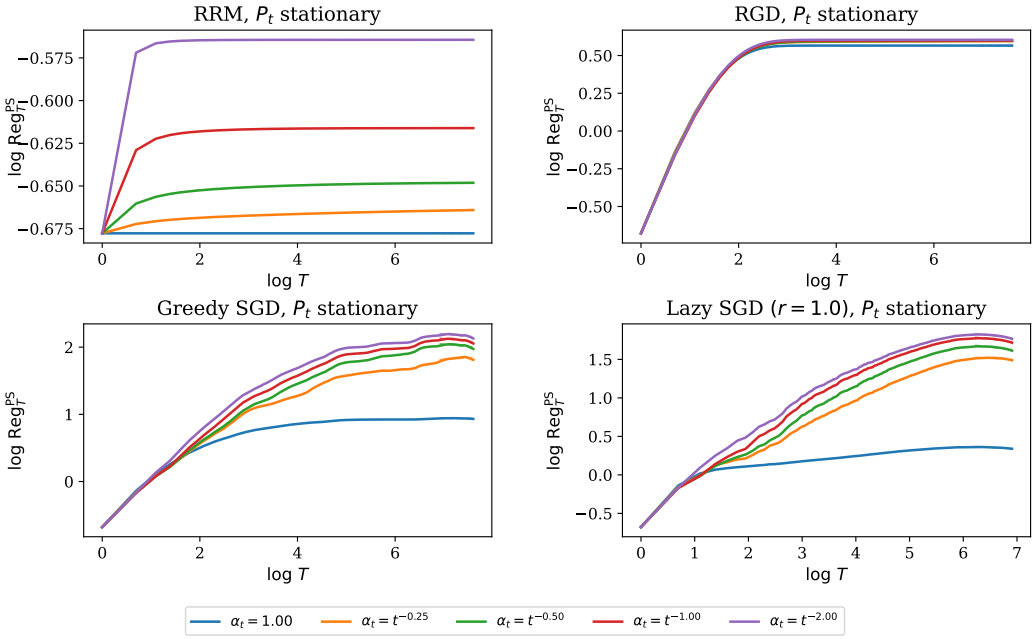

Figure 3: Performative stability regret of $\mathcal{A}_{\mathrm{RRM}}$, $\mathcal{A}_{\mathrm{RGD}}$, $\mathcal{A}_{\mathrm{SGD\text{-}greedy}}$, and $\mathcal{A}_{\mathrm{SGD\text{-}lazy}}$ in the credit scoring experiment, for varying α_t. We consider a stationary exogenous component P_t.

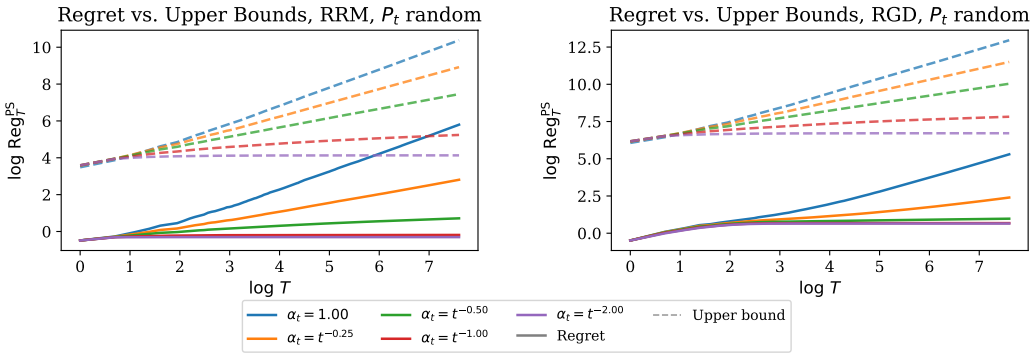

Figure 4: Performative stability regret vs. theoretical upper bounds of $\mathcal{A}_{\mathrm{RRM}}$ (left) and $\mathcal{A}_{\mathrm{RGD}}$ (right) in the credit scoring experiment, for varying α_t. We consider a randomly varying exogenous component P_t.

I.2 Further results on synthetic data

Environment. As a synthetic comparison, we also report a quadratic-Gaussian experiment. We minimize the squared loss $\ell(z, \theta) = \frac{1}{2}\|z - \theta\|^2$ over $\Theta = [-B, B]^d$ with $B = 5$. We use a partially performative setting similar to that of Section I.1.

$$Z_t \sim \mathcal{D}_t(\theta_t) = (1 - \alpha_t)\mathcal{D}(\theta_t) + \alpha_t P_t,$$
$$\mathcal{D}(\theta) = \mathcal{N}(A\theta + m, \Sigma), \qquad P_t = \mathcal{N}(m_t, \Sigma_P).$$

We use $\|A\|_{\mathrm{op}} = \epsilon = 0.3$, $\Sigma = \Sigma^2 I$, $\Sigma_P = \sigma_P^2 I$, and $\Sigma = \sigma_P = 0.5$. The vector m is sampled once from $[-1, 1]^d$, and the matrix A is sampled once and rescaled to have operator norm ϵ. The α_t grid consists of the constant schedule $\alpha_t = 1$ and polynomial decays $\alpha_t \propto t^{-b}$ with $b \in \{0.5, 0.75, 1.0, 1.5, 2.0\}$. For the exogenous sequence, we report both a shared i.i.d. randomly varying sequence in the unit Euclidean ball and the stationary sequence $m_t \equiv m_0$, where m_0 is drawn once from the same bounded construction. The experiments use $d \in \{2, 10\}$, $T_{\mathrm{stab}} = 100000$ for the algorithms $\mathcal{A}_{\mathrm{RRM}}, \mathcal{A}_{\mathrm{RGD}}$, and $\mathcal{A}_{\text{SGD-greedy}}$ and $T_{\mathrm{lazy}} = 1000$ for $\mathcal{A}_{\text{SGD-lazy}}$.

The quadratic loss is 1-strongly convex in θ and 1-smooth in both θ and z. Since $\mathcal{D}(\theta)$ and $\mathcal{D}(\theta')$ have the same covariance for any θ, θ', we have

$$\begin{aligned}\mathcal{W}_1\big(\mathcal{D}(\theta), \mathcal{D}(\theta')\big) &= \mathcal{W}_1\big(\mathcal{N}(A\theta + m, \Sigma), \mathcal{N}(A\theta + m', \Sigma)\big) \\ &= \|A(\theta - \theta')\| \le \epsilon\|\theta - \theta'\|.\end{aligned}$$

Thus $\mathcal{D}$ is ϵ-sensitive, and the partially performative map has sensitivity at most $(1 - \alpha_t)\epsilon$. Since $\epsilon = 0.3 < 1 = \mu/L_z$, the stability condition is satisfied.

As in the previous experiment, because Gaussian distributions have unbounded support, we use slightly conservative constants as a high-probability surrogate for the Lipschitz constants that are only used to compute the theoretical upper bounds.

Computing PS references. For the stable point, the minimizer of the squared loss under a fixed distribution is its mean. Hence θ_t^{PS} solves

$$\theta_t^{\mathrm{PS}} = (1 - \alpha_t)(A\theta_t^{\mathrm{PS}} + m) + \alpha_t m_t,$$

and therefore can be evaluated in closed form as the solution to the linear system:

$$\theta_t^{\mathrm{PS}} = \big(I - (1 - \alpha_t)A\big)^{-1}\big((1 - \alpha_t)m + \alpha_t m_t\big).$$

Results. Figures 5 and 6 gives alpha-schedule ablations for the stability algorithms $\mathcal{A}_{\mathrm{RRM}}$, $\mathcal{A}_{\mathrm{RGD}}$, $\mathcal{A}_{\text{SGD-greedy}}$, and $\mathcal{A}_{\text{SGD-lazy}}$. We can observe a similar trend as in the credit scoring experiment, with a clearer distinction between the randomly varying and stationary cases; Faster decay is better for the randomly varying case and worse for the stationary case, and the effects are more prominent in the deterministic methods. In the stationary case in Fig. 6, $\mathcal{A}_{\mathrm{RGD}}$ clearly follows the same alpha-schedule ordering as the other algorithms, in contrast to its behavior in the credit scoring experiment.

Figure 7 compares the empirical stability regret of $\mathcal{A}_{\mathrm{RRM}}$ and $\mathcal{A}_{\mathrm{RGD}}$ to the finite upper bounds under randomly varying shifts for $\alpha_t \propto t^{-0.5}$ and $\alpha_t \propto t^{-1.0}$. We can see similar trends as in Figure 4, and the difference between slopes of the empirical and upper bound curves are even more pronounced for the Gaussian case.

Figure 8 shows a more visible demonstration of the experiment in two dimensions. For the randomly varying sequence with $\alpha_t \propto t^{-1.0}$, we can visually observe that the algorithm iterates θ_t track the moving target θ_t^{PS} with a slight lag. The log-distance panels on the right show that this tracking error contracts over time but is also repeatedly perturbed by the movement of the reference.

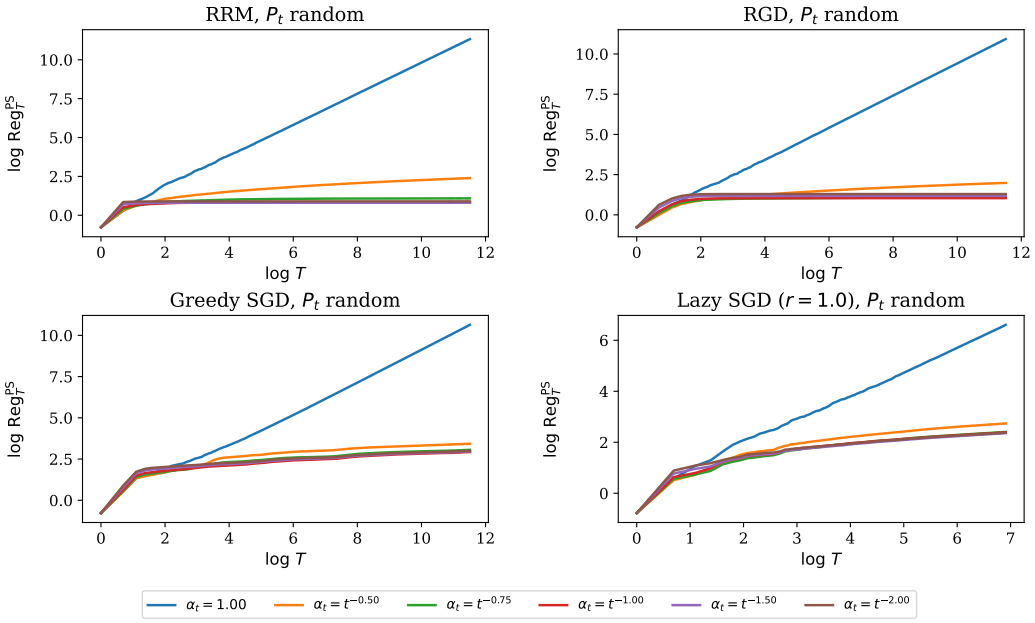

Figure 5: Performative stability regret of $\mathcal{A}_{\mathrm{RRM}}$, $\mathcal{A}_{\mathrm{RGD}}$, $\mathcal{A}_{\mathrm{SGD\text{-}greedy}}$, and $\mathcal{A}_{\mathrm{SGD\text{-}lazy}}$ in the quadratic-Gaussian experiment, for varying α_t. We consider a randomly varying exogenous component P_t.

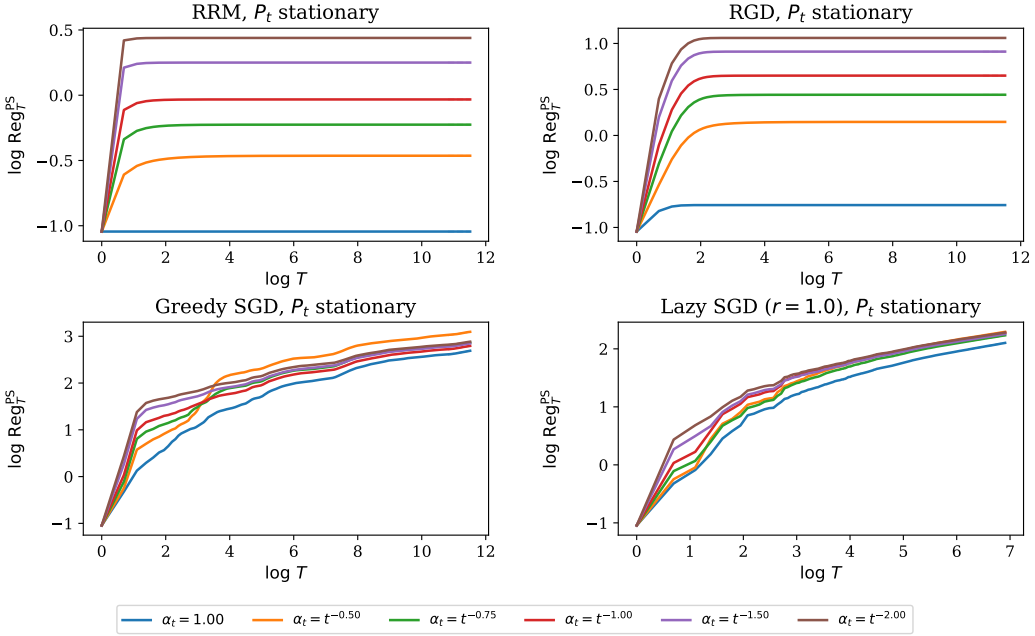

Figure 6: Performative stability regret of $\mathcal{A}_{\mathrm{RRM}}$, $\mathcal{A}_{\mathrm{RGD}}$, $\mathcal{A}_{\mathrm{SGD\text{-}greedy}}$, and $\mathcal{A}_{\mathrm{SGD\text{-}lazy}}$ in the quadratic-Gaussian experiment, for varying α_t. We consider a stationary exogenous component P_t.

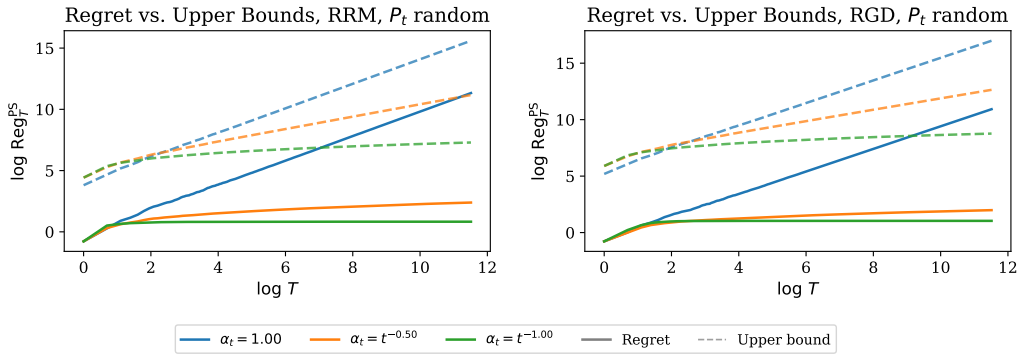

Figure 7: Performative stability regret vs. theoretical upper bounds of $\mathcal{A}_{\mathrm{RRM}}$ (left) and $\mathcal{A}_{\mathrm{RGD}}$ (right) in the quadratic-Gaussian experiment, for varying α_t. We consider a randomly varying exogenous component P_t.

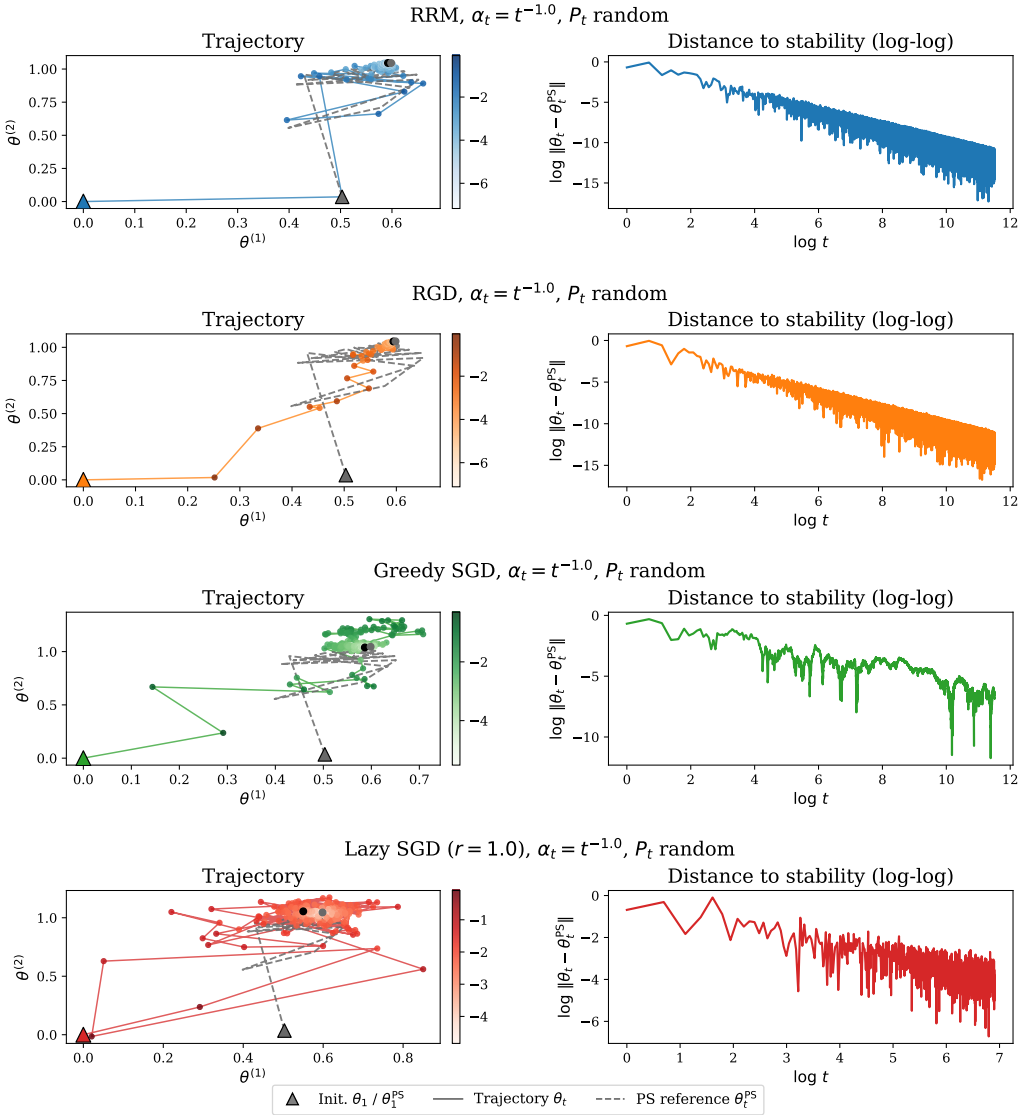

Figure 8: Representative two-dimensional iterate vs. stability trajectories of $\mathcal{A}_{\mathrm{RRM}}$, $\mathcal{A}_{\mathrm{RGD}}$, $\mathcal{A}_{\text{SGD-greedy}}$, and $\mathcal{A}_{\text{SGD-lazy}}$ in the 2D quadratic-Gaussian experiment, for $\alpha_t \propto t^{-1.0}$. We consider a randomly varying exogenous component P_t. The left subfigures show the actual movement of θ_t and θ_t^{PS} in the 2D space and the right subfigures show the log-log plots of the distance $\|\theta_t - \theta_t^{\mathrm{PS}}\|$.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the framework, regret notions, and algorithmic guarantees, which are developed in Sections 2–6.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper states the main regularity and convexity assumptions with the results, and Appendix I notes experimental caveats such as Gaussian-tail bounds and numerical reference computation.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Assumptions are stated in the theorem and lemma statements, with definitions in Section 2 and Appendix A, and proofs in the subsequent appendices.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Section 6 and Appendix I specify the experiment settings, algorithms, and relevant details used for the empirical results.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The supplemental code includes scripts to run and duplicate both experiments presented in the paper. The credit data are loaded from a public source.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 6 and Appendix I specify the experiment settings, algorithms, and relevant details used for the empirical results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: The experiments are simply illustrative demonstrations of the theoretical results and do not require such statistical arguments to understand the results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [No]

   Justification: Experiments are lightweight and numerical. They do not require any GPUs or large memory, and can be run in less than 30 minutes on a M4 Macbook Air.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The work is theoretical with synthetic and public-data experiments, and does not involve human subjects, private data, or high-risk model releases.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The introduction discusses consequential deployment domains and feedback loops, motivating both the benefits of adaptation and risks from endogenous and exogenous distribution shift.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper does not release high-risk models, scraped datasets, or other assets that require misuse safeguards.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The strategic-credit experiment identifies the processed credit dataset, which is open-source and obtainable from the link stated in the paper.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core methods and experiments do not use LLMs as an important, original, or non-standard component.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.